

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

## Qualitative Reachability for Open Interval Markov Chains

### **This is the author's manuscript**

*Original Citation:*

*Availability:*

This version is available <http://hdl.handle.net/2318/1691650> since 2019-02-11T11:40:31Z

*Publisher:*

Springer

*Published version:*

DOI:10.1007/978-3-030-00250-3\_11

*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)

# Qualitative Reachability for Open Interval Markov Chains

Jeremy Sproston

Dipartimento di Informatica, University of Turin, Italy

**Abstract.** Interval Markov chains extend classical Markov chains with the possibility to describe transition probabilities using intervals, rather than exact values. While the standard formulation of interval Markov chains features closed intervals, previous work has considered also open interval Markov chains, in which the intervals can also be open or half-open. In this paper we focus on qualitative reachability problems for open interval Markov chains, which consider whether the optimal (maximum or minimum) probability with which a certain set of states can be reached is equal to 0 or 1. We present polynomial-time algorithms for these problems for both of the standard semantics of interval Markov chains. Our methods do not rely on the closure of open intervals, in contrast to previous approaches for open interval Markov chains, and can characterise situations in which probability 0 or 1 can be attained not exactly but arbitrarily closely.

## 1 Introduction

The development of modern computer systems can benefit substantially from a verification phase, in which a formal model of the system is exhaustively verified in order to identify undesirable errors or inefficiencies. In this paper we consider the verification of probabilistic systems, in which state-to-state transitions are accompanied by probabilities that specify the relative likelihood with which the transitions occur, using model-checking techniques; see [2,11,1] for general overviews of this field. One drawback of classical formalisms for probabilistic systems is that they typically require the specification of exact probability values for transitions: in practice, it is likely that such precise information concerning the probability of system behaviour is not available. A solution to this problem is to associate intervals of probabilities with transitions, rather than exact probability values, leading to *interval Markov chains* (IMCs) or *interval Markov decision processes*. IMCs have been studied in [14,15], and considered in the *qualitative* and *quantitative* model-checking context in [18,5,6]. Qualitative model checking concerns whether a property is satisfied by the system model with probability (equal to or strictly greater than) 0 or (equal to or strictly less than) 1, whereas quantitative model checking considers whether a property is satisfied with probability above or below some threshold in the interval  $[0, 1]$ , and generally involves the computation of the probability of property satisfaction, which is then compared to the aforementioned threshold.




Fig. 1. An open IMC  $\mathcal{O}_1$ .




Fig. 2. An open IMC  $\mathcal{O}_2$ .

In [18,5,6], the intervals associated with transitions are *closed*. This limitation was addressed in [4], which considered the possibility of utilising *open* (and half-open) intervals, in addition to closed intervals. Example of such open IMCs are shown in Figure 1 and Figure 2. In [4], it was shown that the probability of the satisfaction of a property in an open IMC can be approximated arbitrarily closely by a standard, closed IMC obtained by changing all (half-)open intervals featured in the model to closed intervals with the same endpoints. However, although the issue of the determining the existence of exact solutions is mentioned in [4], closing the intervals can involve the loss of information concerning exact solutions. Take, for example, the open IMC in Figure 1: changing the intervals from  $(0, 1)$  to  $[0, 1]$  on both of the transitions means that the minimum probability of reaching the state  $s_1$  after starting in state  $s_0$  becomes 0, whereas the probability of reaching  $s_1$  from  $s_0$  is strictly greater than 0 for all ways of assigning probabilities to the transitions in the original IMC.

In this paper we propose verification methods for qualitative reachability properties of open IMCs. We consider both of the standard semantics for IMCs. The uncertain Markov chain (UMC) semantics associated with an IMC comprises an infinite number of standard Markov chains, each corresponding to a certain choice of probability for each transition. In contrast, the interval Markov decision process (IMDP) semantics associates a single Markov decision process (MDP) with the IMC, where from each state there is available an uncountable number of distributions, each corresponding to one assignment of probabilities belonging to the intervals of the transitions leaving that state. The key difference between the two semantics can be summarised by considering the behaviour from a particular state of the IMC: in the UMC semantics, the same probability distribution over outgoing transitions must always be used from the state, whereas in the IMDP semantics the outgoing probability distribution may change for each visit to the state. We show that we can obtain exact (not approximate) solutions for both semantics in polynomial time in the size of the open IMC.

For the UMC semantics, and for three of the four classes of qualitative reachability problem in the IMDP semantics, the algorithms presented are inspired by methods for finite MDPs. In the case of the IMDP semantics, these algorithms rely on the fact that retaining the memory of previous choices along the behaviour of an IMC is not necessary. A direct method for the construction of a finite MDP that represents an IMC and which can be used for the verification of qualitative properties is the following: the set of states of the finite MDP equals that of the IMC and, for each state  $s$  and each set  $X$  of states, there exists a single distribution from  $s$  in the finite MDP that assigns positive probability to

each state in  $X$  if and only if there exists at least one probability assignment for transitions in the IMC that assigns positive probability to each transition from  $s$  with target state in  $X$ . Intuitively, a distribution associated with  $s$  and  $X$  in the finite MDP can be regarded as the representative distribution of all probability assignments of the IMC that assign positive probability to the transitions from  $s$  to states in  $X$ . However, such a finite MDP construction does not yield polynomial-time algorithms in the size of the open IMC, because the presence of transitions having zero as their left endpoint can result in an exponential number of distributions in the number of IMC transitions. In our methods, apart from considering issues concerning the difference between closed and open intervals and the subsequent implications for qualitative reachability problems, we avoid such an exponential blow up. In particular, we show how the predecessor operations used by some qualitative reachability algorithms for MDPs can be applied directly on the open IMC.

The fourth class of reachability problem in the IMDP semantics concerns determining whether the probability of reaching a certain set of states from the current state is equal to 1 for all schedulers, where a scheduler chooses an outgoing probability distribution from a state on the basis of the choices made so far. For this class of problem, retaining memory of previous choices can be important for showing that the problem is *not* satisfied, i.e., that there exists a scheduler such that the reachability probability is strictly less than 1. As an example, we can take the open IMC in Figure 1. Consider the memoryful scheduler that assigns probability  $\frac{1}{2^i}$  to the  $i$ -th attempt to take a transition from  $s_0$  to  $s_1$ , meaning that the overall probability of reaching  $s_1$  when starting in  $s_0$  under this scheduler is  $\frac{1}{2} + \frac{1}{2}(\frac{1}{4} + \frac{3}{4}(\frac{1}{8} + \dots)) < 1$ . Instead a memoryless scheduler will reach  $s_1$  with probability 1: for any  $\lambda \in (0, 1)$  representing the (constant) probability of taking the transition from  $s_0$  to  $s_1$ , the overall probability of reaching  $s_1$  is  $\lim_{k \rightarrow \infty} 1 - (1 - \lambda)^k = 1$ . Hence our results for this class of reachability problem take the inadequacy of memoryless schedulers into account; indeed, while the algorithms presented for all other classes of problems (and all problems for the UMC semantics) proceed in a manner similar to that introduced in the literature for finite MDPs, for this class we present an *ad hoc* algorithm, based on an adaptation of the classical notion of end components [9].

After introducing open IMCs in Section 2, the algorithms for the UMC semantics and the IMDP semantics are presented in Section 3 and Section 4, respectively. The proofs of the results can be found in the appendix.

*Related work.* Model checking of qualitative properties of Markov chains (see, for example, [20,7]) relies on the fact that transition probability values are fixed throughout the behaviour of the system, and does not require that exact probability values are taken into account during analysis. The majority of work on model checking for IMCs considers the more general quantitative problems: [18,5] present algorithms utilising a finite MDP construction based on encoding within distributions available from a state the extremal probabilities allowed from that state (known as the state's basic feasible solutions). Such a construction results in an exponential blow up, which is also not avoided in [5] for qualitative proper-

ties (when transitions can have 0 as their left endpoint). [6,17] improve on these results to present polynomial-time algorithms for reachability problems based on linear or convex programming. The paper [12] includes polynomial-time methods for computing (maximal) end components, and for computing a single step of value iteration, for interval MDPs. We note that IMCs are a special case of constraint Markov chains [3], and that the UMC semantics of IMCs corresponds to a special case of parametric Markov chains [8,16]. As far as we are aware, only [4] considers open IMCs.

## 2 Open Interval Markov Chains

**Preliminaries.** A (probability) distribution over a finite 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  $Q$ . We use  $\text{support}(\mu) = \{q \in Q \mid \mu(q) > 0\}$  to denote the *support set* of  $\mu$ , i.e., the set of elements assigned positive probability by  $\mu$ , and use  $\{q \mapsto 1\}$  to denote the distribution that assigns probability 1 to the single element  $q$ . 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  $q' \in Q$ . 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  $I$ . The set of closed (respectively, left-open, right-closed; left-closed, right-open; open) intervals in  $\mathcal{I}$  is denoted by  $\mathcal{I}^{[\cdot, \cdot]}$  (respectively,  $\mathcal{I}^{(\cdot, \cdot]}$ ;  $\mathcal{I}^{[\cdot, \cdot)}$ ;  $\mathcal{I}^{(\cdot, \cdot)}$ ). Hence we have  $\mathcal{I} = \mathcal{I}^{[\cdot, \cdot]} \cup \mathcal{I}^{(\cdot, \cdot]}$   $\cup \mathcal{I}^{[\cdot, \cdot)}$   $\cup \mathcal{I}^{(\cdot, \cdot)}$ . Furthermore, we let  $\mathcal{I}^{(+, \cdot)}$  (respectively,  $\mathcal{I}^{(0, \cdot]}$ ;  $\mathcal{I}^{[0, \cdot)}$ ) be the set of intervals in  $\mathcal{I}$  such that the left endpoint is positive (respectively, left-closed intervals with the left endpoint equal to 0; left-open intervals with the left endpoint equal to 0). Finally, let  $\mathcal{I}^{(0, \cdot)} = \mathcal{I}^{[0, \cdot)}$   $\cup \mathcal{I}^{(0, \cdot]}$  be the set of intervals in  $\mathcal{I}$  with left endpoint equal to zero.

A *discrete-time Markov chain* (DTMC)  $\mathcal{D}$  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 matrix*, such that, for each state  $s \in S$ , we have  $\sum_{s' \in S} \mathbf{P}(s, s') = 1$ . Note that  $\mathbf{P}(s, \cdot)$  is a distribution, for each state  $s \in S$ . A *path* of DTMC  $\mathcal{D}$  is a sequence  $s_0 s_1 \dots$  such that  $\mathbf{P}(s_i, s_{i+1}) > 0$  for all  $i \geq 0$ . Given a path  $\rho = s_0 s_1 \dots$  and  $i \geq 0$ , we let  $\rho(i) = s_i$  be the  $(i + 1)$ -th state along  $\rho$ . The set of paths of  $\mathcal{D}$  starting in state  $s \in S$  is denoted by  $\text{Paths}^{\mathcal{D}}(s)$ . In the standard manner (see, for example, [2,11]), given a state  $s \in S$ , we can define a probability measure  $\text{Pr}_s^{\mathcal{D}}$  over  $\text{Paths}^{\mathcal{D}}(s)$ .

A *Markov decision process* (MDP)  $\mathcal{M}$  is a pair  $(S, \Delta)$  where  $S$  is a finite set of *states* and  $\Delta : S \rightarrow 2^{\text{Dist}(S)}$  is a *transition function* such that  $\Delta(s) \neq \emptyset$  for all  $s \in S$ . We say that an MDP is *finite* if  $\Delta(s)$  is finite for all  $s \in S$ .

A (n infinite) path of an MDP  $\mathcal{M}$  is a sequence  $s_0 \mu_0 s_1 \mu_1 \dots$  such that  $\mu_i \in \Delta(s_i)$  and  $\mu_i(s_{i+1}) > 0$  for all  $i \geq 0$ . Given a path  $\rho = s_0 \mu_0 s_1 \mu_1 \dots$  and  $i \geq 0$ , we let  $\rho(i) = s_i$  be the  $(i + 1)$ -th state along  $\rho$ . A finite path is a sequence  $r = s_0 \mu_0 s_1 \mu_1 \dots \mu_{n-1} s_n$  such that  $\mu_i \in \Delta(s_i)$  and  $\mu_i(s_{i+1}) > 0$  for each  $0 \leq i < n$ . Let  $\text{last}(r) = s_n$  denote the final state of  $r$ . Let  $\text{Paths}_*^{\mathcal{M}}$  be the

set of finite paths of the MDP  $\mathcal{M}$ . Let  $Paths^{\mathcal{M}}(s)$  and  $Paths_*^{\mathcal{M}}(s)$  be the sets of infinite paths and finite paths, respectively, of  $\mathcal{M}$  starting in state  $s \in S$ .

A *scheduler* is a mapping  $\sigma : Paths_*^{\mathcal{M}} \rightarrow \text{Dist}(\bigcup_{s \in S} \Delta(s))$  such that  $\sigma(r) \in \text{Dist}(\Delta(\text{last}(r)))$  for each  $r \in Paths_*^{\mathcal{M}}$ . Let  $\Sigma^{\mathcal{M}}$  be the set of schedulers of the MDP  $\mathcal{M}$ . Given a state  $s \in S$  and a scheduler  $\sigma$ , we can define a countably infinite-state DTMC  $\mathcal{D}_s^\sigma$  that corresponds to the behaviour of the scheduler  $\sigma$  from state  $s$ , which in turn can be used to define a probability measure  $\text{Pr}_s^\sigma$  over  $Paths^{\mathcal{M}}(s)$  in the standard manner (see [2,11]). A scheduler  $\sigma \in \Sigma^{\mathcal{M}}$  is *memoryless* if, for finite paths  $r, r' \in Paths_*^{\mathcal{M}}$  such that  $\text{last}(r) = \text{last}(r')$ , we have  $\sigma(r) = \sigma(r')$ . Let  $\Sigma_m^{\mathcal{M}}$  be the set of memoryless schedulers of  $\mathcal{M}$ . Note that, for a memoryless scheduler  $\sigma \in \Sigma_m^{\mathcal{M}}$ , we can construct a finite DTMC  $\tilde{\mathcal{D}}^\sigma = (S, \tilde{\mathbf{P}})$  with  $\tilde{\mathbf{P}}(s, s') = \sum_{\mu \in \Delta(s)} \sigma(s)(\mu) \cdot \mu(s')$ : we call this DTMC the *folded DTMC* of  $\sigma$ . It can be shown that  $\text{Pr}_s^\sigma$  and  $\text{Pr}_s^{\tilde{\mathcal{D}}^\sigma}$  assign the same probabilities to measurable sets of paths, because the state  $s$  of the DTMC  $\mathcal{D}_s^\sigma$  is probabilistic bisimulation equivalent to the state  $s$  of the folded DTMC  $\tilde{\mathcal{D}}^\sigma$  (for a definition of probabilistic bisimulation and more information on this point, see [2, Section 10.4.2]).

**Interval Markov Chains: syntax.** An (open) *interval Markov chain* (IMC)  $\mathcal{O}$  is a pair  $(S, \delta)$ , where  $S$  is a finite set of *states*, and  $\delta : S \times S \rightarrow \mathcal{I}$  is a *interval-based transition function*.

In the following, we refer to *edges* as those state pairs for which the transition function does not assign the probability 0 point interval  $[0, 0]$ . Formally, let the set of edges  $E$  of  $\mathcal{O}$  be defined as  $\{(s, s') \in S \times S \mid \delta(s, s') \neq [0, 0]\}$ . We use edges to define the notion of path for IMCs: a path of an IMC  $\mathcal{O}$  is a sequence  $s_0 s_1 \dots$  such that  $(s_i, s_{i+1}) \in E$  for all  $i \geq 0$ . Given a path  $\rho = s_0 s_1 \dots$  and  $i \geq 0$ , we let  $\rho(i) = s_i$  be the  $(i+1)$ -th state along  $\rho$ . We use  $Paths^{\mathcal{O}}$  to denote the set of paths of  $\mathcal{O}$ ,  $Paths_*^{\mathcal{O}}$  to denote the set of finite paths of  $\mathcal{O}$ , and  $Paths^{\mathcal{O}}(s)$  and  $Paths_*^{\mathcal{O}}(s)$  to denote the sets of paths and finite paths starting in state  $s \in S$ .

Given a state  $s \in S$ , we say that a distribution  $\mathbf{a} \in \text{Dist}(S)$  is an *assignment for  $s$*  if  $\mathbf{a}(s') \in \delta(s, s')$  for each state  $s' \in S$ . We say that the IMC  $\mathcal{O}$  is *well formed* if there exists at least one assignment for each state. Note that an assignment for state  $s \in S$  exists if and only if the following conditions hold: (1a)  $\sum_{s' \in S} \text{left}(\delta(s, s')) \leq 1$ , (1b)  $\sum_{s' \in S} \text{left}(\delta(s, s')) = 1$  implies that  $\delta(s, s')$  is left-closed for all  $s' \in S$ , (2a)  $\sum_{s' \in S} \text{right}(\delta(s, s')) \geq 1$ , and (2b)  $\sum_{s' \in S} \text{right}(\delta(s, s')) = 1$  implies that  $\delta(s, s')$  is right-closed for all  $s' \in S$ . We henceforth consider IMCs that are well formed. We define the *size of an IMC*  $\mathcal{O} = (S, \delta)$  as the size of the representation of  $\delta$ , which is the sum over all states  $s, s' \in S$  of the binary representation of the endpoints of  $\delta(s, s')$ , where rational numbers are encoded as the quotient of integers written in binary.

**Interval Markov Chains: semantics.** IMCs are typically presented with regard to two semantics, which we consider in turn. Given an IMC  $\mathcal{O} = (S, \delta)$ , the *uncertain Markov chain* (UMC) semantics of  $\mathcal{O}$ , denoted by  $[\mathcal{O}]_{\text{U}}$ , is the smallest

set of DTMCs such that  $(S, \mathbf{P}) \in [\mathcal{O}]_{\cup}$  if, for each state  $s \in S$ , the distribution  $\mathbf{P}(s, \cdot)$  is an assignment for  $s$ . The *interval Markov decision process* (IMDP) semantics of  $\mathcal{O}$ , denoted by  $[\mathcal{O}]_{\text{I}}$ , is the MDP  $(S, \Delta)$  where, for each state  $s \in S$ , we let  $\Delta(s)$  be the set of assignments for  $s$ .

**Reachability.** Let  $\mathcal{O} = (S, \delta)$  be an IMC and let  $T \subseteq S$  be a set of states. We define  $\text{Reach}(T) \subseteq \text{Paths}^{\mathcal{O}}$  to be the set of paths of  $\mathcal{O}$  that reach at least one state in  $T$ . Formally,  $\text{Reach}(T) = \{\rho \in \text{Paths}^{\mathcal{O}} \mid \exists i \in \mathbb{N}. \rho(i) \in T\}$ . In the following we assume without loss of generality that states in  $T$  are absorbing in all the IMCs that we consider, i.e.,  $\delta(s, s) = [1, 1]$  for all states  $s \in T$ .

**Edge sets.** Let  $\mathcal{O} = (S, \delta)$  be an IMC. Let  $s \in S$  be a state of  $\mathcal{O}$ , and let  $E(s) = \{(s, s') \in E \mid s' \in S\}$  be the set of edges of  $\mathcal{O}$  with source  $s$ . Let  $\star \in \{[\cdot, \cdot], (\cdot, \cdot), [\cdot, \cdot), (\cdot, \cdot), \langle +, \cdot \rangle, [0, \cdot), (0, \cdot), \langle 0, \cdot \rangle\}$ , and let  $E^{\star} = \{(s, s') \in E \mid \delta(s, s') \in \mathcal{I}^{\star}\}$ . Given  $X \subseteq S$ , and given  $s$  and  $\star$  as defined above, let  $E(s, X) = \{(s, s') \in E(s) \mid s' \in X\}$  and  $E^{\star}(s, X) = E(s, X) \cap E^{\star}$ .

**Valid edge sets.** We are interested in identifying the sets of edges from state  $s \in S$  that result from assignments. Such a set is characterised by two syntactic conditions: the first condition requires that the sum of the upper bounds of the set's edges' intervals is at least 1, whereas the second condition specifies the edges from state  $s$  that are *not* included in the set can be assigned probability 0. Formally, we say that a non-empty subset  $B \subseteq E(s)$  of edges from  $s$  is *large* if either (a)  $\sum_{e \in B} \text{right}(\delta(e)) > 1$  or (b)  $\sum_{e \in B} \text{right}(\delta(e)) = 1$  and  $B \subseteq E^{(\cdot, \cdot]}$ . The set  $B$  is *realisable* if  $E(s) \setminus B \subseteq E^{[0, \cdot)}$ . Then we say that  $B \subseteq E(s)$  is *valid* if it is large and realisable. The following lemma specifies that a valid edge set for state  $s$  characterises exactly the support sets of some assignments for  $s$ .

**Lemma 1.** *Let  $s \in S$  and  $B \subseteq E(s)$ . Then  $B$  is valid if and only if there exists an assignment  $\mathbf{a}$  for  $s$  such that  $\{(s, s') \mid s' \in \text{support}(\mathbf{a})\} = B$ .*

A consequence of Lemma 1 is that, because we consider only well-formed IMCs, there exists at least one valid subset of outgoing edges from each state.

For each state  $s \in S$ , we let  $\text{Valid}(s) = \{B \subseteq E(s) \mid B \text{ is valid}\}$ . Note that, in the worst case (when all edges in  $E(s)$  are associated with intervals  $[0, 1]$ ),  $|\text{Valid}(s)| = 2^{|E(s)|} - 1$ . Let  $\text{Valid} = \bigcup_{s \in S} \text{Valid}(s)$  be the set of valid sets of the IMC. Given a valid set  $B \in \text{Valid}$ , we let  $\text{ValidAssign}(B)$  be the set of assignments  $\mathbf{a}$  that witness Lemma 1, i.e., all assignments  $\mathbf{a}$  such that  $\{(s, s') \mid s' \in \text{support}(\mathbf{a})\} = B$ . A *witness assignment function*  $w : \text{Valid} \rightarrow \text{Dist}(S)$  assigns to each valid set  $B \in \text{Valid}$  an assignment from  $\text{ValidAssign}(B)$ .

*Example 1.* For the state  $s_1$  of the IMC  $\mathcal{O}_2$  of Figure 2, the valid edge sets are  $B_1 = \{(s_1, s_0), (s_1, s_1), (s_1, s_2)\}$  and  $B_2 = \{(s_1, s_0), (s_1, s_2)\}$ , reflecting the intuition that the edge  $(s_1, s_1)$  can be assigned (exactly) probability 0. Note that reducing the right endpoint of  $(s_1, s_0)$  to 0.7 would result in  $B_1$  being the only

valid set associated with  $s_1$ , because  $B_2$  would not be large. An example of a witness assignment function  $w$  for state  $s_1$  of  $\mathcal{O}_2$  is  $w(B_1)(s_0) = 0.7$ ,  $w(B_1)(s_1) = 0.12$  and  $w(B_1)(s_2) = 0.18$ , and  $w(B_2)(s_0) = 0.8$  and  $w(B_2)(s_2) = 0.2$ .

**Qualitative MDP abstractions.** The *qualitative MDP abstraction of  $\mathcal{O}$  with respect to witness assignment function  $w$*  is the MDP  $[\mathcal{O}]_w = (S, \Delta_w)$ , where  $\Delta_w$  is defined by  $\Delta_w(s) = \{w(B) \mid B \in \text{Valid}(s)\}$  for each state  $s \in S$ .

### 3 Qualitative Reachability: UMC semantics

Qualitative reachability problems can be classified into four categories, depending on whether the probability of reaching the target set  $T$  is 0 or 1 for some or for all ways of assigning probabilities to intervals. For the UMC semantics, we consider the computation of the following sets:

- $S_{\forall}^{0,U} = \{s \in S \mid \forall \mathcal{D} \in [\mathcal{O}]_U . \Pr_s^{\mathcal{D}}(\text{Reach}(T)) = 0\}$ ;
- $S_{\exists}^{0,U} = \{s \in S \mid \exists \mathcal{D} \in [\mathcal{O}]_U . \Pr_s^{\mathcal{D}}(\text{Reach}(T)) = 0\}$ ;
- $S_{\exists}^{1,U} = \{s \in S \mid \exists \mathcal{D} \in [\mathcal{O}]_U . \Pr_s^{\mathcal{D}}(\text{Reach}(T)) = 1\}$ ;
- $S_{\forall}^{1,U} = \{s \in S \mid \forall \mathcal{D} \in [\mathcal{O}]_U . \Pr_s^{\mathcal{D}}(\text{Reach}(T)) = 1\}$ .

The remainder of this section is dedicated to showing the following result.

**Theorem 1.** *The sets  $S_{\forall}^{0,U}$ ,  $S_{\exists}^{0,U}$ ,  $S_{\exists}^{1,U}$  and  $S_{\forall}^{1,U}$  can be computed in polynomial time in the size of the IMC.*

**Computation of  $S_{\forall}^{0,U}$ .** The case for  $S_{\forall}^{0,U}$  is straightforward. We compute the state set  $S \setminus S_{\forall}^{0,U} = \{s \in S \mid \exists \mathcal{D} \in [\mathcal{O}]_U . \Pr_s^{\mathcal{D}}(\text{Reach}(T)) > 0\}$ , which reduces to reachability on the graph of the IMC according to the following lemma.

**Lemma 2.** *Let  $s \in S$ . There exists  $\mathcal{D} \in [\mathcal{O}]_U$  such that  $\Pr_s^{\mathcal{D}}(\text{Reach}(T)) > 0$  if and only if there exists a path  $r \in \text{Paths}_*^{\mathcal{O}}(s)$  such that  $\text{last}(r) \in T$ .*

Hence the set  $S_{\forall}^{0,U}$  is equal to the complement of the set of states from which there exists a path reaching  $T$  in the graph of the IMC (that is, the graph  $(S, E)$ ). Given that the latter set of states can be computed in polynomial time, we conclude that  $S_{\forall}^{0,U}$  can be computed in polynomial time.

**Computation of  $S_{\exists}^{0,U}$ .** We show that  $S_{\exists}^{0,U}$  can be obtained by computing the set of states from which there exists a scheduler for which  $T$  is reached with probability 0 in the qualitative MDP abstraction  $[\mathcal{O}]_w = (S, \Delta_w)$  of  $\mathcal{O}$  with respect to some (arbitrary) witness assignment function  $w$ .

First we establish that the set of states of  $[\mathcal{O}]_w$  for which there exists a scheduler such that  $T$  is reached with probability 0 (respectively, probability 1) is equal to the set of states of  $\mathcal{O}$  for which there exists a DTMC in  $[\mathcal{O}]_U$  such that  $T$  is reached with probability 0 (respectively, probability 1).



**Lemma 3.** *Let  $s \in S$ ,  $\bowtie \in \{<, =, >\}$  and  $\lambda \in \{0, 1\}$ . There exists  $\mathcal{D} \in [\mathcal{O}]_{\cup}$  such that  $\text{Pr}_s^{\mathcal{D}}(\text{Reach}(T)) \bowtie \lambda$  if and only if there exists a scheduler  $\sigma \in \Sigma^{[\mathcal{O}]_w}$  such that  $\text{Pr}_s^{\sigma}(\text{Reach}(T)) \bowtie \lambda$ .*

In particular, Lemma 3 allows us to reduce the problem of computing  $S_{\exists}^{0, \cup}$  to that of computing the set  $\{s \in S \mid \exists \sigma \in \Sigma^{[\mathcal{O}]_w} . \text{Pr}_s^{\sigma}(\text{Reach}(T)) = 0\}$  on  $[\mathcal{O}]_w$ . As in the case of standard finite MDP techniques (see [11]), we proceed by computing the *complement* of this set, i.e., we compute the set  $\{s \in S \mid \forall \sigma \in \Sigma^{[\mathcal{O}]_w} . \text{Pr}_s^{\sigma}(\text{Reach}(T)) > 0\}$ . For a set  $X \subseteq S$ , let  $\text{CPre}(X) = \{s \in S \mid \exists \mu \in \Delta_w(s) . \text{support}(\mu) \subseteq X\}$  be the set of states for which there exists a distribution such that all states assigned positive probability by the distribution are in  $X$ . Furthermore, we let  $\overline{\text{CPre}}(X) = \{s \in S \mid \forall \mu \in \Delta_w(s) . \text{support}(\mu) \cap X \neq \emptyset\}$  be the dual of the CPre operator (i.e.,  $\text{CPre}(X) = S \setminus \overline{\text{CPre}}(S \setminus X)$ ), that is the set of states from which it is inevitable to make a transition to  $X$  with positive probability. The standard algorithm for computing the set of states of a finite MDP for which all schedulers are such that a set  $T$  of target states is reached with probability strictly greater than 0 operates in the following way: starting from  $X_0 = T$ , we let  $X_{i+1} = X_i \cup \overline{\text{CPre}}(X_i)$  for progressively larger values of  $i \geq 0$ , until we reach a fixpoint (that is, until we obtain  $X_{i^*+1} = X_{i^*}$  for some  $i^*$ ). However, a direct application of this algorithm to  $[\mathcal{O}]_w$  would result in an exponential-time algorithm, given that the size of the transition function  $\Delta_w$  of  $[\mathcal{O}]_w$  may be exponential in the size of  $\mathcal{O}$ . For this reason, we propose an algorithm that operates directly on the IMC  $\mathcal{O}$ , without needing the explicit construction of  $[\mathcal{O}]_w$ . We proceed by establishing that CPre can be implemented in polynomial time in the size of  $\mathcal{O}$ .

**Lemma 4.** *Let  $s \in S$  and  $X \subseteq S$ . Then  $s \in \text{CPre}(X)$  if and only if (1)  $E(s, S \setminus X) \subseteq E^{[0, \cdot]}$ , and (2)  $E(s, X)$  is large. The set  $\text{CPre}(X)$  can be computed in polynomial time in the size of the IMC  $\mathcal{O}$ .*

The intuition underlying Lemma 4 is that conditions (1) and (2) encode realizability and largeness, i.e., validity, of edge set  $E(s, X)$ . From Lemma 1, their satisfaction means that there exists a distribution in  $\Delta_w(s)$  with support set equal to the set of target states of edges in  $E(s, X)$ . We consider the largest edge set with target states in  $X$ , i.e.,  $E(s, X)$ , because taking smaller edge sets with targets in  $X$  would make the conditions (1) and (2) more difficult to satisfy.

The final part of Lemma 4 follows from the fact that conditions (1) and (2) in Lemma 4 can be checked in polynomial time in the size of  $\mathcal{O}$ . Hence our algorithm avoids the construction of the qualitative MDP abstraction  $[\mathcal{O}]_w$ , and instead consists of direct computation of the sets  $X_0 = T$  and  $X_{i+1} = X_i \cup S \setminus \overline{\text{CPre}}(S \setminus X_i)$  for increasing indices  $i$  until a fixpoint is reached. Given that a fixpoint must be reached within  $|S|$  steps, and the computation of  $\text{CPre}(X_i)$  can be done in polynomial time in the size of  $\mathcal{O}$ , we have that the set  $\{s \in S \mid \forall \sigma \in \Sigma^{[\mathcal{O}]_w} . \text{Pr}_s^{\sigma}(\text{Reach}(T)) > 0\}$  can be computed in polynomial time in the size of  $\mathcal{O}$ . The complement of this set is equal to  $S_{\exists}^{0, \cup}$ , as established by Lemma 3, and hence we can compute  $S_{\exists}^{0, \cup}$  in polynomial time in the size of  $\mathcal{O}$ .

**Computation of  $S_{\exists}^{1,U}$ .** We proceed in a manner analogous to that for the case of  $S_{\exists}^{0,U}$ . First we note that, by Lemma 3, we have that  $S_{\exists}^{1,U}$  is equal to the set of states of  $[\mathcal{O}]_w$  such that there exists a scheduler for which  $T$  is reached with probability 1. Hence, our aim is to compute the set  $\{s \in S \mid \exists \sigma \in \Sigma^{[\mathcal{O}]_w} . \Pr_s^\sigma(\text{Reach}(T)) = 1\}$  on  $[\mathcal{O}]_w$ . We recall the standard algorithm for the computation of this set on finite MDPs [9,10]. Given state sets  $X, Y \subseteq S$ , we let

$$\text{APre}(Y, X) = \{s \in S \mid \exists \mu \in \Delta_w(s) . \text{support}(\mu) \subseteq Y \wedge \text{support}(\mu) \cap X \neq \emptyset\}$$

be the set of states for which there exists a distribution such that all states assigned positive probability by the distribution are in  $Y$  and there exists a state assigned positive probability by the distribution that is in  $X$ . The standard algorithm proceeds by setting  $Y_0 = S$  and  $X_0^0 = T$ . Then the sequence  $X_0^0, X_1^0, \dots$  is computed by letting  $X_{i_0+1}^0 = X_{i_0}^0 \cup \text{APre}(Y_0, X_{i_0}^0)$  for progressively larger indices  $i_0 \geq 0$  until a fixpoint is obtained, that is, until we obtain  $X_{i_0^*+1}^0 = X_{i_0^*}^0$  for some  $i_0^*$ . Next we let  $Y_1 = X_{i_0^*}^0$ ,  $X_0^1 = T$  and compute  $X_{i_1+1}^1 = X_{i_1}^1 \cup \text{APre}(Y_1, X_{i_1}^1)$  for larger  $i_1 \geq 0$  until a fixpoint  $X_{i_1^*}^1$  is obtained. Then we let  $Y_2 = X_{i_1^*}^1$  and  $X_0^2 = T$ , and repeat the process. We terminate the algorithm when a fixpoint is reached in the sequence  $Y_0, Y_1, \dots$ <sup>1</sup> The algorithm requires at most  $|S|^2$  calls to  $\text{APre}$ . In an analogous manner to  $\text{CPre}$  in the case of  $S_{\exists}^{0,U}$ , we show that  $\text{APre}$  can be characterised by efficiently checkable conditions on  $\mathcal{O}$ .

**Lemma 5.** *Let  $s \in S$  and let  $X, Y \subseteq S$ . Then  $s \in \text{APre}(Y, X)$  if and only if (1)  $E(s, X \cap Y) \neq \emptyset$ , (2)  $E(s, S \setminus Y) \subseteq E^{[0, \cdot]}$ , and (3)  $E(s, Y)$  is large. The set  $\text{APre}(Y, X)$  can be computed in polynomial time in the size of the IMC  $\mathcal{O}$ .*

The intuition underlying Lemma 5 is similar to that of Lemma 4.

Hence we obtain an overall polynomial-time algorithm for computing  $\{s \in S \mid \exists \sigma \in \Sigma^{[\mathcal{O}]_w} . \Pr_s^\sigma(\text{Reach}(T)) = 1\}$  which, from Lemma 3, equals  $S_{\exists}^{1,U}$ .

**$S_{\forall}^{1,U}$ .** We recall the standard algorithm for determining the set of states for which all schedulers reach a target set with probability 1 on a finite MDP (see [11]): from the set of states of the MDP, we first remove states from which the target state can be reached with probability 0 (for some scheduler), then successively remove states for which it is possible to reach a previously removed state with positive probability. For each of the remaining states, there exists a scheduler that can reach the target set with probability 1.

We propose an algorithm for IMCs that is inspired by this standard algorithm for finite MDPs. Our aim is to compute the complement of  $S_{\forall}^{1,U}$ , i.e., the state set  $S \setminus S_{\forall}^{1,U} = \{s \in S \mid \exists \mathcal{D} \in [\mathcal{O}]_U . \Pr_s^{\mathcal{D}}(\text{Reach}(T)) < 1\}$ .

**Lemma 6.** *Let  $s \in S$ . There exists  $\mathcal{D} \in [\mathcal{O}]_U$  such that  $\Pr_s^{\mathcal{D}}(\text{Reach}(T)) < 1$  if and only if there exists a path  $r \in \text{Paths}_*^{\mathcal{O}}(s)$  such that  $\text{last}(r) \in S_{\exists}^{0,U}$ .*

<sup>1</sup> Readers familiar with  $\mu$ -calculus will observe that the algorithm can be expressed using the term  $\nu Y . \mu X (T \cup \text{APre}(Y, X))$  [10].

Hence the set  $S_{\forall}^{1,U}$  can be computed by taking the complement of the set of states for which there exists a path to  $S_{\exists}^{0,U}$  in the graph of  $\mathcal{O}$ . Given that  $S_{\exists}^{0,U}$ , and the set of states reaching  $S_{\exists}^{0,U}$ , can be computed in polynomial time, we have obtained a polynomial-time algorithm for computing  $S_{\forall}^{1,U}$ . Together with the cases for  $S_{\forall}^{0,U}$ ,  $S_{\exists}^{0,U}$  and  $S_{\exists}^{1,U}$ , this establishes Theorem 1.

## 4 Qualitative Reachability: IMDP semantics

We now focus on the IMDP semantics, and consider the computation of the following sets:

- $S_{\forall}^{0,I} = \{s \in S \mid \forall \sigma \in \Sigma^{[\mathcal{O}]_I} . \Pr_s^\sigma(\text{Reach}(T)) = 0\}$ ;
- $S_{\exists}^{0,I} = \{s \in S \mid \exists \sigma \in \Sigma^{[\mathcal{O}]_I} . \Pr_s^\sigma(\text{Reach}(T)) = 0\}$ ;
- $S_{\exists}^{1,I} = \{s \in S \mid \exists \sigma \in \Sigma^{[\mathcal{O}]_I} . \Pr_s^\sigma(\text{Reach}(T)) = 1\}$ ;
- $S_{\forall}^{1,I} = \{s \in S \mid \forall \sigma \in \Sigma^{[\mathcal{O}]_I} . \Pr_s^\sigma(\text{Reach}(T)) = 1\}$ .

This section will be dedicated to showing the following result. We note that the cases for  $S_{\forall}^{0,I}$ ,  $S_{\exists}^{0,I}$  and  $S_{\exists}^{1,I}$  proceed in a manner similar to the UMC case (using either graph reachability or reasoning based on the qualitative MDP abstraction); instead the case for  $S_{\forall}^{1,I}$  requires substantially different techniques.

**Theorem 2.** *The sets  $S_{\forall}^{0,I}$ ,  $S_{\exists}^{0,I}$ ,  $S_{\exists}^{1,I}$  and  $S_{\forall}^{1,I}$  can be computed in polynomial time in the size of the IMC.*

**Computation of  $S_{\forall}^{0,I}$ .** As in the case of UMCs, the computation of  $S_{\forall}^{0,I}$  reduces to straightforward reachability analysis on the graph of the IMC  $\mathcal{O}$ . The correctness of the reduction is based on the following lemma.

**Lemma 7.** *Let  $s \in S$ . There exists  $\sigma \in \Sigma^{[\mathcal{O}]_I}$  such that  $\Pr_s^\sigma(\text{Reach}(T)) > 0$  if and only if there exists a path  $r \in \text{Paths}_*^{\mathcal{O}}(s)$  such that  $\text{last}(r) \in T$ .*

Therefore, to obtain  $S_{\forall}^{0,I}$ , we proceed by computing the state set  $S \setminus S_{\forall}^{0,I} = \{s \in S \mid \exists \sigma \in \Sigma^{[\mathcal{O}]_I} . \Pr_s^\sigma(\text{Reach}(T)) > 0\}$ , which reduces to reachability on the graph of the IMC according to Lemma 7, and then taking the complement.

**Computation of  $S_{\exists}^{0,I}$  and  $S_{\exists}^{1,I}$ .** In the following we fix an arbitrary witness assignment function  $w$  of  $\mathcal{O}$ . Lemma 8 establishes that  $S_{\exists}^{0,I}$  (respectively,  $S_{\exists}^{1,I}$ ) equals the set of states of the qualitative MDP abstraction  $[\mathcal{O}]_w$  with respect to  $w$  for which there exists some scheduler such that  $T$  is reached with probability 0 (respectively, probability 1).

**Lemma 8.** *Let  $s \in S$  and  $\lambda \in \{0,1\}$ . There exists  $\sigma \in \Sigma^{[\mathcal{O}]_I}$  such that  $\Pr_s^\sigma(\text{Reach}(T)) = \lambda$  if and only if there exists a scheduler  $\sigma' \in \Sigma^{[\mathcal{O}]_w}$  such that  $\Pr_s^{\sigma'}(\text{Reach}(T)) = \lambda$ .*

Given that we have shown in Section 3 that the set of states of the qualitative MDP abstraction  $[\mathcal{O}]_w$  for which there exists some scheduler such that  $T$  is reached with probability 0 (respectively, probability 1) can be computed in polynomial time in the size of  $\mathcal{O}$ , we obtain polynomial-time algorithms for computing  $S_{\exists}^{0,1}$  (respectively,  $S_{\exists}^{1,1}$ ).

**Computation of  $S_{\forall}^{1,1}$ .** This case is notably different from the other three cases for the IMDP semantics, because schedulers that are *not* memoryless may influence whether a state is included in  $S_{\forall}^{1,1}$ . In particular, we recall the example of the IMC of Figure 1: as explained in Section 1, we have  $s_0 \notin S_{\forall}^{1,1}$ . In contrast, we have  $s_0 \in S_{\forall}^{1,U}$ , and  $s_0$  would be in  $S_{\forall}^{1,1}$  if we restricted the IMDP semantics to memoryless (actually finite-memory, in this case) schedulers. For this reason, a qualitative MDP abstraction is not useful for computing  $S_{\forall}^{1,1}$ , because it is based on the use of witness assignment functions that assign *constant* probabilities to sets of edges available from states: on repeated visits to a state, the (finite) set of available distributions remains the same in a qualitative MDP abstraction. Therefore we require alternative analysis methods that are not based on the qualitative MDP abstraction. Our approach is based on the notion of end components, which is a standard concept in the field of MDP verification [9]. In this section we introduce an alternative notion of end components, defined solely in terms of states of the IMC, which characterises situations in which the IMC can confine its behaviour to certain state sets with positive probability in the IMDP semantics (for example, the IMC of Figure 1 can confine itself to state  $s_0$  with positive probability in the IMDP semantics).

An *IMC-level end component* (ILEC) is a set  $C \subseteq S$  of states that is strongly connected and such that the total probability assigned to edges that have a source state in  $C$  but a target state outside of  $C$  can be made to be arbitrarily small (note that such edges must have an interval with a left endpoint of 0). Formally,  $C \subseteq S$  is an ILEC if, for each state  $s \in C$ , we have (1)  $E^{(+,\cdot)}(s, S \setminus C) = \emptyset$ , (2)  $\sum_{e \in E(s, C)} \text{right}(\delta(e)) \geq 1$ , and (3) the graph  $(C, E \cap (C \times C))$  is strongly connected.

*Example 2.* In the IMC  $\mathcal{O}_1$  of Figure 1, the set  $\{s_0\}$  is an ILEC: for condition (1), the edge  $(s_0, s_1)$  (the only edge in  $E(s_0, S \setminus \{s_0\})$ ) is not in  $E^{(+,\cdot)}$ , and, for condition (2), we have  $\text{right}(\delta(s_0, s_1)) = 1$ . In the IMC  $\mathcal{O}_2$  of Figure 2, the set  $\{s_0, s_1\}$  is an ILEC: for condition (1), the only edge leaving  $\{s_0, s_1\}$  has 0 as its left endpoint, i.e.,  $\delta(s_1, s_2) = (0, 0.2]$ , hence  $E^{(+,\cdot)}(s_0, \{s_2\}) = E^{(+,\cdot)}(s_1, \{s_2\}) = \emptyset$ ; for condition (2), we have  $\text{right}(\delta(s_0, s_0)) + \text{right}(\delta(s_0, s_1)) = 1.6 \geq 1$  and  $\text{right}(\delta(s_1, s_0)) + \text{right}(\delta(s_1, s_1)) = 1.3 \geq 1$ . In both cases, the identified sets clearly induce strongly connected subgraphs, thus satisfying condition (3).

*Remark 1.* Both conditions (1) and (2) are necessary to ensure that the probability of leaving  $C$  in one step can be made arbitrarily small. Consider an IMC with state  $s \in C$  such that  $E(s, C) = \{e_1\}$  and  $E(s, S \setminus C) = \{e_2, e_3\}$ , where  $\delta(e_1) = [0.6, 0.8]$ ,  $\delta(e_2) = [0, 0.2]$  and  $\delta(e_3) = [0, 0.2]$ . Then condition (1) holds

but condition (2) does not: indeed, at least total probability 0.2 must be assigned to the edges ( $e_2$  and  $e_3$ ) that leave  $C$ . Now consider an IMC with state  $s \in C$  such that  $E(s, C) = \{e_1, e_2\}$  and  $E(s, S \setminus C) = \{e_3\}$ , where  $\delta(e_1) = [0, 0.5]$ ,  $\delta(e_2) = [0, 0.5]$  and  $\delta(e_3) = [0.1, 0.5]$ . Then condition (2) holds (because the sum of the right endpoints of the intervals associated with  $e_1$  and  $e_2$  is equal to 1), but condition (1) does not (because the interval associated with  $e_3$  specifies that probability at least 0.1 must be assigned to leaving  $C$ ). Note also that if  $E(s, C) \subseteq E^{[\cdot, \cdot)} \cup E^{(\cdot, \cdot)}$  (all edges in  $E(s, C)$  have right-open intervals) and  $\sum_{e \in E(s, C)} \text{right}(\delta(e)) = 1$ , there must exist a least one edge in  $E(s, S \setminus C)$  by well formedness.

Let  $\mathfrak{J}$  be the set of ILECs of  $\mathcal{O}$ . We say that an ILEC  $C \in \mathfrak{J}$  is *maximal* if there does not exist any  $C' \in \mathfrak{J}$  such that  $C \subset C'$ . For a path  $\rho \in \text{Paths}^{[\mathcal{O}]_1}(s)$ , let  $\text{infst}(\rho) \subseteq S$  be the states that appear infinitely often along  $\rho$ , i.e., for  $\rho = s_0 \mu_0 s_1 \mu_1 \dots$ , we have  $\text{infst}(\rho) = \{s \in S \mid \forall i \in \mathbb{N}. \exists j > i. s_j = s\}$ . We present a result for ILECs that is analogous to the fundamental theorem of end components of [9]: the result specifies that, with probability 1, a scheduler of the IMDP semantics of  $\mathcal{O}$  must confine itself to an ILEC.

**Lemma 9.** *For  $s \in S$  and  $\sigma \in \Sigma^{[\mathcal{O}]_1}$ , we have  $\Pr_s^\sigma(\{\rho \mid \text{infst}(\rho) \in \mathfrak{J}\}) = 1$ .*

We now show that there exists a scheduler that, from a state within an ILEC, can confine the IMC to the ILEC with positive probability. This result is the ILEC analogue of a standard result for end components of finite MDPs that specifies that there exists a scheduler that, from a state of an end component, can confine the MDP to the end component with probability 1 (see [9,2]). In the case of IMCs and ILECs, it is not possible to obtain an analogous result for probability 1; in the example of Figure 1, the singleton set  $\{s_0\}$  is an ILEC, but it is not possible to find a scheduler that remains in  $s_0$  with probability 1, because with each transition the IMC goes to  $s_1$  with positive probability. For our purposes, it is sufficient to have a result stating that, from an ILEC, the IMC can be confined to the ILEC with positive probability.

**Lemma 10.** *Let  $C \in \mathfrak{J}$  and  $s \in C$ . There exists  $\sigma \in \Sigma^{[\mathcal{O}]_1}$  such that  $\Pr_s^\sigma(\{\rho \mid \rho \notin \text{Reach}(S \setminus C) \wedge \text{infst}(\rho) = C\}) > 0$ .*

The key point of the proof of Lemma 10 is the definition of a scheduler that assigns progressively decreasing probability to all edges in  $E^{(0, \cdot)}$  that leave ILEC  $C$ , in such a way as to guarantee that the IMC is confined in  $C$  with positive probability. This is possible because condition (2) of the definition of ILECs specifies that there is no lower bound on the probability that must be assigned to edges that leave  $C$ . Furthermore, the scheduler is defined so that the remaining probability at each step that is assigned between all edges that stay in  $C$  is always no lower than some fixed lower bound; this characteristic of the scheduler, combined with the fact that we remain in  $C$  with positive probability and the fact that  $C$  is strongly connected, means that we visit all states of  $C$  with positive probability under the defined scheduler.

Let  $U_{-T} = \bigcup\{C \in \mathcal{J} \mid C \cap T = \emptyset\}$  be the union of states of ILECs that do not contain states in  $T$ . Using Lemma 9 and Lemma 10 in a standard way, we can show that the existence of a scheduler of  $[\mathcal{O}]_1$  that reaches  $T$  with probability strictly less than 1 is equivalent to the existence of a path in the graph of  $\mathcal{O}$  that reaches  $U_{-T}$ .

**Proposition 1.** *Let  $s \in S$ . There exists  $\sigma \in \Sigma^{[\mathcal{O}]_1}$  such that  $\Pr_s^\sigma(\text{Reach}(T)) < 1$  if and only if there exists a finite path  $r \in \text{Paths}_*^{\mathcal{O}}(s)$  such that  $\text{last}(r) \in U_{-T}$ .*

Hence we identify the set  $S_{\Psi}^{1,I}$  by computing the complement of  $S_{\Psi}^{1,I}$ , i.e., the set  $S \setminus S_{\Psi}^{1,I} = \{s \in S \mid \forall \sigma \in \Sigma^{[\mathcal{O}]_1}. \Pr_s^\sigma(\text{Reach}(T)) < 1\}$ . Using Proposition 1, this set can be computed by considering reachability on the graph of  $\mathcal{O}$  of the set  $U_{-T}$ . The set  $U_{-T}$  can be computed in polynomial time in the size of  $\mathcal{O}$  in a manner similar to the computation of maximal end components of MDPs (see [9,2]). First we compute all strongly connected components  $(C_1, E \cap (C_1 \times C_1)), \dots, (C_m, E \cap (C_m \times C_m))$  of the graph  $(S \setminus T, E \cap ((S \setminus T) \times (S \setminus T)))$  of  $\mathcal{O}$ . Then, for each  $1 \leq i \leq m$ , we remove from  $C_i$  all states for which conditions (1) or (2) in the definition of ILECs do *not* hold with respect to  $C_i$  (these conditions can be checked in polynomial time for each state), to obtain the state set  $C'_i$ . Next, we compute the strongly connected components of the graph  $(C'_i, E \cap (C'_i \times C'_i))$ , and for each of these, repeat the procedure described above. We terminate the algorithm when it is not possible to remove a state (via a failure to satisfy a least one of the conditions (1) and (2) in the definition of ILECs) from any generated strongly connected component. The generated state sets of the strongly connected components obtained will be the maximal ILECs that do not contain any state in  $T$ , and their union is  $U_{-T}$ . Hence the overall algorithm for computing  $S_{\Psi}^{1,I}$  is in polynomial time in the size of  $\mathcal{O}$ .

## 5 Conclusion

We have presented algorithms for qualitative reachability properties for open IMCs. In the context of qualitative properties of system models with *fixed* probabilities on their transitions, probability can be regarded as imposing a fairness constraint, i.e., paths for which a state is visited infinitely often and one of its successors is visited only finitely often have probability 0. In open IMCs, the possibility to make the probability of a transition converge to 0 in the IMDP semantics captures a different phenomenon, which is key for problems concerning the minimum reachability probability being compared to 1. We conjecture that finite-memory strategies are no more powerful than memoryless strategies for this class of problem. For the three other classes of qualitative reachability problems, we have shown that the UMC and IMDP semantics coincide. We note that the algorithms presented in this paper require some numerical computation (a sum and a comparison of the result with 1 in the CPre, APre and ILEC computations), but these operations are simpler than the polynomial-time solutions for quantitative properties of (closed) IMCs in [6,17]. Similarly, the CPre and APre operators are simpler than the polynomial-time step of value iteration

used in the context of quantitative verification of [12]. For the IMDP semantics, our methods give directly a P-complete algorithm for the qualitative fragment of the temporal logic PCTL [13]. Future work could consider quantitative properties and  $\omega$ -regular properties, and applying the results to develop qualitative reachability methods for interval Markov decision processes or for higher-level formalisms such as clock-dependent probabilistic timed automata [19].

## References

1. C. Baier, L. de Alfaro, V. Forejt, and M. Kwiatkowska. Model checking probabilistic systems. In E. M. Clarke, T. A. Henzinger, H. Veith, and R. Bloem, editors, *Handbook of Model Checking*, pages 963–999. Springer, 2018.
2. C. Baier and J.-P. Katoen. *Principles of model checking*. MIT Press, 2008.
3. B. Caillaud, B. Delahaye, K. G. Larsen, A. Legay, M. L. Pedersen, and A. Wasowski. Constraint Markov chains. *Theoretical Computer Science*, 412(34):4373–4404, 2011.
4. 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.
5. K. Chatterjee, K. Sen, and T. A. Henzinger. Model-checking omega-regular properties of interval Markov chains. In R. Amadio, editor, *Proc. FOSSACS 2008*, volume 4962 of *LNCS*, pages 302–317. Springer, 2008.
6. 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.
7. C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, 1995.
8. C. Daws. Symbolic and parametric model checking of discrete-time Markov chains. In Z. Liu and K. Araki, editors, *Proc. ICTAC 2004*, volume 3407 of *LNCS*, pages 280–294. Springer, 2004.
9. L. de Alfaro. *Formal verification of probabilistic systems*. PhD thesis, Stanford University, Department of Computer Science, 1997.
10. L. de Alfaro. Computing minimum and maximum reachability times in probabilistic systems. In J. Baeten and S. Mauw, editors, *Proc. CONCUR 1999*, volume 1664 of *LNCS*, pages 66–81. Springer, 1999.
11. 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.
12. S. Haddad and B. Monmege. Interval iteration algorithm for MDPs and IMDPs. *Theoretical Computer Science*, 735:111–131, 2018.
13. H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6(5):512–535, 1994.
14. B. Jonsson and K. G. Larsen. Specification and refinement of probabilistic processes. In *Proc. LICS 1991*, pages 266–277. IEEE Computer Society, 1991.
15. I. O. Kozine and L. V. Utkin. Interval-valued finite Markov chains. *Reliable Computing*, 8(2):97–113, 2002.
16. R. Lanotte, A. Maggiolo-Schettini, and A. Troina. Parametric probabilistic transition systems for system design and analysis. *Formal Aspects of Computing*, 19(1):93–109, 2007.

17. 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.
18. 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.
19. 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.
20. M. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proc. FOCS 1985*, pages 327–338. IEEE Computer Society, 1985.

## 6 Appendix

### 6.1 Proof of Lemma 1

( $\Rightarrow$ ) Let  $B$  be a valid subset of  $E(s)$ . Firstly, because  $B$  is large, and from well formedness, we have that there exists some function  $f : B \rightarrow [0, 1]$  such that  $f(s, s') \in \delta(s, s')$  for each  $(s, s') \in B$  and  $\sum_{(s, s') \in B} f(s, s') = 1$ . Next, because  $B$  is realisable, we can obtain an assignment  $\mathbf{a}$  for  $s$  by letting  $\mathbf{a}(s') = f(s, s')$  for each  $(s, s') \in B$ , and letting  $\mathbf{a}(s') = 0$  for each  $(s, s') \notin B$ . This definition is clearly such that  $\{(s, s') \mid s' \in \text{support}(\mathbf{a})\} = B$ .

( $\Leftarrow$ ) Let  $\mathbf{a}$  be an assignment for  $s$  such that  $\{(s, s') \mid s' \in \text{support}(\mathbf{a})\} = B$ . We first show that  $B$  is large. By the definition of assignments, we have  $\mathbf{a} \in \delta(s, s')$ , and hence  $\mathbf{a}(s') \leq \text{right}(\delta(s, s'))$  for each state  $s' \in S$ . From this fact, and given that  $\mathbf{a}$  is a distribution (i.e., sums to 1), we have:

$$1 = \sum_{s' \in \text{support}(\mathbf{a})} \mathbf{a}(s') \leq \sum_{s' \in \text{support}(\mathbf{a})} \text{right}(\delta(s, s')) = \sum_{(s, s') \in B} \text{right}(\delta(s, s')).$$

In the case of  $\sum_{(s, s') \in B} \text{right}(\delta(s, s')) > 1$ , we have shown that  $B$  is large. In the case of  $\sum_{(s, s') \in B} \text{right}(\delta(s, s')) = 1$ , then  $\mathbf{a}(s') = \text{right}(\delta(s, s'))$  (i.e.,  $\mathbf{a}(s')$  must be equal to the right endpoint of  $\delta(s, s')$ ) for each  $(s, s') \in B$ . Hence  $B \subseteq E^{(\cdot, \cdot]}$ , and as a consequence  $B$  is large. The fact that  $\{(s, s') \mid s' \in \text{support}(\mathbf{a})\} = B$  implies that  $\mathbf{a}(s') = 0$  for all  $(s, s') \in E(s) \setminus B$ . Hence  $E(s) \setminus B \subseteq E^{[0, \cdot)}$ , and therefore  $B$  is realisable. Because  $B$  is large and realisable,  $B$  is valid.  $\square$

### 6.2 Proofs from Section 3

**Proof of Lemma 2.** ( $\Rightarrow$ ) Let  $\mathcal{D} = (S, \mathbf{P}) \in [\mathcal{O}]_{\cup}$  be a DTMC such that  $\text{Pr}_s^{\mathcal{D}}(\text{Reach}(T)) > 0$ . Because  $\text{Pr}_s^{\mathcal{D}}(\text{Reach}(T)) > 0$ , there exists at least one finite path  $s_0 s_1 \cdots s_n \in \text{Paths}_*^{\mathcal{D}}(s)$  such that  $s_n \in T$ . For each  $i < n$ , we have that  $\mathbf{P}(s_i, s_{i+1}) > 0$ . From this fact, and by the definition of assignment, we must have that  $\delta(s_i, s_{i+1}) \neq [0, 0]$ . This means that  $(s_i, s_{i+1}) \in E$ . Repeating this reasoning for each  $i < n$ , we have that  $s_0 s_1 \cdots s_n \in \text{Paths}_*^{\mathcal{O}}(s)$ . Recalling that  $s_n \in T$ , this direction of the proof is completed.



( $\Leftarrow$ ) Let  $s_0 s_1 \cdots s_n \in \text{Paths}_*^{\mathcal{O}}(s)$  be a finite path of  $\mathcal{O}$  such that  $s_n \in T$ . By the definition of finite paths of  $\mathcal{O}$ , we have  $\delta(s_i, s_{i+1}) \neq [0, 0]$  for each  $i < n$ . This implies that, for each  $i < n$ , there exists an assignment to state  $s_i$  that sets positive probability to  $s_{i+1}$ . In turn, this means that  $\mathbf{P}(s_i, s_{i+1}) > 0$  for each  $i < n$ , for some DTMC  $\mathcal{D} = (S, \mathbf{P}) \in [\mathcal{O}]_{\cup}$ . Given that  $s_n \in T$ , hence we have  $\text{Pr}_s^{\mathcal{D}}(\text{Reach}(T)) > 0$ .  $\square$

**Proof of Lemma 3.** Before we present the proof of Lemma 3, we recall (and specialise to reachability) a technical result from [5] that specifies that the same qualitative properties are satisfied on (finite) DTMCs that have the same graph. Two finite DTMCs  $\mathcal{D}_1 = (S_1, \mathbf{P}_1)$  and  $\mathcal{D}_2 = (S_1, \mathbf{P}_2)$  are *graph equivalent* if (1)  $S_1 = S_2$  and (2)  $\{(s, s') \mid \mathbf{P}_1(s, s') > 0\} = \{(s, s') \mid \mathbf{P}_2(s, s') > 0\}$ .

**Lemma 11 ([5]).** *Given graph equivalent finite DTMCs  $\mathcal{D}_1 = (S_1, \mathbf{P}_1)$  and  $\mathcal{D}_2 = (S_2, \mathbf{P}_2)$ , for any state  $s \in S_1 (= S_2)$ ,  $\bowtie \in \{<, =, >\}$  and  $\lambda \in \{0, 1\}$ , we have  $\text{Pr}_s^{\mathcal{D}_1}(\text{Reach}(T)) \bowtie \lambda$  if and only if  $\text{Pr}_s^{\mathcal{D}_2}(\text{Reach}(T)) \bowtie \lambda$ .*

( $\Rightarrow$ ) Let  $\mathcal{D} = (S, \mathbf{P})$  be a DTMC such that  $\mathcal{D} \in [\mathcal{O}]_{\cup}$  and  $\text{Pr}_s^{\mathcal{D}}(\text{Reach}(T)) = 0$ . Then we define the (memoryless) scheduler  $\sigma_{\mathcal{D}} \in \Sigma^{[\mathcal{O}]_w}$  of  $[\mathcal{O}]_w$  in the following way. Consider a state  $s' \in S$ , and let  $B_{s'} = \{(s', s'') \in E \mid \mathbf{P}(s', s'') > 0\}$  be the set of edges assigned positive probability by  $\mathcal{D}$ . Note that  $B_{s'} \in \text{Valid}(s')$ : this follows from the fact that  $\mathbf{P}(s', \cdot)$  is an assignment for  $s'$  and by Lemma 1. Then, for finite path  $r \in \text{Paths}_*^{[\mathcal{O}]_w}(s)$  ending in state  $s'$  (that is,  $\text{last}(r) = s'$ ), we let  $\sigma_{\mathcal{D}}(r) = \{w(B_{s'}) \mapsto 1\}$ , i.e.,  $\sigma_{\mathcal{D}}$  chooses (with probability 1) the distribution that corresponds to the witness assignment function applied to the edge set  $B_{s'}$  (this is possible because  $\Delta_w(s') = \{w(B) \mid B \in \text{Valid}(s')\}$  by definition). The fact that  $\text{Pr}_s^{\mathcal{D}}(\text{Reach}(T)) \bowtie \lambda$  implies  $\text{Pr}_s^{\sigma_{\mathcal{D}}}(\text{Reach}(T)) \bowtie \lambda$  then follows from the fact that the folded DTMC  $\tilde{\mathcal{D}}^{\sigma_{\mathcal{D}}}$  of  $\sigma_{\mathcal{D}}$  and the DTMC  $\mathcal{D}$  are graph equivalent (because, for all states  $s' \in S$ , we have  $\{(s', s'') \in E \mid \mathbf{P}(s', s'') > 0\} = B_{s'} = \{(s', s'') \mid s' \in \text{support}(w(B_{s'}))\}$ ), and from Lemma 11.

( $\Leftarrow$ ) Assume that there exists a scheduler  $\sigma \in \Sigma^{[\mathcal{O}]_w}$  of  $[\mathcal{O}]_w$  such that  $\text{Pr}_s^{\sigma}(\text{Reach}(T)) = 0$ . Given that  $[\mathcal{O}]_w$  is a finite MDP, from standard results [9], we can assume that  $\sigma$  is (a) pure (that is, for all finite paths  $r \in \text{Paths}_*^{[\mathcal{O}]_w}$ , we have that  $\sigma(r) = \{\mu \mapsto 1\}$  for some  $\mu \in \Delta_w(\text{last}(r))$ ) and (b) memoryless. Then we consider the DTMC  $\mathcal{D} = (S, \mathbf{P})$ , where, for all states  $s' \in S$ , we define  $\mathbf{P}(s', \cdot) = \sigma(s')$ . Given that  $\sigma(s')$  is an assignment (because  $\sigma$  is pure, i.e., assigns probability 1 to a distribution  $w(B)$ , for some  $B \in \text{Valid}(s')$ , and because  $w$  is a witness assignment function, i.e.,  $w(B)$  is an assignment), we have that  $\mathcal{D} \in [\mathcal{O}]_{\cup}$ . The folded DTMC  $\tilde{\mathcal{D}}^{\sigma}$  of  $\sigma$  and  $\mathcal{D}$  are identical. Hence, given that  $\text{Pr}_s^{\sigma}(\text{Reach}(T)) \bowtie \lambda$ , we have  $\text{Pr}_s^{\mathcal{D}}(\text{Reach}(T)) \bowtie \lambda$ .  $\square$

**Proof of Lemma 4.** We start by recalling the definition of CPre:

$$\begin{aligned}
& s \in \text{CPre}(X) \\
& \Leftrightarrow \exists \mu \in \Delta_w(s) . \text{support}(\mu) \subseteq X \\
& \Leftrightarrow \exists B \in E(s) . B \subseteq S \times X \wedge B \in \text{Valid}(s) \quad (\text{definition of } [\mathcal{O}]_w) \\
& \Leftrightarrow \exists B \in E(s) . B \subseteq S \times X \wedge E(s) \setminus B \subseteq E^{[0, \cdot]} \wedge B \text{ is large} \\
& \quad \quad \quad (\text{definition of } \text{Valid}(s)) \\
& \Leftrightarrow E(s, X) \subseteq S \times X \wedge E(s) \setminus E(s, X) \subseteq E^{[0, \cdot]} \wedge E(s, X) \text{ is large} \\
& \quad \quad \quad (\text{substituting } E(s, X) \text{ for } B) \\
& \Leftrightarrow E(s, S \setminus X) \subseteq E^{[0, \cdot]} \wedge E(s, X) \text{ is large}
\end{aligned}$$

where the last step follows from the fact that  $E(s, X) \subseteq S \times X$  holds trivially, and from  $E(s) \setminus E(s, X) = E(s, S \setminus X)$ .  $\square$

**Proof of Lemma 5.** From the definition of APre:

$$\begin{aligned}
& s \in \text{APre}(Y, X) \\
& \Leftrightarrow \exists \mu \in \Delta_w(s) . \text{support}(\mu) \subseteq Y \wedge \text{support}(\mu) \cap X \neq \emptyset \\
& \Leftrightarrow \exists B \in E(s) . B \subseteq S \times Y \wedge B \cap (S \times X) \neq \emptyset \wedge B \in \text{Valid}(s) \\
& \quad \quad \quad (\text{definition of } [\mathcal{O}]_w) \\
& \Leftrightarrow \exists B \in E(s) . B \subseteq S \times Y \wedge B \cap (S \times X) \neq \emptyset \wedge E(s) \setminus B \subseteq E^{[0, \cdot]} \wedge B \text{ is large} \\
& \quad \quad \quad (\text{definition of } \text{Valid}(s)) \\
& \Leftrightarrow E(s, Y) \subseteq S \times Y \wedge E(s, Y) \cap (S \times X) \neq \emptyset \wedge E(s) \setminus E(s, Y) \subseteq E^{[0, \cdot]} \\
& \quad \quad \quad \wedge E(s, Y) \text{ is large} \\
& \quad \quad \quad (\text{substituting } E(s, Y) \text{ for } B) \\
& \Leftrightarrow E(s, X \cap Y) \neq \emptyset \wedge E(s, S \setminus Y) \subseteq E^{[0, \cdot]} \wedge E(s, Y) \text{ is large}
\end{aligned}$$

where the last step follows from the fact that  $E(s, Y) \subseteq S \times Y$  holds trivially, from  $E(s, Y) \cap (S \times X) = E(s, X \cap Y)$ , and from  $E(s) \setminus (E(s, Y)) = E(s, S \setminus Y)$ .  $\square$

**Proof of Lemma 6.** ( $\Rightarrow$ ) Let  $\mathcal{D} = (S, \mathbf{P}) \in [\mathcal{O}]_{\mathcal{U}}$  be a DTMC such that  $\Pr_s^{\mathcal{D}}(\text{Reach}(T)) < 1$ . A bottom strongly connected component (BSCC)  $U \subseteq S$  of  $\mathcal{D}$  is a strongly connected component of the graph of  $\mathcal{D}$  (that is, the graph  $(S, \{(s', s'') \mid \mathbf{P}(s', s'') > 0\})$ ) such that there is no outgoing edge from  $U$  (that is, for all states  $s' \in U$  and  $s'' \in S$ , if  $\mathbf{P}(s', s'') > 0$  then  $s'' \in U$ ). Let  $\mathcal{U} \subseteq 2^S$  be the set of BSCCs of  $\mathcal{D}$ . Given that  $\mathcal{D}$  is a finite DTMC, by standard results for finite DTMCs, we have that BSCCs are reached with probability 1, i.e.,  $\Pr_s^{\mathcal{D}}(\text{Reach}(\bigcup_{U \in \mathcal{U}} U)) = 1$ , and that once a BSCC is entered all of its states are visited with probability 1. Hence  $\Pr_s^{\mathcal{D}}(\text{Reach}(T)) < 1$  implies that there exists some BSCC  $U \in \mathcal{U}$  such that  $U \cap T = \emptyset$  and  $\Pr_s^{\mathcal{D}}(\text{Reach}(U)) > 0$ . At this point we repeat the reasoning of Lemma 2 to show that there exists a finite path  $s_0 s_1 \cdots s_n \in \text{Paths}_*^{\mathcal{O}}(s)$  such that  $s_n \in U$ . Now we show that  $U \subseteq S_{\exists}^{0, \mathcal{U}}$ : for any state  $s' \in U$ , the fact that  $U \cap T = \emptyset$  and from the fact that BSCCs do not

feature outgoing edges, we have that  $\Pr_{s'}^{\mathcal{D}}(\text{Reach}(T)) = 0$ , and hence  $s' \in S_{\exists}^{0,U}$ . Therefore, for the finite path  $s_0 s_1 \cdots s_n \in \text{Paths}_*^{\mathcal{O}}(s)$ , we have  $s_n \in S_{\exists}^{0,U}$ , which completes this direction of the proof.

( $\Leftarrow$ ) Let  $s_0 s_1 \cdots s_n \in \text{Paths}_*^{\mathcal{O}}(s)$  be a finite path of  $\mathcal{O}$  such that  $s_n \in S_{\exists}^{0,U}$ . We assume w.l.o.g. that  $s_i \notin S_{\exists}^{0,U}$  for all  $i < n$ . As in the proof of Lemma 2, we can use this fact to conclude that there exists a DTMC  $\mathcal{D} = (S, \mathbf{P}) \in [\mathcal{O}]_{\cup}$  such that  $\mathbf{P}(s_i, s_{i+1}) > 0$  for each  $i < n$ . Furthermore, for any state  $s' \in S_{\exists}^{0,U}$ , there exists a DTMC  $\mathcal{D}^{s'} = (S, \mathbf{P}_{s'}) \in [\mathcal{O}]_{\cup}$  such that  $\Pr_{s'}^{\mathcal{D}^{s'}}(\text{Reach}(T)) = 0$ . From  $s_n \in S_{\exists}^{0,U}$ , we have  $\Pr_{s_n}^{\mathcal{D}^{s_n}}(\text{Reach}(T)) = 0$ . Note that all states on all paths in  $\text{Paths}_*^{\mathcal{D}^{s_n}}(s_n)$  are in  $S_{\exists}^{0,U}$  (if this was not the case, that is there exists a path in  $\text{Paths}_*^{\mathcal{D}^{s_n}}(s_n)$  featuring a state  $s'' \notin S_{\exists}^{0,U}$ , then we must also have  $s_n \notin S_{\exists}^{0,U}$ , i.e., we have obtained a contradiction). The fact that the set  $\{s_0, s_1, \dots, s_{n-1}\}$  (i.e., the states before  $s_n$  along the finite path) and the set  $S_{\exists}^{0,U}$  are disjoint means that we can combine the DTMCs  $\mathcal{D}$  along the finite path and  $\mathcal{D}^{s_n}$  within  $S_{\exists}^{0,U}$  to obtain a DTMC  $\mathcal{D}' = (S, \mathbf{P}')$ ; formally,  $\mathbf{P}'(s_i, \cdot) = \mathbf{P}(s_i, \cdot)$  for  $i < n$ ,  $\mathbf{P}'(s', \cdot) = \mathbf{P}_{s_n}(s', \cdot)$  for  $s' \in S_{\exists}^{0,U}$ , and  $\mathbf{P}'(s', \cdot)$  can be defined by arbitrary assignments for all other states, i.e.,  $s' \in S \setminus (\{s_0, s_1, \dots, s_{n-1}\} \cup S_{\exists}^{0,U})$ . Given that  $\mathcal{D}, \mathcal{D}^{s_n} \in [\mathcal{O}]_{\cup}$ , we have that  $\mathcal{D}' \in [\mathcal{O}]_{\cup}$ . Furthermore, in  $\mathcal{D}'$ , from  $s$  there exists a finite path to states  $S_{\exists}^{0,U}$  in which the DTMC is subsequently confined, and (by the trivial fact that  $S_{\exists}^{0,U} \subseteq S \setminus T$ ) hence  $\Pr_s^{\mathcal{D}'}(\text{Reach}(T)) < 1$ , completing this direction of the proof.  $\square$

### 6.3 Proofs from Section 4

**Proof of Lemma 7.** ( $\Rightarrow$ ) Let  $\sigma \in \Sigma^{[\mathcal{O}]_I}$  be a scheduler of  $[\mathcal{O}]_I$  such that  $\Pr_s^{\sigma}(\text{Reach}(T)) > 0$ . Recall that the DTMC  $\mathcal{D}_s^{\sigma}$  is defined as  $(\text{Paths}_*^{\sigma}(s), \mathbf{P})$ , i.e., the state set comprises the finite paths resulting from choices of  $\sigma$  from state  $s$ , and, for finite paths  $r, r' \in \text{Paths}_*^{\sigma}(s)$ , we have  $\mathbf{P}(r, r') > 0$  if  $r' = r\mu s'$  such that  $\mu \in \Delta(\text{last}(r))$ ,  $\sigma(\mu) > 0$  and  $\mu(s') > 0$ , otherwise  $\mathbf{P}(r, r') = 0$ . Furthermore, we say that a path  $\rho$  of  $\mathcal{D}_s^{\sigma}$  is in  $\text{Reach}(T)$  if the last state of a finite path that is a prefix of  $\rho$  is in  $T$ . For a path  $\rho$  of  $\mathcal{D}_s^{\sigma}$ , we can obtain a path  $\rho'$  of  $\mathcal{O}$  simply by extracting the sequence of the final states for all finite prefixes of  $\rho$ : that is, from the path  $s_0(s_0\mu_0s_1)(s_0\mu_0s_1\mu_1s_2)\cdots$  of  $\mathcal{D}_s^{\sigma}$  we can obtain the path  $s_0s_1s_2\cdots$  of  $\mathcal{O}$ , because  $\mu_i(s_{i+1}) > 0$  implies that  $\delta(s_i, s_{i+1}) \neq [0, 0]$  and hence  $(s_i, s_{i+1}) \in E$  for all  $i \in \mathbb{N}$ . Observe that  $\rho \in \text{Reach}(T)$  implies that  $\rho' \in \text{Reach}(T)$ . Hence  $\Pr_s^{\sigma}(\text{Reach}(T)) > 0$  implies that there exists  $\rho \in \text{Paths}_*^{\mathcal{D}_s^{\sigma}}(s)$  such that  $\rho \in \text{Reach}(T)$ , which in turn implies that there exists  $\rho' \in \text{Paths}_*^{\mathcal{O}}(s)$  such that  $\rho' \in \text{Reach}(T)$ , i.e., there exists a finite prefix  $r$  of  $\rho$  (and therefore  $r \in \text{Paths}_*^{\mathcal{O}}(s)$ ) such that  $\text{last}(r) \in T$ .

( $\Leftarrow$ ) Let  $s_0 s_1 \cdots s_n \in \text{Paths}_*^{\mathcal{O}}(s)$  be a finite path of  $\mathcal{O}$  such that  $s_n \in T$ . By the definition of finite paths of  $\mathcal{O}$ , we have  $\delta(s_i, s_{i+1}) \neq [0, 0]$  for each  $i < n$ , which implies in turn that there exists an assignment  $\mathbf{a}_i$  to  $s_i$  such that  $\mathbf{a}_i(s_{i+1}) > 0$ . We consider a scheduler  $\sigma \in \Sigma^{[\mathcal{O}]_I}$  of  $[\mathcal{O}]_I$  such that, for any finite path  $r \in \text{Paths}_*^{[\mathcal{O}]_I}(s)$ , if  $\text{last}(r) = s_i$  then  $\sigma(r) = \mathbf{a}_i$  for  $i < n$  ( $\sigma$  can

be defined in an arbitrary manner for all other finite paths). Observe that the finite path  $r = s_0(s_0\mathbf{a}_0s_1)(s_0\mathbf{a}_0s_1\mathbf{a}_1s_2) \cdots (s_0\mathbf{a}_0 \cdots \mathbf{a}_{n-1}s_n)$  is a finite path of  $\mathcal{D}_s^\sigma$  such that all infinite paths that have  $r$  as a prefix are in  $\text{Reach}(T)$ , and hence  $\Pr_s^\sigma(\text{Reach}(T)) > 0$ .  $\square$

**Proof of Lemma 8.** The proof consists of two parts, depending on whether  $\lambda = 0$  or  $\lambda = 1$ .

*Case:  $\lambda = 0$ .* ( $\Rightarrow$ ) Let  $\sigma \in \Sigma^{[\mathcal{O}]_I}$  be a scheduler of  $[\mathcal{O}]_I$  such that  $\Pr_s^\sigma(\text{Reach}(T)) = 0$ . We show how we can define a scheduler  $\mathfrak{f}(\sigma) \in \Sigma^{[\mathcal{O}]_w}$  such that  $\Pr_s^{\mathfrak{f}(\sigma)}(\text{Reach}(T)) = 0$ .

For any finite path  $r \in \text{Paths}_*^{[\mathcal{O}]_I}(s)$ , let  $B_r^\sigma$  be the set of edges assigned positive probability by  $\sigma$  after  $r$ . Formally:

$$B_r^\sigma = \{(s, s') \in E \mid s = \text{last}(r) \wedge \exists \mu \in \Delta(s). (\sigma(r)(\mu) > 0 \wedge \mu(s') > 0)\}.$$

We now define the partial function  $\mathfrak{g} : \text{Paths}_*^{[\mathcal{O}]_w}(s) \rightarrow \text{Paths}_*^\sigma(s)$ , which associates with a finite path of  $[\mathcal{O}]_w$  some finite path of  $[\mathcal{O}]_I$  (more precisely, of  $\sigma$ ) that features the same sequence of states. For  $r = s_0\mu_0s_1\mu_1 \cdots \mu_{n-1}s_n \in \text{Paths}_*^{[\mathcal{O}]_w}(s)$ , let  $\mathfrak{g}(r) = s'_0\mu'_0s'_1\mu'_1 \cdots \mu'_{n-1}s'_n$  such that:

- $s'_i = s_i$  for all  $i \leq n$ ;
- $\sigma(\mathfrak{g}(s_0\mu_0s_1\mu_1 \cdots \mu_{i-1}s_i))(\mu'_i) > 0$  and  $\mu'_i(s'_{i+1}) > 0$  for all  $i < n$ ;
- $\mathfrak{g}(s_0\mu_0s_1\mu_1 \cdots \mu_{i-1}s_i) = s'_0\mu'_0s'_1\mu'_1 \cdots \mu'_{i-1}s'_i$  for all  $i < n$ .

The first two conditions and  $s_0 = s$  ensure that  $\mathfrak{g}(r) \in \text{Paths}_*^\sigma(s)$ ; if these conditions do not hold, then  $\mathfrak{g}(r)$  is undefined. The third condition ensures that  $\mathfrak{g}$  maps paths of  $\text{Paths}_*^{[\mathcal{O}]_w}(s)$  to paths of  $\text{Paths}_*^\sigma(s)$  in a consistent manner (for example, the condition ensures that, if there are two finite paths  $r_1$  and  $r_2$  of  $[\mathcal{O}]_w$  with a common prefix of length  $i$ , then  $\mathfrak{g}(r_1)$  and  $\mathfrak{g}(r_2)$  will also have a common prefix of length  $i$ ). We are now in a position to define the scheduler  $\mathfrak{f}(\sigma)$  of  $[\mathcal{O}]_w$ : let  $\mathfrak{f}(\sigma)(r) = \{w(B_{\mathfrak{g}(r)}^\sigma) \mapsto 1\}$  for each finite path  $r \in \text{Paths}_*^{[\mathcal{O}]_w}(s)$ .

Next we need to show that  $\Pr_s^{\mathfrak{f}(\sigma)}(\text{Reach}(T)) = 0$ . Note that  $\Pr_s^\sigma(\text{Reach}(T)) = 0$  if and only if all paths  $\rho \in \text{Paths}_*^\sigma(s)$  are such that  $\rho \notin \text{Reach}(T)$ , and similarly  $\Pr_s^{\mathfrak{f}(\sigma)}(\text{Reach}(T)) = 0$  if and only if all paths  $\rho \in \text{Paths}_*^{\mathfrak{f}(\sigma)}(s)$  are such that  $\rho \notin \text{Reach}(T)$ . Hence we show that the existence of a path  $\rho \in \text{Paths}_*^{\mathfrak{f}(\sigma)}(s)$  such that  $\rho \in \text{Reach}(T)$  implies the existence of a path  $\rho' \in \text{Paths}_*^\sigma(s)$  such that  $\rho' \in \text{Reach}(T)$ . Let  $\rho \in \text{Paths}_*^{\mathfrak{f}(\sigma)}(s)$  be such that  $\rho \in \text{Reach}(T)$ . We can identify a path  $\rho' \in \text{Paths}_*^\sigma(s)$  that visits the same states as  $\rho$ , and hence is such that  $\rho' \in \text{Reach}(T)$ . Let  $\rho = s_0\mu_0s_1\mu_1 \cdots$ , and let  $r_i$  be the  $i$ -th prefix of  $\rho$  (i.e.,  $r_i = s_0\mu_0s_1\mu_1 \cdots \mu_{i-1}s_i$ ). We show the existence of  $\rho'$  by induction, by considering prefixes  $r'_i$  of  $\rho'$  of increasing length.

(Base case.) Let  $r'_0 = s_0$ .

(Inductive step.) Assume that we have constructed the finite path  $r'_i = s_0\mu'_0s_1\mu'_1 \cdots \mu'_{i-1}s_i$ . We show how to extend  $r'_i$  with one transition to obtain the finite path  $r'_{i+1}$ . The transition  $s_i\mu_i s_{i+1}$  along  $\rho$  implies that

$(s_i, s_{i+1}) \in B_{\mathbf{g}(r_i)}^\sigma$ . This fact then implies that there exists some  $\mu' \in \Delta(s_i)$  such that  $\sigma(\mathbf{g}(r_i))(\mu') > 0$  and  $\mu'(s_{i+1}) > 0$ . We then let  $r'_{i+1} = r'_i \mu' s_{i+1}$ .

From this construction of  $\rho'$ , we can see that  $\rho' \in \text{Paths}^\sigma(s)$  (because  $\rho'$  follows the definition of a path of scheduler  $\sigma$ , and because the choice of  $\mu'$  after  $r'_i$  must be consistent with the choices made along all prefixes of  $r'_i$ ). Furthermore, we have  $\rho' \in \text{Reach}(T)$  (because  $\rho$  and  $\rho'$  feature the same states, and  $\rho \in \text{Reach}(T)$ ). Hence we have shown that  $\Pr_s^\sigma(\text{Reach}(T)) = 0$  implies  $\Pr_s^{\mathbf{f}(\sigma)}(\text{Reach}(T)) = 0$ , concluding this direction of the proof.

( $\Leftarrow$ ) Let  $\sigma \in \Sigma^{[\mathcal{O}]_w}$  such that  $\Pr_s^\sigma(\text{Reach}(T)) = 0$ . Given that  $[\mathcal{O}]_w$  is a finite MDP, we can assume that  $\sigma$  is memoryless and pure. We now define  $\sigma' \in \Sigma^{[\mathcal{O}]_1}$  and show that  $\Pr_s^{\sigma'}(\text{Reach}(T)) = 0$ . For a finite path  $r \in \text{Paths}_*^{[\mathcal{O}]_1}(s)$ , let  $\sigma'(r) = \sigma(\text{last}(r))$  (recall that, for all states  $s' \in S$ ,  $\sigma(s') = \{w(B) \mapsto 1\}$  for some  $B \in \text{Valid}(s')$ , where  $w(B)$  is an assignment for  $s'$ , hence  $\sigma'$  is well defined). Then we have that the DTMCs  $\mathcal{D}_s^\sigma$  and  $\mathcal{D}_s^{\sigma'}$  are identical, and hence  $\Pr_s^\sigma(\text{Reach}(T)) = 0$  implies that  $\Pr_s^{\sigma'}(\text{Reach}(T)) = 0$ .

Hence the part of the proof for  $\lambda = 0$  is concluded.

Case:  $\lambda = 1$ . ( $\Rightarrow$ ) Let  $\sigma \in \Sigma^{[\mathcal{O}]_1}$  be such that  $\Pr_s^\sigma(\text{Reach}(T)) = 1$ . We use the construction of the scheduler  $\mathbf{f}(\sigma)$  from the case of  $S_{\exists}^{0,1}$  (i.e., the case  $\lambda = 0$  of this proof), and show that  $\Pr_s^{\mathbf{f}(\sigma)}(\text{Reach}(T)) = 1$ . We proceed to show that  $\Pr_s^\sigma(\text{Reach}(T)) = 1$  implies  $\Pr_s^{\mathbf{f}(\sigma)}(\text{Reach}(T)) = 1$  by showing the contrapositive, i.e.,  $\Pr_s^{\mathbf{f}(\sigma)}(\text{Reach}(T)) < 1$  implies  $\Pr_s^\sigma(\text{Reach}(T)) < 1$ . To show this property, we use properties of *end components*. An end component of finite MDP  $[\mathcal{O}]_w = (S, \Delta_w)$  is a pair  $(C, D)$  where  $C \subseteq S$  and  $D : C \rightarrow 2^{\text{Dist}(S)}$  is such that (1)  $\emptyset \neq D(s') \subseteq \Delta_w(s')$  for all  $s' \in C$ , (2)  $\text{support}(\mu) \subseteq C$  for all  $s' \in C$  and  $\mu \in D(s')$ , and (3) the graph  $(C, \{(s', s'') \in C \times C \mid \exists \mu \in D(s') . \mu(s'') > 0\})$  is strongly connected. Let  $\mathcal{E}$  be the set of end components of  $[\mathcal{O}]_w$ . For an end component  $(C, D)$ , let  $sa(C, D) = \{(s', \mu) \mid s \in C \wedge \mu \in D(s')\}$  be the set of state-action pairs associated with  $(C, D)$ . For a path  $\rho \in \text{Paths}^{[\mathcal{O}]_w}(s)$ , we let  $\text{inf}(\rho) = \{(s, \mu) \mid s \text{ and } \mu \text{ appear infinitely often along } \rho\}$ . The fundamental theorem of end components [9] specifies that, for any scheduler  $\sigma' \in \Sigma^{[\mathcal{O}]_w}$ , we have that  $\Pr_s^{\sigma'}(\{\rho \mid \exists (C, D) \in \mathcal{E} . \text{inf}(\rho) = sa(C, D)\}) = 1$ .

Now, given that  $[\mathcal{O}]_w$  is a finite MDP, and from  $\Pr_s^{\mathbf{f}(\sigma)}(\text{Reach}(T)) < 1$  and the fact that states in  $T$  are absorbing, there exists an end component  $(C, D) \in \mathcal{E}$  such that  $C \cap T = \emptyset$  and  $\Pr_s^{\mathbf{f}(\sigma)}(\{\rho \mid \text{inf}(\rho) = sa(C, D)\}) > 0$ . We observe the following property: for any finite path  $r \in \text{Paths}_*^{\mathbf{f}(\sigma)}(s)$ , we have that  $\{(s', s'') \mid w(B_{\mathbf{g}(r)}^\sigma)(s'') > 0\} = B_{\mathbf{g}(r)}^\sigma$ . From the definition of  $B_{\mathbf{g}(r)}^\sigma$ , we can see that  $\sigma$  assigns positive probability only to those distributions that assign positive probability to states that are targets of edges in  $B_{\mathbf{g}(r)}^\sigma$ . Putting these two facts together, we conclude that, for any set  $X \subseteq S$ , we have  $\text{support}(w(B_{\mathbf{g}(r)}^\sigma)) \subseteq X$  if and only if  $\{s' \in S \mid \exists \mu \in \Delta(\text{last}(\mathbf{g}(\rho))) . (\sigma(\mathbf{g}(r))(\mu) > 0 \wedge \mu(s') > 0)\} \subseteq X$ . This means that the existence of a finite path  $r \in \text{Paths}_*^{\mathbf{f}(\sigma)}(s)$  such that all suffixes of  $r$  generated by  $\mathbf{f}(\sigma)$  visit only states in  $C$  implies that all suffixes of  $\mathbf{g}(r)$  (which we recall is a finite path of  $\sigma$ , i.e.,  $\mathbf{g}(r) \in \text{Paths}_*^\sigma(s)$ ) generated by

$\sigma$  visit only states in  $C$ . Given that  $\Pr_s^{f(\sigma)}(\{\rho \mid \text{inf}(\rho) = \text{sa}(C, D)\}) > 0$ , such a finite path in  $\text{Paths}_*^{f(\sigma)}(s)$  exists. Then, given that there exists a finite path in  $\text{Paths}_*^\sigma(s)$  such that all suffixes generated by  $\sigma$  visit only states in  $C$ , and from the fact that  $C \cap T = \emptyset$ , we have that  $\Pr_s^\sigma(\text{Reach}(T)) < 1$ . Hence we have shown that  $\Pr_s^{f(\sigma)}(\text{Reach}(T)) < 1$  implies that  $\Pr_s^\sigma(\text{Reach}(T)) < 1$ . This means that  $\Pr_s^\sigma(\text{Reach}(T)) = 1$  implies that  $\Pr_s^{f(\sigma)}(\text{Reach}(T)) = 1$ . Hence, for  $\sigma \in \Sigma^{[O]_1}$  such that  $\Pr_s^\sigma(\text{Reach}(T)) = 1$  there exists  $\sigma' \in \Sigma^{[O]_w}$  such that  $\Pr_s^{\sigma'}(\text{Reach}(T)) = 1$ .

( $\Leftarrow$ ) Let  $\sigma \in \Sigma^{[O]_w}$  be such that  $\Pr_s^\sigma(\text{Reach}(T)) = 1$ . As in the case of the analogous direction of the case for  $\lambda = 0$ , we can show the existence of  $\sigma' \in \Sigma^{[O]_1}$  such that  $\mathcal{D}_s^\sigma$  and  $\mathcal{D}_s^{\sigma'}$  are identical, and hence  $\Pr_s^\sigma(\text{Reach}(T)) = 1$  implies that  $\Pr_s^{\sigma'}(\text{Reach}(T)) = 1$ .  $\square$

**Proof of Lemma 9.** The proof is structured in the same manner as that for classical end components in [9]. Consider  $C \subseteq S$  such that  $C \notin \mathfrak{J}$ . Our aim is to show that  $\Pr_s^\sigma(\{\rho \mid \text{infst}(\rho) = C\}) = 0$ . Given that  $\mathfrak{J}$  is a finite set, the required result then follows from this.

First suppose that the condition (1) in the definition of ILECs does not hold, i.e., there exists  $(s', s'') \in E(s)$  such that  $s' \in C$ ,  $(s', s'') \in E^{(+, \cdot)}$  and  $s'' \notin C$ . Observe that  $\text{left}(\delta(s', s'')) > 0$ , which means that the probability of remaining in  $C$  when visiting  $s'$ , i.e.,  $1 - \text{left}(\delta(s', s''))$ , is strictly less than 1. Given that  $s' \in \text{infst}(\rho)$  for every  $\rho$  such that  $\text{infst}(\rho) = C$ , we have that  $\Pr_s^\sigma(\{\rho \mid \text{infst}(\rho) = C\}) \leq \lim_{k \rightarrow \infty} (1 - \text{left}(\delta(s', s'')))^k = 0$ .

Suppose that condition (2) in the definition of ILECs does not hold, i.e., there exists  $s' \in C$  such that  $\sum_{e \in E(s', C)} \text{right}(\delta(e)) < 1$ . Therefore, for all  $\mu \in \Delta(s')$ , we must have  $\sum_{s'' \in C} \mu(s'') \leq \sum_{e \in E(s', C)} \text{right}(\delta(e)) < 1$ . Hence, the probability of remaining in  $C$  when visiting  $s'$  is strictly less than 1. Then, as in the case of the first condition in the definition of ILECs, we conclude that  $\Pr_s^\sigma(\{\rho \mid \text{infst}(\rho) = C\}) \leq \lim_{k \rightarrow \infty} (\sum_{e \in E(s', C)} \text{right}(\delta(e)))^k = 0$ .

Now suppose that the condition (3) in the definition of ILECs does not hold, i.e., there is no path from state  $s'$  to state  $s''$  in the graph induced by  $C$ , where all states along the path (including  $s'$  and  $s''$ ) belong to  $C$ . Given that  $S$  is finite, from some point onwards, any path  $\rho$  features only those states from  $\text{infst}(\rho)$ . Then for any occurrence of  $s'$  in the suffix of  $\rho$  that features only states from  $\text{infst}(\rho)$ , if we have  $\text{infst}(\rho) = C$ , there cannot be a subsequent occurrence of  $s''$  along the path. Hence we must have  $\text{infst}(\rho) \neq C$ , and thus  $\Pr_s^\sigma(\{\rho \mid \text{infst}(\rho) = C\}) = 0$ .  $\square$

**Proof of Lemma 10.** As mentioned in Section 4, the intuition of the proof is that we define a scheduler  $\sigma \in \Sigma^{[O]_1}$  that depends only on the current state and the number of transitions done so far, that assigns progressively decreasing probability to all edges in  $E^{(0, \cdot)}$  that leave  $C$ . To guarantee that all states in  $C$  are visited infinitely often, we also define a constant probability that is assigned to states in  $E^{(0, \cdot)}$  that remain in  $C$ .

In the following, we assume that there is at least one edge  $(s', s'') \in E$  such that  $s', s'' \in C$  and  $\delta(s', s'')$  is non-singular (otherwise the graph  $(C, E \cap (C \times C))$  is a bottom strongly connected component, i.e., there are no edges leaving  $C$ , with fixed probabilities on each edge, and hence  $\Pr_s^\sigma(\{\rho \mid \rho \notin \text{Reach}(S \setminus C) \wedge \text{infst}(\rho) = C\}) = 1$ ). Similarly, we assume that there is at least one state  $s' \in C$  for which  $E^{(0, \cdot)}(s', S \setminus C)$  is non-empty; if this is not the case, from the definition of ILECs, there is no edge that leaves  $C$ . It then remains to define the scheduler  $\sigma$  as selecting only assignments that assign fixed, positive probability to edges in  $E^{(0, \cdot)}$  that have source states in  $C$  (note that these edges remain in  $C$ ) in a similar manner as in the following proof, and we omit the details.

First, we define the constant probability that is assigned by the scheduler  $\sigma$  to edges in  $E^{(0, \cdot)}$  that have source and target states in  $C$ . For  $s' \in C$ , we let:

$$\kappa_{s'} = \min\left\{ \min_{e \in E^{(0, \cdot)}(s', C)} \text{right}(\delta(e)), 1 - \sum_{e \in E(s', C)} \text{left}(\delta(e)) \right\}.$$

Intuitively,  $\kappa_{s'}$  is the supremum probability that can be assigned to edges with left endpoint 0 from state  $s'$  that remain in  $C$ . We let  $\bar{\kappa} = \frac{1}{2} \min_{s' \in C} \kappa_{s'}$ , which will be the probability assigned to edges in  $E^{(0, \cdot)}(s', C)$  (if it is non-empty) by the distribution chosen by  $\sigma$  from paths ending in state  $s'$ .

Next, for  $s' \in C$ , we let:

$$\lambda_{s'} = \min\left\{ \min_{e \in E^{(0, \cdot)}(s', S \setminus C)} \text{right}(\delta(e)), 1 - \bar{\kappa} - \sum_{e \in E(s', C)} \text{left}(\delta(e)) \right\}.$$

Intuitively,  $\lambda_{s'}$  is the supremum probability that can be assigned to edges with left endpoint 0 from state  $s'$  to a state *not*  $C$ , assuming that probability  $\bar{\kappa}$  is assigned to edges with left endpoint 0 that remain in  $C$  (note that this assumption is conservative, because in some states there may be no outgoing edges with left endpoint 0 that remain in  $C$ ). Let  $\bar{\lambda} = \frac{1}{2} \min_{s' \in C} \lambda_{s'}$ , and let  $i$  equal the minimum index  $i \geq 2$  such that  $\frac{1}{2^i} < \bar{\lambda}$ .

Given  $s' \in C$  and  $i \in \mathbb{N}$ , let  $\mathbf{a}_{s'}^i$  be an assignment for  $s'$  defined in the following way:

- If  $E^{(0, \cdot)}(s', S \setminus C) \neq \emptyset$ , then  $\mathbf{a}_{s'}^i(s'') = \frac{1}{2^{i+i}} \cdot \frac{1}{|E^{(0, \cdot)}(s', S \setminus C)|}$  for each  $s'' \in S \setminus C$  such that  $(s', s'') \in E^{(0, \cdot)}(s', S \setminus C)$ .
- If  $E^{(0, \cdot)}(s', C) \neq \emptyset$ , then  $\mathbf{a}_{s'}^i(s'') = \bar{\kappa} \cdot \frac{1}{|E^{(0, \cdot)}(s', C)|}$  for each  $s'' \in C$  such that  $(s', s'') \in E^{(0, \cdot)}(s', C)$ .

The definitions of  $\bar{\kappa}$  and  $\bar{\lambda}$  allow us to complete the definition of  $\mathbf{a}_{s'}^i$ , for states not considered in these two points above, so that it  $\mathbf{a}_{s'}^i$ , an assignment for  $s'$ . The intuition here is as follows: condition (2) of the definition of ILECs specifies that obtaining an assignment that obeys the above constraints and which sums to 1 is not problematic; furthermore, the definitions of  $\bar{\kappa}$  and  $\bar{\lambda}$  guarantee that enough probability mass is available to edges that do *not* have a left endpoint of 0 to

equal or exceed their left endpoints, and that the probability that is assigned to edges that *do* have a left endpoint of 0 does not exceed their right endpoint.

We now formalise the definition of  $\sigma$ . Let  $r \in \text{Paths}_*^{[\mathcal{O}]1}(s)$ , where  $r = s_0\mu_0s_1\mu_1\cdots\mu_{n-1}s_n$  and  $s_i \in C$  for all  $i \leq n$ . We let  $\sigma(r) = \{\mathbf{a}_{s_n}^n \mapsto 1\}$ .

It remains to show that  $\Pr_s^\sigma(\{\rho \mid \rho \notin \text{Reach}(S \setminus C) \wedge \text{infst}(\rho) = C\}) > 0$ . For a path  $\rho \in \text{Paths}^{[\mathcal{O}]1}(s)$ , let  $\rho(i)$  be the  $(i+1)$ -th state along the path. Given state set  $X \subseteq S$  and  $k \in \mathbb{N}$ , we let  $\text{Reach}^{\leq k}(X)$  be the set of paths that reach  $X$  within  $k$  transitions; formally,  $\text{Reach}^{\leq k}(X) = \{\rho \in \text{Paths}^{[\mathcal{O}]1}(s) \mid \exists i \leq k. \rho(i) \in X\}$ . By definition, we have  $\text{Reach}(X) = \bigcup_{k \in \mathbb{N}} \text{Reach}^{\leq k}(X)$ , and  $\Pr_s^\sigma(\text{Reach}(X)) = \lim_{k \rightarrow \infty} \Pr_s^\sigma(\text{Reach}^{\leq k}(X))$ . From the definition of  $\sigma$  above, we have  $\Pr_s^\sigma(\text{Reach}^{\leq 0}(S \setminus C)) = 0$  (from  $s \in C$ ) and, for  $k \geq 1$ , we have  $\Pr_s^\sigma(\text{Reach}^{\leq k}(S \setminus C)) \leq \frac{1}{4} + \sum_{i=1}^k \frac{1}{2^{i+2}} \prod_{j=0}^{i-1} (1 - \frac{1}{2^{j+1}}) \leq \frac{1}{4} + \sum_{i=1}^k \frac{1}{2^{i+2}} = \sum_{i=0}^k \frac{1}{2^{i+2}}$ . We note that  $\sum_{i=0}^k \frac{1}{2^{i+2}} \leq \frac{1}{2}$  for all  $k \in \mathbb{N}$ . Hence  $\Pr_s^\sigma(\text{Reach}(S \setminus C)) = \lim_{k \rightarrow \infty} \Pr_s^\sigma(\text{Reach}^{\leq k}(S \setminus C)) \leq \lim_{k \rightarrow \infty} \sum_{i=0}^k \frac{1}{2^{i+2}} \leq \frac{1}{2}$ . Therefore we have shown that  $\Pr_s^\sigma(\{\rho \mid \rho \notin \text{Reach}(S \setminus C)\}) \geq \frac{1}{2} > 0$ .

Because the assignments used by  $\sigma$  dedicate a probability value to all edges with source and target states in  $C$  that is no lower than some fixed lower bound, and because the graph  $(C, E \cap (C \times C))$  is strongly connected, we have that  $\Pr_s^\sigma(\{\rho \mid \rho \notin \text{Reach}(S \setminus C) \wedge \text{infst}(\rho) = C\}) = \Pr_s^\sigma(\{\rho \mid \rho \notin \text{Reach}(S \setminus C)\})$ . Hence we have shown that  $\Pr_s^\sigma(\{\rho \mid \rho \notin \text{Reach}(S \setminus C) \wedge \text{infst}(\rho) = C\}) > 0$ .  $\square$

*Remark 2.* Note that we can confine the IMC to an ILEC with probability 1 if the only edges leaving the ILEC belong to  $E^{[0,\cdot]}$ ; however, we do not require that result, and above we settle for a simplified scheduler construction that does not distinguish between edges in  $E^{[0,\cdot]}$  and  $E^{(0,\cdot)}$ , and hence allows some progressively decreasing probability of exiting from the ILEC even if the outgoing edges of the ILEC are only from  $E^{[0,\cdot]}$ .

**Proof of Proposition 1.** ( $\Rightarrow$ ) Let  $\sigma \in \Sigma^{[\mathcal{O}]1}$  be such that  $\Pr_s^\sigma(\text{Reach}(T)) < 1$ . Then, by duality, we have  $\Pr_s^\sigma(\{\rho \mid \rho \notin \text{Reach}(T)\}) > 0$ . From this fact, and from Lemma 9, we then have that  $\Pr_s^\sigma(\{\rho \mid \rho \notin \text{Reach}(T) \wedge \text{infst}(\rho) \in \mathfrak{J}\}) > 0$ . For any path  $\rho \in \text{Paths}^\sigma(s)$ , we have that  $\rho \notin \text{Reach}(T)$  and  $\text{infst}(\rho) \in \mathfrak{J}$  implies that  $\rho \in \text{Reach}(U_{-T})$  from the definition of  $U_{-T}$ . Hence  $\Pr_s^\sigma(\{\rho \mid \rho \in \text{Reach}(U_{-T})\}) > 0$ . From Lemma 7, we then have that there exists a finite path  $r \in \text{Paths}_*^{\mathcal{O}}(s)$  such that  $\text{last}(r) \in U_{-T}$ , completing this direction of the proof.

( $\Leftarrow$ ) The existence of a path  $r \in \text{Paths}_*^{\mathcal{O}}(s)$  such that  $\text{last}(r) \in U_{-T}$  implies that there exists  $\sigma \in \Sigma^{[\mathcal{O}]1}$  such that  $\Pr_s^\sigma(\{\rho \mid \rho \in \text{Reach}(U_{-T})\}) > 0$  by Lemma 7. We define  $\sigma' \in \Sigma^{[\mathcal{O}]1}$  such that  $\Pr_s^{\sigma'}(\text{Reach}(T)) < 1$  in the following way: the scheduler  $\sigma'$  behaves as  $\sigma$  until a state in  $U_{-T}$  is reached; once such a state has been reached, the scheduler then behaves as a scheduler defined in Lemma 10, which ensures that the states ILEC that does not contain any state from  $T$  are henceforth visited infinitely often with positive probability. The construction of the scheduler is standard, and we describe it here for completeness.

Let  $s' \in U_{-T}$ , and let  $C_{s'} \in \mathfrak{J}$  be the maximal ILEC such that  $s' \in C_{s'}$  and  $C_{s'} \cap T = \emptyset$  (which exists by definition of  $U_{-T}$ ). Now let  $\sigma_{s'} \in \Sigma^{[\mathcal{O}]1}$  be such that



$\Pr_s^\sigma(\{\rho \mid \rho \notin \text{Reach}(S \setminus C_{s'}) \wedge \text{infst}(\rho) = C_{s'}\}) > 0$ , which exists by Lemma 10. For finite paths  $r = s_0\mu_0s_1\mu_1 \cdots \mu_{n-1}s_n \in \text{Paths}^{[\mathcal{O}]_1}(s)$  such that  $s_i \notin U_{-T}$  for all  $i \leq n$ , we let  $\sigma'(r) = \sigma(r)$ . For a finite path  $r = s_0\mu_0s_1\mu_1 \cdots \mu_{n-1}s_n \in \text{Paths}^{[\mathcal{O}]_1}(s)$  such that there exists  $i \leq n$  for which  $s_j \notin U_{-T}$  for all  $j < i$  and  $s_j \in U_{-T}$  for all  $j \geq i$ , we let  $\sigma'(r) = \sigma_{s_i}(s_i\mu_i \cdots \mu_{n-1}s_n)$ . For other finite paths, the definition of  $\sigma'$  can be arbitrary. Given that  $\Pr_s^\sigma(\{\rho \mid \rho \in \text{Reach}(U_{-T})\}) > 0$ , we have  $\Pr_s^{\sigma'}(\{\rho \mid \rho \in \text{Reach}(U_{-T})\}) > 0$ , and we know that there exists a finite path  $r \in \text{Paths}_*^{\sigma'}(s)$  (with positive probability) such that  $\text{last}(r) \in U_{-T}$ . Given the definition of the behaviour of  $\sigma'$  after finite path  $r$ , we have  $\Pr_s^{\sigma'}(\{\rho \mid r \text{ is a prefix of } \rho \wedge \rho \notin \text{Reach}(S \setminus C_{\text{last}(r)}) \wedge \text{infst}(\rho) = C_{\text{last}(r)}\}) > 0$ . This then implies that  $\Pr_s^{\sigma'}(\{\rho \mid \rho \notin \text{Reach}(S \setminus C_{\text{last}(r)}) \wedge \text{infst}(\rho) = C_{\text{last}(r)}\}) > 0$ . Because we have assumed that all states in  $T$  are absorbing, the set  $T$  is not reached along  $r$  (otherwise  $U_{-T}$  could not be the final state of  $r$ ); then, given that  $C_{\text{last}(r)} \cap T = \emptyset$ , we have that  $\Pr_s^{\sigma'}(\{\rho \mid \rho \notin \text{Reach}(T)\}) > 0$ , and by duality  $\Pr_s^{\sigma'}(\text{Reach}(T)) < 1$ . Hence this direction of the proof is completed.  $\square$