

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

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

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1767366> since 2023-12-04T13:32:05Z

Publisher:

Springer

Published version:

DOI:10.1007/978-3-030-50086-3_9

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

PROBABILISTIC TIMED AUTOMATA WITH ONE CLOCK AND INITIALISED CLOCK-DEPENDENT PROBABILITIES

JEREMY SPROSTON

Dipartimento di Informatica, University of Turin, Italy
e-mail address: sproston@di.unito.it

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. Model checking [13, 3] is an example of a formal verification technique, and comprises the automatic verification of a system model 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 probabilistic systems formally, 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, 15] 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* [17, 26, 29], 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, in automotive and aeronautic contexts, the probability of certain reactions of human operators may depend on factors

Key words and phrases: Timed automata, interval Markov chains, probabilistic model checking.




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

such as fatigue, which can increase over time (see, for example, [14]); 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 must have rational-numbered endpoints). This limitation led to the development of an extension of probabilistic timed automata called *clock-dependent probabilistic timed automata* [33], 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 restart the task from the beginning (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 (which is a finite-state abstraction used extensively for timed automata) can be employed to approximate optimal reachability probabilities [33]. 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, [27, 8, 24, 5, 6, 1]), we restrict our attention to clock-dependent probabilistic timed automata with a *single* clock variable. As in [33], 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 [22], 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, i.e., at least one natural-numbered value, 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 [31, 21, 18]), and also in other contexts, such as planning [16] and control [28, 37]. 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, 34], 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 [23, 25, 32] 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, [34] 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. This paper extends the conference version [35] with full proofs of the results.

2. INTERVAL MARKOV DECISION PROCESSES

2.1. 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, 15]), 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 $s_0 a_0 s_1 a_1 \dots$ such that $a_i \in A(s_i)$ and $\Delta(s_i, a_i)(s_{i+1}) > 0$ for all $i \geq 0$. Given an infinite path $\mathbf{r} = s_0 a_0 s_1 a_1 \dots$ and $i \geq 0$, we let $\mathbf{r}(i) = s_i$ be the $(i+1)$ -th state along \mathbf{r} . Let $\text{Paths}^{\mathcal{M}}$ be the set of infinite paths of \mathcal{M} . A finite path is a sequence $r = s_0 a_0 s_1 a_1 \dots a_{n-1} s_n$ such that $a_i \in A(s_i)$ and $\Delta(s_i, a_i)(s_{i+1}) > 0$ for all $0 \leq i < n$. Let $\text{last}(r) = s_n$ denote the final state of r . For $a \in A(s_n)$ and $s \in S$ such that $\Delta(s_n, a)(s) > 0$, we use ras to denote the finite path $s_0 a_0 s_1 a_1 \dots a_{n-1} s_n a s$. Let $\text{Paths}_*^{\mathcal{M}}$ be the set of finite paths of the MDP \mathcal{M} . Let $\text{Paths}^{\mathcal{M}}(s)$ and $\text{Paths}_*^{\mathcal{M}}(s)$ be the sets of infinite paths and finite paths, respectively, of \mathcal{M} starting in state $s \in S$.

A *scheduler* is a function $\sigma : \text{Paths}_*^{\mathcal{M}} \rightarrow \bigcup_{s \in S} \text{Dist}(A(s))$ such that $\sigma(r) \in \text{Dist}(A(\text{last}(r)))$ for all $r \in \text{Paths}_*^{\mathcal{M}}$.¹ Let $\Sigma^{\mathcal{M}}$ be the set of schedulers of the MDP \mathcal{M} . We say

¹From [19, 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 $r \in \text{Paths}_*^{\mathcal{M}}$.

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 σ : we let $\mathcal{C}^\sigma = (Paths_*^\sigma, \mathbf{P})$, where, for $r, r' \in Paths_*^\sigma$, we have $\mathbf{P}(r, r') = \sigma(r)(a) \cdot \Delta(\text{last}(r), a)(s)$ if $r' = ras$ and $a \in A(\text{last}(r))$, 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 $i \leq n$. Given $s \in S$ and $r \in Paths_*^\sigma$, 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, 15]).

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 paths of \mathcal{M} such that some state of T is visited along the path. 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 $\text{Pr}_s^\sigma(\Diamond T) = 0$; $(\exists 1)$ decide whether there exists $\sigma \in \Sigma^\mathcal{M}$ such that $\text{Pr}_s^\sigma(\Diamond T) = 1$; $(\forall 1)$ decide whether $\text{Pr}_s^\sigma(\Diamond T) = 1$ for all $\sigma \in \Sigma^\mathcal{M}$.

2.2. 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 .

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




Figure 2: Example fragments of IMDP \mathfrak{M} (left) and the constructed IMC $\mathfrak{C}[\mathfrak{M}]$ (right).

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, [32, 11, 12]), 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 [32, 12]): 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 [34]. 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.

2.3. 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 $\mathfrak{A}(s) = \{a \in \mathfrak{A} \mid \mathfrak{D}(s, a) \neq \perp\}$, 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 $a \in \mathfrak{A}(s)$ is chosen, then (2) an assignment α for $\mathfrak{D}(s, a)$ 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 $A(s) = \{(a, \alpha) \in \mathfrak{A} \times \mathfrak{Dist}(S) \mid a \in \mathfrak{A}(s) \text{ and } \alpha \in \mathfrak{G}(\mathfrak{D}(s, a))\}$ for each state $s \in S$, and $\Delta(s, (a, \alpha)) = \alpha$ for each state $s \in S$ and action/assignment pair $(a, \alpha) \in A(s)$. Note that (as in, for example, [31, 21, 18]) 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 in polynomial time such that the maximal and minimal reachability probabilities for the IMDP and the constructed IMC coincide, and furthermore qualitative properties agree on the IMDP and the constructed 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 $S \otimes \mathfrak{A} = \bigcup_{s \in S} \{(s, a) \in S \times \mathfrak{A} \mid a \in \mathfrak{A}(s)\}$;
- for $s \in S$ and $a \in \mathfrak{A}(s)$, let $\tilde{\mathfrak{P}}(s, (s, a)) = [0, 1]$, and let $\tilde{\mathfrak{P}}((s, a), \cdot) = \mathfrak{D}(s, a)$.

Example 2.1. In Figure 2 we illustrate a fragment of an IMDP (left), and give the corresponding fragment of the constructed IMC $\mathfrak{C}[\mathfrak{M}]$ (right). Note that the nondeterministic choice between actions a_1 and a_2 from state s of \mathfrak{M} is replaced in $\mathfrak{C}[\mathfrak{M}]$ by a choice between states (s, a_1) and (s, a_2) ; more precisely, the choice of a scheduler of \mathfrak{M} from a path ending in state s as to probabilities to assign to actions a_1 and a_2 is reflected by the choice of assignment made in $\mathfrak{C}[\mathfrak{M}]$ over states (s, a_1) and (s, a_2) (note that this latter choice is unconstrained by the interval $[0, 1]$ used for these transitions).

The following proposition states the correctness of the construction of $\mathfrak{C}[\mathfrak{M}]$ with respect to quantitative and qualitative problems.

Proposition 2.2. *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:*

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

In order to show Proposition 2.2, we consider two lemmata that show that a scheduler of \mathfrak{M} can be matched by a “mimicking” scheduler of $\mathfrak{C}[\mathfrak{M}]$ (Lemma 2.3), and vice versa (Lemma 2.4), such that the probabilities of reaching a set of target states for the two schedulers coincide. Together, these lemmata suffice to establish Proposition 2.2.

Lemma 2.3. *Let $s \in S$ and $T \subseteq S$. Then, for each $\sigma \in \Sigma^{[\mathfrak{M}]}$, there exists $\pi \in \Sigma^{[\mathfrak{C}[\mathfrak{M}]]}$ such that $\Pr_s^\sigma(\Diamond T) = \Pr_s^\pi(\Diamond T)$.*

Proof. Let $\mathfrak{M} = (S, \mathfrak{A}, \mathfrak{D})$. To avoid ambiguity, we henceforth refer to elements of \mathfrak{A} as \mathfrak{M} -actions. Recall that $[\mathfrak{M}] = (S, A, \Delta_{[\mathfrak{M}]})$ where $A(s) = \{(\mathfrak{a}, \alpha) \in \mathfrak{A} \times \text{Dist}(S) \mid \mathfrak{a} \in \mathfrak{A}(s) \text{ and } \alpha \in \mathfrak{G}(\mathfrak{D}(s, \mathfrak{a}))\}$ and $\Delta_{[\mathfrak{M}]}(s, (\mathfrak{a}, \alpha)) = \alpha$ for each $s \in S$ and $(\mathfrak{a}, \alpha) \in A(s)$. Note that an action $(\mathfrak{a}, \alpha) \in A$ of $[\mathfrak{M}]$ comprises an \mathfrak{A} -action \mathfrak{a} and an assignment α from $\text{Dist}(S)$. We can observe that finite paths of $[\mathfrak{M}]$ (i.e., elements of $\text{Paths}_*^{[\mathfrak{M}]}$) have the form $s_0(\mathfrak{a}_0, \alpha_0)s_1(\mathfrak{a}_1, \alpha_1) \cdots (\mathfrak{a}_{n-1}, \alpha_{n-1})s_n$, where $s_i \in S$ for all $0 \leq i \leq n$, and $(\mathfrak{a}_j, \alpha_j) \in \mathfrak{A} \times \text{Dist}(S)$ for all $0 \leq j \leq n$.

Recall that $\mathfrak{C}[\mathfrak{M}] = (S \cup (S \otimes \mathfrak{A}), \tilde{\mathfrak{P}})$, where $\tilde{\mathfrak{P}}(s, (s, \mathfrak{a})) = [0, 1]$ and $\tilde{\mathfrak{P}}((s, \mathfrak{a}), \cdot) = \mathfrak{D}(s, \mathfrak{a})$ for $s \in S$ and $\mathfrak{a} \in \mathfrak{A}(s)$, i.e., $\mathfrak{C}[\mathfrak{M}]$ alternates between states from S and $S \otimes \mathfrak{A}$, where the probability of transitions from S to $S \otimes \mathfrak{A}$ are unconstrained, and the probability of transitions from $S \otimes \mathfrak{A}$ to S are constrained by intervals defined by \mathfrak{D} . Finally, we also recall that $[\mathfrak{C}[\mathfrak{M}]] = (S \cup (S \otimes \mathfrak{A}), \mathfrak{G}(\tilde{\mathfrak{P}}), \Delta_{[\mathfrak{C}[\mathfrak{M}]])$, where $\mathfrak{G}(\tilde{\mathfrak{P}}) = \bigcup_{\tilde{s} \in S \cup (S \otimes \mathfrak{A})} \mathfrak{G}(\tilde{\mathfrak{P}}(\tilde{s}, \cdot))$ and $\Delta_{[\mathfrak{C}[\mathfrak{M}]]}(\tilde{s}, \alpha) = \alpha$ for all $\tilde{s} \in S \cup (S \otimes \mathfrak{A})$ and $\alpha \in \mathfrak{G}(\tilde{\mathfrak{P}}(\tilde{s}, \cdot))$. Note that the actions of $[\mathfrak{C}[\mathfrak{M}]]$, i.e., elements of the set $\mathfrak{G}(\tilde{\mathfrak{P}})$, are themselves the assignments that are used to determine the next state: given the aforementioned alternation between S and $S \otimes \mathfrak{A}$ in $\mathfrak{C}[\mathfrak{M}]$, from states in S , only assignments in $\text{Dist}(S \otimes \mathfrak{A})$ are available; similarly, from states in $S \otimes \mathfrak{A}$, only assignments in $\text{Dist}(S)$ are available. We partition $\text{Paths}_*^{[\mathfrak{C}[\mathfrak{M}]]}$ into two sets, on the basis of whether the final state of a path is in S or $S \otimes \mathfrak{A}$: let $\text{Paths}_{*,S}^{[\mathfrak{C}[\mathfrak{M}]]} \subseteq \text{Paths}_*^{[\mathfrak{C}[\mathfrak{M}]]}$ be the set of finite paths of the form $s_0\beta_0(s_0, \mathfrak{a}_0)\gamma_0s_1\beta_1(s_1, \mathfrak{a}_1)\gamma_1 \cdots s_{n-1}\beta_{n-1}(s_{n-1}, \mathfrak{a}_{n-1})\gamma_{n-1}s_n$, and let $\text{Paths}_{*,S \otimes \mathfrak{A}}^{[\mathfrak{C}[\mathfrak{M}]]} \subseteq \text{Paths}_*^{[\mathfrak{C}[\mathfrak{M}]]}$ be the set of finite paths of the form $s_0\beta_0(s_0, \mathfrak{a}_0)\gamma_0s_1\beta_1(s_1, \mathfrak{a}_1)\gamma_1 \cdots s_m\beta_m(s_m, \mathfrak{a}_m)$.

Let $s \in S$, $T \subseteq S$ and $\sigma \in \Sigma^{\llbracket \mathfrak{M} \rrbracket}$. The proof of the lemma consists of constructing $\pi \in \Sigma^{\llbracket \mathfrak{C}[\mathfrak{M}] \rrbracket}$ from σ , and showing that $\Pr_s^\sigma(\Diamond T) = \Pr_s^\pi(\Diamond T)$. Intuitively, the construction of scheduler π proceeds in the following manner: each choice made by σ (which we recall, from the definition of $\llbracket \mathfrak{M} \rrbracket$, is a distribution over \mathfrak{M} -action/assignment pairs) is mimicked by a sequence of *two* choices made by π , the first of which mimics the choice of σ over \mathfrak{M} -actions, and where the second mimics the choice of σ over assignments.

To describe formally the construction of π , we first introduce the function $f : \text{Paths}_{*,S}^{\llbracket \mathfrak{C}[\mathfrak{M}] \rrbracket}(s) \rightarrow \text{Paths}_*^{\llbracket \mathfrak{M} \rrbracket}(s)$ that associates, for each finite path in $\text{Paths}_{*,S}^{\llbracket \mathfrak{C}[\mathfrak{M}] \rrbracket}(s)$, a finite path in $\text{Paths}_*^{\llbracket \mathfrak{M} \rrbracket}(s)$ that (1) visits the same states of S , and (2) features the same actions from \mathfrak{A} and assignments from $\text{Dist}(S)$. Formally, for $r \in \text{Paths}_{*,S}^{\llbracket \mathfrak{C}[\mathfrak{M}] \rrbracket}(s)$ such that $r = s_0\beta_0(s_0, \mathbf{a}_0)\gamma_0s_1\beta_1(s_1, \mathbf{a}_1)\gamma_1 \cdots s_{n-1}\beta_{n-1}(s_{n-1}, \mathbf{a}_{n-1})\gamma_{n-1}s_n$, we let $f(r) = s_0(\mathbf{a}_0, \gamma_0)s_1(\mathbf{a}_1, \gamma_1) \cdots (\mathbf{a}_{n-1}, \gamma_{n-1})s_n$. That is, $f(r)$ retains fully the subsequence $s_0s_1 \cdots s_n$ of states from S , the subsequence $\mathbf{a}_0\mathbf{a}_1 \cdots \mathbf{a}_{n-1}$ of actions of \mathfrak{M} from \mathfrak{A} , and the subsequence $\gamma_0\gamma_1 \cdots \gamma_{n-1}$ of assignments from $\text{Dist}(S)$, but not $\beta_0\beta_1 \cdots \beta_{n-1}$ of assignments from $\text{Dist}(S \otimes \mathfrak{A})$. We note that $f(r)$ is an element of $\text{Paths}_*^{\llbracket \mathfrak{M} \rrbracket}(s)$ by the following reasoning: for each $0 \leq i < n$, we require that (a) $\mathbf{a}_i \in \mathfrak{A}(s_i)$ and (b) $\gamma_i \in \mathfrak{G}(\mathfrak{D}(s_i, \mathbf{a}_i))$; both (a) and (b) follow from the definition of $\mathfrak{C}[\mathfrak{M}]$, with (a) following because $(s_i, \mathbf{a}_i) \in S \otimes \mathfrak{A}$ implies that $\mathbf{a}_i \in \mathfrak{A}(s_i)$, and (b) following because $\tilde{\mathfrak{P}}((s_i, \mathbf{a}_i), \cdot) = \mathfrak{D}(s_i, \mathbf{a}_i)$, and hence $\gamma_i \in \mathfrak{G}(\tilde{\mathfrak{P}}((s_i, \mathbf{a}_i), \cdot))$ implies that $\gamma_i \in \mathfrak{G}(\mathfrak{D}(s_i, \mathbf{a}_i))$.

Consider finite path $r \in \text{Paths}_{*,S}^{\llbracket \mathfrak{C}[\mathfrak{M}] \rrbracket}(s)$. We define the choice of π after r by mimicking the choice of σ after $f(r)$. Recall that $\sigma(f(r)) \in \text{Dist}(A)$, and that $A \subseteq \mathfrak{A} \times \text{Dist}(S)$, i.e., $\text{support}(\sigma(f(r)))$ contains pairs of the form (\mathbf{a}, α) for $\mathbf{a} \in A$ and $\alpha \in \text{Dist}(S)$. We assume w.l.o.g. that, for $\mathbf{a} \in \mathfrak{A}(\text{last}(r))$, there exists at most one $\alpha \in \mathfrak{G}(\mathfrak{D}(\text{last}(r), \mathbf{a}))$ such that $(\mathbf{a}, \alpha) \in \text{support}(\sigma(f(r)))$. The assumption can be made w.l.o.g. because the set of assignments $\mathfrak{G}(\mathfrak{D}(\text{last}(r), \mathbf{a}))$ is closed under convex combinations (as noted in, for example, [21]). In the following, for $\mathbf{a} \in \mathfrak{A}(\text{last}(r))$ such that $(\mathbf{a}, \alpha) \in \text{support}(\sigma(f(r)))$ for some $\alpha \in \text{Dist}(S)$, we let $\text{pol}(r, \mathbf{a}) = \alpha$, i.e., $\text{pol}(r, \mathbf{a})$ denotes the unique assignment such that $(\mathbf{a}, \text{pol}(r, \mathbf{a})) \in \text{support}(\sigma(f(r)))$. For each $\mathbf{a} \in \mathfrak{A}(\text{last}(r))$, let $\beta_{(\text{last}(r), \mathbf{a})} : S \otimes \mathfrak{A} \rightarrow \{0, 1\}$ be the function such that, for each $(s', \mathbf{a}') \in S \otimes \mathfrak{A}$, we have $\beta_{(\text{last}(r), \mathbf{a})}(s', \mathbf{a}') = 1$ if $(s', \mathbf{a}') = (\text{last}(r), \mathbf{a})$ and $\beta_{(\text{last}(r), \mathbf{a})}(s', \mathbf{a}') = 0$ otherwise. Note that $\beta_{(\text{last}(r), \mathbf{a})} \in \mathfrak{G}(\tilde{\mathfrak{P}}(\text{last}(r), \cdot))$, i.e., $\beta_{(\text{last}(r), \mathbf{a})}$ is an assignment for $\tilde{\mathfrak{P}}(\text{last}(r), \cdot)$ (recall that, for any $\mathbf{a} \in \mathfrak{A}(\text{last}(r))$, we have $\tilde{\mathfrak{P}}(s, (\text{last}(r), \mathbf{a})) = [0, 1]$). We then let $\pi(r)$ be such that $\pi(r)(\beta_{(\text{last}(r), \mathbf{a})}) = \sigma(f(r))(\mathbf{a}, \text{pol}(r, \mathbf{a}))$ for each $\mathbf{a} \in \mathfrak{A}(\text{last}(r))$ such that $(\mathbf{a}, \text{pol}(r, \mathbf{a})) \in \text{support}(\sigma(f(r)))$.²

Now consider finite path $r' \in \text{Paths}_{*,S \otimes \mathfrak{A}}^{\llbracket \mathfrak{C}[\mathfrak{M}] \rrbracket}(s)$, where $r' = r\beta(\text{last}(r), \mathbf{a})$ for some $r \in \text{Paths}_{*,S}^{\llbracket \mathfrak{C}[\mathfrak{M}] \rrbracket}(s)$, $\beta \in \mathfrak{G}(\tilde{\mathfrak{P}}(\text{last}(r), \cdot))$ and $(\text{last}(r), \mathbf{a}) \in S \otimes \mathfrak{A}$ (note that we use the notation $r\beta(\text{last}(r), \mathbf{a})$ to denote the finite path with prefix r and suffix $\text{last}(r)\beta(\text{last}(r), \mathbf{a})$; we use similar notation throughout this and subsequent proofs). We define the choice of π after r' as follows. Consider the case in which $(\mathbf{a}, \text{pol}(r, \mathbf{a})) \in \text{support}(\sigma(f(r)))$; then we let $\pi(r')$ be such that $\pi(r')(\text{pol}(r, \mathbf{a})) = 1$. Instead, in the case in which there does not exist any $(\mathbf{a}, \alpha) \in \text{support}(\sigma(f(r)))$, we let $\pi(r')$ be an arbitrary distribution.

²We note that the scheduler $\pi(r)$ can be defined in an alternative way, as follows. Define $\beta : S \otimes \mathfrak{A} \rightarrow [0, 1]$ to be the function such that $\beta(\text{last}(r), \mathbf{a}) = \sigma(f(r))(\mathbf{a}, \text{pol}(r, \mathbf{a}))$ for each $\mathbf{a} \in \mathfrak{A}(\text{last}(r))$ such that $(\mathbf{a}, \text{pol}(r, \mathbf{a})) \in \text{support}(\sigma(f(r)))$, and let $\pi(r)(\beta) = 1$, i.e., the probabilistic choice of the distribution $\sigma(f(r))$ is encoded in β , which is chosen by $\pi(r)$ with probability 1. We do not use this alternative definition of π .

Recall that $Paths_*^\sigma(s)$ and $Paths_*^\pi(s)$ are the sets of finite paths from state s induced by σ and π , respectively. Let $Paths_{*,S}^\pi(s) = Paths_*^\pi(s) \cap Paths_{*,S}^{\llbracket \mathcal{C}[\mathcal{M}] \rrbracket}(s)$ and $Paths_{*,S \otimes \mathcal{A}}^\pi(s) = Paths_*^\pi(s) \cap Paths_{*,S \otimes \mathcal{A}}^{\llbracket \mathcal{C}[\mathcal{M}] \rrbracket}(s)$. Let $\Pr_{*,s}^\sigma$ (respectively, \Pr_s^π) be the probability measure over finite paths induced by σ (respectively, π), defined in the standard manner. In particular, we note that, for finite paths $r, r' \in Paths_{*,S}^\pi(s)$, if $r' = r\beta(\text{last}(r), \mathbf{a})\gamma s'$, then $\Pr_{*,s}^\pi(r') = \Pr_{*,s}^\pi(r) \cdot \pi(r)(\beta) \cdot \beta(\text{last}(r), \mathbf{a}) \cdot \pi(r\beta(\text{last}(r), \mathbf{a}))(\gamma) \cdot \gamma(s')$ (where $\Pr_{*,s}^\pi(s) = 1$). Similarly, for finite paths $r, r' \in Paths_*^\sigma(s)$, if $r' = r(\mathbf{a}, \alpha)s'$, then $\Pr_{*,s}^\sigma(r') = \Pr_{*,s}^\sigma(r) \cdot \sigma(\mathbf{a}, \alpha) \cdot \alpha(s')$ (where $\Pr_{*,s}^\sigma(s) = 1$). We show that $\Pr_{*,s}^\pi(r) = \Pr_{*,s}^\sigma(f(r))$, for $r \in Paths_{*,S}^\pi(s)$, by induction on the length of paths, where the length of a finite path of $Paths_*^\sigma(s)$ or $Paths_{*,S}^\pi(s)$ is the number of states *from* S along the path (although paths $Paths_{*,S}^\pi(s)$ visit states from S and from $S \otimes \mathcal{A}$, when considering the length of such a path, we only consider states from S). For the base case, i.e., for the path of length 1, which comprises state s only, we note that $f(s) = s$, and that $\Pr_{*,s}^\pi(s) = \Pr_{*,s}^\sigma(f(s)) = 1$. Now consider $r, r' \in Paths_{*,S}^\pi(s)$ such that $r' = r\beta(\text{last}(r), \mathbf{a})\gamma s'$, and assume that we have already established that $\Pr_{*,s}^\pi(r) = \Pr_{*,s}^\sigma(f(r))$. Furthermore (given that $r' \in Paths_{*,S}^\pi(s)$), we have $\beta(\text{last}(r), \mathbf{a}) = 1$ and $\gamma = \text{pol}(r, \mathbf{a})$. Recalling that $f(r') = f(r)(\mathbf{a}, \text{pol}(r, \mathbf{a}))s'$ and, by construction, $\pi(r)(\beta) = \sigma(f(r))(\mathbf{a}, \text{pol}(r, \mathbf{a}))$ and $\pi(r\beta(\text{last}(r), \mathbf{a}))(\text{pol}(r, \mathbf{a})) = 1$, we can establish the following:

$$\begin{aligned} \Pr_{*,s}^\pi(r') &= \Pr_{*,s}^\pi(r) \cdot \pi(r)(\beta) \cdot \beta(\text{last}(r), \mathbf{a}) \cdot \pi(r\beta(\text{last}(r), \mathbf{a}))(\gamma) \cdot \gamma(s') \\ &= \Pr_{*,s}^\sigma(f(r)) \cdot \sigma(f(r))(\mathbf{a}, \text{pol}(r, \mathbf{a})) \cdot \text{pol}(r, \mathbf{a})(s') \\ &= \Pr_{*,s}^\sigma(f(r')) . \end{aligned}$$

The fact that $\Pr_{*,s}^\pi$ and $\Pr_{*,s}^\sigma$ assign the same probability to finite paths related by f , together with the fact that finite paths related by f visit the same states from S , means that $\Pr_{*,s}^\pi$ and $\Pr_{*,s}^\sigma$ assign the same probability to sets of finite paths that visit T . By standard reasoning, this means that $\Pr_s^\sigma(\Diamond T) = \Pr_s^\pi(\Diamond T)$. \square

Lemma 2.4. *Let $s \in S$ and $T \subseteq S$. Then, for each $\pi \in \Sigma^{\llbracket \mathcal{C}[\mathcal{M}] \rrbracket}$, there exists $\sigma \in \Sigma^{\llbracket \mathcal{M} \rrbracket}$ such that $\Pr_s^\sigma(\Diamond T) = \Pr_s^\pi(\Diamond T)$.*

Proof. We proceed by constructing σ from π by using the “reverse” of the construction used for Lemma 2.3: each sequence of two choices of π , the first determining a transition from a state in S to a state in $S \otimes \mathcal{A}$, the second determining a transition from a state in $S \otimes \mathcal{A}$ to a state in S , is mimicked by a single choice of σ .

In the following, we assume w.l.o.g. that, for each finite path $r \in Paths_*^{\llbracket \mathcal{C}[\mathcal{M}] \rrbracket}$, there exists some $\beta \in \mathfrak{G}(\mathfrak{P}(\text{last}(r), \cdot))$ such that $\pi(r)(\beta) = 1$. The assumption is w.l.o.g. because, as in the proof of Lemma 2.3, the set $\mathfrak{G}(\mathfrak{P}(\text{last}(r), \cdot))$ is closed under convex combinations.

As in the proof of Lemma 2.3, we introduce a function from the set of finite paths of $Paths_*^{\llbracket \mathcal{M} \rrbracket}(s)$ to the set of finite paths of $Paths_*^{\llbracket \mathcal{C}[\mathcal{M}] \rrbracket}$; however, in contrast to the proof of Lemma 2.3, the range of the function includes only finite paths that are generated by π , with \perp used to represent cases in which a finite path of $Paths_*^{\llbracket \mathcal{M} \rrbracket}(s)$ has no corresponding finite path in $Paths_*^\pi(s)$. Let $g : Paths_*^{\llbracket \mathcal{M} \rrbracket}(s) \rightarrow Paths_*^\pi(s) \cup \{\perp\}$ be the function defined as follows: for $r \in Paths_*^{\llbracket \mathcal{M} \rrbracket}(s)$, where $r = s_0(\mathbf{a}_0, \alpha_0)s_1(\mathbf{a}_1, \alpha_1) \cdots (\mathbf{a}_{n-1}, \alpha_{n-1})s_n$, if there exists sequence $\beta_0\beta_1 \cdots \beta_{n-1}$ for which the finite path

$$r' = s_0\beta_0(s_0, \mathbf{a}_0)\alpha_0s_1\beta_1(s_1, \mathbf{a}_1)\alpha_1 \cdots s_{n-1}\beta_{n-1}(s_{n-1}, \mathbf{a}_{n-1})\alpha_{n-1}s_n$$

is such that $r' \in \text{Paths}_*^\pi(s)$, then $g(r) = r'$, otherwise $g(r) = \perp$. Note that, from the assumption on π made in the previous paragraph, if such a sequence $\beta_0\beta_1 \cdots \beta_{n-1}$ exists, it is unique, and hence the function g is well-defined.

Let $r \in \text{Paths}_*^{\llbracket \mathfrak{M} \rrbracket}(s)$. If $g(r) = \perp$, then we define $\sigma(r)$ to be an arbitrary distribution. Otherwise we derive $\sigma(r)$ from $\pi(g(r))$. First, recall that $\pi(g(r))(\beta) = 1$ for some $\beta \in \mathfrak{G}(\tilde{\mathfrak{P}}(\text{last}(r), \cdot))$. Furthermore, recall that $\beta \in \text{Dist}(S \otimes \mathfrak{A})$ and note that, by construction, for all $(s', \mathfrak{a}) \in \text{support}(\beta)$ we have $s' = \text{last}(r)$. Consider some $(\text{last}(r), \mathfrak{a}) \in \text{support}(\beta)$. For the resulting finite path $g(r)\beta(\text{last}(r), \mathfrak{a})$, we now consider $\pi(g(r)\beta(\text{last}(r), \mathfrak{a}))$. As above, we can assume w.l.o.g. that $\pi(g(r)\beta(\text{last}(r), \mathfrak{a}))(\gamma) = 1$ for some $\gamma \in \mathfrak{G}(\tilde{\mathfrak{P}}((\text{last}(r), \mathfrak{a}), \cdot))$. We can now define $\sigma(r)(\mathfrak{a}', \alpha)$ for each $(\mathfrak{a}', \alpha) \in A(s)$: if $(\mathfrak{a}', \alpha) = (\mathfrak{a}, \gamma)$, then $\sigma(r)(\mathfrak{a}', \alpha) = \beta(\text{last}(r), \mathfrak{a})$ and $\sigma(r)(\mathfrak{a}', \alpha) = 0$ otherwise.

Next, we show that $\text{Pr}_{*,s}^\sigma(r) = \text{Pr}_{*,s}^\pi(g(r))$, for $r \in \text{Paths}_*^\sigma(s)$ by induction on the length of paths. For the base case, given that $g(s) = s$, we have $\text{Pr}_{*,s}^\sigma(s) = \text{Pr}_{*,s}^\pi(g(s)) = 1$. Now consider $r, r' \in \text{Paths}_*^\sigma(s)$ such that $r' = r(\mathfrak{a}, \alpha)s'$, and assume that we have already established that $\text{Pr}_{*,s}^\sigma(r) = \text{Pr}_{*,s}^\pi(g(r))$. From the definition of g , we have that $g(r') = g(r)\beta(\text{last}(r), \mathfrak{a})\alpha s'$, where $\beta \in \mathfrak{G}(\tilde{\mathfrak{P}}(\text{last}(r), \cdot))$ such that $\pi(g(r))(\beta) = 1$ (and hence $g(r') \in \text{Paths}_*^\pi(s)$). Recall that we have constructed σ such that $\sigma(r)(\mathfrak{a}, \alpha) = \beta(\text{last}(r), \mathfrak{a})$; furthermore, we have assumed w.l.o.g. $\pi(g(r))(\beta) = 1$ and $\pi(g(r)\beta(\text{last}(r), \mathfrak{a}))(\alpha) = 1$. Then we have:

$$\begin{aligned} \text{Pr}_{*,s}^\sigma(r') &= \text{Pr}_{*,s}^\sigma(r) \cdot \sigma(r)(\mathfrak{a}, \alpha) \cdot \alpha(s') \\ &= \text{Pr}_{*,s}^\pi(g(r)) \cdot \pi(r)(\beta) \cdot \beta(\text{last}(r), \mathfrak{a}) \cdot \pi(g(r)\beta(\text{last}(r), \mathfrak{a}))(\alpha) \cdot \alpha(s') \\ &= \text{Pr}_{*,s}^\pi(g(r')) . \end{aligned}$$

As in the proof of Lemma 2.3, we conclude that $\text{Pr}_s^\sigma(\Diamond T) = \text{Pr}_s^\pi(\Diamond T)$. \square

3. CLOCK-DEPENDENT PROBABILISTIC TIMED AUTOMATA WITH ONE CLOCK

In this section, we recall the formalism of clock-dependent probabilistic timed automata. The definition of clock-dependent probabilistic timed automata of [33] features an arbitrary number of clock variables. In contrast, we consider models with only *one* clock variable. This clock variable 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) $\mathcal{P} = (L, inv, prob)$ comprises the following components:

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

A probabilistic edge $(l, g, \wp) \in prob$ 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 [17, 26, 24]: 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 nondeterministic; 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 probabilistic timed automata, the key novelty of 1c-cdPTAs is that the distribution used to make the aforementioned 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 $v \models inv(l)$. 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 $inv(l)$ remains satisfied continuously while time passes. The resulting state after the elapse of time is $(l, v+t)$. A probabilistic edge $(l', g, \wp) \in prob$ can then be chosen from state $(l, v+t)$ 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 $\wp[v+t]$. For example, in the case of the 1c-cdPTA of Figure 1, from state $(W, 0)$ (i.e., the location is W and the value of clock x is equal to 0), a time delay $t \in (1, 3)$ elapses, increasing the value of x to t , before the probabilistic edge leaving W is traversed. Then the resulting state will be (S, t) with probability $\frac{3t-3}{8}$, (T, t) with probability $\frac{11t-3}{16}$, and (F, t) with probability $\frac{11t-3}{16}$.

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 upper bound and impose no lower bound: formally, for each location $l \in L$, we have that $inv(l)$ 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 $\mathcal{P} = (L, \text{inv}, \text{prob})$ is the MDP $\llbracket \mathcal{P} \rrbracket = (S, A, \Delta)$ where:

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

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

Note that the summation in the first case of the definition of Δ is performed because the valuation $v' = 0$ may be obtained from $\tilde{v} = 0$ either by resetting the clock to 0 (summand $\wp[\tilde{v}](\{x\}, l')$) or by not resetting the clock to 0 (summand $\wp[\tilde{v}](\emptyset, l')$).

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 $(l, 0)$ 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} , $F \subseteq L$, $l \in L$, $\trianglelefteq \in \{\leq, <\}$ and $\lambda \in [0, 1]$ is to decide whether $\mathbb{P}_{\llbracket \mathcal{P} \rrbracket, (l, 0)}^{\min}(\Diamond T_F) \trianglelefteq \lambda$. 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 $\sigma \in \Sigma^{\llbracket \mathcal{P} \rrbracket}$.

3.1. 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 $p = (l, g, \wp) \in \text{prob}$, let I^p be the set of clock valuations in which p is enabled, i.e., $I^p = \{v \in \mathbb{R}_{\geq 0} \mid v \models g \wedge \text{inv}(l)\}$. 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 $v \in \overline{I^p}$. The following technical fact will be useful in subsequent sections: 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 $d_e^p \neq 0$ and the fact that I is open, would




Figure 3: Overview of the use of the IMDP construction in the overall solution process.

mean that there exists $v' \in I$ such that $\wp[v'](e) < 0$, which contradicts the definition of distribution templates).

3.2. 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 $(l_0, g_0, \wp_0)(X_0, l_1)(l_1, g_1, \wp_1)(X_1, l_2) \cdots (l_n, g_n, \wp_n) \in (prob \times (2^{\{x\}} \times L))^+ \times prob$ 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 non-constant probabilistic edges requires that the clock takes a natural numbered value at some point along the path fragment, 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) $X_i = \{x\}$ for some $0 \leq i < n$ 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 [33], 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 [27, 24].

We summarise our overall approach in Figure 3. From the 1c-cdPTA \mathcal{P} , we construct an IMDP $\mathfrak{M}[\mathcal{P}]$, from which we can obtain in turn the IMC $\mathfrak{C}[\mathfrak{M}[\mathcal{P}]]$ according to the construction of Section 2 (top line, from left to right). Furthermore, by Lemma 2.3 and Lemma 2.4, for any scheduler of $\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket$, we can find an equivalent (in terms of assigning the same probabilities for reaching a set of locations) scheduler of $\llbracket \mathfrak{C}[\mathfrak{M}[\mathcal{P}]] \rrbracket$, and vice versa (as indicated by the lower right dashed line). In this section, in Lemma 4.5 and Lemma 4.7, we will also show that analogous results allowing us to find, for any scheduler of $\llbracket \mathcal{P} \rrbracket$, an equivalent scheduler of $\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket$, and vice versa (lower left dashed line). Overall, this allows us to relate the quantitative and qualitative problems defined at the level of 1c-cdPTA to analogous problems defined at the level of IMCs, for which there exist efficient solution algorithms [31, 12, 9, 34].

Let $\mathcal{P} = (L, inv, prob)$ be a 1c-cdPTA. Let $Cst(\mathcal{P})$ be the set of constants that are used in the guards of probabilistic edges and invariants of \mathcal{P} , and let $\mathbb{B} = Cst(\mathcal{P}) \cup \{0\}$. 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

$\mathcal{I}_{\mathbb{B}} = \{[b_0, b_0], (b_0, b_1), [b_1, b_1], \dots, [b_k, b_k]\}$, i.e., $\mathcal{I}_{\mathbb{B}}$ is a partition of the interval $[0, b_k]$ into subintervals with endpoints in \mathbb{B} , and where each element of \mathbb{B} has a corresponding closed interval in $\mathcal{I}_{\mathbb{B}}$ comprising only that element. 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 $\text{le}(B) = b$ and $\text{re}(B) = b'$ refer to the left- and right-endpoints of B . For a closed interval $[b, b] \in \mathcal{I}_{\mathbb{B}}$, we let $\text{le}(B) = \text{re}(B) = b$.

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 $B \subseteq \{v \in \mathbb{R}_{\geq 0} \mid v \models \psi\}$ (representing the fact that all valuations of B satisfy ψ).

Example 4.1. 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 $(4, 5) \models (4 < x < 5)$.

4.1. \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 a *single* action (\tilde{v}, p) such that $\tilde{v} \in \tilde{B}$. Formally, we say that a scheduler $\sigma \in \Sigma[\mathcal{P}]$ of $[\mathcal{P}]$ is \mathbb{B} -minimal if, for all finite paths $r \in \text{Paths}_*^{\mathcal{P}}$, 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 \tilde{v}_1 and \tilde{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 $\Sigma_{\mathbb{B}}^{\mathcal{P}}$ be the set of schedulers of $\Sigma[\mathcal{P}]$ that are \mathbb{B} -minimal. The lemma allows us to consider only \mathbb{B} -minimal schedulers in the sequel, permitting us to obtain a close correspondence between the schedulers of $[\mathcal{P}]$ and the schedulers of the IMDP that we describe how to construct in the next subsection.

Lemma 4.2. *Let $(l, v) \in S_{\mathcal{P}}$ and $F \subseteq L$. Then, for each $\sigma \in \Sigma[\mathcal{P}]$, there exists $\pi \in \Sigma_{\mathbb{B}}^{\mathcal{P}}$ such that $\Pr_{(l,v)}^{\sigma}(\Diamond T_F) = \Pr_{(l,v)}^{\pi}(\Diamond T_F)$.*

Proof. Let $(l, v) \in S_{\mathcal{P}}$, $F \subseteq L$ and $\sigma \in \Sigma[\mathcal{P}]$. We proceed by describing the construction of $\pi \in \Sigma_{\mathbb{B}}^{\mathcal{P}}$, and then show that $\Pr_{(l,v)}^{\sigma}(\Diamond T_F) = \Pr_{(l,v)}^{\pi}(\Diamond T_F)$. The overall approach to the construction of π is that each finite path of π is associated with a number of finite paths of σ ; then a choice of π after a finite path r is based on a weighted average of the corresponding choices made by σ after the finite paths associated with r . The association between finite paths of σ and finite paths of π is defined on the basis of the locations that are visited, the intervals from $\mathcal{I}_{\mathbb{B}}$ that the value of the clock passes through, and the probabilistic edges that are taken along those finite paths. To define the choice of transition of π after a finite path r , we consider a pair $(\tilde{B}, p) \in \mathcal{I}_{\mathbb{B}} \times \text{prob}$, and, if σ assigns positive probability to at least one transition (\tilde{v}, p) such that $\tilde{v} \in \tilde{B}$, we define a unique transition (\tilde{v}', p) that is assigned positive probability by π and for which $\tilde{v}' \in \tilde{B}$. There are two principal cases for the transitions assigned positive probability by π : if \tilde{B} is a closed interval, then there is only one possible choice for \tilde{v}' (i.e., if $\tilde{B} = [b, b]$ then $\tilde{v}' = b$); instead, if \tilde{B} is an open interval, the

definition of the π is substantially more complicated, and the choice of $\tilde{v}' \in \tilde{B}$ takes into account the “recent” history of the scheduler. More precisely, while the scheduler π chooses transitions for which the value of the clock remains within \tilde{B} , the scheduler passes between three phases, where the number of transitions in the first and third phases is arbitrary, and where the second phase consists of at most one transition. The first phase consists of transitions derived from constant probabilistic edges, where the value of the clock is kept “low” (equal to the minimum of the value of the clock at the end of the corresponding finite paths of σ). The second phase consists of a single transition derived from a non-constant probabilistic edge, where the value of the clock is chosen carefully so as to replicate exactly the probability of taking the same probabilistic edge by σ . The third phase consists of transitions derived from constant probabilistic edges, and the value of the clock can be chosen in an arbitrary manner. Note that the assumption of initialisation guarantees that we cannot have more than one non-constant probabilistic edge between points at which the clock is equal to a value from \mathbb{B} , and hence at most one non-constant probabilistic edge can be taken while the value of the clock remains in the open interval \tilde{B} .

We first introduce the following notation. Given a probabilistic edge $(l, g, \wp) \in \text{prob}$, we let $\text{source}(l, g, \wp) = l$ and $\text{guard}(l, g, \wp) = g$. Furthermore, given interval $\tilde{B} \in \mathcal{I}_{\mathbb{B}}$, we say that $(l', B') \in L \times \mathcal{I}_{\mathbb{B}}$ is a *successor* of $(\tilde{B}, (l, g, \wp))$ if either $B' = [0, 0]$ and $\wp[v](\{x\}, l') > 0$ for some $v \in \tilde{B}$, or $B' = \tilde{B}$ and $\wp[v](\emptyset, l') > 0$ for some $v \in \tilde{B}$.

A (finite) \mathbb{B} -path is a sequence $(l_0, B_0)(\tilde{B}_0, p_0)(l_1, B_1)(\tilde{B}_1, p_1) \cdots (\tilde{B}_{n-1}, p_{n-1})(l_n, B_n)$ where (1) $l_i \in L$, $B_i \in \mathcal{I}_{\mathbb{B}}$ and $B_i \models \text{inv}(l_i)$ for all $0 \leq i \leq n$, and (2) $\text{source}(p_i) = l_i$, $\tilde{B}_i \in \mathcal{I}_{\mathbb{B}}$, $\tilde{B}_i \geq B_i$, $\tilde{B}_i \models \text{guard}(p_i) \wedge \text{inv}(l_i)$ and (l_{i+1}, B_{i+1}) is a successor of (\tilde{B}_i, p_i) , for all $0 \leq i < n$.

Consider a finite path $r = (l_0, v_0)(\tilde{v}_0, p_0)(l_1, v_1)(\tilde{v}_1, p_1) \cdots (\tilde{v}_{n-1}, p_{n-1})(l_n, v_n)$ of $\llbracket \mathcal{P} \rrbracket$ (i.e., $r \in \text{Paths}_*^{\llbracket \mathcal{P} \rrbracket}$), and a \mathbb{B} -path $\rho = (l'_0, B_0)(\tilde{B}_0, p'_0)(l'_1, B_1)(\tilde{B}_1, p'_1) \cdots (\tilde{B}_{m-1}, p'_{m-1})(l'_m, B_m)$. Then r *corresponds to* ρ if (1) $n = m$, (2) $l_i = l'_i$ and $v_i \in B_i$ for all $0 \leq i \leq n$, and (3) $\tilde{v}_i \in \tilde{B}_i$ and $p_i = p'_i$ for all $0 \leq i < n$. Given a scheduler $\sigma \in \Sigma^{\llbracket \mathcal{P} \rrbracket}$ and a \mathbb{B} -path ρ , we let $\Lambda^\sigma(\rho) \subseteq \text{Paths}_*^{\llbracket \mathcal{P} \rrbracket}$ be the set of finite paths of σ that correspond to ρ .

Note that, for a \mathbb{B} -minimal scheduler $\pi \in \Sigma_{\mathbb{B}}^{\llbracket \mathcal{P} \rrbracket}$ and a \mathbb{B} -path

$$\rho = (l_0, B_0)(\tilde{B}_0, p_0)(l_1, B_1)(\tilde{B}_1, p_1) \cdots (\tilde{B}_{n-1}, p_{n-1})(l_n, B_n) ,$$

there exists *at most one* finite path $r \in \Lambda^\pi(\rho)$ (i.e., $\Lambda^\pi(\rho)$ is either empty or a singleton). The reasoning underlying this fact is as follows: for each $i < n$, letting $\rho_i = (l_0, B_0)(\tilde{B}_0, p_0)(l_1, B_1)(\tilde{B}_1, p_1) \cdots (\tilde{B}_{i-1}, p_{i-1})(l_i, B_i)$, and assuming that $\Lambda^\pi(\rho_i)$ contains a unique path denoted by r , then there exists at most one valuation $\tilde{v} \in \tilde{B}_i$ such that $\pi(r)(\tilde{v}, p_i) > 0$. Furthermore, there exists at most one valuation $v' \in B_{i+1}$ such that $\Delta_{\mathcal{P}}(\text{last}(r), (\tilde{v}, p_i))(l_{i+1}, v') > 0$. Conversely, by using similar reasoning, for any finite path $r \in \text{Paths}_*^{\llbracket \mathcal{P} \rrbracket}$, there exists exactly one \mathbb{B} -path ρ such that $r \in \Lambda^\pi(\rho)$.

We now construct $\pi \in \Sigma_{\mathbb{B}}^{\llbracket \mathcal{P} \rrbracket}$, proceeding by induction on the length of paths. Given a finite path $r = (l_0, v_0)(\tilde{v}_0, p_0)(l_1, v_1)(\tilde{v}_1, p_1) \cdots (\tilde{v}_{n-1}, p_{n-1})(l_n, v_n)$, we consider the problem of defining $\pi(r)$ (assuming that π has been defined for all prefixes of r). Let $\rho = (l_0, B_0)(\tilde{B}_0, p_0)(l_1, B_1)(\tilde{B}_1, p_1) \cdots (\tilde{B}_{n-1}, p_{n-1})(l_n, B_n)$ be the unique \mathbb{B} -path such that r corresponds to ρ . Now consider the extension of ρ with (\tilde{B}, p) such that $\text{source}(p) = l_n$, $\tilde{B} \in \mathcal{I}_{\mathbb{B}}$, $\tilde{B} \geq B_n$, and $\tilde{B} \models \text{guard}(p) \wedge \text{inv}(l_n)$. Let $\Xi_{\rho, (\tilde{B}, p)}$ be the set of clock valuations in \tilde{B} that are featured with p in actions that are assigned positive probability by σ after

paths corresponding to ρ ; formally:

$$\Xi_{\rho,(\tilde{B},p)} = \{\tilde{v} \in \tilde{B} \mid (\tilde{v}, p) \in \text{support}(\sigma(r')) \text{ and } r' \in \Lambda^\sigma(\rho)\}.$$

In order to define $\pi(r)$, we consider a number of cases that depend on r and the pair (\tilde{B}, p) used to extend ρ . In all of the cases, we identify a clock valuation $\tilde{v}_{\rho,(\tilde{B},p)}^*$ that depends on the case, and define:

$$\pi(r)(\tilde{v}_{\rho,(\tilde{B},p)}^*, p) = \frac{\sum_{r' \in \Lambda^\sigma(\rho)} \Pr_{*,(l,v)}^\sigma(r') \cdot \sum_{\tilde{v} \in \Xi_{\rho,(\tilde{B},p)}} \sigma(r')(\tilde{v}, p)}{\Pr_{*,(l,v)}^\sigma(\Lambda^\sigma(\rho))},$$

where $\Pr_{*,(l,v)}^\sigma(\Lambda^\sigma(\rho)) = \sum_{r' \in \Lambda^\sigma(\rho)} \Pr_{*,(l,v)}^\sigma(r')$.

First we consider the case in which \tilde{B} is a closed interval, i.e., $\tilde{B} = [b, b]$ for $b \in \mathbb{B}$, which corresponds to the first phase described above. In this case, we simply let $\tilde{v}_{\rho,(\tilde{B},p)}^* = b$. Note that $\Xi_{\rho,(\tilde{B},p)}$ is a singleton, i.e., $\Xi_{\rho,(\tilde{B},p)} = \{\tilde{v}_{\rho,(\tilde{B},p)}^*\}$.

Next, we consider the case in which \tilde{B} is an open interval. Recall the description of the three phases given above. Let $k \leq n$ be the maximum index for which $B_k \neq \tilde{B}$ (and let $k = 0$ if no such index exists). Note that $\tilde{B}_k = B_{k+1} = \tilde{B}_{k+1} = \dots = \tilde{B}_{n-1} = B_n = \tilde{B}$. First we consider the subcase in which all probabilistic edges in the sequence $p_k \dots p_n$ are constant, and p is also constant: this corresponds to the first phase described above. Note that there may be multiple choices made by σ that correspond to (\tilde{B}, p) (i.e., from each finite path in $\Lambda^\sigma(\rho)$, the scheduler σ may assign positive probability to *multiple* time delays, each of which corresponds to a clock valuation in \tilde{B}). In order to replicate these choices in π we consider a time delay that results in a clock valuation that is equal to the *minimum* clock valuation assigned positive probability after any finite path in $\Lambda^\sigma(\rho)$, with the motivation that taking such a minimum gives π sufficient freedom in the third phase. Hence we let $\tilde{v}_{\rho,(\tilde{B},p)}^* = \min \Xi_{\rho,(\tilde{B},p)}$.

Now consider the subcase in which \tilde{B} is an open interval, and all probabilistic edges in the sequence $p_k \dots p_n$ are constant, but p is non-constant, i.e., the transition corresponds to the second phase described above. The definition of π is similar to that of the first phase, although the choice of clock valuation by π corresponds to the *weighted average* of the choice of clock valuations made by σ , where the weights refer to the probabilities of the finite paths of σ multiplied by the probability assigned by σ to the particular clock valuation. Formally, let:

$$\tilde{v}_{\rho,(\tilde{B},p)}^* = \frac{\sum_{r' \in \Lambda^\sigma(\rho)} \Pr_{*,(l,v)}^\sigma(r') \cdot \sum_{\tilde{v} \in \Xi_{\rho,(\tilde{B},p)}} \sigma(r')(\tilde{v}, p) \cdot \tilde{v}}{\sum_{r' \in \Lambda^\sigma(\rho)} \Pr_{*,(l,v)}^\sigma(r') \cdot \sum_{\tilde{v} \in \Xi_{\rho,(\tilde{B},p)}} \sigma(r')(\tilde{v}, p)}.$$

Finally, we consider the subcase in which \tilde{B} is an open interval, there exists $i \in \{k, \dots, n\}$ such that p_i is non-constant, and for all other $j \in \{k, \dots, n\}$ we have that p_j is constant, and also p is constant. This corresponds to the third phase described above. For this subcase, the clock valuation can be arbitrary: for simplicity we retain the same clock valuation as was used in the final state of r . Formally, if $\text{last}(r)$ is equal to (l', v') , we let $\tilde{v}_{\rho,(\tilde{B},p)}^* = v'$.

We repeat this process for all (\tilde{B}, p) such that $\text{source}(p) = l_n$, $\tilde{B} \in \mathcal{I}_{\mathbb{B}}$, $\tilde{B} \geq B_i$, and $\tilde{B} \models \text{guard}(p) \wedge \text{inv}(l_n)$. This suffices to define comprehensively the distribution $\pi(r)$. For example, consider the case in which the choice of π , to replicate the choices of σ , assigns

positive probability to actions with clock valuations in both closed and open intervals, and with probabilistic edges that are both constant and non-constant. Let Υ be the set of pairs (\tilde{B}, p) such that $\text{source}(p) = l_n$, $\tilde{B} \in \mathcal{I}_{\mathbb{B}}$, $\tilde{B} \geq B_i$, and $\tilde{B} \models \text{guard}(p) \wedge \text{inv}(l_n)$. Now observe that, for any $r' \in \Lambda^\sigma(\rho)$, we have:

$$\sum_{(\tilde{B}, p) \in \Upsilon} \sum_{\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}} \sigma(r')(\tilde{v}, p) = 1$$

(because, for each $(\tilde{v}, p) \in \text{support}(\sigma(r'))$ there exists $(\tilde{B}, p) \in \Upsilon$ such that $\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}$). We now show that $\sum_{(\tilde{B}, p) \in \Upsilon} \pi(r)(\tilde{v}_{\rho, (\tilde{B}, p)}^*, p) = 1$, i.e., $\pi(r)$ is a distribution:

$$\begin{aligned} & \sum_{(\tilde{B}, p) \in \Upsilon} \pi(r)(\tilde{v}_{\rho, (\tilde{B}, p)}^*, p) \\ &= \sum_{(\tilde{B}, p) \in \Upsilon} \left(\frac{\sum_{r' \in \Lambda^\sigma(\rho)} \text{Pr}_{*, (l, v)}^\sigma(r') \cdot \sum_{\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}} \sigma(r')(\tilde{v}, p)}{\text{Pr}_{*, (l, v)}^\sigma(\Lambda^\sigma(\rho))} \right) \\ &= \frac{\sum_{r' \in \Lambda^\sigma(\rho)} \text{Pr}_{*, (l, v)}^\sigma(r') \cdot \sum_{(\tilde{B}, p) \in \Upsilon} \sum_{\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}} \sigma(r')(\tilde{v}, p)}{\text{Pr}_{*, (l, v)}^\sigma(\Lambda^\sigma(\rho))} \\ &= \frac{\sum_{r' \in \Lambda^\sigma(\rho)} \text{Pr}_{*, (l, v)}^\sigma(r')}{\text{Pr}_{*, (l, v)}^\sigma(\Lambda^\sigma(\rho))} \\ &= 1. \end{aligned}$$


We now proceed to show that $\text{Pr}_{*, (l, v)}^\sigma(\Diamond T_F) = \text{Pr}_{*, (l, v)}^\pi(\Diamond T_F)$. It suffices to show that $\text{Pr}_{*, (l, v)}^\sigma(\Lambda^\sigma(\rho)) = \text{Pr}_{*, (l, v)}^\pi(\Lambda^\pi(\rho))$ for all \mathbb{B} -paths ρ (because all paths in $\Lambda^\sigma(\rho)$ reach a location in F if and only if the unique path in $\Lambda^\pi(\rho)$ reaches a location in F). We proceed by induction on the length of \mathbb{B} -paths. Let ρ be a \mathbb{B} -path, and assume that we have established that $\text{Pr}_{*, (l, v)}^\sigma(\Lambda^\sigma(\rho)) = \text{Pr}_{*, (l, v)}^\pi(\Lambda^\pi(\rho))$. Now consider the \mathbb{B} -path $\rho(\tilde{B}, p)(l', B')$ that extends ρ with one transition. Our aim is to show that $\text{Pr}_{*, (l, v)}^\sigma(\Lambda^\sigma(\rho(\tilde{B}, p)(l', B')))) = \text{Pr}_{*, (l, v)}^\pi(\Lambda^\pi(\rho(\tilde{B}, p)(l', B'))))$. In the sequel, we overload the notation $\Delta_{\mathcal{P}}$ in the following way: given $\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}$, we let $\Delta_{\mathcal{P}}(\text{last}(r'), (\tilde{v}, p))(l', B') = 0$ if $\Delta_{\mathcal{P}}(\text{last}(r'), (\tilde{v}, p))(l', v') = 0$ for all $v' \in B'$, and $\Delta_{\mathcal{P}}(\text{last}(r'), (\tilde{v}, p))(l', B') = \Delta_{\mathcal{P}}(\text{last}(r'), (\tilde{v}, p))(l', v')$ for the unique clock valuation $v' \in B'$ such that $\Delta_{\mathcal{P}}(\text{last}(r'), (\tilde{v}, p))(l', v') > 0$ (note that uniqueness follows from the fact that either $v' = 0$ or $v' = \tilde{v}$). In the following we consider the most involved case, namely that concerning \tilde{B} being open and p being non-constant, i.e., the second phase of the case of open \tilde{B} described above. Assume that $\Delta_{\mathcal{P}}(\text{last}(r'), (\tilde{v}, p))(l', B') > 0$, i.e., there is a clock valuation $v' \in B'$ such that $\Delta_{\mathcal{P}}(\text{last}(r'), (\tilde{v}, p))(l', v') > 0$. Furthermore, assume that $v' = \tilde{v}$, which means that $\Delta_{\mathcal{P}}(\text{last}(r'), (\tilde{v}, p))(l', B') = \wp[\tilde{v}](\emptyset, l')$ (the case in which $B' = [0, 0]$, and hence $v' = 0$ and $\Delta_{\mathcal{P}}(\text{last}(r'), (\tilde{v}, p))(l', B') = \wp[\tilde{v}](\{x\}, l')$, is similar). Let $(c_{(\emptyset, l')}, d_{(\emptyset, l')}) \in \mathbb{Q}^2$ be the pair associated with p and (\emptyset, l') defining the clock dependency; in order to simplify notation in the sequel, we let $c = c_{(\emptyset, l')}^p$ and $d = d_{(\emptyset, l')}^p$. Furthermore, we let $\Lambda^\pi(\rho) = r$, and write Pr_*^σ and Pr_*^π rather than $\text{Pr}_{*, (l, v)}^\sigma$ and $\text{Pr}_{*, (l, v)}^\pi$, respectively. In the following, we use \wp to denote the distribution template of probabilistic edge p . Then

we have:

$$\begin{aligned}
& \Pr_*^\sigma(\Lambda^\sigma(\rho(\tilde{B}, p)(l', B')))) \\
= & \sum_{r' \in \Lambda^\sigma(\rho(\tilde{B}, p)(l', B'))} \Pr_*^\sigma(r') \\
& \quad \text{(by definition of } \Lambda^\sigma(\rho(\tilde{B}, p)(l', B'))\text{)} \\
= & \sum_{r' \in \Lambda^\sigma(\rho)} \sum_{\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}} \Pr_*^\sigma(r') \cdot \sigma(r')(\tilde{v}, p) \cdot \Delta_{\mathcal{P}}(\text{last}(r'), (\tilde{v}, p))(l', B') \\
& \quad \text{(by definition of } \Pr_*^\sigma\text{)} \\
= & \sum_{r' \in \Lambda^\sigma(\rho)} \Pr_*^\sigma(r') \cdot \sum_{\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}} \sigma(r')(\tilde{v}, p) \cdot \wp[\tilde{v}](\emptyset, l') \\
& \quad \text{(rearranging and by definition of } \Delta_{\mathcal{P}}\text{)} \\
= & \sum_{r' \in \Lambda^\sigma(\rho)} \Pr_*^\sigma(r') \cdot \sum_{\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}} \sigma(r')(\tilde{v}, p) \cdot (c + d \cdot \tilde{v}) \\
& \quad \text{(by definition of non-constant probabilistic edges)} \\
= & c \left(\sum_{r' \in \Lambda^\sigma(\rho)} \Pr_*^\sigma(r') \cdot \sum_{\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}} \sigma(r')(\tilde{v}, p) \right) \\
& + d \left(\sum_{r' \in \Lambda^\sigma(\rho)} \Pr_*^\sigma(r') \cdot \sum_{\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}} \sigma(r')(\tilde{v}, p) \cdot \tilde{v} \right) \\
& \quad \text{(rearranging)} \\
= & \left(\sum_{r' \in \Lambda^\sigma(\rho)} \Pr_*^\sigma(r') \cdot \sum_{\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}} \sigma(r')(\tilde{v}, p) \right) \\
& \left(c + d \frac{\sum_{r' \in \Lambda^\sigma(\rho)} \Pr_*^\sigma(r') \cdot \sum_{\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}} \sigma(r')(\tilde{v}, p) \cdot \tilde{v}}{\sum_{r' \in \Lambda^\sigma(\rho)} \Pr_*^\sigma(r') \cdot \sum_{\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}} \sigma(r')(\tilde{v}, p)} \right) \\
& \quad \text{(rearranging)} \\
= & \left(\sum_{r' \in \Lambda^\sigma(\rho)} \Pr_*^\sigma(r') \cdot \sum_{\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}} \sigma(r')(\tilde{v}, p) \right) (c + d \cdot \tilde{v}_{\rho, (\tilde{B}, p)}^*) \\
& \quad \text{(by definition of } \tilde{v}_{\rho, (\tilde{B}, p)}^*\text{)} \\
= & \left(\sum_{r' \in \Lambda^\sigma(\rho)} \Pr_*^\sigma(r') \cdot \sum_{\tilde{v} \in \Xi_{\rho, (\tilde{B}, p)}} \sigma(r')(\tilde{v}, p) \right) \wp[\tilde{v}_{\rho, (\tilde{B}, p)}^*](\emptyset, l') \\
& \quad \text{(by definition of non-constant probabilistic edges)} \\
= & \Pr_*^\sigma(\Lambda^\sigma(\rho)) \cdot \pi(r)(\tilde{v}_{\rho, (\tilde{B}, p)}^*, p) \cdot \Delta_{\mathcal{P}}(\text{last}(r), (\tilde{v}_{\rho, (\tilde{B}, p)}^*, p))(l', B') \\
& \quad \text{(from the construction of } \pi \text{ and by the overloaded definition of } \Delta_{\mathcal{P}}\text{)} \\
= & \Pr_*^\pi(r) \cdot \pi(r)(\tilde{v}_{\rho, (\tilde{B}, p)}^*, p) \cdot \Delta_{\mathcal{P}}(\text{last}(r), (\tilde{v}_{\rho, (\tilde{B}, p)}^*, p))(l', B') \\
& \quad \text{(by induction)} \\
= & \Pr_*^\pi(\Lambda^\pi(\rho)(\tilde{B}, p)(l', B')) \\
& \quad \text{(by definition of } \Pr_*^\pi\text{)}.
\end{aligned}$$

We remark briefly that the cases for the second and fourth phases are simpler, because the probabilities $\wp[\cdot](\emptyset, l')$ and $\wp[\cdot](\{x\}, l')$ are constant (this means that the choice of clock valuation made by π , as long as it is in the open region denoted by \tilde{B} above, is arbitrary). The first phase is even more straightforward, because there is only *one* choice of clock valuation when considering choices of σ and π that correspond to ρ extended with (\tilde{B}, p) for which \tilde{B} is closed. \square

Example 4.3. 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 $\sigma \in \Sigma[\mathcal{P}]$, 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 $\pi \in \Sigma_{\mathbb{B}}[\mathcal{P}]$ 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

Figure 4: Interval Markov decision process $\mathfrak{M}[\mathcal{P}]$ obtained from \mathcal{P} of Figure 1.

$r = (W, 0)(\frac{5}{4}, p_W)(F, \frac{5}{4})$ 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 $(\frac{17}{4}, p_F)$ and probability 1 to $(\frac{19}{4}, p_F)$. 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 .

4.2. 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 in the sequence 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, an endpoint of the interval chosen in (ii). The second IMDP transition in the sequence represents the probabilistic choice made according to the chosen probabilistic edge, interval and endpoint chosen in the first transition of the sequence.

Example 4.4. The IMDP construction, applied to the example of Figure 1, is shown in Figure 4 (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, $((1, 3), p_W, \text{le})$ and $((1, 3), p_W, \text{re})$, where the third component of the state (le or re) denotes whether the state refers to the left- or right-endpoint of $(1, 3)$. Each of the states $((1, 3), p_W, \text{le})$ and $((1, 3), p_W, \text{re})$ is associated with the probability interval $(0, 1)$, referring to the probability of making a transition to those states from $(W, [0, 0])$. 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 $((1, 3), p_W, \text{le})$ and $\frac{1}{4}$ with $((1, 3), p_W, \text{re})$ (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 $(1, 3)$, obtaining the weighted combination $\frac{3}{4} \cdot 1 + \frac{1}{4} \cdot 3 = \frac{3}{2}$). Then, from both $((1, 3), p_W, \text{le})$ and $((1, 3), p_W, \text{re})$, there is a probabilistic choice regarding the target IMDP state to make the subsequent transition to, i.e., the transitions from $((1, 3), p_W, \text{le})$ and $((1, 3), p_W, \text{re})$ 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.³ The probabilities of the transitions from $((1, 3), p_W, \text{le})$ and $((1, 3), p_W, \text{re})$ 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 $((1, 3), p_W, \text{le})$ or $((1, 3), p_W, \text{re})$, 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 $((1, 3), p_W, \text{le})$ and $\frac{1}{4}$ to the transition to $((1, 3), p_W, \text{re})$, 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 $\mathfrak{M}[\mathcal{P}] = (S_{\mathfrak{M}[\mathcal{P}]}, \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}, \mathfrak{D}_{\mathfrak{M}[\mathcal{P}]})$. 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:

$$\begin{aligned} S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}} &= \{(l, B) \in L \times \mathcal{I}_{\mathbb{B}} \mid B \models \text{inv}(l)\} \\ S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}} &= \{(\tilde{B}, (l, g, \wp), \text{ep}) \in \mathcal{I}_{\mathbb{B}} \times \text{prob} \times \{\text{le}, \text{re}\} \mid \tilde{B} \models g \wedge \text{inv}(l)\} . \end{aligned}$$

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 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 $\mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(l, B) = \{(\tilde{B}, (l', g, \wp)) \in \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]} \mid l = l' \text{ and } \tilde{B} \geq B\}$.⁴ For each $(\tilde{B}, p, \text{ep}) \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}}$, let $\mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(\tilde{B}, p, \text{ep}) = \{\tau\}$. 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 follows:

³Given that there is only one action available from states such as $((1, 3), p_W, \text{le})$ and $((1, 3), p_W, \text{re})$, we omit both the action and the usual black box from the figure.

⁴Note 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).

- 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 $\mathfrak{D}_{\mathfrak{M}[\mathcal{P}]}((l, B), (\tilde{B}, p))(\tilde{B}, p, \text{le}) = (0, 1)$, $\mathfrak{D}_{\mathfrak{M}[\mathcal{P}]}((l, B), (\tilde{B}, p))(\tilde{B}, p, \text{re}) = (0, 1)$, and $\mathfrak{D}_{\mathfrak{M}[\mathcal{P}]}((l, B), (\tilde{B}, p))(s) = [0, 0]$ for all $s \in S_{\mathfrak{M}[\mathcal{P}]} \setminus \{(\tilde{B}, p, \text{le}), (\tilde{B}, p, \text{re})\}$.
- For each $(\tilde{B}, (l, g, \wp), \text{ep}) \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}}$ and $(l', B') \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}}$, let:

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

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

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

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

4.3. Correctness of the IMDP construction. 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 a transition 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 .

Lemma 4.5. *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 a scheduler $\hat{\pi} \in \Sigma^{\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket}$ such that $\Pr_{(l, 0)}^{\pi}(\diamond T_F) = \Pr_{\text{reg}(l, 0)}^{\hat{\pi}}(\diamond \mathfrak{T}_F)$.*

Before presenting the formal proof of Lemma 4.5, we sketch the overall approach for the construction of a scheduler $\hat{\pi}$ of $\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket$ from a \mathbb{B} -minimal scheduler π of $\llbracket \mathcal{P} \rrbracket$, and give an example. 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 *le* and *re* components used in endpoint indicators. Furthermore, $\hat{\pi}$ replicates exactly the choice of π made after r in terms of interval of $\mathcal{I}_{\mathbb{B}}$ and probabilistic edge chosen 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 4.4 above: more precisely, the choice of action (\tilde{v}, p) by π , where \tilde{B} is

the unique open interval such that $\tilde{v} \in \tilde{B}$, is mimicked by $\hat{\pi}$ choosing the action $((\tilde{B}, p), \alpha)$ for which $\alpha(\tilde{B}, p, \text{le}) = \frac{\text{re}(\tilde{B}) - \tilde{v}}{\text{re}(\tilde{B}) - \text{le}(\tilde{B})}$, and $\alpha(\tilde{B}, p, \text{re}) = 1 - \alpha(\tilde{B}, p, \text{le}) = \frac{\tilde{v} - \text{le}(\tilde{B})}{\text{re}(\tilde{B}) - \text{le}(\tilde{B})}$.

Example 4.6. Consider the 1c-cdPTA of Figure 1. Let $\pi \in \Sigma_{\mathbb{B}}^{\llbracket \mathcal{P} \rrbracket}$ be a \mathbb{B} -minimal 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 $\alpha((1, 3), p_W, \text{le}) = \frac{3}{4}$ and $\alpha((1, 3), p_W, \text{re}) = \frac{1}{4}$ (observe that $\alpha((1, 3), p_W, \text{le}) = \frac{3 - \frac{3}{2}}{2}$ and that $\alpha((1, 3), p_W, \text{re}) = 1 - \alpha((1, 3), p_W, \text{le}) = \frac{\frac{3}{2} - 1}{2}$). 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 $r' = (W, [0, 0])((1, 3), p_W, \alpha)((1, 3), p_W, \text{le})(F, (1, 3))$ and $r'' = (W, [0, 0])((1, 3), p_W, \alpha)((1, 3), p_W, \text{re})(F, (1, 3))$. 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 $\alpha'((4, 5), p_F, \text{le}) = \frac{1}{2}$ and $\alpha'((4, 5), p_F, \text{re}) = \frac{1}{2}$ (note that $\alpha'((4, 5), p_F, \text{le}) = 5 - \frac{9}{2}$ and $\alpha'((4, 5), p_F, \text{re}) = 1 - \alpha'((4, 5), p_F, \text{le}) = \frac{9}{2} - 4$). Hence, regardless of whether $((1, 3), p_W, \text{le})$ or $((1, 3), p_W, \text{re})$ was visited, scheduler $\hat{\pi}$ makes the *same* choice to mimic $\pi(r)$.

Proof of Lemma 4.5. Let $l \in L$, $F \subseteq L$ and $\pi \in \Sigma_{\mathbb{B}}^{\llbracket \mathcal{P} \rrbracket}$. We proceed by describing the construction of scheduler $\hat{\pi} \in \Sigma^{\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket}$, then show that $\Pr_{(l, 0)}^{\pi}(\Diamond T_F) = \Pr_{\text{reg}(l, 0)}^{\hat{\pi}}(\Diamond \mathfrak{T}_F)$.

Before describing the construction of $\hat{\pi} \in \Sigma^{\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket}$, we first simplify the notation for transitions along paths of $\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket$ to remove redundant elements, in the following way: a sequence $(l, B)((\tilde{B}, p), \alpha)(\tilde{B}, p, \text{ep})(\tau, \beta)(l', B')$ of two transitions between regions will be simplified to $(l, B)((\tilde{B}, p), \alpha)(\text{ep})(l', B')$ (note that, for the endpoint indicator $(\tilde{B}, p, \text{ep})$, the interval \tilde{B} and the probabilistic edge p have been featured in the previous position in the sequence, and that, from $(\tilde{B}, p, \text{ep})$, only one action is available, namely (τ, β) , and hence we can omit it from the sequence).

Recall that, from the proof of Lemma 4.2, each finite path $r \in \text{Paths}_{*}^{\pi}$ of π corresponds to a \mathbb{B} -path. Furthermore, for a given \mathbb{B} -path, we can identify a set of finite paths of $\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket$ that agree with the \mathbb{B} -path in terms of locations, intervals, and probabilistic edges. Formally, for a \mathbb{B} -path ρ which equals

$$(l_0, B_0)(\tilde{B}_0, p_0)(l_1, B_1)(\tilde{B}_1, p_1) \cdots (\tilde{B}_{n-1}, p_{n-1})(l_n, B_n)$$

and a finite path r which equals

$(l'_0, B'_0)((\tilde{B}'_0, p'_0), \alpha_0)(\text{ep}_0)(l'_1, B'_1)((\tilde{B}'_1, p'_1), \alpha_1)(\text{ep}_1) \cdots ((\tilde{B}'_{m-1}, p'_{m-1}), \alpha_{m-1})(\text{ep}_{m-1})(l'_m, B'_m)$ of $\mathfrak{M}[\mathcal{P}]$, we say that r *corresponds to* ρ if (1) $n = m$, (2) $l_i = l'_i$, $B_i = B'_i$ for all $0 \leq i \leq n$, and (3) $\tilde{B}_i = \tilde{B}'_i$ and $p_i = p'_i$ for all $0 \leq i < n$. Given a scheduler $\hat{\pi} \in \Sigma^{\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket}$ and a \mathbb{B} -path ρ , we let $\Lambda^{\hat{\pi}}(\rho) \subseteq \text{Paths}_{*}^{\hat{\pi}}$ be the set of finite paths of $\hat{\pi}$ that correspond to ρ .

Recall that $\mathfrak{M}[\mathcal{P}]$ alternates between regions (elements of $S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}}$) and endpoint indicators (elements of $S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}}$). We partition $\text{Paths}_{*}^{\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket}(\text{reg}(l, 0))$ into the set $\text{Paths}_{*, S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}}}^{\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket}(\text{reg}(l, 0))$ of finite paths of the form

$$(l_0, B_0)((\tilde{B}_0, p_0), \alpha_0)(\text{ep}_0)(l_1, B_1)((\tilde{B}_1, p_1), \alpha_1)(\text{ep}_1) \cdots ((\tilde{B}_{n-1}, p_{n-1}), \alpha_{n-1})(\text{ep}_{n-1})(l_n, B_n)$$

that end with a region, and the set $Paths_{*, S_{\mathfrak{M}[\mathcal{P}]}}^{\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket}(\text{reg}(l, 0))$ of finite paths of the form

$$(l_0, B_0)((\tilde{B}_0, p_0), \alpha_0)(\text{ep}_0)(l_1, B_1)((\tilde{B}_1, p_1), \alpha_1)(\text{ep}_1) \cdots ((\tilde{B}_m, p_m), \alpha_m)(\text{ep}_m)$$

that end with an endpoint indicator (note that ep_m , in the context of the last transition of the finite path, denotes the endpoint indicator $(\tilde{B}_m, p_m, \text{ep}_m)$).

The construction of $\hat{\pi}$ is based on the following principles: for any two finite paths that correspond to the same \mathbb{B} -path, the choice made by $\hat{\pi}$ after those finite paths will be the same; furthermore, for any action of $\mathfrak{M}[\mathcal{P}]$ of the form (\tilde{B}, p) , the scheduler $\hat{\pi}$ assigns positive probability to at most one action of $\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket$ of the form $((\tilde{B}, p), \alpha)$, where α is used to represent the actual clock valuation chosen by π within the interval \tilde{B} .

We now describe the formal construction of $\hat{\pi}$. As elsewhere in this paper, we proceed by induction on the length of paths. Let $r \in Paths_{*, S_{\mathfrak{M}[\mathcal{P}]}}^{\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket}(\text{reg}(l, 0))$ be a finite path of $\mathfrak{M}[\mathcal{P}]$ that ends in a region, and assume that we have already defined the choices of $\hat{\pi}$ along all of the prefixes of r . Let ρ be the unique \mathbb{B} -path to which r corresponds (it is obtained by simply removing the components denoted by α and ep), and let $r' \in Paths_{*}^{\pi}$ be the unique path of π that corresponds to ρ . Consider $(\tilde{v}, p) \in \text{support}(\pi(r'))$; we aim to mimic this choice in the construction of $\hat{\pi}(r)$. Let $\tilde{B} \in \mathcal{I}_{\mathbb{B}}$ be the unique interval for which $\tilde{v} \in \tilde{B}$. By definition, $(\tilde{B}, p) \in \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(\text{last}(r))$. Then we let $\hat{\pi}(r)((\tilde{B}, p), \alpha) = \pi(r')(\tilde{v}, p)$, where α is defined as follows: let $\alpha(\tilde{B}, p, \text{le}) = \frac{\text{re}(\tilde{B}) - \tilde{v}}{\text{re}(\tilde{B}) - \text{le}(\tilde{B})}$, and let $\alpha(\tilde{B}, p, \text{re}) = 1 - \alpha(\tilde{B}, p, \text{le}) = \frac{\tilde{v} - \text{le}(\tilde{B})}{\text{re}(\tilde{B}) - \text{le}(\tilde{B})}$. This completes the construction of $\hat{\pi}$, because only the action τ is available in the final states of finite paths in $Paths_{*, S_{\mathfrak{M}[\mathcal{P}]}}^{\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket}(\text{reg}(l, 0))$, and because $\hat{\pi}$ can be defined in an arbitrary manner for paths not starting in $\text{reg}(l, 0)$.

Next we show that $\Pr_{(l, 0)}^{\pi}(\diamond T_F) = \Pr_{\text{reg}(l, 0)}^{\hat{\pi}}(\diamond \mathfrak{T}_F)$. Our approach is to show that $\Pr_{(l, 0)}^{\pi}(\Lambda^{\pi}(\rho)) = \Pr_{\text{reg}(l, 0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho))$ for all \mathbb{B} -paths ρ (this is sufficient to show $\Pr_{(l, 0)}^{\pi}(\diamond T_F) = \Pr_{\text{reg}(l, 0)}^{\hat{\pi}}(\diamond \mathfrak{T}_F)$ because the unique path in $\Lambda^{\pi}(\rho)$ reaches a location in F if and only if all paths in $\Lambda^{\hat{\pi}}(\rho)$ reach a location in F). We proceed by induction on the length of \mathbb{B} -paths. Consider the \mathbb{B} -path $\rho(\tilde{B}, p)(l', B')$, and assume that it has already been established that $\Pr_{(l, 0)}^{\pi}(\Lambda^{\pi}(\rho)) = \Pr_{\text{reg}(l, 0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho))$. Our aim is now to show that $\Pr_{(l, 0)}^{\pi}(\Lambda^{\pi}(\rho(\tilde{B}, p)(l', B')))) = \Pr_{\text{reg}(l, 0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho(\tilde{B}, p)(l', B'))))$. As in the construction of $\hat{\pi}$, also in the following we denote $\Lambda^{\pi}(\rho)$ by r' . Given that π is a \mathbb{B} -minimal scheduler, there is at most one clock valuation $\tilde{v} \in \tilde{B}$ such that $(\tilde{v}, p) \in \text{support}(\sigma(r'))$. For the case in which no such $\tilde{v} \in \tilde{B}$ exists, i.e., $\{\tilde{v} \mid (\tilde{v}, p) \in \text{support}(\sigma(r'))\} \cap \tilde{B} = \emptyset$, then $\Pr_{(l, 0)}^{\pi}(\Lambda^{\pi}(\rho(\tilde{B}, p)(l', B')))) = 0$ and $\Pr_{\text{reg}(l, 0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho(\tilde{B}, p)(l', B')))) = 0$, and we are done. In the remainder of this proof, we consider the case in which there exists $\tilde{v} \in \tilde{B}$ such that $(\tilde{v}, p) \in \text{support}(\pi(r'))$. As in the proof of Lemma 4.2, we let $\Delta_{\mathcal{P}}(\text{last}(r'), (\tilde{v}, p))(l', B')$ be the probability of making a transition to a state in (l', B') from state $\text{last}(r')$ with action (\tilde{v}, p) . We also concentrate on the case in which \tilde{B} is open, p is non-constant, and $B' = \tilde{B}$, which corresponds to the outcome (\emptyset, l') (i.e., the clock is not reset), with the other cases being similar. We write c and d rather than $c_{(\emptyset, l')}$ and $d = d_{(\emptyset, l')}$, respectively. Furthermore, we write $\alpha(\text{ep})$ rather than $\alpha(\tilde{B}, p, \text{ep})$, for $\text{ep} \in \{\text{le}, \text{re}\}$.

First we observe the following (where the second step is obtained from $\tilde{v} - \text{le}(\tilde{B}) = \alpha(\text{re})(\text{re}(\tilde{B}) - \text{le}(\tilde{B}))$, and the fourth and fifth steps also use the fact that $\alpha(\text{le}) + \alpha(\text{re}) = 1$):

$$\begin{aligned}
c + d \cdot \tilde{v} &= c + d \cdot (\text{le}(\tilde{B}) + \tilde{v} - \text{le}(\tilde{B})) \\
&= c + d \cdot (\text{le}(\tilde{B}) + \alpha(\text{re})(\text{re}(\tilde{B}) - \text{le}(\tilde{B}))) \\
&= c + d \cdot (\text{le}(\tilde{B}) + \alpha(\text{re}) \cdot \text{re}(\tilde{B}) - \alpha(\text{re}) \cdot \text{le}(\tilde{B})) \\
&= c + d \cdot (\text{le}(\tilde{B}) + \alpha(\text{re}) \cdot \text{re}(\tilde{B}) - \text{le}(\tilde{B}) + \alpha(\text{le}) \cdot \text{le}(\tilde{B})) \\
&= (\alpha(\text{le}) + \alpha(\text{re})) \cdot c + d \cdot \alpha(\text{re}) \cdot \text{re}(\tilde{B}) + d \cdot \alpha(\text{le}) \cdot \text{le}(\tilde{B}) \\
&= \alpha(\text{le})(c + d \cdot \text{le}(\tilde{B})) + \alpha(\text{re})(c + d \cdot \text{re}(\tilde{B})) .
\end{aligned}$$

The fact that $c + d \cdot \tilde{v} = \alpha(\text{le})(c + d \cdot \text{le}(\tilde{B})) + \alpha(\text{re})(c + d \cdot \text{re}(\tilde{B}))$ is then used in the following derivation (where \wp denotes the distribution template of probabilistic edge p):

$$\begin{aligned}
&\Pr_{(l,0)}^\pi(\Lambda^\pi(\rho(\tilde{B}, p)(l', B'))) \\
&= \Pr_{(l,0)}^\pi(\Lambda^\pi(\rho) \cdot \pi(r')(\tilde{v}, p) \cdot \Delta_{\mathcal{P}}(\text{last}(r'), (\tilde{v}, p))(l', B')) \\
&= \Pr_{(l,0)}^\pi(\Lambda^\pi(\rho) \cdot \pi(r')(\tilde{v}, p) \cdot \wp[\tilde{v}](\emptyset, l')) \\
&= \Pr_{(l,0)}^\pi(\Lambda^\pi(\rho) \cdot \pi(r')(\tilde{v}, p) \cdot (c + d \cdot \tilde{v})) \\
&= \Pr_{(l,0)}^\pi(\Lambda^\pi(\rho) \cdot \pi(r')(\tilde{v}, p) \cdot \left(\alpha(\text{le})(c + d \cdot \text{le}(\tilde{B})) + \alpha(\text{re})(c + d \cdot \text{re}(\tilde{B})) \right)) \\
&= \Pr_{(l,0)}^\pi(\Lambda^\pi(\rho) \cdot \pi(r')(\tilde{v}, p) \cdot \left(\alpha(\text{le}) \cdot \wp[\text{le}(\tilde{B})](\emptyset, l') + \alpha(\text{re}) \cdot \wp[\text{re}(\tilde{B})](\emptyset, l') \right)) .
\end{aligned}$$

For the subsequent steps, we require the following notation. Recall that, for any given $(\tilde{B}, p, \text{ep}) \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}}$, the interval distribution $\mathfrak{D}_{\mathfrak{M}[\mathcal{P}]}((\tilde{B}, (l, g, \wp), \text{ep}), \tau)$ assigns *singleton* intervals to each $s \in S_{\mathfrak{M}[\mathcal{P}]}$. Hence, there is only one action associated with $(\tilde{B}, p, \text{ep})$ in $\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket$ which we denote as $(\tau, \alpha^{\tilde{B}, p, \text{ep}})$, where $\alpha^{\tilde{B}, p, \text{ep}}$ is the assignment such that $\alpha^{\tilde{B}, p, \text{ep}}(s) = \lambda_s^{(\tilde{B}, p, \text{ep})}$ for each $s \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}}$. Recall that we have assumed $\Pr_{(l,0)}^\pi(\Lambda^\pi(\rho)) = \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho))$, and have $\hat{\pi}(r)((\tilde{B}, p), \alpha) = \pi(r')(\tilde{v}, p)$ for all finite paths $r \in \Lambda^{\hat{\pi}}(\rho)$. In the following, we let $\hat{r} \in \Lambda^{\hat{\pi}}(\rho)$ be an arbitrary finite path from $\Lambda^{\hat{\pi}}(\rho)$. Then by construction:

$$\begin{aligned}
&\Pr_{(l,0)}^\pi(\Lambda^\pi(\rho) \cdot \pi(r')(\tilde{v}, p) \cdot \left(\alpha(\text{le}) \cdot \wp[\text{le}(\tilde{B})](\emptyset, l') + \alpha(\text{re}) \cdot \wp[\text{re}(\tilde{B})](\emptyset, l') \right)) \\
&= \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho) \cdot \hat{\pi}(\hat{r})((\tilde{B}, p), \alpha) \\
&\quad \times \left(\alpha(\text{le}) \cdot \Delta_{\mathfrak{M}[\mathcal{P}]}((\tilde{B}, p, \text{le}), (\tau, \alpha^{\tilde{B}, p, \text{le}}))(l', B') \right. \\
&\quad \left. + \alpha(\text{re}) \cdot \Delta_{\mathfrak{M}[\mathcal{P}]}((\tilde{B}, p, \text{re}), (\tau, \alpha^{\tilde{B}, p, \text{re}}))(l', B') \right) \\
&= \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho) \cdot \hat{\pi}(\hat{r})((\tilde{B}, p), \alpha) \cdot \alpha(\text{le}) \cdot \Delta_{\mathfrak{M}[\mathcal{P}]}((\tilde{B}, p, \text{le}), (\tau, \alpha^{\tilde{B}, p, \text{le}}))(l', B') \\
&\quad + \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho) \cdot \hat{\pi}(\hat{r})((\tilde{B}, p), \alpha) \cdot \alpha(\text{re}) \cdot \Delta_{\mathfrak{M}[\mathcal{P}]}((\tilde{B}, p, \text{re}), (\tau, \alpha^{\tilde{B}, p, \text{re}}))(l', B')) \\
&= \sum_{r \in \Lambda^{\hat{\pi}}(\rho)} \sum_{\text{ep} \in \{\text{le}, \text{re}\}} \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha) \cdot \alpha(\text{ep}) \\
&\quad \times \Delta_{\mathfrak{M}[\mathcal{P}]}((\tilde{B}, p, \text{ep}), (\tau, \alpha^{\tilde{B}, p, \text{ep}}))(l', B') \\
&= \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho(\tilde{B}, p)(l', B')))) .
\end{aligned}$$

This completes showing that $\Pr_{(l,0)}^\pi(\Lambda^\pi(\rho(\tilde{B},p)(l',B')))) = \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho(\tilde{B},p)(l',B'))))$ when $B' = \tilde{B}$. The case for $B' = [0,0]$ follows similarly. \square

The following lemma considers the converse direction, namely that (starting from a given location with the clock equal to 0) any scheduler of $\llbracket \mathcal{M}[\mathcal{P}] \rrbracket$ 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 4.7. *Let $l \in L$ be a location and let $F \subseteq L$ be a set of locations. Given a scheduler $\hat{\pi} \in \Sigma^{\llbracket \mathcal{M}[\mathcal{P}] \rrbracket}$, there exists a \mathbb{B} -minimal scheduler $\pi \in \Sigma_{\mathbb{B}}^{\llbracket \mathcal{P} \rrbracket}$, such that $\Pr_{(l,0)}^\pi(\Diamond T_F) = \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(\Diamond \mathfrak{T}_F)$.*

Proof. Let $l \in L$, $F \subseteq L$ and $\hat{\pi} \in \Sigma^{\llbracket \mathcal{M}[\mathcal{P}] \rrbracket}$. We describe the construction of scheduler $\pi \in \Sigma_{\mathbb{B}}^{\llbracket \mathcal{P} \rrbracket}$, and then show that $\Pr_{(l,0)}^\pi(\Diamond T_F) = \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(\Diamond \mathfrak{T}_F)$.

Before presenting the details, we sketch the overall approach of the proof. The proof proceeds in a similar manner to that of Lemma 4.2: here, as for the proof of Lemma 4.2, our aim is to obtain a (\mathbb{B} -minimal) scheduler π of $\llbracket \mathcal{P} \rrbracket$, i.e., a scheduler for which there is a one-to-one relationship between its finite paths and (a subset of) \mathbb{B} -paths. The principles underlying the construction of π are the same as those underlying the analogous construction of Lemma 4.2: for a finite path r of $\llbracket \mathcal{P} \rrbracket$, in order to define $\pi(r)$, we consider extensions of the unique \mathbb{B} -path ρ that corresponds to r ; for each of those extensions that correspond to a set of finite paths of $\hat{\pi}$, we define a choice made by the distribution $\pi(r)$ that mimics the final transition of the aforementioned set of finite paths of $\hat{\pi}$. The probabilities of the distribution $\pi(r)$ are obtained as weighted averages of the choices of $\hat{\pi}$. Showing that $\Pr_{(l,0)}^\pi(\Diamond T_F) = \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(\Diamond \mathfrak{T}_F)$ will be done in a similar manner to the analogous part of the proof of Lemma 4.5.

We first describe the construction of $\pi \in \Sigma_{\mathbb{B}}^{\llbracket \mathcal{P} \rrbracket}$. Consider the finite path

$$r = (l_0, v_0)(\tilde{v}_0, p_0)(l_1, v_1)(\tilde{v}_1, p_1) \cdots (\tilde{v}_{n-1}, p_{n-1})(l_n, v_n)$$

of $\llbracket \mathcal{P} \rrbracket$. Let

$$\rho = (l_0, B_0)(\tilde{B}_0, p_0)(l_1, B_1)(\tilde{B}_1, p_1) \cdots (\tilde{B}_{n-1}, p_{n-1})(l_n, B_n)$$

be the unique \mathbb{B} -path such that r corresponds to ρ . Similarly, recall that $\Lambda^{\hat{\pi}}(\rho)$ is the set of finite paths of $\hat{\pi}$ that correspond to ρ (where the notion of correspondence of finite paths of $\hat{\pi}$ to a \mathbb{B} -path is given in the proof of Lemma 4.5). Consider the extension of ρ with (\tilde{B}, p) such that $\text{source}(p) = l_n$, $\tilde{B} \in \mathcal{I}_{\mathbb{B}}$, $\tilde{B} \geq B_n$, and $\tilde{B} \models \text{guard}(p) \wedge \text{inv}(l_n)$.

We recall an assumption that we can make without loss of generality, as already described in the proof of Lemma 2.3: given $r' \in \Lambda^{\hat{\pi}}(\rho)$, we assume that there exists at most one $\alpha \in \mathfrak{G}(\mathfrak{D}_{\mathcal{M}[\mathcal{P}]}(\text{last}(r'), (\tilde{B}, p)))$ such that $((\tilde{B}, p), \alpha) \in \text{support}(\hat{\pi}(r'))$ (this follows from the fact that the set of assignments $\mathfrak{G}(\mathfrak{D}_{\mathcal{M}[\mathcal{P}]}(\text{last}(r'), (\tilde{B}, p)))$ is closed under convex combinations). If such $\alpha \in \mathfrak{G}(\mathfrak{D}_{\mathcal{M}[\mathcal{P}]}(\text{last}(r'), (\tilde{B}, p)))$ exists, we denote it by $\alpha_{(\tilde{B}, p)}^{r'}$.

As in the proof of Lemma 4.2, to define $\pi(r)$, for each pair (\tilde{B}, p) used to extend ρ , we identify a clock valuation $\tilde{v}_{\rho, (\tilde{B}, p)}^*$, and then let:

$$\pi(r)(\tilde{v}_{\rho, (\tilde{B}, p)}^*, p) = \frac{\sum_{r' \in \Lambda^{\hat{\pi}}(\rho)} \Pr_{*, \text{reg}(l,0)}^{\hat{\pi}}(r') \cdot \hat{\pi}(r')((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^{r'})}{\Pr_{*, \text{reg}(l,0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho))},$$

where $\Pr_{*,\text{reg}(l,0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho)) = \sum_{r' \in \Lambda^{\hat{\pi}}(\rho)} \Pr_{*,(l,0)}^{\hat{\pi}}(r')$, and where $\hat{\pi}(r')((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^{r'}) = 0$ if there does not exist any $\alpha \in \mathfrak{G}(\mathfrak{D}_{\mathfrak{M}[\mathcal{P}]}(\text{last}(r'), (\tilde{B}, p)))$ such that $((\tilde{B}, p), \alpha) \in \text{support}(\hat{\pi}(r'))$.

For the case in which \tilde{B} is a closed interval, i.e., $\tilde{B} = [b, b]$ for $b \in \mathbb{B}$, we let $\tilde{v}_{\rho, (\tilde{B}, p)}^* = b$.

We now consider the case in which \tilde{B} is an open interval. The approach we take is analogous to that taken in the proof of Lemma 4.2: intuitively, if p is a constant probabilistic edge, and no non-constant probabilistic edge has been taken while the clock value remains in \tilde{B} , then the scheduler π maintains the clock as having a value that is a (small) amount greater than the left endpoint of \tilde{B} ; if p is instead a non-constant probabilistic edge, then the value of the clock is advanced to a particular value in the interval \tilde{B} , in order to mimic the choice of particular assignments in $\hat{\pi}$; finally, if p is a constant probabilistic edge, and a non-constant probabilistic edge *has* been taken previously while the clock value remains in \tilde{B} , then the scheduler π can choose the value of the clock from \tilde{B} in an arbitrary manner. The construction more complicated than that of the proof of Lemma 4.2, because the clock valuations must be chosen by π to reflect *assignments* chosen by $\hat{\pi}$ (instead, for Lemma 4.2, the clock valuations were chosen in the construction of π to reflect *clock valuations* chosen by σ , which was simpler conceptually). As in the proof of Lemma 4.2, let $k \leq n$ be the maximum index for which $B_k \neq \tilde{B}$ (with $k = 0$ if no such index exists), noting that $\tilde{B}_k = B_{k+1} = \tilde{B}_{k+1} = \dots = \tilde{B}_{n-1} = B_n = \tilde{B}$. We consider the following function $h_{\tilde{B}} : \text{Dist}(\{(\tilde{B}, -, \text{le}), (\tilde{B}, -, \text{re})\}) \rightarrow \tilde{B}$ that maps an assignment (over the target endpoint indicators $(\tilde{B}, -, \text{le}), (\tilde{B}, -, \text{re})$ associated with (\tilde{B}, p)) to a clock valuation in \tilde{B} (where the probabilistic edge in the endpoint indicators $(\tilde{B}, -, \text{le})$ and $(\tilde{B}, -, \text{re})$ is omitted because it is irrelevant to the definition of the function). The function $h_{\tilde{B}}$ is defined as follows: given $\alpha \in \text{Dist}(\{(\tilde{B}, -, \text{le}), (\tilde{B}, -, \text{re})\})$, we let $h_{\tilde{B}}(\alpha) = \text{re}(\tilde{B}) - \alpha(\text{le}) \cdot (\text{re}(\tilde{B}) - \text{le}(\tilde{B}))$ (note that, as in the proof of Lemma 4.5, we write $\alpha(\text{le})$ rather than $\alpha(\tilde{B}, -, \text{le})$ for simplicity; in the sequel, we will also usually write $\alpha(\text{re})$ rather than $\alpha(\tilde{B}, -, \text{re})$). The intuition underlying the definition of $h_{\tilde{B}}$ is that $h_{\tilde{B}}(\alpha)$ is a clock valuation that represents faithfully the assignment α , in the sense that the greater the probability of $\alpha(\text{le})$, the closer the clock valuation is to the left endpoint of \tilde{B} .

First we consider the subcase in which all probabilistic edges in $p_k \dots p_n$ are constant, and p is also constant. While the clock must be set to some value in \tilde{B} , it is important that the value chosen is no greater than the value needed in the future to replicate the probabilities corresponding to a non-constant probabilistic edge. We now identify such values. We first define the set $\Omega_{\rho, (\tilde{B}, p)}$, which contains finite suffixes of paths in $\Lambda^{\hat{\pi}}(\rho)$, where the suffixes (1) start with a transition derived from (\tilde{B}, p) , (2) feature only \tilde{B} (apart from the suffixes' first state), and (3) features only constant probabilistic edges until the final transition, which features a non-constant probabilistic edge. We now define formally $\Omega_{\rho, (\tilde{B}, p)}$. For a finite path $r' \in \text{Paths}_{*, S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}}}^{\hat{\pi}}(\text{reg}(l, 0))$, let $\text{SuffFrom}_{*}^{\hat{\pi}}(r')$ be the set of (finite) suffixes of r' generated by $\hat{\pi}$ that end in a state in $S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}}$; formally, $\text{SuffFrom}_{*}^{\hat{\pi}}(r')$ is the smallest set such that, if $r'' \in \text{Paths}_{*, S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}}}^{\hat{\pi}}(\text{reg}(l, 0))$, where $r'' = r'r'''$, then $r''' \in \text{SuffFrom}_{*}^{\hat{\pi}}(r')$. We let $\text{Suffixes}_{*}^{\hat{\pi}}(\text{reg}(l, 0))$ be the set of all suffixes of finite paths of $\hat{\pi}$, starting in a state of $S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}}$

and terminating in a state in $S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}}$, that is:

$$\text{Suffixes}_*^{\hat{\pi}}(\text{reg}(l, 0)) = \bigcup_{r' \in \text{Paths}_{*, S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}}}^{\hat{\pi}}(\text{reg}(l, 0))} \text{SuffFrom}_*^{\hat{\pi}}(r').$$

Consider $r' \in \text{Suffixes}_*^{\hat{\pi}}(\text{reg}(l, 0))$, where $r' = (l_0, B_0)((\tilde{B}_0, p_0), \alpha_0)(\text{ep}_0) \cdots ((\tilde{B}_m, p_m), \alpha_m)(\text{ep}_m)$. We let $\text{FA}(r') = \alpha_m$ be the assignment featured in the final transition of r' . Furthermore, we say that:

- r' starts with (\tilde{B}, p) if $\tilde{B}_0 = \tilde{B}$ and $p_0 = p$;
- r' remains in \tilde{B} if $\tilde{B}_i = \tilde{B}$ for all $0 \leq i \leq m$ and $B_i = \tilde{B}$ for all $0 < i \leq m$;
- r' terminates with a non-constant probabilistic edge if p_m is a non-constant probabilistic edge and p_i is a constant probabilistic edge for all $0 \leq i < m$.

We say that $r' \in \text{Suffixes}_*^{\hat{\pi}}(\text{reg}(l, 0))$ is a (\tilde{B}, p) -critical path if r' starts with (\tilde{B}, p) , remains in \tilde{B} , and terminates with a non-constant probabilistic edge. Let $\text{Critical}_{(\tilde{B}, p)}^{\hat{\pi}}$ be the set of (\tilde{B}, p) -critical paths. We can now let:

$$\Omega_{\rho, (\tilde{B}, p)} = \text{Critical}_{(\tilde{B}, p)}^{\hat{\pi}} \cap \bigcup_{r' \in \Lambda^{\hat{\pi}}(\rho)} \text{SuffFrom}_*^{\hat{\pi}}(r').$$

Then we let $\Psi_{\rho, (\tilde{B}, p)}$ be the set of assignments featured in the final transition of finite paths in $\Omega_{\rho, (\tilde{B}, p)}$: formally $\Psi_{\rho, (\tilde{B}, p)} = \{\text{FA}(r') \mid r' \in \Omega_{\rho, (\tilde{B}, p)}\}$. Finally we let $\tilde{v}_{\rho, (\tilde{B}, p)}^* = \min\{h_{\tilde{B}}(\alpha) \mid \alpha \in \Psi_{\rho, (\tilde{B}, p)}\}$.

Now consider the subcase in which \tilde{B} is an open interval, and all probabilistic edges in the sequence $p_k \cdots p_n$ are constant, but p is non-constant. As in the analogous case of the proof of Lemma 4.2, the choice of clock valuation by π corresponds to a weighted average from a set of clock valuations. In the case of this proof, the set of clock valuations used is that obtained from assignments used by $\hat{\pi}$ after finite paths corresponding to ρ extended with (\tilde{B}, p) , using a similar “transformation from assignment to clock valuation” approach as used in the previous paragraph. We now give the formal details. Recall that we have assumed w.l.o.g. that, given $r' \in \Lambda^{\hat{\pi}}(\rho)$, there exists at most one $\alpha \in \text{Dist}(S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}})$ such that $\hat{\pi}(r')((\tilde{B}, p), \alpha) > 0$, and that we denote such an assignment α by $\alpha_{(\tilde{B}, p)}^{r'}$. Then we let:

$$\tilde{v}_{\rho, (\tilde{B}, p)}^* = \frac{\sum_{r' \in \Lambda^{\hat{\pi}}(\rho)} \text{Pr}_{*, (\text{reg}(l, 0))}^{\hat{\pi}}(r') \cdot \hat{\pi}(r')((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^{r'}) \cdot h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^{r'})}{\sum_{r' \in \Lambda^{\hat{\pi}}(\rho)} \text{Pr}_{*, (\text{reg}(l, 0))}^{\hat{\pi}}(r') \cdot \hat{\pi}(r')((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^{r'})}.$$

Finally, we consider the subcase in which \tilde{B} is an open interval, there exists $i \in \{k, \dots, n\}$ such that p_i is non-constant, and for all other $j \in \{k, \dots, n\}$ we have that p_j is constant, and also p is constant. As in the proof of Lemma 4.2, for this subcase, the clock valuation can be arbitrary; for simplicity we retain the same clock valuation as was used in the final state of r , i.e., if $\text{last}(r)$ is equal to (l', v') , we let $\tilde{v}_{\rho, (\tilde{B}, p)}^* = v'$.

Repeating this overall process for all $(\tilde{B}, p) \in \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(l, B)$ suffices to define the distribution $\pi(r)$. Following similar reasoning to the analogous part of Lemma 4.2, we can show that $\pi(r)$ is indeed a distribution, i.e., $\sum_{(\tilde{B}, p) \in \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(\text{reg}(\text{last}(r)))} \pi(r)(\tilde{v}_{\rho, (\tilde{B}, p)}^*, p) = 1$ (recall that $\text{reg}(\text{last}(r)) = (l_n, B_n)$). From the construction of π , and from the fact that

$\sum_{(\tilde{B}, p) \in \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(l_n, B_n)} \hat{\pi}(r')((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^{r'}) = 1$, we have the following:

$$\begin{aligned}
& \sum_{(\tilde{B}, p) \in \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(l_n, B_n)} \pi(r)(\tilde{v}_{\rho, (\tilde{B}, p)}^*, p) \\
&= \sum_{(\tilde{B}, p) \in \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(l_n, B_n)} \left(\frac{\sum_{r' \in \Lambda^{\hat{\pi}}(\rho)} \text{Pr}_{*, \text{reg}(l, 0)}^{\hat{\pi}}(r') \cdot \hat{\pi}(r')((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^{r'})}{\text{Pr}_{*, \text{reg}(l, 0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho))} \right) \\
&= \frac{\sum_{r' \in \Lambda^{\hat{\pi}}(\rho)} \text{Pr}_{*, \text{reg}(l, 0)}^{\hat{\pi}}(r') \cdot \sum_{(\tilde{B}, p) \in \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(l_n, B_n)} \hat{\pi}(r')((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^{r'})}{\text{Pr}_{*, \text{reg}(l, 0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho))} \\
&= \frac{\sum_{r' \in \Lambda^{\hat{\pi}}(\rho)} \text{Pr}_{*, \text{reg}(l, 0)}^{\hat{\pi}}(r')}{\text{Pr}_{*, \text{reg}(l, 0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho))} \\
&= 1.
\end{aligned}$$

The next step of the proof is to show that $\text{Pr}_{(l, 0)}^{\pi}(\diamond T_F) = \text{Pr}_{\text{reg}(l, 0)}^{\hat{\pi}}(\diamond \mathfrak{T}_F)$. As in previous proofs, it is sufficient to show that $\text{Pr}_{(l, 0)}^{\pi}(\Lambda^{\pi}(\rho)) = \text{Pr}_{\text{reg}(l, 0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho))$ for all \mathbb{B} -paths ρ . Proceeding by induction on the length of \mathbb{B} -paths, consider the \mathbb{B} -path $\rho(\tilde{B}, p)(l', B')$, and assume that $\text{Pr}_{(l, 0)}^{\pi}(\Lambda^{\pi}(\rho)) = \text{Pr}_{\text{reg}(l, 0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho))$. Our aim is to show that

$$\text{Pr}_{(l, 0)}^{\pi}(\Lambda^{\pi}(\rho(\tilde{B}, p)(l', B')))) = \text{Pr}_{\text{reg}(l, 0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho(\tilde{B}, p)(l', B')))).$$

We consider the case in which \tilde{B} is open, p is non-constant and $B' = \tilde{B}$, i.e., the target region (l', B') is obtained using the outcome (\emptyset, l') (which does not reset the clock, and hence the interval B' of the target region is equal to the interval \tilde{B} obtained just before the probabilistic edge p is taken). Other cases are dealt with similarly; in particular, the cases for constant probabilistic edges are more straightforward. Our approach is slightly different to that used in the proof of Lemma 4.5 because, in the scheduler of $\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket$ constructed in the proof of Lemma 4.5, the choice made after all finite paths corresponding to a particular \mathbb{B} -path was the same; instead, this property does not necessarily hold for $\hat{\pi}$. Let $r \in \Lambda^{\hat{\pi}}(\rho)$. Now consider the finite paths in $\Lambda^{\hat{\pi}}(\rho(\tilde{B}, p)(l', B'))$ that have r as a prefix: these finite paths are $r((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r)(\text{le})(l', B')$ and $r((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r)(\text{re})(l', B')$. From the construction of $\hat{\pi}$ and by definition (where \wp denotes the distribution template of probabilistic edge p),

we have the following:

$$\begin{aligned}
& \sum_{\text{ep} \in \{\text{le}, \text{re}\}} \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r)(\text{ep})(l', B')) \\
&= \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \\
&\quad \times \sum_{\text{ep} \in \{\text{le}, \text{re}\}} \alpha_{(\tilde{B}, p)}^r(\tilde{B}, p, \text{ep}) \cdot \Delta_{\mathfrak{M}[P]}((\tilde{B}, p, \text{ep}), (\tau, \alpha^{\tilde{B}, p, \text{ep}}))(l', B') \\
&= \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \cdot \sum_{\text{ep} \in \{\text{le}, \text{re}\}} \alpha_{(\tilde{B}, p)}^r(\tilde{B}, p, \text{ep}) \cdot \wp[\text{ep}(\tilde{B})](\emptyset, l') \\
&= \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \cdot \sum_{\text{ep} \in \{\text{le}, \text{re}\}} \alpha_{(\tilde{B}, p)}^r(\tilde{B}, p, \text{ep}) \cdot (c + d \cdot \text{ep}(\tilde{B})) \\
&= c \cdot \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \cdot \sum_{\text{ep} \in \{\text{le}, \text{re}\}} \alpha_{(\tilde{B}, p)}^r(\tilde{B}, p, \text{ep}) \\
&\quad + d \cdot \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \cdot \sum_{\text{ep} \in \{\text{le}, \text{re}\}} \alpha_{(\tilde{B}, p)}^r(\tilde{B}, p, \text{ep}) \cdot \text{ep}(\tilde{B}) .
\end{aligned}$$

Note that $\sum_{\text{ep} \in \{\text{le}, \text{re}\}} \alpha_{(\tilde{B}, p)}^r(\tilde{B}, p, \text{ep}) = 1$ (because $\alpha_{(\tilde{B}, p)}^r$ is a distribution over the set $\{(\tilde{B}, p, \text{le}), (\tilde{B}, p, \text{re})\}$). Furthermore, recalling the definition of $h_{\tilde{B}}$, we observe that:

$$\begin{aligned}
\alpha_{(\tilde{B}, p)}^r(\tilde{B}, p, \text{le}) &= \frac{\text{re}(\tilde{B}) - h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^r)}{\text{re}(\tilde{B}) - \text{le}(\tilde{B})} \\
\alpha_{(\tilde{B}, p)}^r(\tilde{B}, p, \text{re}) &= \frac{h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^r) - \text{le}(\tilde{B})}{\text{re}(\tilde{B}) - \text{le}(\tilde{B})} .
\end{aligned}$$

Hence we have:

$$\begin{aligned}
& \sum_{\text{ep} \in \{\text{le}, \text{re}\}} \alpha_{(\tilde{B}, p)}^r(\tilde{B}, p, \text{ep}) \cdot \text{ep}(\tilde{B}) \\
&= \frac{\text{re}(\tilde{B}) - h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^r)}{\text{re}(\tilde{B}) - \text{le}(\tilde{B})} \cdot \text{le}(\tilde{B}) + \frac{h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^r) - \text{le}(\tilde{B})}{\text{re}(\tilde{B}) - \text{le}(\tilde{B})} \cdot \text{re}(\tilde{B}) \\
&= \frac{\text{re}(\tilde{B}) \cdot \text{le}(\tilde{B}) - h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^r) \cdot \text{le}(\tilde{B}) + h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^r) \cdot \text{re}(\tilde{B}) - \text{le}(\tilde{B}) \cdot \text{re}(\tilde{B})}{\text{re}(\tilde{B}) - \text{le}(\tilde{B})} \\
&= \frac{h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^r) \cdot \text{re}(\tilde{B}) - h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^r) \cdot \text{le}(\tilde{B})}{\text{re}(\tilde{B}) - \text{le}(\tilde{B})} \\
&= h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^r) .
\end{aligned}$$

Therefore we have:

$$\begin{aligned}
& c \cdot \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \cdot \sum_{\text{ep} \in \{\text{le}, \text{re}\}} \alpha_{(\tilde{B}, p)}^r(\tilde{B}, p, \text{ep}) \\
& + d \cdot \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \cdot \sum_{\text{ep} \in \{\text{le}, \text{re}\}} \alpha_{(\tilde{B}, p)}^r(\tilde{B}, p, \text{ep}) \cdot \text{ep}(\tilde{B}) \\
= & c \cdot \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \\
& + d \cdot \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \cdot h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^r) .
\end{aligned}$$

We now use this equality to establish that $\Pr_{(l,0)}^{\pi}(\Lambda^{\pi}(\rho)) = \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho))$. First observe the following:

$$\begin{aligned}
& \Pr_{(l,0)}^{\hat{\pi}}(\Lambda^{\pi}(\rho(\tilde{B}, p))(l', B')) \\
= & \sum_{r \in \Lambda^{\hat{\pi}}(\rho(\tilde{B}, p))(l', B')} \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \\
= & \sum_{r \in \Lambda^{\hat{\pi}}(\rho)} \sum_{\text{ep} \in \{\text{le}, \text{re}\}} \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r)(\text{ep})(l', B') \\
= & \sum_{r \in \Lambda^{\hat{\pi}}(\rho)} c \cdot \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \\
& + \sum_{r \in \Lambda^{\hat{\pi}}(\rho)} d \cdot \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \cdot h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^r)
\end{aligned}$$

Recalling the definition of $\tilde{v}_{\rho, (\tilde{B}, p)}^*$ (for the case of open \tilde{B} and non-constant p), we have that $\sum_{r \in \Lambda^{\hat{\pi}}(\rho)} \Pr_{*, (\text{reg}(l,0))}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \cdot h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^r) = \tilde{v}_{\rho, (\tilde{B}, p)}^* \cdot \sum_{r \in \Lambda^{\hat{\pi}}(\rho)} \Pr_{*, (\text{reg}(l,0))}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r)$. Hence we then obtain:

$$\begin{aligned}
& \sum_{r \in \Lambda^{\hat{\pi}}(\rho)} c \cdot \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \\
& + \sum_{r \in \Lambda^{\hat{\pi}}(\rho)} d \cdot \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \cdot h_{\tilde{B}}(\alpha_{(\tilde{B}, p)}^r) \\
= & c \cdot \sum_{r \in \Lambda^{\hat{\pi}}(\rho)} \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \\
& + d \cdot \tilde{v}_{\rho, (\tilde{B}, p)}^* \cdot \sum_{r \in \Lambda^{\hat{\pi}}(\rho)} \Pr_{*, (\text{reg}(l,0))}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) .
\end{aligned}$$

Let \check{r} denote the unique path in $\Lambda^{\pi}(\rho)$. By the definition of $\pi(\check{r})$, and from the fact that $\Pr_{*, (l,0)}^{\pi}(\Lambda^{\pi}(\rho)) = \Pr_{*, \text{reg}(l,0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho))$ by induction, we have:

$$\sum_{r \in \Lambda^{\hat{\pi}}(\rho)} \Pr_{*, \text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) = \Pr_{*, (l,0)}^{\pi}(\Lambda^{\pi}(\rho)) \cdot \pi(\check{r})(\tilde{v}_{\rho, (\tilde{B}, p)}^*, p) .$$

We then obtain:

$$\begin{aligned}
& c \cdot \sum_{r \in \Lambda^{\hat{\pi}}(\rho)} \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \\
& + d \cdot \tilde{v}_{\rho, (\tilde{B}, p)}^* \cdot \sum_{r \in \Lambda^{\hat{\pi}}(\rho)} \Pr_{*, (\text{reg}(l,0))}^{\hat{\pi}}(r) \cdot \hat{\pi}(r)((\tilde{B}, p), \alpha_{(\tilde{B}, p)}^r) \\
& = c \cdot \Pr_{*, (l,0)}^{\pi}(\Lambda^{\pi}(\rho)) \cdot \pi(\check{r})(\tilde{v}_{\rho, (\tilde{B}, p)}^*, p) \\
& + d \cdot \tilde{v}_{\rho, (\tilde{B}, p)}^* \cdot \Pr_{*, (l,0)}^{\pi}(\Lambda^{\pi}(\rho)) \cdot \pi(\check{r})(\tilde{v}_{\rho, (\tilde{B}, p)}^*, p) \\
& = \Pr_{*, (l,0)}^{\pi}(\Lambda^{\pi}(\rho)) \cdot \pi(\check{r})(\tilde{v}_{\rho, (\tilde{B}, p)}^*, p) \cdot (c + d \cdot \tilde{v}_{\rho, (\tilde{B}, p)}^*) \\
& = \Pr_{*, (l,0)}^{\pi}(\Lambda^{\pi}(\rho)) \cdot \pi(\check{r})(\tilde{v}_{\rho, (\tilde{B}, p)}^*, p) \cdot \wp[\tilde{v}_{\rho, (\tilde{B}, p)}^*](\emptyset, l') \\
& = \Pr_{*, (l,0)}^{\pi}(\Lambda^{\pi}(\rho)) \cdot \pi(\check{r})(\tilde{v}_{\rho, (\tilde{B}, p)}^*, p) \cdot \Delta_{\mathcal{P}}(\text{last}(\check{r}), (\tilde{v}_{\rho, (\tilde{B}, p)}^*, p))(l', B') \\
& = \Pr_{(l,0)}^{\pi}(\Lambda^{\pi}(\rho(\tilde{B}, p)(l', B'))).
\end{aligned}$$

This establishes that $\Pr_{(l,0)}^{\pi}(\Lambda^{\pi}(\rho(\tilde{B}, p)(l', B')))) = \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(\Lambda^{\hat{\pi}}(\rho(\tilde{B}, p)(l', B')))$. We note that the cases for other outcomes are similar, and for constant probabilistic edges and/or closed intervals are more straightforward. \square

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 4.8. *Quantitative and qualitative problems for 1c-cdPTA can be solved in polynomial time.*

The theorem follows from Lemma 4.2, Lemma 4.5, Lemma 4.7, Proposition 2.2, the fact that the IMDP defined in this section can be constructed in polynomial time, 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 [31, 12, 9, 34]. We add that the quantitative and qualitative problems for 1c-cdPTAs are PTIME-hard, following from the PTIME-hardness of the corresponding problems for MDPs [30, 10].

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 \tilde{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 \tilde{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 [33] 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 [20, 36]. 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, L. Doyen, and T. A. Henzinger. Qualitative analysis of partially-observable Markov decision processes. In P. Hlinený and A. Kucera, editors, *Proc. MFCS 2010*, volume 6281 of *LNCs*, pages 258–269. Springer, 2010.
- [11] 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.
- [12] 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.
- [13] E. M. Clarke, O. Grumberg, and D. A. Peled. *Model checking*. MIT Press, 2001.
- [14] L. Feng, C. Wilsche, 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.

- [15] 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.
- [16] R. Givan, S. M. Leach, and T. L. Dean. Bounded-parameter Markov decision processes. *Artificial Intelligence*, 122(1-2):71–109, 2000.
- [17] 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.
- [18] S. Haddad and B. Monmege. Interval iteration algorithm for MDPs and IMDPs. *Theoretical Computer Science*, 735:111–131, 2018.
- [19] E. M. Hahn. *Model checking stochastic hybrid systems*. PhD thesis, Universität des Saarlandes, 2013.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] B. Jonsson and K. G. Larsen. Specification and refinement of probabilistic processes. In *Proc. LICS 1991*, pages 266–277. IEEE Computer Society, 1991.
- [24] 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.
- [25] I. O. Kozine and L. V. Utkin. Interval-valued finite Markov chains. *Reliable Computing*, 8(2):97–113, 2002.
- [26] 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.
- [27] 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.
- [28] A. Nilim and L. El Ghaoui. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.
- [29] G. Norman, D. Parker, and J. Sproston. Model checking for probabilistic timed automata. *Formal Methods in System Design*, 43(2):164–190, 2013.
- [30] C. H. Papadimitriou and J. N. Tsitsiklis. The complexity of Markov decision processes. *Mathematics of Operations Research*, 12(3):441–450, 1987.
- [31] 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.
- [32] 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.
- [33] 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.
- [34] 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.
- [35] J. Sproston. Probabilistic timed automata with one clock and initialised clock-dependent probabilities. In A. Gotsman and A. Sokolova, editors, *Proc. FORTE 2020*, volume 12136 of *LNCS*, pages 150–168. Springer, 2020.
- [36] T. Winkler, S. Junges, G. A. Pérez, and J. Katoen. On the complexity of reachability in parametric Markov decision processes. In W. Fokink and R. van Glabbeek, editors, *Proc. CONCUR 2019*, volume 140 of *LIPIcs*, pages 14:1–14:17. Leibniz-Zentrum für Informatik, 2019.
- [37] 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.