



HAL
open science

Long-run Satisfaction of Path Properties

Christel Baier, Nathalie Bertrand, Jakob Piribauer, Ocan Sankur

► **To cite this version:**

Christel Baier, Nathalie Bertrand, Jakob Piribauer, Ocan Sankur. Long-run Satisfaction of Path Properties. LICS 2019 - 34th Annual ACM/IEEE Symposium on Logic in Computer Science, Jun 2019, Vancouver, Canada. pp.1-31, 10.1109/LICS.2019.8785672 . hal-02349456

HAL Id: hal-02349456

<https://hal.science/hal-02349456v1>

Submitted on 20 Nov 2019

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

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

Long-run Satisfaction of Path Properties

(extended version)

Christel Baier¹, Nathalie Bertrand², Jakob Piribauer¹, Ocan Sankur²

¹Technische Universität Dresden, Germany

²Univ Rennes, Inria, CNRS, IRISA, France

Abstract—The paper introduces the concepts of long-run frequency of path properties for paths in Kripke structures, and their generalization to long-run probabilities for schedulers in Markov decision processes. We then study the natural optimization problem of computing the optimal values of these measures, when ranging over all paths or all schedulers, and the corresponding decision problem when given a threshold. The main results are as follows. For (repeated) reachability and other simple properties, optimal long-run probabilities and corresponding optimal memoryless schedulers are computable in polynomial time. When it comes to constrained reachability properties, memoryless schedulers are no longer sufficient, even in the non-probabilistic setting. Nevertheless, optimal long-run probabilities for constrained reachability are computable in pseudo-polynomial time in the probabilistic setting and in polynomial time for Kripke structures. Finally for co-safety properties expressed by NFA, we give an exponential-time algorithm to compute the optimal long-run frequency, and prove the PSPACE-completeness of the threshold problem.

I. INTRODUCTION

While the standard semantics of temporal logics relies on Boolean truth values for formulas over system models, several approaches have been studied to quantify how well a system model satisfies a temporal formula. This includes work on the robust satisfaction of temporal specifications [31], [36], vacuity and coverage semantics [15], [16], [30], [32], robustness distances [10] and the more general model-measurement semantics based on automatic distance functions of [26]. Another direction attempts to measure the degree to which a specification is satisfied when evolving over time. This includes, *e.g.*, the work on frequency LTL [7] where a quantitative variant $aU_q b$ of the until operator relaxes the standard meaning of $aU b$ by requiring that a holds at a fraction q or more of the positions before b holds. Other variants of frequency LTL [23], [24] allow only a quantitative variant \Box_q of the globally operator. The

The authors are partly supported by the DFG through the Collaborative Research Centers CRC 912 (HAEC) and TRR 248 (see <https://perspicuous-computing.science>, project ID 389792660), the Cluster of Excellence EXC 2050/1 (CeTI, project ID 390696704, as part of Germany's Excellence Strategy), the Research Training Groups QuantLA (GRK 1763) and RoSI (GRK 1907), the DFG-project BA-1679/11-1, and the DFG-project BA-1679/12-1. The collaboration is supported by the Inria associated team QuantProb.

semantics here is that $\Box_q \phi$ holds on a path if the long-run average of the frequency of positions at which ϕ holds is at least q . Alternatively, averaging LTL [8] rather than truth values, assigns quantities to pairs of paths and formula. It is based on a quantitative labeling function for atomic propositions and inductively defines the semantics of $\Box \phi$ as the average of the value of ϕ along the path. A notable similarity of these two quantitative extensions of LTL is the undecidability of the model checking problem of the full logics [7], [8]. Decidable fragments of frequency LTL can be obtained by restricting the nesting of temporal operators or the allowed frequency thresholds [7], [23], [24].

Following the spirit of quantifying the validity of a property along a path, we introduce the notion of *long-run frequencies* for ω -regular properties. Phrased in averaging LTL words, no nesting of the averaging operators is allowed, and the labeling function is Boolean. As the name suggests, long-run frequencies measure in the long-run how frequently a property holds. For finite-state Kripke structures (KS), we study the optimization over all paths of the long-run frequency of a given property.




Fig. 1. Kripke structures requiring memory to maximize the long-run frequency of $aU b$.

Fig. 1 gives two examples of KS on which one wants to evaluate the long-run frequency of an until property. Here a, b, c stand for atomic propositions. For the KS on the left, the long run frequency of $aU b$ along *e.g.* the path $(abc)^\omega$ is $\frac{2}{3}$. The maximal long-run frequency is 1, which is achieved, *e.g.*, by the infinite path $abca^2bca^4bca^8bc\dots$ that successively doubles the number of times the self-loop at state a is taken. However, there is no finite-memory strategy for generating an infinite path where the long-run frequency for $aU b$ is 1. The KS on the right illustrates, that, even when infinite-memory is not needed, memoryless is not

enough: for aUb , the maximal long-run frequency is achieved by alternating between the two simple cycles and amounts to $\frac{4}{9}$, which is indeed more than $\frac{2}{5}$ the long-run frequency of iterating the bottom cycle only.

When turning to the probabilistic world, we introduce the corresponding concept of *long-run probabilities*. On Markov chains, long-run probabilities are limit-average probabilities for path properties, indicating the probability for a property to hold on the suffix of a path after many steps. They can, among others, serve to provide refined measures for the system availability, understood as the proportion of time a system is functioning under “normal” operating conditions (after the initialization phase). For finite Markov decision processes (MDP), the corresponding optimization problem is to compute the optimal long-run probability of a given property, when ranging over all schedulers, or to decide how this value compares to a threshold.

To illustrate the notion of long-run probability, consider the MDP \mathcal{N}_k shown in Fig. 2, the only non-determinism is between actions α and β , and α yields a uniform distribution over the three successors.




Fig. 2. MDP with labels indicated by the state names requiring counting to maximize the long-run probability of aUb .

Under the memoryless scheduler \mathfrak{S}_α that always picks action α , the probability of aUb in the a -state is $\frac{1}{2}$, and its frequency is $\frac{3}{5}$. The state b_1 has frequency $\frac{1}{5}$ and from there the probability of aUb is 1. We thus compute the long-run probability under \mathfrak{S}_α to be $\frac{1}{2}$. Similarly, the steady-state probability of the states a and b_2 under the memoryless scheduler \mathfrak{S}_β are $\frac{1}{k+2}$, and the probability that aUb holds from there is 1. The long-run probability of aUb under \mathfrak{S}_β equals $\frac{2}{k+2}$. Observe that the satisfaction probability at the a -state depends on the scheduler (it is $\frac{1}{2}$ for \mathfrak{S}_α and 1 for \mathfrak{S}_β). The intricate interaction of satisfaction probability and frequency of each state makes the optimization of long-run probability particularly challenging. Here, we will see that counting the number of consecutive visits to the a -state allows one to derive a scheduler that achieves a higher long-run probability than the two memoryless ones.

Contributions: Beyond the introduction of the notions of long-run frequency and long-run probability, in this paper, we establish complexity bounds for the computation of the value (optimal long-run frequency or probability) and the associated threshold problem. These are summarized in Table I, split depending on

the type of properties and the models (KS or MDP). In particular, computing the optimal long-run probability for simple properties (such as reachability, invariant, Rabin or Streett conditions) can be done in polynomial time for MDP, in which case optimal memoryless schedulers exist. This entails the same complexity upper bound for the particular case of computing long-run frequency of these simple properties in KS. Moreover, the computation of the optimal long-run frequency in KS for constrained reachability properties (expressed by until formulas) can also be done in polynomial time, although, as explained already, infinite memory can be necessary. Our main contribution for non-probabilistic structures, is an exponential time algorithm for computing optimal long-run frequency for regular co-safety properties (specified by an NFA). It is obtained by reducing to the computation of the optimal mean-payoff in an exponentially large weighted KS. We also prove the PSPACE-completeness of the corresponding threshold problem.

In comparison, the probabilistic setting is substantially harder, already for constrained reachability properties, expressed by until formulas. As in our illustrating example, and contrary to the case of the simple properties mentioned above, maximizing in each state the probability that an until property holds does not yield the maximal long-run probability. Also when finite-memory schedulers are optimal, as opposed to the non-probabilistic case where memory with two modes suffices, a counter up to some bound that depends on the size of the description of the MDP is needed. Proving the existence of this *saturation point* (the ideal number of consecutive visits to a -states for the property aUb) is the crux to derive our pseudo-polynomial time algorithm for computing maximal or minimal long-run probabilities. The corresponding threshold problem is shown to be NP-hard. These two results certainly constitute the most involved contribution of the paper. We also show that the corresponding questions for qualitative threshold problems (*e.g.*, whether the maximal long-run probability for an until property is positive, or is 1) are solvable in polynomial time.

Related work: We mentioned the quantitative semantics of LTL and the decidable fragments of frequency LTL with a quantitative globally operator \Box_q [23], [24] which are the closest to our work. Frequency LTL, however, is a logic to specify quantitative measures for the satisfaction of properties along paths using the \Box_q -modality, while long-run probabilities are a quantitative measure across behaviors. For finite strongly connected Markov chains, the probabilities for \Box_q -formulas are 0 or 1, while long-run probabilities can be strictly between 0 and 1. There is still a connection as for each finite, strongly connected Markov chain \mathcal{M} , $\Box_q(aUb)$ holds in \mathcal{M} with probability 1 iff the long-run probability of aUb is at least q . Nevertheless,

	<i>non-probabilistic case</i> Kripke structures	<i>probabilistic case</i> Markov decision processes
reachability/invariant/ Rabin/Streect conditions	value computable in polynomial-time (special case of Theorem IV.3)	value computable in polynomial-time (Theorem IV.3)
constrained reachability ($\alpha \cup b$)	value computable in polynomial-time (Corollary III.3)	qualitative decision problems in polynomial-time (Lemma IV.5) value computable in pseudo-polynomial time (Theorem IV.10) NP-hard threshold problem (Theorem IV.11)
regular co-safety (NFA)	value computable in exponential time (Corollary III.6) PSPACE-complete threshold problem (Theorem III.7)	computability of the value: open PSPACE-hard threshold problem (consequence of Theorem III.7)

TABLE I
SUMMARY OF THE MAIN RESULTS.

the contribution for MDPs in [23], [24] are orthogonal to ours. On the one hand, they can treat much more complex properties with nested \Box_q -formulas. On the other hand, they cannot deal with formulas of the type $\Box_q(\alpha \cup b)$ for $q < 1$. The results in [23] only apply to $q = 1$. The fragment in [24] can deal with \Box_q -modalities for arbitrary q , but imposes the constraint that no until operator occurs in the scope of the \Box_q -modality.

Despite many works on long-run properties in MDPs (*e.g.*, mean payoff [9], [11], [27] and other cost objectives [20] or ratios [19], [37]), we are not aware that long-run probabilities for MDPs have been studied before. Long-run probabilities can be seen as mean-payoff, where the weights are the satisfaction probabilities. A crucial difference however with mean-payoff and other long-run properties is that, for long-run probabilities, the “weights” along a path are not fixed *a priori*, but do depend on the scheduler. In this aspect, there is some conceptual relation to dynamic Markov processes [34] where cost or transition probabilities depend on previously made decisions, or the stochastic variant of the Canadian traveler problem [25]. These problems, however, are concerned with finite-horizon objectives; moreover, their weights are affected by the past, whereas our “weights” (satisfaction probabilities) are induced by the future scheduler.

Outline: Section II summarizes the notations used in the paper. Our results for non-probabilistic systems are presented in Section III, while Section IV discusses long-run probabilities in MDPs. We conclude in Section V. Proofs can be found in the appendix.

II. PRELIMINARIES

We suppose familiarity with linear temporal logic (LTL), Kripke structures, finite automata, and basic concepts of discrete Markovian models, and only provide a summary of the notations used in the paper. Details can be found in textbooks, *e.g.*, [3], [17], [35].

Nondeterministic finite automata (NFA): An NFA is a tuple $\mathcal{A} = (Q, \Sigma, \delta, Q_0, F)$ where Q is a finite set of states, Σ an alphabet, $\delta \subseteq S \times \Sigma \times S$ the transition relation, $Q_0 \subseteq Q$ the set of initial states and $F \subseteq Q$ the set of final states. $\mathcal{L}(\mathcal{A})$ is the accepted language of \mathcal{A} .

Kripke structures (KSs): A KS is a tuple $\mathcal{T} = (S, \Delta, AP, L)$ where (S, Δ) is a finite directed graph, AP a finite set of atomic propositions and $L : S \rightarrow 2^{AP}$ a labeling function. The *trace* of a path $\pi = s_0 s_1 s_2 \dots$ is the word $L(\pi) = L(s_0)L(s_1)L(s_2)\dots$ over 2^{AP} obtained by projecting states to their labels. If $\pi = s_0 s_1 s_2 \dots$ is a path then we write $first(\pi)$ for its first state s_0 , $\pi_{[i]}$ for the $(i+1)$ -st state s_i , and $\pi_{[i..j]}$ for the path fragment $s_i s_{i+1} \dots s_j$. Likewise, $\pi_{[0..i]}$ and $\pi_{[i..]}$ stand for the prefix ending in state s_i resp. the suffix starting from s_i . If $T \subseteq S$, a T -state is a state in T , and a T -cycle is a cycle consisting of T -states.

We shall use LTL-like notations to denote path properties. For instance, if T is a set of states then $\Diamond T$ stands for the event “eventually reaching T ”, $\Box T$ for the event “always T ” and $\Box \Diamond T$ stands for Büchi condition “infinitely often T ”. The modality U stands for the standard until operator. Likewise, CTL-like notations are used for state properties, *e.g.*, $s \models \exists \Diamond T$ indicates that a T -state is reachable from state s .

Weighted structures: A *weighted KS* extends a plain KS \mathcal{T} as above by a weight function $wgt : \Delta \rightarrow \mathbb{Q}$ that assigns rational values to transitions. We also use weight functions on states (rather than transitions), which can be seen as a special case of transition-based weight functions. Given an infinite path $\pi = s_0 s_1 \dots$ in a weighted KS \mathcal{T} , the mean payoff $mp(\pi)$ is defined by:

$$mp(\pi) = \liminf_{n \rightarrow \infty} \frac{1}{n+1} \cdot \sum_{i=0}^n wgt(s_i) .$$

The maximal mean payoff from state s , $\mathbb{MPP}_{\mathcal{T},s}^{\max}$, is $\sup_{\pi} mp(\pi)$ where π ranges over all infinite paths starting at s . For $I \subseteq S$ we define $\mathbb{MPP}_{\mathcal{T},I}^{\max} = \max_{s \in I} \mathbb{MPP}_{\mathcal{T},s}^{\max}$. Analogous notations are used for the minimal mean payoff.

Markov decision processes (MDPs): An MDP is a tuple $\mathcal{M} = (S, Act, P)$ where S is a finite state space, Act a finite nonempty set of actions and $P : S \times Act \times S \rightarrow [0, 1] \cap \mathbb{Q}$ the transition probability function satisfying $\sum_{t \in S} P(s, \alpha, t) \in \{0, 1\}$ for all $(s, \alpha) \in S \times Act$. The triples (s, α, t) with $P(s, \alpha, t) > 0$ are called transitions of \mathcal{M} . The actions enabled in s form $Act(s) = \{\alpha \in Act : \sum_{t \in S} P(s, \alpha, t) = 1\}$. We often refer to the pairs (s, α) with $\alpha \in Act(s)$ as the state-action pairs of \mathcal{M} . MDPs can be augmented with atomic propositions labeling states and weight functions, as for KSs. The *size* of an MDP is the number of states and actions plus the sum of the logarithmic lengths of the transition probabilities. For a weighted MDP, we add the sum of the logarithmic lengths of the (numerator and denominator of) all weights.

Intuitively, when \mathcal{M} is at a state s , then an action α of $Act(s)$ is selected nondeterministically; afterwards the next state is obtained by probabilistically choosing one of the potential successor states according to the probability distribution $P(s, \alpha, \cdot)$. Paths in MDP are alternating sequences of states and actions: $\pi = s_0 \alpha_0 s_1 \alpha_1 \dots$ where $\alpha_i \in Act(s_i)$ and $P(s_i, \alpha_i, s_{i+1}) > 0$ for all $i \geq 0$. The notations $first(\pi)$, $\pi_{[i]}$, $\pi_{[i..j]}$ are defined as for paths in KS. *End components* (ECs) of an MDP are strongly connected sub-MDPs: they are formed of sets of state-action pairs where the induced graph is strongly connected. *Maximal ECs* (MECs) are ECs not contained in other ECs.

A *scheduler* for \mathcal{M} is a function \mathfrak{S} that maps a finite path ω to a probability distribution over $Act(last(\omega))$ (where $last(\omega)$ is the last state of ω). Given a finite path ω , then $\mathfrak{S} \uparrow \omega$ denotes the *residual scheduler* defined by $(\mathfrak{S} \uparrow \omega)(\omega') = \mathfrak{S}(\omega; \omega')$ for each finite path ω' starting in $last(\omega)$. \mathfrak{S} is *deterministic* if $\mathfrak{S}(\omega)$ is a Dirac distribution for each finite path, in which case \mathfrak{S} can be viewed as a function that maps finite paths to actions. *Finite-memory* schedulers are those that can be realized using a finite-state automaton whose states, called *modes*, serve as memory cells (see e.g. [3] for the formal definition). The decisions of *memoryless* schedulers only depend on the last state. They can be viewed as functions from states to distributions over actions. We write HR for the full class of (history-dependent randomized) schedulers, FMR or briefly FM (resp. FMD) for the class of finite-memory randomized (resp. deterministic) schedulers and MD for the class of memoryless deterministic schedulers.

$\Pr_{\mathcal{M}, s}^{\mathfrak{S}}$ denotes the probability measure induced by \mathfrak{S} , when s is the initial state. It is well-known that all ω -regular path properties φ are measurable and there exist FMD-schedulers maximizing or minimizing the probability for φ . This justifies the notations $\Pr_{\mathcal{M}, s}^{\max}(\varphi)$ and $\Pr_{\mathcal{M}, s}^{\min}(\varphi)$ for ω -regular properties. We consider the maximal (resp. minimal) expected mean payoff, denoted $E_{\mathcal{M}, s}^{\max}(\text{MIP})$ and $E_{\mathcal{M}, s}^{\min}(\text{MIP})$, where

the extrema are taken over all schedulers. It is well-known that optimal MD-schedulers for mean payoff objectives exist and are computable in polynomial time (see e.g. [35]). States belonging to the same MEC have the same maximal (resp. minimal) probability for reachability and prefix-independent properties and the same extremal expected mean payoff. This justifies notations like $\Pr_{\mathcal{E}}^{\max}(\diamond b)$ or $E_{\mathcal{E}}^{\max}(\text{MIP})$ for ECs \mathcal{E} . As shown in [19], for each scheduler \mathfrak{S} and each state s , the limit of almost all infinite paths constitutes an EC.

Remark II.1 (Traps). When studying long-run frequencies and long-run probabilities, we assume the given KSs and MDPs have no traps, *i.e.*, states with no outgoing transitions. This ensures the existence of infinite paths from every state. The presented reductions, however, can generate structures with traps, in which case optimal strategies in the constructed structures need to avoid these (trap) states.

III. LONG-RUN FREQUENCIES IN NON-PROBABILISTIC SYSTEMS

Let $\mathcal{T} = (S, \Delta, AP, L)$ be a KS and φ a path property. The *long-run frequency* for φ along an infinite path π of \mathcal{T} is defined as:

$$lrf_{\varphi}(\pi) = \liminf_{n \rightarrow \infty} \frac{1}{n+1} \cdot \sum_{i=0}^n \mathbb{1}_{\pi_{[i..]} \models \varphi}$$

where $\mathbb{1}_{\pi_{[i..]} \models \varphi}$ is 1 if $\pi_{[i..]} \models \varphi$ and 0 otherwise. The problem we address is how to compute the *maximal long-run frequency* for φ given by

$$\mathbb{L}\mathbb{F}_{\mathcal{T}, s}^{\max}(\varphi) = \sup_{\pi} lrf_{\varphi}(\pi)$$

where $s \in S$ and π ranges over all infinite paths starting in state s , and the analogously defined *minimal long-run frequency* $\mathbb{L}\mathbb{F}_{\mathcal{T}, s}^{\min}(\varphi)$.

Obviously, the value $lrf_{\varphi}(\pi)$ is prefix-independent, and all states belonging to the same strongly connected component (SCC) of \mathcal{T} have the same extremal values. It thus suffices to determine the optimal values for the SCC of \mathcal{T} . The optimal value for a given state s of \mathcal{T} is then the maximum or minimum of the optimal values for the SCCs reachable from s . In the sequel we therefore assume \mathcal{T} is strongly connected, and simply write $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\max}(\varphi)$ and $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\min}(\varphi)$.

As a consequence of a result established later for the probabilistic setting (see Theorem IV.3), the extremal long-run frequencies for invariants, reachability, Rabin and Streett conditions are computable in polynomial time. For KS, these techniques essentially require to identify “good” cycles ξ , where the property under consideration holds from all states along ξ .

Reasoning about long-run frequencies becomes more challenging when considering properties that are not

prefix-independent and where a classification of cycles into good and bad ones is not sufficient. The example from Fig. 1 left, illustrates this phenomenon as already for apparently simple properties such as constrained reachability, maximizing the long-run frequency may require infinite memory. As a stepping stone in this direction, we consider until properties, and regular co-safety properties where satisfaction is witnessed by “good” prefixes. We first give a polynomial time algorithm for computing optimal long-run frequencies for until properties. The main contribution for long-run frequencies in KSs, is then the PSPACE-completeness for co-safety properties specified by NFA.

A. Long-run frequencies for until properties

Before presenting the polynomial-time computation scheme for extremal long-run frequencies of until properties, we introduce useful notations, and identify easy particular cases. Fix a strongly connected KS \mathcal{T} , and the until property $\varphi = aUb$ with $a, b \in AP$. Let

$$\begin{aligned} A &= \{s \in S \mid s \models \exists(aUb), s \not\models \forall(aUb)\}, \\ B &= \{s \in S \mid s \models \forall(aUb)\}, \quad C = S \setminus (A \cup B). \end{aligned}$$

If $B = \emptyset$, $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\max}(aUb) = \mathbb{L}\mathbb{F}_{\mathcal{T}}^{\min}(aUb) = 0$, thus we assume $B \neq \emptyset$. The until properties aUb and AUB are equivalent: for each infinite path π of \mathcal{T} , $\pi \models aUb$ iff $\pi \models AUB$, and therefore $lrf_{aUb}(\pi) = lrf_{AUB}(\pi)$.

Using these sets, we first characterize easy cases:

Lemma III.1. $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\max}(aUb) = 1$ iff \mathcal{T} has a AUB -cycle. $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\min}(aUb) = 0$ iff \mathcal{T} has a $A \cup C$ -cycle.

As a consequence of Lemma III.1 (see Lemma A.1), if \mathcal{T} contains an A -cycle then the values of $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\max}(aUb)$ and $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\min}(aUb)$ are respectively 1 and 0. This observation permits the additional assumption that \mathcal{T} has no A -cycles.

The computation of $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\max}(aUb)$ and $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\min}(aUb)$ reduces to the computation of the extremal mean-payoff in a polynomial size weighted KS. The idea of the reduction is to deal with two copies of A -states. Intuitively, when entering an A -state there is a nondeterministic choice between mode 0 and mode 1 where the former expects AUB to be violated, while the latter expects some B -state will be visited after consecutive A -states. From any state $(s, 0)$ with $s \in A$, the transitions to B -states are removed. Likewise, from states $(s, 1)$ with $s \in A$ there are no transitions to C -states. Since \mathcal{T} has no A -cycles, the sub-graph consisting of states in $A \times \{0, 1\}$ is acyclic. The weight function in \mathcal{T}' is state-based and assigns weight 1 to the states in $B \cup A \times \{1\}$ and weight 0 to the states in $C \cup A \times \{0\}$. The construction is illustrated on Fig. 3 for the KS right of Fig. 1. The soundness of the construction is stated in the following lemma (see Appendix, Lemma A.2 for its proof).




Fig. 3. Weighted Kripke structures obtained for the KS right of Fig. 1: gray states have weight 1, while others have weight 0.

Lemma III.2. Suppose \mathcal{T} has no A -cycles. Then, $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\max}(aUb) = \mathbb{M}\mathbb{P}_{\mathcal{T}'}^{\max}$ and $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\min}(aUb) = \mathbb{M}\mathbb{P}_{\mathcal{T}'}^{\min}$.

Thus, the extremal long-run frequencies for until properties are computable using standard techniques for weighted KS with mean payoff objectives. This yields:

Corollary III.3. The values $\mathbb{L}\mathbb{F}_{\mathcal{T},s}^{\max}(aUb)$ and $\mathbb{L}\mathbb{F}_{\mathcal{T},s}^{\min}(aUb)$ are computable in polynomial time.

Remark III.4 (Memory requirements). The example from Fig. 1 left in the introduction shows that infinite memory can be necessary to optimize the long-frequency for until-properties. If, however, \mathcal{T} does not contain A -cycles, then two memory cells suffice for optimizing the long-run frequency for AUB . (This follows from Lemma III.2 and the well-known fact that optimal memoryless strategies for mean payoff objectives exist.) This stands in great contrast with the probabilistic case, as we will see in Section IV.

B. Regular co-safety properties

We now address extremal long-run frequencies in KS for regular co-safety properties φ . Fix \mathcal{T} a strongly connected KS, and let $\mathcal{A} = (Q, \Sigma, \delta, Q_0, F)$ be an NFA over the alphabet $\Sigma = 2^{AP}$ representing φ , i.e., an infinite path of \mathcal{T} satisfies φ iff it has a prefix accepted by \mathcal{A} . Hence, $lrf_{\varphi}(\pi)$, also denoted $lrf_{\mathcal{A}}(\pi)$, is the long-run average of positions in π where a word in $\mathcal{L}(\mathcal{A})$ starts, and we write $\mathbb{L}\mathbb{F}_{\mathcal{T}}^*(\mathcal{A})$ rather than $\mathbb{L}\mathbb{F}_{\mathcal{T}}^*(\varphi)$. We show that the computation of $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\max}(\mathcal{A})$ and $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\min}(\mathcal{A})$ reduces to determine the maximal and minimal mean-payoff in a weighted KS \mathcal{G} with a generalized Büchi side condition, whose size is exponential in the size of \mathcal{T} .

For simplicity, we suppose here that $Q_0 = \{q_0\}$ is a singleton and that q_0 is not accessible from any other state in \mathcal{A} . We also assume that $q_0 \notin F$ (otherwise \mathcal{A} accepts the empty word and the long-run frequency of φ along any infinite path is 1). We fix an arbitrary state $s_0 \in S$ which we treat as a starting state of \mathcal{T} . (Since \mathcal{T} is strongly connected, the extremal long-run frequencies in \mathcal{T} do not depend on the choice of the starting state.) We define a weighted KS \mathcal{G} as follows.

Let $\ell = |Q|$ denote the number of states in the NFA \mathcal{A} . Then, the state space $S_{\mathcal{G}}$ of \mathcal{G} is equal to

$$S \times \{\{\emptyset\} \cup \{\text{merged}\} \times \{0, \dots, \ell\} \cup Q \times \{\text{true}, \text{false}\}\}^{2\ell+1}.$$

Each state of \mathcal{G} tracks a state of \mathcal{T} and in addition a vector of $2\ell+1$ elements, which we explain now. For a state $(s, f) \in S_{\mathcal{G}}$, $f(i)$ denotes the i -th component and is called the i -th *track*. Intuitively, at every step of the execution, the construction guesses whether φ holds from the current position and creates a new track to check this guess. Hence, the symbol \emptyset is used for tracks that are not in use; (q, true) means that the track is at state q of \mathcal{A} and that the final state of \mathcal{A} will be eventually reached (confirming the guess that φ holds), while (q, false) means that no final state of \mathcal{A} should be reached at this track (expressing the guess that φ does not hold). The symbol *merged* is used to merge tracks that arrive at the same states of \mathcal{A} with the same guess. The formal details follow.

The set $I_{\mathcal{G}}$ of initial states in \mathcal{G} is given by:

$$I_{\mathcal{G}} = \left\{ (s_0, f) \mid f(0) \in \{(q_0, \text{true}), (q_0, \text{false})\} \text{ and } f(i) = \emptyset \text{ for } 1 \leq i \leq 2\ell \right\}$$

That is, a single track is created at initial state q_0 of \mathcal{A} that is either true or false. The non-determinism in the choice of the initial state allows the execution to guess whether φ holds at this position. The definition of transitions also uses non-determinism for two purposes: (i) to guess whether φ holds at the given step and (ii) to guess a successor state in \mathcal{A} . Note that for each track of the form (q, true) , the construction will guess *some* successor q' of q in \mathcal{A} in an attempt to prove the guess correct, while for each track of the form (q, false) , it will check *all* successors of q in order to prove that the word that is being read is not accepted by \mathcal{A} .

We define the transitions in two steps. The first step consists in allowing the tracks to make progress non-deterministically, while in the second step we merge tracks at identical states and add the new track q_0 for the current position.

For any state (s, f) in \mathcal{G} and each state s' of \mathcal{T} , we define relation \rightsquigarrow_1 as follows: we have $(s, f) \rightsquigarrow_1 (s', f', H)$ where $H = \bigcup_{q: \exists i. f(i) = (q, \text{false})} \delta(q, L(s))$, if, and only if the following holds for all $0 \leq i \leq 2\ell$,

- If $f(i) \in \{\emptyset, (\text{merged}, j)\}$ for some j , then $f'(i) = \emptyset$,
- If $f(i) = (q, b)$ for some $(q, b) \in (Q \setminus F) \times \{\text{true}, \text{false}\}$, then $f'(i) = (q', b)$ where $q' \in \delta(q, L(s))$.
- If $f(i) = (q, \text{true})$ for some $q \in F$, then $f'(i) = \emptyset$.
- If $f(i) = (q, \text{false})$ for some $q \in F$, then $f'(i) = (q, \text{false})$.

Here, the additional component H is the set of all successors in \mathcal{A} of states q which appear in f as (q, false) . In fact, these must be checked for non-acceptance, and

in the next step, new tracks (q, false) will be created for each $q \in H$ unless they already exist or they contradict another guess (*i.e.* if some track with (q, true) exists).

The intermediary state (s', f', H) is called *inconsistent* if there exists $0 \leq i, j \leq 2\ell$ and $q \in Q$ such that one of the conditions hold:

- 1) $f'(i) = (q, \text{false})$ with $q \in F$,
- 2) $f'(i) = (q, \text{true})$ and $f'(j) = (q, \text{false})$,
- 3) $f'(i) = (q, \text{true})$ for some $q \in H$.

Assume (s', f', H) is consistent. Let us write $H' = \{q \in H \mid \forall i. f'(i) \neq (q, \text{false})\} = \{q'_1, \dots, q'_k\}$. We define the relation \rightsquigarrow_2 by $(s', f', H) \rightsquigarrow_2 (s'', f'')$ if, and only if, the following conditions hold for all $0 \leq i \leq 2\ell$:

Case $f'(i) = \emptyset$: if i is the first such position then $f''(i) = (q_0, b)$ for some $b \in \{\text{true}, \text{false}\}$, if it is the $l+1$ -th such position, we let $f''(i) = q'_1$, and otherwise $f''(i) = \emptyset$.

Case $f'(i) = (q, b)$ for some $(q, b) \in Q \times \{\text{true}, \text{false}\}$: if for all $0 \leq j < i$, $f'(j) \neq (q, b)$, then $f''(i) = (q, b)$; and otherwise, $f''(i) = (\text{merged}, j)$ where j is the least index of a track containing (q, b) .

Hence, the second step is used to stop all redundant tracks; for each pair (q, b) only the first track survives, and others are stopped and labeled by (merged, j) with j being the position of the surviving track. Making the positions j visible in the construction is not necessary but will be useful for the proofs. In addition, all states $q \in H$ that do not already appear in some track are added as a new track (q, false) , and a new track is started from q_0 .

The transitions of \mathcal{G} are defined as follows. We have $((s, f), (s', f'')) \in \Delta_{\mathcal{G}}$ if and only if $(s, s') \in \Delta$ and there exists a consistent state (s', f', H) such that $(s, f) \rightsquigarrow_1 (s', f', H) \rightsquigarrow_2 (s'', f'')$.

The weight function of \mathcal{G} is state-based and given by $wgt(s, f) = 1$ if there exists some i with $f(i) = (q_0, \text{true})$, and $wgt(s, f) = 0$ otherwise.

For each $0 \leq i \leq 2\ell$, let us define the following labels. Let F_i , false_i , merged_i , and \emptyset_i denote the states (s, f) satisfying, respectively, $f(i) = (q, \text{true})$ for some $q \in F$, $f(i) = (q, \text{false})$ for some $q \in Q$, $f(i) = (\text{merged}, j)$ for some $0 \leq j < i$, or $f(i) = \emptyset$.

We now show that the extremal long-run frequencies for co-safety properties are computable using the above construction in combination with techniques for one-player mean payoff games with generalized Büchi conditions. For this, we define Φ denote the following generalized Büchi condition:

$$\Phi = \bigwedge_{i=0}^{2\ell} \square \diamond (\text{false}_i \vee \text{merged}_i \vee F_i \vee \emptyset_i).$$

Let $\text{MIP}_{\mathcal{G}}^{\max}(\Phi)$ denote the maximal mean payoff in \mathcal{G} starting in one of the two initial states in $I_{\mathcal{G}}$ while sat-

isfying Φ . Similarly, $\text{MIP}_{\mathcal{G}}^{\min}(\Phi)$ denotes the minimal payoff in \mathcal{G} while satisfying Φ . The soundness of our construction is stated as follows.

Theorem III.5. *For \mathcal{G} and Φ defined as above, $\text{LRF}_{\mathcal{T}}^{\max}(\mathcal{A}) = \text{MIP}_{\mathcal{G}}^{\max}(\Phi)$ and $\text{LRF}_{\mathcal{T}}^{\min}(\mathcal{A}) = \text{MIP}_{\mathcal{G}}^{\min}(\Phi)$.*

We already explained that the tracks in \mathcal{G} guess and check the satisfaction or violation of φ from each position. The additional Büchi condition is used to make sure that these checks are conclusive: each track of the form (q, true) must eventually either end in an accepting state (F_i) or merge with a track with strictly lower index (which will itself reach an accepting state or merge with even smaller indexed-track). Other tracks may stay in (\cdot, false) or not be used (\emptyset_i). A detailed proof of the above theorem is given in Theorem A.4 in the Appendix.

Using [6], [28], one can compute extremal mean payoff values in one-player games in polynomial time in the size of the game. We apply this to \mathcal{G} which yields the following result (See also Appendix B):

Corollary III.6. *$\text{LRF}_{\mathcal{T},s}^{\max}(\varphi)$ and $\text{LRF}_{\mathcal{T},s}^{\min}(\varphi)$ are computable in time exponential in the size of \mathcal{A} and polynomial in the size of \mathcal{T} .*

We now establish the complexity of the decision problem associated with the maximization of long-run frequency for co-safety properties. Formally, given a KS \mathcal{T} , an NFA \mathcal{A} and a rational threshold ϑ , the threshold problem asks whether \mathcal{T} admits an infinite path π with $\text{lrf}_{\mathcal{A}}(\pi) \geq \vartheta$.

Theorem III.7. *The threshold problem for co-safety properties in KS is PSPACE-complete.*

Proof Sketch. The PSPACE upper bound is obtained by providing a nondeterministic guess-&-check algorithm that guesses a reachable state $s_{\mathcal{G}}$ in \mathcal{G} and two cycles ξ_{Φ} and ξ_{MP} containing $s_{\mathcal{G}}$ and satisfying certain length restrictions and $\xi_{\Phi} \models \Phi$ and $\text{mp}(\xi_{\text{MP}}) \geq \vartheta$. In fact, one can construct the desired path by alternating between ξ_{Φ} and ξ_{MP} by making sure the former appears infinitely often but with frequency 0.

The PSPACE lower bound follows by a polynomial reduction from the intersection problem for deterministic finite automata (DFA): given k DFA $\mathcal{D}_1, \dots, \mathcal{D}_k$ over the same alphabet Σ , is the intersection language $\mathcal{L}(\mathcal{D}_1) \cap \dots \cap \mathcal{L}(\mathcal{D}_k)$ nonempty? This problem is known to be PSPACE-complete [29]. A detailed proof is given in Theorem A.5 in the Appendix. \square

Proposition III.8 (Qualitative thresholds).

Deciding the existence of an infinite path π in \mathcal{T} with

- $\text{lrf}_{\mathcal{A}}(\pi) > 0$ is in PTIME;

- $\text{lrf}_{\mathcal{A}}(\pi) = 0$ (resp. $= 1$) is NP-hard;
- $\text{lrf}_{\mathcal{A}}(\pi) < 1$ is PSPACE-hard.

Proof sketch. The existence of an infinite path π in \mathcal{T} with $\text{lrf}_{\mathcal{A}}(\pi) > 0$ is equivalent to the existence of a finite path in (an SCC of) \mathcal{T} that is accepted by \mathcal{A} . The latter can be checked by performing a graph analysis of the product $\mathcal{T} \otimes \mathcal{A}$.

To prove the two NP-hardness results, we provide a polynomial reduction from 3SAT. The construction of the KS \mathcal{T} is the same in both reductions: it runs through each clause one after the other and selects a literal. The constructed automata slightly differ in the two cases. For example, for the existence of a path with long-run frequency 0, the automaton accepts behavior where the choice of literals to be true is conflicting.

For the last hardness result, we reduce the universality of NFA to the problem whether for all infinite paths π in a KS $\text{lrf}_{\mathcal{A}}(\pi) = 1$, and use the fact that PSPACE is closed under complementation. Given an NFA \mathcal{B} over the alphabet Σ , we construct an NFA \mathcal{A} over the alphabet $\Sigma \cup \{\#\}$ for a new symbol $\#$ such that \mathcal{A} accepts $\Sigma \cup \{\#w\# \mid w \in \mathcal{L}(\mathcal{B})\}$. The Kripke structure that generates all possible sequences of letters over $\Sigma \cup \{\#\}$ then has a path π with $\text{lrf}_{\mathcal{A}}(\pi) < 1$ iff \mathcal{B} is non-universal. See Lemma A.8 for the full proofs. \square

IV. LONG-RUN PROBABILITIES IN MDP

When turning to probabilistic models, long-run probabilities generalize long-run frequencies, and the objective is to compute the optimal values of these long-run averages, when ranging over all schedulers. Our first contribution is the identification of efficiently solvable instances, including prefix-independent properties (such as Rabin or Streett conditions) where the satisfaction only depends on the states that are visited infinitely often and that can be treated by a polynomial-time analysis of end components. Our second contribution, and main result of this section is a pseudo-polynomial algorithm for computing extremal long-run probabilities for constrained reachability (until properties $a \cup b$). This result can be seen as a first step towards the treatment of more general co-safety properties in MDP and serves to illustrate which major extra difficulties arise when switching from long-run frequencies in KS to long-run probabilities in MDP.

It is important to emphasize that the computation of optimal long-run probabilities for constrained reachability does not easily reduce to reachability via a pre-processing of the MDP, as it typically does for most verification problems. Also, the traditional reduction to the case of a Rabin condition for the treatment of arbitrary ω -regular properties fails here. These highlight the challenge and specificity in computing long-run probabilities.

Throughout this section, we suppose that we are given an MDP $\mathcal{M} = (S, Act, P, AP, L)$ and a path property φ . The *long-run probability* for φ of an infinite path π under a scheduler \mathfrak{S} for \mathcal{M} is defined as the long-run average of the probabilities for φ in all positions of π with respect to the residual schedulers $\mathfrak{S} \uparrow \pi_{[0..i]}$ (see Section II for the definition of these):

$$lrp_{\varphi}^{\mathfrak{S}}(\pi) = \liminf_{n \rightarrow \infty} \frac{1}{n+1} \cdot \sum_{i=0}^n \Pr_{\mathcal{M}, \pi_{[i]}}^{\mathfrak{S} \uparrow \pi_{[0..i]}}(\varphi).$$

The long-run probability for property φ under scheduler \mathfrak{S} from state s , denoted $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(\varphi)$, is defined as the expectation of the random variable $\pi \mapsto lrp_{\varphi}^{\mathfrak{S}}(\pi)$ under \mathfrak{S} with starting state s . We now address the task to compute the extremal long-run probabilities for φ :

$$\begin{aligned} \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\max}(\varphi) &= \sup_{\mathfrak{S}} \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(\varphi) \\ \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\min}(\varphi) &= \inf_{\mathfrak{S}} \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(\varphi) \end{aligned}$$

where \mathfrak{S} ranges over all schedulers for \mathcal{M} , and refer to this task as the *lrp-problem*. In contrast to classical optimization problems for MDPs, the random variable whose expectation we aim to optimize, namely $lrp_{\varphi}^{\mathfrak{S}}$, depends on the scheduler \mathfrak{S} .

Remark IV.1 (Long-run frequencies and probabilities). Computing extremal long-run frequencies for path properties in KSSs is a special case of the lrp-problem: KSSs can be viewed as MDPs where all distributions are Dirac. For such degenerated MDPs, if π is an infinite \mathfrak{S} -path then $\Pr_{\mathcal{M}, \pi_{[i]}}^{\mathfrak{S} \uparrow \pi_{[0..i]}}(\varphi) = \mathbb{1}_{\pi_{[i..]} \models \varphi}$. Hence, $lrf_{\varphi}(\pi) = lrp_{\varphi}^{\mathfrak{S}}(\pi)$, and therefore $\mathbb{L}\mathbb{F}_{\mathcal{M},s}^{\max}(\varphi) = \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\max}(\varphi)$ and $\mathbb{L}\mathbb{F}_{\mathcal{M},s}^{\min}(\varphi) = \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\min}(\varphi)$.

Analogously to extremal long-run frequencies in KSSs, we can assume MDPs to be strongly connected. Indeed for a general MDP \mathcal{M} without traps, one can compute the MECs $\mathcal{E}_1, \dots, \mathcal{E}_k$ in polynomial-time [13], [19]. Given a path formula φ , we define a weighted MDP \mathcal{N} , obtained from \mathcal{M} by collapsing each MEC \mathcal{E}_i into a single state, say s_i , which is duplicated into s'_i a trap state. The copy s'_i is reachable only from s_i by a Dirac action on a fresh action symbol. The weight function assigns $\mathbb{L}\mathbb{P}_{\mathcal{E}_i}^{\max}(\varphi)$ to state s'_i and 0 to all other states. With this construction, $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\max}(\varphi)$ equals the maximal expected accumulated weight in \mathcal{N} under all proper schedulers, *i.e.*, schedulers that reach the trap states almost surely. This value is computable in polynomial-time using standard algorithms for the stochastic shortest path problem [1], [4], [21].

When \mathcal{M} is strongly connected, the optimal long-run probabilities do not depend on the starting state and we simply write $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\varphi)$ and $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\min}(\varphi)$.

A. Efficiently solvable lrp instances

We first investigate special cases for which one can obtain efficient algorithms to compute optimal long-run probabilities: first, we explain the case of Markov chains, and then we identify restricted classes of properties for MDPs.

Remark IV.2 (Special case of Markov chains). A Markov chain can be seen as a degenerated MDP where each state has a unique enabled action, in which case the concept of schedulers is irrelevant and we simply write $lrp_{\varphi}(\pi)$ for each infinite path π . If φ is an ω -regular property, for each bottom strongly connected component (BSCC) \mathcal{B} of the Markov chain \mathcal{M} , the long-run probability for all states in \mathcal{B} is the same:

$$\mathbb{L}\mathbb{P}_{\mathcal{B}}(\varphi) = \sum_{t \in \mathcal{B}} lrf_{\mathcal{B}}(t) \cdot \Pr_{\mathcal{M},t}(\varphi)$$

where $lrf_{\mathcal{B}}(t)$ denotes the steady-state probability (defined as the long-run frequency) of state t in \mathcal{B} . Thus, $\mathbb{L}\mathbb{P}_{\mathcal{B}}(\varphi)$ equals the probability for φ in \mathcal{B} viewed as a Markov chain where the initial distribution is given by the long-run frequencies in \mathcal{B} , which again coincides with the expected mean payoff in \mathcal{B} when $\Pr_{\mathcal{M},t}(\varphi)$ is viewed as weight for state t . The long-run frequencies inside the BSCC are computable in polynomial-time using a linear equation system. The values $\Pr_{\mathcal{M},t}(\varphi)$ for the states inside the BSCC are computable using standard techniques for the analysis of Markov chains against ω -regular properties (see *e.g.* [3]). The complexity depends on the type and representation of φ : for instance, exponential-time algorithms exist for LTL formulas [18]. Thus, long-run probabilities for LTL-properties in Markov chains are computable in exponential time. Moreover, $\mathbb{L}\mathbb{P}_{\mathcal{M},s}(\varphi)$ is computable in polynomial time for those properties φ where $\Pr_{\mathcal{M},t}(\varphi)$ is computable in polynomial time, such as until, Rabin or Streett properties.

Alternatively, the computation of long-run probabilities as expected mean payoff when dealing with the weight function that assigns weight $wgt(s) = \Pr_{\mathcal{B},s}(\varphi)$ to each state s can also be written as a quotient of expectations as follows. Let s be an arbitrary state in \mathcal{B} , called *reference state*. Then, the long-run probability for φ in \mathcal{B} equals the quotient of the expected accumulated weight along paths of length at least 1 from s until returning to s and the expected return time (*i.e.*, expected number of steps) along such paths from s to s . (Strong connectivity ensures that both expectations are finite.) That is,

$$\mathbb{L}\mathbb{P}_{\mathcal{B}}(\varphi) = \frac{\mathbb{E}_{\mathcal{B},s}(\text{“weight until } s\text{”})}{\mathbb{E}_{\mathcal{B},s}(\text{“steps until } s\text{”})} \quad (\star)$$

Finally, if \mathfrak{B} denotes the set of all BSCCs of \mathcal{M} then

for each state s in \mathcal{M} :

$$\mathbb{L}\mathbb{P}_{\mathcal{M},s}(\varphi) = \sum_{\mathcal{B} \in \mathfrak{B}} \Pr_{\mathcal{M},s}(\diamond \mathcal{B}) \cdot \mathbb{L}\mathbb{P}_{\mathcal{B}}(\varphi).$$

We now identify classes of path formulas, for which the lrp-problem is solvable in polynomial time. This applies to (repeated) reachability and other properties φ where efficiently computable schedulers \mathfrak{S} exist that maximize the probabilities for φ from every visited state in the sense that

$$\Pr_{\mathcal{M},\pi_{[i]}^{\mathfrak{S}}}(\varphi) = \Pr_{\mathcal{M},\pi_{[i]}}^{\max}(\varphi)$$

for each infinite \mathfrak{S} -path π and each position $i \in \mathbb{N}$. Obviously, such a scheduler \mathfrak{S} maximizes the long-run probability for φ . (See Theorem A.9 and its proof.)

Theorem IV.3 (Efficiently solvable lrp-instances). *The values $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\max}(\varphi)$ and $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\min}(\varphi)$ are computable in polynomial-time if φ has the form $\diamond b$ (reachability), $\square b$ (invariant), $\bigwedge_{i=1}^n \bigvee_{j=1}^{\ell_i} (\square \diamond b_{i,j} \wedge \diamond \square a_{i,j})$ (generalized Rabin condition), or $\bigwedge_{i=1}^n (\square \diamond a_i \rightarrow \square \diamond b_i)$ (Streett condition). Moreover, FMD-schedulers are optimal for all these cases, and even MD-schedulers for reachability, invariant, Büchi and co-Büchi conditions.*

An analogous result for much richer classes of properties cannot be expected given that, already in the non-probabilistic setting, infinite-memory can be necessary for until properties (see Fig. 1 left), and co-safety properties yield PSPACE-hardness (see Theorem III.7).

B. Long-run probabilities for until properties

We now address the case of an until property aUb . While long-run frequencies for aUb in KSSs can be handled fairly easily using a reduction to classical mean payoff objectives, the probabilistic setting adds major extra challenges. Recall that in strongly connected KSSs, if no infinite-memory is required, then optimal long-run frequencies can be achieved by strategies operating in just two modes. This stands in contrast with the probabilistic setting where, in those cases where no infinite-memory is needed, optimal schedulers can need a counter for the number of consecutive a -states up to some bound that depends on \mathcal{M} (rather than a constant). This phenomenon is illustrated in the example from Fig. 2, with which we carry on now.

Example IV.4. For the MDP \mathcal{N}_k from in Fig. 2, we already argued that the two MD-schedulers \mathfrak{S}_α and \mathfrak{S}_β achieve long-run probabilities of $1/2$ and $\frac{2}{k+2}$ for the until property aUb . Thus, as soon as $k \geq 3$, \mathfrak{S}_α achieves the higher long-run probability, whereas \mathfrak{S}_β maximizes the probability of aUb from the a -state. However, none of them is optimal. To see this, consider the FMD-schedulers \mathfrak{T}_n for $n \geq 1$, that use a

counter for the number of consecutive visits to the a -state, starting with counter value 1 when entering that state via the transitions from the other states. When in the a -state, \mathfrak{T}_n schedules action α if the counter value is at most n and β otherwise.

We compute $\mathbb{L}\mathbb{P}_{\mathcal{N}_k}^{\mathfrak{T}_n}(aUb)$ via the quotient representation shown in (\star) in Remark IV.2 using that the Markov chain \mathcal{C}_n induced by \mathfrak{T}_n is strongly connected and consists of the states b_1, c_1, b_2, c_2 and the states $(a, 1), (a, 2), \dots, (a, n+1)$, where (a, i) means state a with counter value i . We pick $(a, 1)$ as reference state.

Let us first compute the denominator. The expected return time from $(a, 1)$ to $(a, 1)$ can be written as the sum of the expected frequencies $ef(s)$ of the states s in \mathcal{C}_n along the return paths from and to $(a, 1)$. These values are: $ef(a, i) = \frac{1}{3^{i-1}}$ for $i = 1, \dots, n+1$, $ef(b_1) = ef(c_1) = (1 - \frac{1}{3^n}) \cdot \frac{1}{2}$, $ef(b_2) = \frac{1}{3^n}$, and $ef(c_2) = \frac{k}{3^n}$.

Note that each of the states (a, i) , b_1 , c_1 and b_2 occurs exactly once on each return path from $(a, 1)$ to $(a, 1)$. Thus, for these states s , the expected frequency equals the probability of reaching s from the reference state $(a, 1)$. For state c_2 , we take into account that the self-loop is taken an expected $k-1$ -times. We conclude:

$$\begin{aligned} & \mathbb{E}_{\mathcal{C}_n, (a, 1)}(\text{“steps until } (a, 1)\text{”}) \\ &= \sum_{i=1}^{n+1} \frac{1}{3^{i-1}} + (1 - \frac{1}{3^n}) \cdot \frac{1}{2} \cdot 2 + \frac{1}{3^n} + \frac{k}{3^n} \\ &= \frac{1}{4} \cdot (10 + (4k-2) \cdot \frac{1}{3^n}) \end{aligned}$$

We now compute the expected accumulated weight along the return paths from and to $(a, 1)$ under scheduler \mathfrak{T}_n . This value can be computed as the sum of the expected frequency of every state s multiplied with its weight. In our case, the weights are the probabilities for aUb in \mathcal{C}_n . That is:

$$\begin{aligned} & \mathbb{E}_{\mathcal{C}_n, (a, 1)}(\text{“weight until } (a, 1)\text{”}) \\ &= \sum_s ef(s) \cdot \Pr_{\mathcal{C}_n, s}(aUb) \end{aligned}$$

where s ranges over all states in the Markov chain \mathcal{C}_n induced by \mathfrak{T}_n . The probability values are as follows: $\Pr_{\mathcal{C}_n, (a, i)}(aUb) = \frac{1}{2} \cdot (1 + \frac{1}{3^{n-i+1}})$ for $i = 1, \dots, n+1$, $\Pr_{\mathcal{C}_n, b_1}(aUb) = \Pr_{\mathcal{C}_n, b_2}(aUb) = 1$, and $\Pr_{\mathcal{C}_n, c_1}(aUb) = \Pr_{\mathcal{C}_n, c_2}(aUb) = 0$. So, we get:

$$\begin{aligned} & \mathbb{E}_{\mathcal{C}_n, (a, 1)}(\text{“weight until } (a, 1)\text{”}) \\ &= \sum_{i=1}^{n+1} \frac{1}{3^{i-1}} \cdot \frac{1}{2} \cdot (1 + \frac{1}{3^{n-i+1}}) + (1 - \frac{1}{3^n}) \cdot \frac{1}{2} \cdot 1 \\ &= \frac{1}{4} \cdot (5 + (2n+3) \cdot \frac{1}{3^n}), \text{ and} \\ & \mathbb{L}\mathbb{P}_{\mathcal{N}_k}^{\mathfrak{T}_n}(aUb) = \frac{5 + (2n+3) \cdot \frac{1}{3^n}}{10 + (4k-2) \cdot \frac{1}{3^n}}. \end{aligned}$$

Thus, \mathfrak{T}_n achieves a higher long-run probability than \mathfrak{S}_α if $n > k-2$. To determine which scheduler is optimal among the schedulers \mathfrak{T}_n with $n \in \mathbb{N}$, we determine the least natural number n such that $\mathbb{L}\mathbb{P}_{\mathcal{N}_k}^{\mathfrak{T}_n}(\mathfrak{a}U\mathfrak{b}) > \mathbb{L}\mathbb{P}_{\mathcal{N}_k}^{\mathfrak{T}_{n+1}}(\mathfrak{a}U\mathfrak{b})$. Treating the computed expression for $\mathbb{L}\mathbb{P}_{\mathcal{N}_k}^{\mathfrak{T}_n}(\mathfrak{a}U\mathfrak{b})$ as a real function in n , one can check that the derivative has only one root and that hence the obtained value n indeed yields the optimal scheduler among these schedulers. For $k \geq 2$, we obtain $n = k-1$.

We will see later that the maximal long-run probability of \mathcal{N}_k is indeed achieved by \mathfrak{T}_n for this n (see Remark IV.12 below). Note that \mathcal{N}_k has 5 states and its size is in $\mathcal{O}(\log k)$. So, this example illustrates that even in cases where finite memory is sufficient, the memory requirements of optimal schedulers can grow exponentially with the size of the MDP. The same applies to the logarithmic length of the optimal values. To see this, we observe that in

$$\mathbb{L}\mathbb{P}_{\mathcal{N}_k}^{\mathfrak{T}_{k-1}}(\mathfrak{a}U\mathfrak{b}) = \frac{5 \cdot 3^{k-1} + 2k + 1}{10 \cdot 3^{k-1} + 4k - 2}$$

the greatest common divisor of numerator and denominator is at most 4 (note that $2(5 \cdot 3^{k-1} + (2k+1)) - (10 \cdot 3^{k-1} + (4k-2)) = 4$). Therefore, the binary representation of the optimal value requires exponentially many bits in the size of \mathcal{N}_k . \square

After this example that highlighted the challenge in computing optimal long-run probabilities, in the remainder of this section, we provide a pseudo-polynomial time algorithm for computing extremal long-run probabilities for $\mathfrak{a}U\mathfrak{b}$. We first fix useful notations. Given a state s of \mathcal{M} , let $p_s^{\max} = \Pr_{\mathcal{M},s}^{\max}(\mathfrak{a}U\mathfrak{b})$, $p_s^{\min} = \Pr_{\mathcal{M},s}^{\min}(\mathfrak{a}U\mathfrak{b})$ and:

$$\begin{aligned} A &= \{s \in S \mid p_s^{\max} > 0 \text{ and } p_s^{\min} < 1\}, \\ B &= \{s \in S \mid p_s^{\min} = 1\}, \quad C = S \setminus (A \cup B). \end{aligned}$$

Then, $\Pr_{\mathcal{M},s}^{\mathfrak{S}}(\mathfrak{a}U\mathfrak{b}) = \Pr_{\mathcal{M},s}^{\mathfrak{S}}(A \cup B)$ for every s and \mathfrak{S} . Hence, we may safely assume that the labeling function fulfills $\mathfrak{a} \in L(s)$ iff $s \in A$ and $\mathfrak{b} \in L(s)$ iff $s \in B$. For $T \subseteq S$, a T-EC denotes an end component consisting of T-states.

Let us first consider the four qualitative variants of the lrp-problem where the objective is to decide the existence of a scheduler \mathfrak{S} such that $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(\mathfrak{a}U\mathfrak{b})$ equals 1, 0, is strictly less than 1 or is positive.

Lemma IV.5 (Qualitative lrp-problems). *The qualitative variants of the lrp-problem for MDPs and until properties are decidable in PTIME.*

Proof sketch. Positivity check is trivial as (in strongly connected MDPs): $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(\mathfrak{a}U\mathfrak{b}) > 0$ for some scheduler \mathfrak{S} iff $B \neq \emptyset$. For the value 1 problem, we can

rely on: $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\mathfrak{a}U\mathfrak{b}) = 1$ iff $\exists \mathfrak{S}. \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(\mathfrak{a}U\mathfrak{b}) = 1$ iff \mathcal{M} has a $(A \cup B)$ -EC \mathcal{E} with $\mathcal{E} \cap B \neq \emptyset$ or \mathcal{M} has an A-EC \mathcal{E} with $\Pr_{\mathcal{E}}^{\max}(\mathfrak{a}U\mathfrak{b}) = 1$. The arguments for the other two problems are analogous. For details see Lemma A.10. \square

We observe that we may concentrate on maximal long-run probabilities for until-properties since (see Lemma A.11):

$$\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\min}(\mathfrak{a}U\mathfrak{b}) = 1 - \mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(A \cup C)$$

if \mathcal{M} has no A-EC. Otherwise $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\min}(\mathfrak{a}U\mathfrak{b}) = 0$.

For the rest of this section, we suppose that \mathcal{M} is strongly connected and does not have $(A \cup B)$ -EC \mathcal{E} with $\mathcal{E} \cap B \neq \emptyset$, as otherwise $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\mathfrak{a}U\mathfrak{b})$ is 1.

The first important result is that $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\mathfrak{a}U\mathfrak{b})$ can be approximated by finite-memory schedulers. This is proved using Fatou's lemma (see Lemma A.12).

Lemma IV.6. $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\mathfrak{a}U\mathfrak{b}) = \sup_{\mathfrak{S} \in \text{FM}} \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(\mathfrak{a}U\mathfrak{b})$.

The main idea behind our algorithm is that particular FM-schedulers suffice to achieve the extremal long-run probabilities. In fact, we prove the existence of a *saturation point* $K \in \mathbb{N}$, such that the only kind of memory schedulers need is to count the number of consecutive \mathfrak{a} -states up to the current instant. We establish this by proving that each FM-scheduler \mathfrak{S} can be “improved” into another FM-scheduler if it does not maximize the probability for $\mathfrak{a}U\mathfrak{b}$ (which does not require memory) after having generated a sequence of K consecutive \mathfrak{a} -states. Here, improving means increasing the long-run probability of $\mathfrak{a}U\mathfrak{b}$. We will show that such a saturation point K is computable in time polynomial in the size of \mathcal{M} .

For $\alpha \in \text{Act}(s)$, let $p_{s,\alpha} = \sum_{t \in S} P(s, \alpha, t) \cdot p_t^{\max}$ and we write $\text{Act}^{\max}(s)$ for the set of *maximizing actions*, i.e. actions $\alpha \in \text{Act}(s)$ with $p_{s,\alpha} = p_s^{\max}$.

For a given bound K , define the class $\text{FM}(K)$ of FM-schedulers \mathfrak{S} for \mathcal{M} such that

- 1) $\Pr_{\mathcal{M}, \text{last}(\omega)}^{\mathfrak{S} \uparrow \omega}(\mathfrak{a}U\mathfrak{b}) = p_{\text{last}(\omega)}^{\max}$ for each finite \mathfrak{S} -path ω that has a suffix consisting of K or more consecutive A-states, and
- 2) the Markov chain induced by \mathfrak{S} has a single BSCC.

By definition, any $\mathfrak{S} \in \text{FM}(K)$ only schedules maximizing actions (in $\text{Act}^{\max}(\cdot)$) for paths ending with K or more consecutive A-states (by 1)); moreover, all states in \mathcal{M} have the same long-run probability under \mathfrak{S} , written $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{S}}(\mathfrak{a}U\mathfrak{b})$, (by 2)). We now strengthen the statement of Lemma IV.6 as follows:

Lemma IV.7 (Saturation point). *There exists K computable in polynomial time and satisfying:*

$$\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\alpha U b) = \sup_{\mathfrak{S} \in \text{FM}(K)} \mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{S}}(\alpha U b) .$$

Proof sketch. Let us explain the choice of K and its computation. For the soundness see Lemma A.13. Let

$$K = \max\{|A|, \lceil e/\delta \rceil\} \quad (*)$$

where e and δ are defined as follows. Let

$$\delta_s = \min\{p_s^{\max} - p_{s,\alpha} : \alpha \in \text{Act}(s) \setminus \text{Act}^{\max}(s)\}$$

with the convention that $\min \emptyset = \infty$. The value δ is then set as $\delta = \min_{s \in A} \delta_s$ if there exists at least one A -state s with $\delta_s < \infty$ and $\delta = 1$ otherwise. Obviously, δ is computable in polynomial time in the size of \mathcal{M} . Intuitively, each δ_s is the minimum probability loss when the first action of an optimal strategy for $\alpha U b$ is replaced by a non-optimal action.

Let us now give a construction which allows us to define e and also to compute long-run probabilities under schedulers of $\text{FM}(K)$. Consider an MDP which arises from \mathcal{M} by adding a counter representing the number of consecutive visits to A -states. When the counter value exceeds a given threshold, only maximizing actions (in $\text{Act}^{\max}(\cdot)$) are enabled. Formally, given a positive integer n , we define the MDP $\mathcal{M}_n = (S_n, \text{Act}, P_n)$ as follows. $S_n = B \cup C \cup (A \times \{1, \dots, n, \top\})$, and the transition probability function is defined by:

- If $s \in B \cup C$ then:
 - $P_n(s, \alpha, (s', 1)) = P(s, \alpha, s')$ if $s' \in A$
 - $P_n(s, \alpha, s') = P(s, \alpha, s')$ if $s' \in B \cup C$
- If $(s, i) \in A \times \{1, \dots, n-1\}$ then:
 - $P_n((s, i), \alpha, s') = P(s, \alpha, s')$ if $s' \in B \cup C$
 - $P_n((s, i), \alpha, (s', i+1)) = P(s, \alpha, s')$ if $s' \in A$
- If $s \in A$, $k \in \{n, \top\}$ and $\alpha \in \text{Act}^{\max}(s)$ then:
 - $P_n((s, k), \alpha, s') = P(s, \alpha, s')$ if $s' \in B \cup C$
 - $P_n((s, k), \alpha, (s', \top)) = P(s, \alpha, s')$ if $s' \in A$
- $P_n(\cdot) = 0$ in all remaining cases.

For every state $s \in S$ of \mathcal{M} , let s_n denote the ‘‘corresponding’’ state in \mathcal{M}_n . That is, $s_n = s$ if $s \in B \cup C$. For $s \in A$, we pick an arbitrary state (s, i) in \mathcal{M}_n that is reachable from $B \cup C$ in \mathcal{M}_n and put $s_n = (s, i)$. In the following, we identify \mathcal{M}_n with its fragment consisting of the states reachable from the set $\{s_n : s \in S\}$. As \mathcal{M} , \mathcal{M}_n is strongly connected, provided that $n \geq |A|$. Letting $N = |A|$, we define

$$e_{t,s} = \mathbb{E}_{\mathcal{M}_N, t_N}^{\min}(\text{‘‘steps until } s_N\text{’’})$$

as the minimal expected number of steps from state t_N to s_N in \mathcal{M}_N , and define $e = \max_{s, t \in S} e_{t,s}$. The values $e_{t,s}$ and a corresponding MD-scheduler $\mathfrak{R}_{N,s}$ for \mathcal{M}_N that minimizes the expected number of steps to state

s_N from every state t_N are computable in polynomial-time using standard linear programming techniques for stochastic shortest paths [4], [21]. The size of \mathcal{M}_N is at most quadratic in the size of \mathcal{M} .

Intuitively, K is chosen such that after a path of K consecutive α -states to s , a scheduler choosing $\beta \notin \text{Act}^{\max}(s)$ can be improved by instead choosing an action in $\text{Act}^{\max}(s)$ and hence increasing the probability to satisfy $\alpha U b$ in the K -many α -states and afterwards returning to s via a memoryless scheduler. \square

We define the MDP $\mathcal{K} = \mathcal{M}_K$ extended by the following transition-based weight function:

- transitions from B -states have weight 1,
- transitions from states in $A \times \{k\}$ with $k \in \{1, \dots, K-1\}$ leading to a B -state have weight k ,
- transitions from states (s, K) where $s \in A$ have weight $K \cdot p_s^{\max}$,
- transitions from states (s, \top) where $s \in A$ have weight p_s^{\max} , and

the weight of all other transitions is 0.

We reduce the computation of maximal long-run probabilities in \mathcal{M} to computing the maximal expected mean payoff in \mathcal{K} (see also Corollary A.18):

Lemma IV.8. $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\alpha U b) = \mathbb{E}_{\mathcal{K}}^{\max}(\text{MIP})$.

Proof sketch. Clearly, all schedulers for \mathcal{K} correspond to schedulers for \mathcal{M} . Vice versa, the behavior of $\text{FM}(K)$ -schedulers for \mathcal{M} can be mimicked by schedulers for \mathcal{K} . Thanks to Lemma IV.7, we deduce that $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\alpha U b) = \mathbb{L}\mathbb{P}_{\mathcal{K}}^{\max}(\alpha U b)$.

Now, for FM -schedulers \mathfrak{S} in \mathcal{K} that have no BSCC consisting of A -states, the expected mean payoff agrees with the long-run probability for $\alpha U b$, while otherwise $\mathbb{L}\mathbb{P}_{\mathcal{K}, s}^{\mathfrak{S}}(\alpha U b) \leq \mathbb{E}_{\mathcal{K}}^{\mathfrak{S}}(\text{MIP})$. Also, each FM -scheduler \mathfrak{S} in \mathcal{K} that has a BSCC consisting of A -states can be transformed into an infinite-memory scheduler \mathfrak{T} for \mathcal{K} such that $\mathbb{E}_{\mathcal{K}}^{\mathfrak{S}}(\text{MIP}) \leq \mathbb{L}\mathbb{P}_{\mathcal{K}, s}^{\mathfrak{T}}(\alpha U b)$. Therefore $\mathbb{E}_{\mathcal{K}}^{\max}(\text{MIP}) = \mathbb{L}\mathbb{P}_{\mathcal{K}}^{\max}(\alpha U b)$. \square

Example IV.9. Consider the MDP \mathcal{N}_k with $k = 2$ in Example IV.4 which is depicted in Fig. 2. We call this MDP \mathcal{M} and use the notation \mathcal{M}_n as well as the other notations as above. First, we compute the saturation point K : $\delta = p_a^{\max} - p_{a,\alpha} = \frac{1}{3}$. For the computation of e , we consider the MDP \mathcal{M}_1 . As we only consider the states reachable from $(a, 1)$ and only action β is available in this state in \mathcal{M}_1 , we see that e equals the expected number of steps from the left b state to $(a, 1)$. So, $e = k + 1 = 3$. Therefore, the saturation point is given by $K = \frac{e}{\delta} = 9$. We obtain the MDP $\mathcal{K} = \mathcal{M}_9$ extended with the weight structure described above.

For $0 \leq n \leq 8$, let \mathfrak{T}_n be the scheduler which chooses α in state (a, i) if $i \leq n$ and β otherwise. All MD-schedulers for \mathcal{K} are of this form. By summation over all paths from $(a, 1)$ to $(a, 1)$, we compute

$$\begin{aligned} \mathbb{E}_{\mathcal{K}}^{\mathfrak{T}_n}(\text{MIP}) &= \frac{\sum_{\ell=1}^n \frac{1}{3^\ell} (\ell+1) + \frac{1}{3^n} (n+2)}{2 \sum_{\ell=1}^n \frac{1}{3^\ell} (\ell+1) + \frac{1}{3^n} (n+3)} \\ &= \frac{5 + (2n+3) \frac{1}{3^n}}{10 + \frac{2}{3^n}}. \end{aligned}$$

The maximum is obtained for $n = 1$. As some MD-scheduler maximizes the expected mean payoff in \mathcal{K} , we conclude that $\mathbb{E}_{\mathcal{K}}^{\max}(\text{MIP}) = \mathbb{E}_{\mathcal{K}}^{\mathfrak{T}_1}(\text{MIP}) = \frac{5}{9}$. This value matches the value $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{T}_1}(aUb)$ computed in Example IV.4. \square

The length of the binary representation of K is polynomial in the size of \mathcal{M} , so that the size of $\mathcal{K} = \mathcal{M}_{\mathcal{K}}$ is pseudo-polynomial in the size of \mathcal{M} . We thus derive a complexity upper-bound for the computation of optimal long-run probabilities of until properties:

Theorem IV.10 (Computing optimal values). *The values $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(aUb)$ and $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\min}(aUb)$ are computable in pseudo-polynomial time.*

Example IV.4 illustrates that optimal schedulers can need a counter for the number of consecutive a -states up to a saturation point that can grow exponentially in the size of the MDP. Moreover, the logarithmic length of the maximal long-run probability can also be exponential in size of the MDP. This indicates that one cannot expect a polynomial-time algorithm to compute optimal schedulers or their values. Even the threshold problem, asking whether $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(aUb) \geq \vartheta$, is hard:

Theorem IV.11. *The threshold problem for until properties in MDP is NP-hard.*

Proof sketch. We sketch the main idea here. For the full proof see Theorem A.19. We provide a polynomial reduction from the intersection problem for unary DFA, i.e., DFA over a one-letter alphabet. This problem is known to be NP-complete [5]. Let $\mathcal{D}_1, \dots, \mathcal{D}_k$ be unary DFAs, and let ℓ be the product of their numbers of states. Then, $\mathcal{L}(\mathcal{D}_1) \cap \dots \cap \mathcal{L}(\mathcal{D}_k)$ is nonempty if and only if $\mathcal{L}(\mathcal{D}_1) \cap \dots \cap \mathcal{L}(\mathcal{D}_k)$ contains a word of length at most ℓ .

We construct an MDP \mathcal{M} with the disjoint union \mathcal{A} of these DFAs as a substructure and a threshold ϑ such that a scheduler has to attempt to follow a path of the same length to an accepting state in each of the DFAs in order to surpass the threshold. The states in \mathcal{M} are the \mathcal{A} -states and four additional states *init*, a , b , c . States in \mathcal{A} and a are labeled with a and b is labeled with b . In *init* and a , actions *pump* leading to a and *enter* randomly leading to one of the initial states in \mathcal{A} are available. In each state in \mathcal{A} an action α following

the transition in the DFAs is available. All of these actions may fail with a small probability of $\frac{1}{r}$. Failure leads back to *init*. In the final states in \mathcal{A} , an action β leading to b is available. In b and c an action τ is available which leads to *init* with probability $\frac{1}{r'}$ and to c otherwise.

The values r , r' , and ϑ are chosen such that a scheduler \mathfrak{S} achieves $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{S}}(aUb) \geq \vartheta$ if and only if it chooses β exactly in the moment it visited ℓ consecutive a -states. For this, \mathfrak{S} can pump $m \leq \ell$ times before entering \mathcal{A} , but then has to follow the transitions in a randomly chosen DFA \mathcal{D}_i for exactly $\ell - m$ steps before choosing β . But this is only possible if there is a word of length $\leq \ell$ which is accepted by all the DFAs. By the definition of ℓ , this is in turn equivalent to the non-emptiness of the intersection $\bigcap_i \mathcal{L}(\mathcal{D}_i)$. The main difficulty is to make sure that the chosen ϑ is computable in polynomial time. \square

Remark IV.12 (Memory requirements). If \mathcal{M} does not have an A-EC then there are FMD-schedulers achieving the optimal long-run frequency for aUb that behave memoryless in (BUC)-states and that use a counter for the number of consecutive A-states up to the saturation point. For MDPs with A-EC, optimal schedulers can need infinite-memory depending on whether all MD-schedulers for \mathcal{K} with maximal expected mean payoff have a BSCC consisting of states in $A \times \{\top\}$. This, e.g., applies to cases where each EC of \mathcal{M} containing some B-state also contains a C-state and there is an A-EC \mathcal{E} where $\Pr_{\mathcal{E}}^{\max}(aUb) > p_s^{\max}$ for all A-states s that do not belong to some A-EC. \square

Proposition IV.13 (Efficiently solvable cases). $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(aUb)$ can be computed in polynomial time in each of the following particular cases

- if \mathcal{M} has no A-cycle, or
- if for every scheduler \mathfrak{S} and every A-state s , $\Pr_{\mathcal{M},s}^{\mathfrak{S}}(aUb) = p_s^{\max}$.

Proof. In the first case indeed, we can let $K = |A|$ as saturation point, in which case the size of $\mathcal{K} = \mathcal{M}_{\mathcal{K}}$ is only polynomial in the size of \mathcal{M} . In the second case, $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(aUb)$ is simply the maximal mean payoff in \mathcal{M} when the weight of state s is p_s^{\max} . \square

V. CONCLUSION

The results of the paper illustrate that the introduced notions of long-run frequencies and probabilities for reasoning about (long-run) average satisfaction of path properties lead to computable measures in various cases. While efficient and simple algorithms exist for reachability, invariance, Streett and Rabin properties, the complexity raises to PSPACE for co-safety. The

probabilistic setting adds extra difficulties as illustrated for the case of until properties in MDPs.

There are various interesting open problems. The major open problem in the non-probabilistic setting is the computability of long-run frequencies in KSs for ω -regular properties specified by a Büchi automaton. Future directions in the probabilistic setting include the precise complexity of the threshold problems for until properties and algorithms for computing long-run probabilities for co-safety or ω -regular properties and long-run expectations of, *e.g.*, accumulated or discounted weights.

REFERENCES

- [1] Christel Baier, Nathalie Bertrand, Clemens Dubschlaff, Daniel Gburek, and Ocan Sankur. Stochastic shortest paths and weight-bounded properties in Markov decision processes. In *33rd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 86–94. ACM, 2018.
- [2] Christel Baier, Marcus Größer, and Frank Ciesinski. Quantitative analysis under fairness constraints. In *7th International Symposium on Automated Technology for Verification and Analysis (ATVA)*, volume 5799 of *Lecture Notes in Computer Science*, pages 135–150. Springer, 2009.
- [3] Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT Press, 2008.
- [4] Dimitri P. Bertsekas and John N. Tsitsiklis. An analysis of stochastic shortest path problems. *Mathematics of Operations Research*, 16(3):580–595, 1991.
- [5] Michael Blondin, Andreas Krebs, and Pierre McKenzie. The complexity of intersecting finite automata having few final states. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:90, 2012.
- [6] Udi Boker, Krishnendu Chatterjee, Thomas A. Henzinger, and Orna Kupferman. Temporal specifications with accumulative values. In *Proc. of the 26th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 43–52. IEEE Computer Society, 2011.
- [7] Benedikt Bollig, Normann Decker, and Martin Leucker. Frequency linear-time temporal logic. In *6th International Symposium on Theoretical Aspects of Software Engineering (TASE)*, pages 85–92. IEEE Computer Society, 2012.
- [8] Patricia Bouyer, Nicolas Markey, and Raj Mohan Matteplackel. Averaging in LTL. In *25th International Conference on Concurrency Theory (CONCUR)*, volume 8704 of *Lecture Notes in Computer Science*, pages 266–280. Springer, 2014.
- [9] Tomáš Brázdil, Václav Brozek, Krishnendu Chatterjee, Vojtech Forejt, and Antonín Kucera. Two views on multiple mean-payoff objectives in markov decision processes. *Logical Methods in Computer Science*, 10(1), 2014.
- [10] Pavol Cerný, Thomas A. Henzinger, and Arjun Radhakrishna. Simulation distances. *Theoretical Computer Science*, 413(1):21–35, 2012.
- [11] Krishnendu Chatterjee and Laurent Doyen. Energy and mean-payoff parity Markov decision processes. In *36th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 6907 of *Lecture Notes in Computer Science*, pages 206–218. Springer, 2011.
- [12] Krishnendu Chatterjee, Andreas Gaiser, and Jan Kretínský. Automata with generalized Rabin pairs for probabilistic model checking and LTL synthesis. In *25th International Conference on Computer Aided Verification (CAV)*, volume 8044 of *Lecture Notes in Computer Science*, pages 559–575. Springer, 2013.
- [13] Krishnendu Chatterjee and Monika Henzinger. Faster and dynamic algorithms for maximal end-component decomposition and related graph problems in probabilistic verification. In *22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1318–1336. SIAM, 2011.
- [14] Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdzinski. Mean-payoff parity games. In *20th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 178–187, June 2005.
- [15] Hana Chockler, Orna Kupferman, and Moshe Y. Vardi. Coverage metrics for temporal logic model checking*. *Formal Methods in System Design*, 28(3):189–212, 2006.
- [16] Hana Chockler and Ofer Strichman. Before and after vacuity. *Formal Methods in System Design*, 34(1):37–58, 2009.
- [17] Edmund Clarke, Orna Grumberg, and Doron Peled. *Model Checking*. MIT Press, 2000.
- [18] Costas Courcoubetis and Mihalis Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, 1995.
- [19] Luca de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, Department of Computer Science, 1997.
- [20] Luca de Alfaro. How to specify and verify the long-run average behavior of probabilistic systems. In *13th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 454–465. IEEE Computer Society, 1998.
- [21] Luca de Alfaro. Computing minimum and maximum reachability times in probabilistic systems. In *10th International Conference on Concurrency Theory (CONCUR)*, volume 1664 of *Lecture Notes in Computer Science*, pages 66–81, 1999.
- [22] Andrzej Ehrenfeucht and Jan Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, Jun 1979.
- [23] Vojtech Forejt and Jan Krcál. On frequency LTL in probabilistic systems. In Luca Aceto and David de Frutos-Escrig, editors, *26th International Conference on Concurrency Theory (CONCUR)*, volume 42 of *LIPICs*, pages 184–197. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.
- [24] Vojtech Forejt, Jan Krcál, and Jan Kretínský. Controller synthesis for MDPs and frequency $LTL_{\setminus GU}$. In Martin Davis, Ansgar Fehnker, Annabelle McIver, and Andrei Voronkov, editors, *20th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR)*, volume 9450 of *Lecture Notes in Computer Science*, pages 162–177. Springer, 2015.
- [25] Dror Fried, Solomon Eyal Shimony, Amit Benbassat, and Cenny Wenner. Complexity of Canadian traveler problem variants. *Theoretical Computer Science*, 487:1–16, 2013.
- [26] Thomas A. Henzinger and Jan Otop. From model checking to model measuring. In *24th International Conference on Concurrency Theory (CONCUR)*, volume 8052 of *Lecture Notes in Computer Science*, pages 273–287. Springer, 2013.
- [27] Lodewijk C.M. Kallenberg. Linear programming and finite Markovian control problems. *Mathematical Center Tracts*, 148, 1983.

- [28] Richard M. Karp. A characterization of the minimum cycle mean in a digraph. *Discrete Mathematics*, 23(3):309 – 311, 1978.
- [29] Dexter Kozen. Lower bounds for natural proof systems. In *18th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 254–266. IEEE Computer Society, 1977.
- [30] Orna Kupferman, Wenchao Li, and Sanjit A. Seshia. A theory of mutations with applications to vacuity, coverage, and fault tolerance. In *Formal Methods in Computer-Aided Design (FMCAD)*, pages 1–9. IEEE Computer Society, 2008.
- [31] Orna Kupferman and Moshe Y. Vardi. Robust satisfaction. In *10th International Conference on Concurrency Theory (CONCUR)*, volume 1664 of *Lecture Notes in Computer Science*, pages 383–398. Springer, 1999.
- [32] Orna Kupferman and Moshe Y. Vardi. Vacuity detection in temporal model checking. *STTT*, 4(2):224–233, 2003.
- [33] Albert R. Meyer and Larry J. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. In *13th Annual Symposium on Switching and Automata Theory (SWAT/FOCS)*, pages 125–129. IEEE Computer Society, 1972.
- [34] Christos H. Papadimitriou. Games against nature. *Journal of Computer and System Science*, 31(2):288–301, 1985.
- [35] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, 1994.
- [36] Paulo Tabuada and Daniel Neider. Robust linear temporal logic. In *25th EACSL Annual Conference on Computer Science Logic (CSL)*, volume 62 of *LIPICs*, pages 10:1–10:21. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
- [37] Christian von Essen and Barbara Jobstmann. Synthesizing systems with optimal average-case behavior for ratio objectives. In *International Workshop on Interactions, Games and Protocols (iWIGP)*, volume 50 of *EPTCS*, pages 17–32, 2011.

A. Long-run frequencies for until properties in KS

Lemma A.1 (see Lemma III.1).

- (a) *The following three statements are equivalent:*
- (1) $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\max}(a \cup b) = 1$
 - (2) \mathcal{T} has an infinite path π with $\text{lrf}_{a \cup b}(\pi) = 1$.
 - (3) \mathcal{T} has a $A \cup B$ -cycle.
- (b) *The following three statements are equivalent:*
- (1) $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\min}(a \cup b) = 0$
 - (2) \mathcal{T} has an infinite path π with $\text{lrf}_{a \cup b}(\pi) = 0$.
 - (3) \mathcal{T} has a $A \cup C$ -cycle.

Proof. We only prove the equivalence of (1), (2) and (3) for maximal long-run frequencies (item (a)).

“(1) \implies (3)”: suppose that \mathcal{T} does not have $A \cup B$ -cycles. Then, each cycle in \mathcal{T} contains a state that is not contained in $A \cup B$. If $n = |S|$ is the total of number of states then along every path latest after n steps a state in $S \setminus (A \cup B)$ will be reached. Hence, the long-run frequency of each infinite path is less or equal $(n-1)/n$. But then $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\max}(a \cup b) \leq (n-1)/n < 1$

“(3) \implies (2)”: suppose that \mathcal{T} has a cycle ξ , consisting of states in $A \cup B$. If one of the states in ξ is a B-state then the long-run frequency of the infinite path ξ^ω is 1. Suppose now that all states in ξ are contained in A. Let $\xi = s_0 s_1 \dots s_n$ and let $\omega = t_0 t_1 \dots t_m t_{m+1} \dots t_k$ be a path with $t_0 = s_0 = s_n = t_k$, $t_1, \dots, t_{m-1} \in A$, $t_m \in B$. (Recall that by definition of A we have $s \models \exists(A \cup B)$ for all states $s \in A$.) Then, the long-run frequency of the infinite path $\xi; \omega; \xi^2; \omega; \xi^4; \omega; \xi^8; \omega; \dots$ is 1.

“(2) \implies (3)”: obvious.

The argument for statement (b) is analogous and omitted here. (Note that the absence of A-cycle yields that $s \models \exists(A \cup C)$ for all states $s \in A$.) \square

We give the precise definition of the transition function for the weighted KS constructed in Section III-A. We switch from \mathcal{T} to the weighted KS \mathcal{T}' with state space $S' = B \cup C \cup A \times \{0, 1\}$ and the following transitions:

- If $(s, s') \in \Delta$ is a transition in \mathcal{T} where $s \in B \cup C$ then \mathcal{T}' contains the following transitions:
 - $s \rightarrow (s', 0)$ and $s \rightarrow (s', 1)$ if $s' \in A$
 - $s \rightarrow s'$ if $s' \in B \cup C$
- If $(s, s') \in \Delta$ is a transition in \mathcal{T} where $s \in A$ then \mathcal{T}' has the following transitions:
 - $(s, 0) \rightarrow s'$ if $s' \in C$
 - $(s, 1) \rightarrow s'$ if $s' \in B$
 - $(s, 0) \rightarrow (s', 0)$ and $(s, 1) \rightarrow (s', 1)$ if $s' \in A$

Recall that the weight function in \mathcal{T}' is state-based and assigns weight 1 to the states in $B \cup A \times \{1\}$ and weight 0 to the states in $C \cup A \times \{0\}$.

We now introduce the following additional notation. For $s \in S$ let $I(s) = \{s\}$ if $s \in B \cup C$ and $I(s) = \{(s, 0), (s, 1)\}$ if $s \in A$. For the following lemma, the assumption that \mathcal{T} is strongly connected, is irrelevant.

Lemma A.2 (See Lemma III.2). *Suppose \mathcal{T} has no A-cycles and let s be a state of \mathcal{T} . Then, $\mathbb{L}\mathbb{F}_{\mathcal{T}, s}^{\max}(a \cup b) = \text{MIP}_{\mathcal{T}', I(s)}^{\max}$ and $\mathbb{L}\mathbb{F}_{\mathcal{T}, s}^{\min}(a \cup b) = \text{MIP}_{\mathcal{T}', I(s)}^{\min}$.*

In particular, if \mathcal{T} is strongly connected then $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\max}(a \cup b) = \text{MIP}_{\mathcal{T}'}^{\max}$ and $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\min}(a \cup b) = \text{MIP}_{\mathcal{T}'}^{\min}$ as stated in Lemma III.2.

Proof. We show that there is a bijection ι between the infinite paths in \mathcal{T} and in \mathcal{T}' such that for each infinite path π in \mathcal{T} : $\text{first}(\iota(\pi)) \in I(\text{first}(\pi))$ and $\text{lrf}_{a \cup b}(\pi) = \text{mp}(\iota(\pi))$.

Given an infinite path $\pi = s_0, s_1 s_2 \dots$ in \mathcal{T} , we define $\iota(\pi) = s'_0 s'_1 s'_2 \dots$ as follows. If $s_i \in B \cup C$ then $s'_i = s_i$. If $i = 0 \vee s_{i-1} \in B \cup C$ and $s_i s_{i+1} \dots s_k \in A^* B$ then $s'_i s'_{i+1} \dots s'_k = (s_i, 1) (s_{i+1}, 1) \dots (s_k, 1)$. Likewise, if $i = 0 \vee s_{i-1} \in B \cup C$ and $s_i s_{i+1} \dots s_k \in A^* C$ then $s'_i s'_{i+1} \dots s'_k = (s_i, 0) (s_{i+1}, 0) \dots (s_k, 0)$. It is easy to see that $\iota(\pi)$ is indeed a path in \mathcal{T}' satisfying $\text{first}(\iota(\pi)) \in I(\text{first}(\pi))$ and $\text{lrf}_{a \cup b}(\pi) = \text{mp}(\iota(\pi))$. Moreover, ι is a bijection as \mathcal{T} has no A-cycles (and hence cannot stay forever in the sub-structures consisting of states in $A \times \{0, 1\}$) and as $A \times \{0\}$ can only be left via a transition to a C-state and $A \times \{1\}$ can only be left via a transition to a B-state. \square

B. Long-run frequencies for co-safety properties in KS

 Soundness of the weighted KS \mathcal{G}

Before providing the proof of Theorem III.5, we show that there are enough empty tracks to insert tracks at states H' and one at q_0 at the second step.

Lemma A.3. *For all paths $(s_k, f_k)_{k \geq 0}$ of \mathcal{G} , for all $k \geq 0$, the following conditions hold:*

- *There are at most ℓ positions $0 \leq i \leq 2\ell$ such that $f(i) \in Q \times \{\text{true}, \text{false}\}$,*
- *there exists $0 \leq i \leq 2\ell$ such that $f_k(i) \in \{q_0\} \times \{\text{true}, \text{false}\}$,*
- *if $k > 0$, then for all $0 \leq i \leq 2\ell$ such that $f_{k-1}(i) = (q, \text{false})$ for some $q \in Q$, for all $q' \in \delta(q, L(s_k))$, there exists $0 \leq j \leq 2\ell$ such that $f_k(j) = (q', \text{false})$.*

Proof. We prove the three properties by induction on the length of the paths of \mathcal{G} .

Consider any path $(s_k, f_k)_{k \geq 0}$. The properties are true in the initial state (s_0, f_0) by definition. Assume this holds at (s_k, f_k) . Consider (s_{k+1}, f'_k, H) such that

$$(s_k, f_k) \rightsquigarrow_1 (s_{k+1}, f'_k, H) \rightsquigarrow_2 (s_{k+1}, f_{k+1}),$$

so that (s_{k+1}, f'_k, H) is consistent. Observe that by construction, the second and third properties hold whenever there are enough empty tracks in f'_k . In f'_k , all tracks different than $Q \times \{\text{true}, \text{false}\}$ in f_k are set to \emptyset ; while each track of the form $(q, b) \in Q \times \{\text{true}, \text{false}\}$ is replaced by (q', b) for some successor q' of q in \mathcal{A} . Thus, $|\{0 \leq i \leq 2\ell \mid f'_k(i) \neq \emptyset\}| \leq \ell$. All properties then follow from this observation. In fact, in f_{k+1} , the track (q_0, true) or (q_0, false) is added and its occurrence is unique since q_0 there is no edge entering q_0 in \mathcal{A} . Furthermore, a track (q, false) is added for each $q \in H$ unless another track with the same state q already exists. Then duplicate tracks are removed and labeled by $\text{merged}(\cdot)$. Thus, no state $q \in Q$ appears twice in (s_{k+1}, f_{k+1}) . \square

Theorem A.4 (See Theorem III.5). *Let \mathcal{T} be a strongly connected KS, \mathcal{A} an NFA encoding a co-safety property φ and \mathcal{G} the weighted KS defined as above. Then, $\text{LFP}_{\mathcal{T}}^{\text{max}}(\varphi) = \text{MIP}_{\mathcal{G}}^{\text{max}}(\Phi)$ and $\text{LFP}_{\mathcal{T}}^{\text{min}}(\varphi) = \text{MIP}_{\mathcal{G}}^{\text{min}}(\Phi)$.*

Proof. We prove the following even stronger result:

- (a) For each infinite path π in \mathcal{T} starting in s_0 there exists an infinite path π' in \mathcal{G} with $\pi' \models \Phi$ and $\text{lrf}_{\varphi}(\pi) = \text{mp}(\pi')$.
- (b) For each initial infinite path π' in \mathcal{G} with $\pi' \models \Phi$ there is an infinite path π in \mathcal{T} starting in s_0 such that $\text{lrf}_{\varphi}(\pi) = \text{mp}(\pi')$

Proof of statement (a).: Consider any infinite path $\pi = (s_k)_{k \geq 0}$ of \mathcal{T} that starts at s_0 . We will construct a uniquely defined path $(s_k, f_k)_{k \geq 0}$ in \mathcal{G} satisfying Φ and whose mean payoff is exactly $\text{lrf}_{\varphi}(\pi)$.

The path follows π in its first component. The only non-determinism in the choice of the second component is in the choice of (q_0, b) with $b \in \{\text{true}, \text{false}\}$ that is introduced at each step, and in the choice of a successor for each (q, true) . In fact, the rest of the components are uniquely determined by the construction.

We are going to construct a path $(s_k, f'_k)_{k \geq 0}$ with an additional information: at each step k , tracks of the form (q, true) are replaced by (q, true, ρ) where ρ is a path witnessing that the trace $L(s_k)L(s_{k+1})\dots$ is accepted by \mathcal{A} from state q . The final path $(s_k, f_k)_{k \geq 0}$ is then obtained by removing this additional component.

We define f'_0 by choosing (q_0, true, ρ) if and only if $\pi, 0 \models \phi$, where ρ is some witness accepting path of \mathcal{A} from q_0 .

Assume that the path is constructed up to step $k-1$. Then f'_k is defined with the following resolution of the non-determinism.

- We add the track (q_0, true, ρ) if $\pi, k \models \phi$, with ρ a witness path, and add (q_0, false) otherwise.

- For any $0 \leq i \leq 2\ell$ with $f'_{k-1}(i) = (q, \text{true}, \rho)$, if $q \in F$ then ρ is the empty path and $f'_k(i) = \emptyset$ (there is no non-determinism to resolve in this case). Otherwise, let us write $\rho = qq'\rho'$ where q, q' are first two states of ρ and ρ' is the rest of the path. We let $f'_k(i) = (q', \text{true}, q'\rho')$.

This sequence thus constructed never produces inconsistent states. This is because all tracks are chosen according to the satisfaction of the property φ at given position with given witness accepting paths in \mathcal{A} . Now, the mean payoff of $(s_k, f_k)_{k \geq 0}$ is obtained as the average of the number of positions in which a track (q_0, true) is created. This is the case exactly when φ holds from a given position. Thus, the mean payoff of $(s_k, f_k)_{k \geq 0}$ is equal to $\text{lrf}_{\varphi}(\pi)$.

It remains to show that $(s_k, f_k)_{k \geq 0}$ satisfies Φ . Consider any position $0 \leq i \leq 2\ell$, and any $k \geq 0$. Let us show that $(\text{false}_i \vee \text{merged}_i \vee F_i \vee \emptyset_i)$ is satisfied at some position $l \geq k$. This is clear if $f_k(i) \models \text{false}_i \vee \text{merged}_i \vee F_i \vee \emptyset_i$, so assume this is not the case. The only remaining case is $f_k(i) = (q, \text{true})$ for $q \in Q \setminus F$. Consider the witness path ρ above. By construction, there is $l > k$ such that either $f_l(i) \in F \times \{\text{true}\}$ or $f_l(i) \in \{\text{merged}\} \times \mathbb{N}$.

Proof of statement (b).: Conversely, consider a path $(s_k, f_k)_{k \geq 0}$ of \mathcal{G} satisfying Φ . Let $\pi = (s_k)_{k \geq 0}$ be the corresponding path in \mathcal{T} . We will prove that for all positions k such that f_k contains a track with (q_0, true) , $\pi, k \models \phi$, and for all other positions k , $\pi, k \not\models \phi$. This shows that $\text{lrf}_{\varphi}(\pi)$ is equal to the mean payoff of the path $(s_k, f_k)_{k \geq 0}$.

We prove a more general statement: for all states (s_k, f_k) containing some track (q, true) , the suffix $L(\pi_{[k..]})$ has a prefix that is accepted by some path from q to F in \mathcal{A} . We proceed by induction on the position i of the track containing (q, true) , with $0 \leq i \leq 2\ell + 1$.

- Case $i = 0$. Consider $k \geq 0$ such that $f_k(0) = (q, \text{true})$. We claim that there exists $k' > k$ and states $q_l \in Q$ for all $k \leq l \leq k'$ such that $f_l(i) = (q_l, \text{true})$, $q_k = q$ and $q_{k'} \in F$, and $q_k q_{k+1} \dots q_{k'}$ is a path of \mathcal{A} . In fact, since $(s_k, f_k)_{k \geq 0}$ satisfies Φ , the track i cannot be of the form (q', true) indefinitely from position k . Moreover, it can never be of the form (merged, j) since this would imply $j < 0$. It cannot be either of the form (q, false) before visiting some (q_f, true) with $q_f \in F$ since this is not possible in the construction. Thus, in order to satisfy Φ , there must exist such a position k' as above. Then $q_k q_{k+1} \dots q_{k'}$ is a path of \mathcal{A} by construction of \mathcal{G} .
- Consider $i > 0$ and $k \geq 0$ such that $f_k(i) = (q, \text{true})$. We claim that one of the two conditions hold:

- 1) there exists $k' > k$ and states $q_l \in Q$ for all $k \leq l \leq k'$ such that $f_l(i) = (q_l, \text{true})$, $q_k = q$ and $q_{k'} \in F$, and $q_k q_{k+1} \dots q_{k'}$ is a path of \mathcal{A} .
- 2) or, there exists $k' > k$ and states $q_l \in Q$ for all $k \leq l \leq k' - 1$ such that $f_l(i) = (q_l, \text{true})$, $q_k = q$ and $q_l q_{l+1} \dots q_{k'-1}$ is a path of \mathcal{A} and $f_{k'}(i) = (\text{merged}, j)$ with some $j < i$.

The result follows immediately in the first case. Assume the first case does not hold. By Φ , there must exist some position $k' > k$ satisfying $(\text{false}_i \vee \text{merged}_i \vee F_i \vee \emptyset_i)$. Consider the least such position k' . By construction, we cannot have $(s_{k'}, f_{k'}) \models \text{false}_i$ or $(s_{k'}, f_{k'}) \models \emptyset_i$ since all tracks of the form $Q \times \{\text{true}\}$ are stopped by some $F \times \{\text{true}\}$ or $\text{merged} \times \mathbb{N}$. The former does not hold by hypothesis so there must exist $j < i$ such that $f_{k'}(i) = (\text{merged}, j)$. By construction, there exists $q, q' \in Q$ such that $f_{k'-1}(i) = (q, \text{true})$ and $f_{k'}(j) = (q, \text{true})$ with $q' \in \delta(q, L(s_{k'-1}))$. By induction hypothesis there exists a path ρ from q' to F along a prefix of the trace $L(s_{k'})L(s_{k'+1})\dots$. Now if $q_k q_{k+1} \dots q_{k'-1}$ denotes the path of \mathcal{A} defined by the track i from positions k to $k' - 1$, then $q_k q_{k+1} \dots q_{k'-1} \cdot \rho$ is an accepting path of \mathcal{A} on a prefix of $L(\pi_{[k..]})$. \square

Computation of the extremal values

Let us explain the exponential-time algorithm of Corollary III.6 for computing $\mathbb{L}\mathbb{F}_{\mathcal{T},s}^{\max}(\varphi)$ and $\mathbb{L}\mathbb{F}_{\mathcal{T},s}^{\min}(\varphi)$.

We start by building \mathcal{G} which has size $|S|(3\ell + 1)^{2\ell+1}$, which is linear in the size of \mathcal{T} and exponential in that of \mathcal{A} . We then compute the strongly connected components of \mathcal{G} and keep those SCCs which contain at least one accepting state per each Büchi condition of Φ . In fact, any infinite path satisfying Φ must have a suffix that belongs to such an SCC. We then use the algorithm of [28] to compute the cycle that maximizes the mean payoff in time polynomial in the size of \mathcal{G} . Then, as done in [6], one can build an infinite path by repeating this cycle, and interleaving an infinite number of visits to accepting states with frequency that goes to 0. A detailed construction of this idea is also given below, in the proof of Theorem A.5.

PSPACE-completeness of the threshold problem

We now turn to the proof of Theorem III.7. Let us first recall the result:

Theorem A.5 (see Theorem III.7). *The threshold problem “given a KS \mathcal{T} , an NFA \mathcal{A} and a rational threshold ϑ , check whether \mathcal{T} has an infinite path π with $\text{lr}_f(\pi) \geq \vartheta$ ” is PSPACE-complete.*

The polynomial-space upper bound will be shown in Lemma A.6 and PSPACE-hardness will be shown in Lemma A.7.

Lemma A.6. *Given KS \mathcal{T} , co-safety property φ described by NFA \mathcal{A} , and a rational number ϑ , a comparison operator $\triangleright \in \{\geq, >, <, \leq\}$ one can check in polynomial space whether $\mathbb{L}\mathbb{F}_{\mathcal{T}}^{\max}(\varphi) \triangleright \vartheta$.*

Proof. As PSPACE equals NPSpace (Savitch’s theorem) it suffices to provide a non-deterministic polynomially space-bounded procedure to check whether $\mathbb{MIP}_{\mathcal{G}}^{\max}(\Phi) \triangleright \vartheta$. The idea is to use the weighted KS \mathcal{G} , without constructing \mathcal{G} explicitly. To simplify the following argument we concentrate on the case $\triangleright = \geq$. The algorithm for strict lower bounds and upper bounds are similar.

The algorithm starts by guessing a state $s_{\mathcal{G}}$ in $S_{\mathcal{G}}$ and checks whether $s_{\mathcal{G}}$ is reachable from one of the initial states in $I(s_0)$. This check can be done in polynomial space. Then, it checks whether

- (1) there exists a cycle ξ_{Φ} of length $2(2\ell + 1)|S_{\mathcal{G}}|$ containing $s_{\mathcal{G}}$ which satisfies each conjunct of Φ ,
- (2) there exists a simple cycle ξ_{MP} containing $s_{\mathcal{G}}$ with mean payoff of at least ϑ .

The first check can be done in polynomial space by guessing the path and using a counter up to $2\ell|S_{\mathcal{G}}|$, which can be represented in polynomial space. The second check can be done similarly by keeping the sum of the guessed cycle which can also be represented in polynomial space, since it is bounded by $|S_{\mathcal{G}}|$ as all weights are 0 or 1.

We first show: If the above nondeterministic algorithm returns “yes”, then \mathcal{G} has an infinite path ρ with $\text{mp}(\rho) \geq \vartheta$. This construction follows the more general results on mean payoff parity games of [14]. Let ξ_{Φ} and ξ_{MP} denote the two cycles computed above, the first one satisfying Φ , and the second one with mean payoff at least ϑ . The path we construct first reaches $s_{\mathcal{G}}$. From here, it runs in phases alternating between ξ_{Φ} and ξ_{MP} . At phase $i \geq 1$, it follows 2^i times ξ_{MP} and then once ξ_{Φ} . This path clearly satisfies Φ since the cycle ξ_{Φ} is seen infinitely often. Let π denote the suffix of ρ starting at the first occurrence of $s_{\mathcal{G}}$. We calculate the mean payoff of π , which is the same as the mean payoff of ρ . Consider any prefix $\pi_{[0..K]}$ of π of length K . Let us write $a = |\xi_{\Phi}|$, $b = |\xi_{\text{MP}}|$ and $\text{avg}(\pi_{[0..K]})$ for the average payoff of the first K states of π , i.e.,

$$\text{avg}(\pi_{[0..K]}) = \frac{1}{K} \cdot \sum_{i=0}^{K-1} \text{wgt}(\pi_{[i]}).$$

If $K \geq a + 2b$, then there exists a unique integer $i \geq 0$ and some integer r with $0 \leq r < a + 2^{i+1}b$ such that:

$$\begin{aligned} K &= a + 2^1b + a + 2^2b + \dots + a + 2^ib + r \\ &= ia + b(2^{i+1} - 1) + r \end{aligned}$$

If $r < a + b$, then:

$$\text{avg}(\pi_{[0\dots K]}) \geq \frac{\vartheta(2^{i+1} - 1)}{2^{i+1} - 1 + (i+1)a + b}$$

If $r \geq a + b$, there exists $m < 2^{i+1}$ such that:

$$\text{avg}(\pi_{[0\dots K]}) \geq \frac{\vartheta(m + 2^{i+1} - 1)}{m + 2^{i+1} - 1 + ia}$$

Then, for all $\varepsilon > 0$, there exists K_ε such that for all $K \geq K_\varepsilon$, $\text{avg}(\pi_{[0\dots K]}) \geq \vartheta - \varepsilon$. This yields:

$$\text{mp}(\pi) = \liminf_{K \rightarrow \infty} \text{avg}(\pi_{[0\dots K]}) \geq \vartheta$$

As the mean payoff of ρ coincides with the mean payoff of each of its suffixes, we get:

$$\text{mp}(\rho) = \text{mp}(\pi) \geq \vartheta$$

This also shows that the limit of the average payoffs of finite prefixes exists, so the limsup and liminf variants of the mean payoff have also the same value for ρ .

Conversely, assume that there exists an infinite path ρ of \mathcal{G} satisfying Φ , starting in a state of $I(s_0)$ and with $\text{mp}(\rho) \geq \vartheta$. The task is to show that the sketched nondeterministic algorithm has a computation returning the answer “yes”. Pick some state $s_{\mathcal{G}}$ that occurs infinitely in ρ . Then, $s_{\mathcal{G}}$ is reachable from $I(s_0)$. Let us write Φ_i for the Büchi condition $\square\Diamond(\text{false}_i \vee \text{merged}_i \vee F_i \vee \emptyset_i)$. Then, $\Phi = \bigwedge_{i=0}^{2\ell} \Phi_i$ and ρ has a fragment that constitutes a cycle in \mathcal{G} containing state $s_{\mathcal{G}}$ and satisfying all conjuncts Φ_i of Φ . We now sketch how to construct new a cycles ξ_Φ and ξ_{MP} satisfying the constraints of (1) and (2). For each $i \in \{0, 1, \dots, 2\ell\}$, ρ contains a path from $s_{\mathcal{G}}$ to some state t_i with $t_i \models \Phi_i$, and a path from t_i to $s_{\mathcal{G}}$. A cycle ξ_i of length $\leq 2|S_{\mathcal{G}}|$ containing $s_{\mathcal{G}}$ and t_i can then be constructed by concatenating simple paths from $s_{\mathcal{G}}$ to t_i and from t_i to $s_{\mathcal{G}}$ in \mathcal{G} . Concatenating the cycles $\xi_0, \xi_1, \dots, \xi_{2\ell}$ for all conjuncts Φ_i of Φ yields the desired cycle ξ_Φ . Furthermore, it is known that in mean payoff automata, the maximal mean payoff is achieved on simple cycles [22]. Thus, the algorithm has a computation returning the answer “yes”. \square

Lemma A.7. *The threshold problem “given a KS \mathcal{T} , an NFA \mathcal{A} and a rational threshold ϑ , decide whether $\mathbb{L}_{\mathcal{T}}^{\text{max}}(\varphi) \geq \vartheta$ ” is PSPACE-hard.*

Proof. The PSPACE lower bound follows by a polynomial reduction from the intersection problem for deterministic finite automata (DFA): given k DFA $\mathcal{D}_1, \dots, \mathcal{D}_k$ over the same alphabet Σ , is the intersection

language $\mathcal{L}(\mathcal{D}_1) \cap \dots \cap \mathcal{L}(\mathcal{D}_k)$ nonempty? This problem is known to be PSPACE-complete [29].

To provide a polynomial reduction from the intersection problem for DFA, we suppose we are given DFA $\mathcal{D}_1, \dots, \mathcal{D}_k$ over some alphabet Σ . W.l.o.g. we may assume that $k \geq 2$ and that the empty word is not included in any of the languages $\mathcal{L}(\mathcal{D}_i)$. Let Q_i be the state space of \mathcal{D}_i , $\ell_i = |Q_i|$ and $\ell = \ell_1 \cdot \dots \cdot \ell_k$. Then, $\mathcal{L}(\mathcal{D}_1) \cap \dots \cap \mathcal{L}(\mathcal{D}_k)$ is nonempty if and only if there is a word $w \in \Sigma^*$ of length at most ℓ such that $w \in \mathcal{L}(\mathcal{D}_i)$ for $i = 1, \dots, k$.

Let $\$, \dots, \$, \#$ be pairwise distinct fresh letters (not contained in Σ), and let $\Gamma = \Sigma \cup \{\$, \dots, \$, \#\}$.

Given a finite word $w = \sigma_1 \sigma_2 \dots \sigma_n \in \Sigma^+$, let \hat{w} denote the word over $\Sigma \cup \{\#\}$ that arises from w by inserting $(k-1)$ -times the symbol $\#$ after each letter σ_j . That is,

$$\hat{w} = \sigma_1 \#^{k-1} \sigma_2 \#^{k-1} \dots \sigma_n \#^{k-1}$$

For $i = 1, \dots, k$, one can easily construct in time $\mathcal{O}(k^2 + k \cdot \text{size}(\mathcal{D}_i))$ a new DFA \mathcal{B}_i over the alphabet Γ such that:

$$\mathcal{L}(\mathcal{B}_i) = \{ \$^j \$_{i+1}^{k-1} \dots \$_k^{k-1} \hat{w} : w \in \mathcal{L}(\mathcal{D}_i), 1 \leq j < k \}$$

Furthermore, we can construct in time linear in the sizes of $\mathcal{B}_1, \dots, \mathcal{B}_k$ an NFA \mathcal{A} over the alphabet Γ with:

$$\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{B}_1) \cup \dots \cup \mathcal{L}(\mathcal{B}_k) \cup \{\#^i : i \geq 1\}$$

Note that \mathcal{A} does not accept the empty word and no word starting with a letter in Σ . Likewise, we can construct in time polynomial in k a strongly connected KS \mathcal{T} with the following states:

- $s_{i,j}$ for $i = 1, \dots, k$ and $j = 1, \dots, k-1$,
- t_1, \dots, t_{k-1} and
- u_σ for each symbol $\sigma \in \Sigma$.

We treat the symbols in Γ as atomic propositions and identify the singletons $\{\gamma\}$ with γ , where γ ranges over all symbols of the alphabet Γ . The labeling function of \mathcal{T} is then given by:

$$L(s_{i,j}) = \$_i, L(t_j) = \# \text{ and } L(u_\sigma) = \sigma.$$

The transition relation of \mathcal{T} is depicted in Figure 4.

Thus, the words generated by \mathcal{T} are the substrings of the infinite words $y_1 y_2 y_3 \dots$ where each word y_i has the form

$$\$_1^{k-1} \$_2^{k-1} \dots \$_k^{k-1} \hat{w}$$

for some $w \in \Sigma^+$. Let

$$\vartheta = \frac{k(k-1) + (k-1)\ell}{k(k-1) + k\ell}$$

Clearly, $\mathcal{T}, \mathcal{A}, \vartheta$ can be constructed in time polynomial in the size of the DFA $\mathcal{D}_1, \dots, \mathcal{D}_k$. It remains to show that \mathcal{T} has an infinite path π with long-run frequency

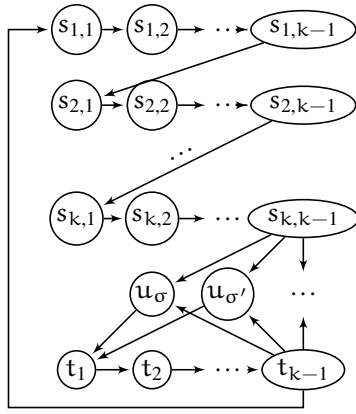


Fig. 4. The Kripke structure \mathcal{T} in the reduction of Theorem III.7.

$lrf_{\mathcal{A}}(\pi)$ at least ϑ if and only if the intersection language of the \mathcal{D}_i 's is nonempty.

It remains to prove that the following equivalence: \mathcal{T} has an infinite path π with long-run frequency $lrf_{\mathcal{A}}(\pi)$ at least ϑ if and only if the intersection language of the \mathcal{D}_i 's is nonempty.

Let us recall that, formally, the relation of \mathcal{T} consists of the following transitions:

$$\begin{aligned} s_{i,1} &\rightarrow s_{i,2} \rightarrow \dots \rightarrow s_{i,k-1} && \text{for } i = 1, \dots, k \\ s_{i,k-1} &\rightarrow s_{i+1,1} && \text{for } i = 1, \dots, k-1 \\ s_{k,k-1} &\rightarrow u_{\sigma} \rightarrow t_1 && \text{for } \sigma \in \Sigma \\ t_1 &\rightarrow t_2 \rightarrow \dots \rightarrow t_{k-1} \\ t_{k-1} &\rightarrow s_{1,1} \text{ and } t_{k-1} \rightarrow u_{\sigma} && \text{for } \sigma \in \Sigma \end{aligned}$$

Suppose first that there is some word $w \in \Sigma^*$ accepted by each of the DFSs $\mathcal{D}_1, \dots, \mathcal{D}_k$. As stated before, we then can safely assume that $|w| \leq \ell$. \mathcal{T} has a cycle ξ generating the word $v = s_1^{k-1} s_2^{k-1} \dots s_k^{k-1} \hat{w}$. We then have:

$$|v| = k(k-1) + k|w|$$

The word v contains exactly $k(k-1) + (k-1)|w|$ positions from which a word accepted by \mathcal{A} starts. This follows from the following two observations:

- The suffixes $s_i^j s_{i+1}^{k-1} \dots s_k^{k-1} \hat{w}$ of v are accepted by \mathcal{B}_i , and therefore by \mathcal{A} (for $j = 1, \dots, k-1$).
- \hat{w} contains exactly $(k-1) \cdot |w|$ positions from which a subword contained in $\#^+$ starts.

Thus, the long-run frequency of the infinite path ξ^ω that repeats this cycle ad infinity is:

$$lrf_{\mathcal{A}}(\xi^\omega) = f(|w|)$$

where $f: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is the following function:

$$f(x) = \frac{k(k-1) + (k-1)x}{k(k-1) + kx}$$

Function f is monotonically decreasing.¹ Therefore, $f(|w|) \geq f(\ell)$. As $f(\ell) = \vartheta$, we conclude that \mathcal{T} has an infinite path with long-run frequency at least ϑ .

We assume now that \mathcal{T} has an infinite path π with long-run frequency at least ϑ . We first observe that π must visit $s_{1,1}$ infinitely often as otherwise π would have an infinite suffix consisting of t - and u -states, in which case the long-run frequency would be smaller or equal than $(k-1)/k$, and therefore strictly smaller than ϑ . Suppose by contradiction $\mathcal{L}(\mathcal{D}_1) \cap \dots \cap \mathcal{L}(\mathcal{D}_k)$ is empty. Then, the average weight obtained by each cycle $s_{1,1} s_{1,2} \dots s_{1,k-1} \dots s_{k,1} s_{k,2} \dots s_{k,k-1} \omega s_{1,1}$ where ω consists of t - and u -states contained in π in \mathcal{T} is less or equal

$$\frac{(k-1)^2 + (k-1)y}{k(k-1) + ky} = \frac{k-1}{k}$$

where y is the number of u -states in ω , in which case the number of t -states in ω is $(k-1)y$. But then again, the long-run frequency of π would be bounded by $(k-1)/k$, and therefore strictly smaller than ϑ . Contradiction. We conclude that $\mathcal{L}(\mathcal{D}_1) \cap \dots \cap \mathcal{L}(\mathcal{D}_k)$ must be nonempty if \mathcal{T} has an infinite path with long-run frequency at least ϑ . \square

Qualitative decision problems

Lemma A.8. Given a state s of a KS \mathcal{T} and NFA \mathcal{A} , the problem to check whether there is an infinite path π of \mathcal{T} starting in s with

- $lrf_{\mathcal{A}}(\pi) > 0$ is in P ,
- $lrf_{\mathcal{A}}(\pi) < 1$ is PSPACE-hard,
- $lrf_{\mathcal{A}}(\pi) = 0$ is NP-hard,
- $lrf_{\mathcal{A}}(\pi) = 1$ is NP-hard.

Proof. We first consider statement (a). As stated before, it suffices to consider the case where \mathcal{T} is strongly connected, in which case the starting state s is irrelevant. We build the synchronous product $\mathcal{T} \otimes \mathcal{A}$ and treat it as an NFA where $S \times Q_0$ is the set of initial states and $S \times F$ the set of final states. We now show: \mathcal{T} has an infinite path π with $lrf_{\mathcal{A}}(\pi) > 0$ if and only if the language of $\mathcal{T} \otimes \mathcal{A}$ is nonempty.

- For the implication " \Leftarrow " we pick a finite word w over $2^{\mathcal{AP}}$ that is accepted by $\mathcal{T} \otimes \mathcal{A}$. As $Q_0 \cap F = \emptyset$, w is non-empty. Pick an accepting run $(s_0, q_0)(s_1, q_1) \dots (s_n, q_n)$ for w in $\mathcal{T} \otimes \mathcal{A}$. We then pick a finite path $t_0 t_1 \dots t_m$ from $s_n = t_0$ to $s_0 = t_m$, and regard the infinite path π arising by the infinite repetition of the cycle $s_0 s_1 \dots s_n t_1 \dots t_m$. Obviously, we then have $lrf_{\mathcal{A}}(\pi) > 0$.

¹Each rational function $h(x) = (a + cx)/(b + dx)$ with $cb < ad$ is decreasing. This is a consequence of the fact the the first derivative is strictly negative. Note that $h'(x) = (cb - ad)/(b + dx)^2$, which is strictly negative if $cb < ad$.

- For the implication “ \implies ”, we suppose that we are given an infinite path π with $lrf_{\mathcal{A}}(\pi) > 0$. Then, there is a pair (i, j) of integers with $i < j$ such that the word induced by the path fragment $\pi_{[i..j]}$ is accepted by \mathcal{A} . This path fragment can be lifted to an accepting run in $\mathcal{T} \otimes \mathcal{A}$. Thus, the language of $\mathcal{T} \otimes \mathcal{A}$ is nonempty.

We prove statement (b) via a polynomial reduction from the universality of finite automata, which is known to be PSPACE-complete [33], to the problem that takes as input a KS \mathcal{T} and an NFA \mathcal{A} and asks whether $lrf_{\mathcal{A}}(\pi) = 1$ for all infinite paths π in \mathcal{T} . Because PSPACE is closed under complement, this implies that checking the existence of an infinite path π such that $lrf_{\mathcal{A}}(\pi) < 1$ is PSPACE-hard. For the reduction from the universality problem for NFA we may restrict to NFA over the alphabet $\Sigma = 2^{AP}$ for some fixed set AP of atomic propositions.

Given an NFA $\mathcal{A} = (Q, \Sigma, \delta, Q_0, F)$, we construct a new NFA \mathcal{B} over the alphabet $\Sigma' = \Sigma \cup \{\#\}$ for a new symbol $\{\#\}$ as follows. The states of \mathcal{B} are Q together with two new states s_{init} and acc . The transitions in \mathcal{B} are given by δ and a new $\#$ -transition from s_{init} to each state $q \in Q_0$, a new $\#$ -transition from each state $f \in F$ to acc , and a new transition from s_{init} to acc for each symbol $\sigma \in \Sigma$. The only accepting state in \mathcal{B} is acc and the only initial state is s_{init} . So, we get that $\mathcal{L}(\mathcal{B}) = \Sigma \cup \{\#w\# \mid w \in \mathcal{L}(\mathcal{A})\}$.

We define a KS \mathcal{T} with state space $S = \{s_\sigma : \sigma \in \Sigma'\}$, the obvious labeling function $L(s_\sigma) = \sigma$, and with transitions between all pairs of states. Thus, all words on Σ' can be generated by \mathcal{T} .

Assume \mathcal{A} is universal. Consider some infinite path π of \mathcal{T} . If π contains only finitely many $\#$ s then $lrf_{\mathcal{B}}(\pi) = 1$, as all suffixes not starting with $\#$ have a prefix accepted by \mathcal{B} in one step. If π contains infinitely many $\#$ s, then $lrf_{\mathcal{B}}(\pi) = 1$ as all suffixes not starting with $\#$ have a prefix of length 1 which is accepted by \mathcal{B} and all suffixes starting with $\#$ have a prefix of the form $\#w\#$ for some word w . As \mathcal{A} is universal, $w \in \mathcal{L}(\mathcal{A})$ and hence $\#w\# \in \mathcal{L}(\mathcal{B})$.

Conversely, assume that there is a word $w \notin \mathcal{L}(\mathcal{A})$. Then, the path $\pi = (\#w)^\omega$ in \mathcal{T} satisfies $lrf_{\mathcal{B}}(\pi) = \frac{|w|}{|w|+1} < 1$. As π starts with $\#$, the only prefix that could be accepted by \mathcal{B} is $\#w\#$. But this word is rejected by construction. So, every $|w| + 1$ th suffix of π has no prefix accepted by \mathcal{B} .

Hence, we showed that there is a path π in \mathcal{T} with $lrf_{\mathcal{B}}(\pi) < 1$ if and only if \mathcal{A} is non-universal.

For the proof of statement (c), we describe a polynomial reduction from 3SAT. Let $\psi = c_1 \vee \dots \vee c_m$ be a 3CNF formula with m clauses using n Boolean variables, say x_1, \dots, x_n . Let \mathcal{T} be a KS containing one state per clause c_i , and for each c_i , a fresh state

per literal of c_i . The initial state is c_1 . At state c_i , there are three successors which are the literals of c_i , which all go directly to c_{i+1} . We create an additional state $c_{m+1} = \perp$ whose only outgoing transition goes to c_1 . The automaton \mathcal{A} reads literals and guesses any contradiction in the input. It accepts iff for some i , the word contains both x_i and $\neg x_i$. Then, if ψ is satisfiable, there exists an infinite path π in \mathcal{T} such that $lrf_{\mathcal{A}}(\pi) = 0$. If ψ is not satisfiable, for all π in \mathcal{T} , at each cycle, some x_i and $\neg x_i$ must be read. Thus, $lrf_{\mathcal{A}}(\pi) \geq \frac{1}{n+m+1}$.

The proof of statement (d) is also a reduction from 3SAT and uses the same KS \mathcal{T} as in the previous construction. Consider 3CNF formula $\psi = c_1 \vee \dots \vee c_m$ over variables x_1, \dots, x_n , and let us write $c_i = l_1^i \vee l_2^i \vee l_3^i$ where l_k^i are literals. Let L denote the set of all literals, and $C = \{c_1, \dots, c_m\}$. We define \mathcal{A} as follows. From the initial state, one goes to an accepting sink state by reading any letter of $L \cup \{\perp\}$. By reading $c_i \in C$, one goes to s_{c_i} from which the automaton accepts when \perp is read iff one of the literals of c_i was seen in the mean time, but not its negation. That is, if none of the literals of c_i are seen, or some literal l_k^i and its negation are both seen, then the automaton ends in a rejecting sink state. This part of the automaton can be constructed using 3 extra bits in order to store which literals have been seen.

If ψ is satisfiable, then there is a word $w = c_1 l_1^1 c_2 l_2^2 \dots c_m l_m^m \perp$ where $l_k^k \in c_k$ for each k , and such that for all k, k' , $l_k^k \neq \neg l_{k'}^{k'}$. Consider the word $\pi = w^\omega$. The property \mathcal{A} is satisfied from all positions. This is trivial from positions containing $X \cup \{\perp\}$. At any other position containing c_i , the letter is followed by an actual literal of c_i , and the rest of the word does not contain its negation. Thus, $lrf_{\mathcal{A}}(\pi) = 1$.

Assume now that ψ is not satisfiable. Consider any word π of \mathcal{T} which can be written as $\pi = w_1 \perp w_2 \perp \dots$ where each w_i has the form $c_1 l_1 c_2 l_2 \dots c_m l_m$ with $l_k \in c_k$ for each k . Since ψ is not satisfiable, each w_i must contain l_j and l_k with $l_j = \neg l_k$. That is, at least once every $m+n+1$ positions, the property is violated. Thus, $lrf_{\mathcal{A}}(\pi) < 1$. \square

C. Proofs for the probabilistic case

Additional notations: End components (ECs) have been introduced in Section II as sets of state-action pairs where the induced graph is strongly connected. Occasionally, we shall use a representation of an end component \mathcal{E} as a pair (E, \mathfrak{A}) where E is a set of states and $\mathfrak{A} : E \rightarrow Act$, namely $E = \{s \in S : \exists \alpha \in Act. (s, \alpha) \in \mathcal{E}\}$ and $\mathfrak{A}(s) = \{\alpha \in Act : (s, \alpha) \in \mathcal{E}\}$. With this representation is mind and identifying \mathcal{E} with E , we sometimes use notations like $s \in \mathcal{E}$ or $T \cap \mathcal{E}$ for $s \in S$ and $T \subseteq S$.

If \mathcal{G} is a scheduler and $\pi = s_0\alpha_0s_1\alpha_1\dots$ a path then π is said to be a \mathcal{G} -path if $\mathcal{G}(s_0\alpha_0\dots\alpha_{n-1}s_n)(\alpha_n) > 0$ for all $n \geq 0$. Many proofs will rely on de Alfaro's observation [19] stating that for each scheduler of a (finite-state) MDP, the limit of almost all infinite \mathcal{G} -paths constitutes an end component. Here, the limit $\text{Lim}(\pi)$ of an infinite path π denotes the set of state-action pairs that occur infinitely often in π .

Efficiently solvable instances of the lrp-problem

Theorem A.9 (see **Theorem IV.3**). *The values $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\max}(\varphi)$ and $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\min}(\varphi)$ are computable in polynomial-time if φ is a condition of one the following types:*

- reachability $\diamond b$,
- invariance $\square b$,
- generalized Rabin $\bigwedge_{i=1}^n \bigvee_{j=1}^{\ell_i} (\square \diamond b_{i,j} \wedge \diamond \square a_{i,j})$
- or Streett $\bigwedge_{i=1}^n (\square \diamond a_{i,j} \rightarrow \square \diamond b_{i,j})$.

In all these cases, optimal FMD-scheduler exist. Moreover, optimal MD-schedulers exist for reachability, invariances, Büchi and co-Büchi conditions.

Proof. We provide the argument for $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\max}(\varphi)$. The argument for $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\min}(\varphi)$ is analogous and omitted here. As stated above, we may assume that \mathcal{M} is strongly connected.

It is well-known [2], [12], [19] that for all properties listed in Theorem IV.3 there is an FMD-scheduler \mathcal{G} that maximizes the probability for φ from every visited state in the following sense:

$$\Pr_{\mathcal{M},\pi_{[i]}^{\mathcal{G}}}^{\uparrow \pi_{[0..i]}}(\varphi) = \Pr_{\mathcal{M},\pi_{[i]}}^{\max}(\varphi)$$

for each infinite \mathcal{G} -path π and each position $i \in \mathbb{N}$. For reachability, invariances, Büchi and co-Büchi conditions, we may even suppose that \mathcal{G} is an MD-scheduler with a single BSCC \mathcal{B} .

If φ is a reachability, generalized Rabin or Streett condition then $\Pr_{\mathcal{M},s}^{\max}(\varphi) = \Pr_{\mathcal{M},t}^{\max}(\varphi)$ for all states s, t in \mathcal{M} . Moreover, this value is either 0 or 1. But then \mathcal{G} obviously achieves the maximal long-run probability from every state.

The states in \mathcal{M} can have different maximal probabilities for invariances $\varphi = \square b$. However, for invariances we either have $\max_{s \in \mathcal{S}} \Pr_{\mathcal{M},s}^{\max}(\varphi) = 0$, in which case $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\max}(\varphi) = 0$ for all states s , or the unique BSCC \mathcal{B} of \mathcal{G} consists of b -states. In the latter case, $\Pr_{\mathcal{M},s}^{\mathcal{G}}(\square b) = \Pr_{\mathcal{M},s}^{\max}(\square b) = 1$ for all states s in \mathcal{B} . Let now \mathcal{T} be the following MD-scheduler:

- From the states not in \mathcal{E} , \mathcal{T} mimics an MD-scheduler maximizing the probability to reach \mathcal{B} (which is 1 as \mathcal{M} is strongly connected).

- For the state inside \mathcal{B} , \mathcal{T} behaves as \mathcal{G} .

We then have $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathcal{T}}(\square b) = 1$ for all states s in \mathcal{M} , which is obviously maximal. \square

Qualitative lrp-problems

Recall that the task of the qualitative lrp-problems is to decide the existence of a scheduler \mathcal{G} such that $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathcal{G}}(\varphi)$ is positive, equals 1, is strictly less than 1 or equals 0.

Lemma A.10. *The four qualitative lrp-problems for MDPs and until properties are decidable in polynomial time. Moreover, if \mathcal{M} is strongly connected and s a state of \mathcal{M} then:*

$$\begin{aligned} \mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(aUb) = 1 & \quad \text{iff} \quad \exists \mathcal{G}. \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathcal{G}}(\varphi) = 1 \\ \mathbb{L}\mathbb{P}_{\mathcal{M}}^{\min}(aUb) = 0 & \quad \text{iff} \quad \exists \mathcal{G}. \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathcal{G}}(\varphi) = 0 \end{aligned}$$

Proof. Let $\varphi = aUb$ and \mathcal{M} be a strongly connected MDP. Polynomial-time decidability of the four qualitative lrp-problems is a direct consequence of the following observations:

- (a) $\exists \mathcal{G}. \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathcal{G}}(\varphi) > 0$ iff $\mathcal{B} \neq \emptyset$
- (b) $\exists \mathcal{G}. \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathcal{G}}(\varphi) = 1$ iff \mathcal{M} has a $(A \cup B)$ -EC \mathcal{E} with $\mathcal{E} \cap \mathcal{B} \neq \emptyset$ or \mathcal{M} has an A-EC \mathcal{E} with $\Pr_{\mathcal{E}}^{\max}(\varphi) = 1$ (or both).
- (c) $\exists \mathcal{G}. \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathcal{G}}(\varphi) < 1$ iff \mathcal{M} has an A-EC or $\mathcal{C} \neq \emptyset$
- (d) $\exists \mathcal{G}. \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathcal{G}}(\varphi) = 0$ iff \mathcal{M} has a $(A \cup C)$ -EC

Statement (a) is obvious. For the proof of statement (b), we show the equivalence of the following three statements:

- (i) $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\varphi) = 1$
- (ii) $\exists \mathcal{G}. \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathcal{G}}(\varphi) = 1$
- (iii) \mathcal{M} has a $(A \cup B)$ -EC \mathcal{E} with $\mathcal{E} \cap \mathcal{B} \neq \emptyset$ or \mathcal{M} has an A-EC \mathcal{E} with $\Pr_{\mathcal{E}}^{\max}(\varphi) = 1$ (or both).

“(iii) \implies (ii)”: The claim is obvious if \mathcal{M} has a $(A \cup B)$ -EC \mathcal{E} with $\mathcal{E} \cap \mathcal{B} \neq \emptyset$. Suppose now that \mathcal{M} has an A-EC \mathcal{E} with $\Pr_{\mathcal{E}}^{\max}(\varphi) = 1$. An infinite-memory scheduler with long-run probability 1 can be obtained by using the A-EC \mathcal{E} to have longer and longer portions of A-states, before satisfying almost-surely aUb (by following an MD-scheduler maximizing the probability for aUb) and returning to \mathcal{E} from the reached B-state (by following an MD-scheduler maximizing the probability for reaching \mathcal{E} from every state in \mathcal{M}).

“(i) \implies (iii)”: by contradiction. Suppose that none of the two alternative condition holds. Then, C-states are seen “often” on average in the following sense: if n denotes the number of states then $\Pr_{\mathcal{M},s}^{\min}(\diamond^{\leq n} C) > 0$ for every state s . As $\Pr_{\mathcal{M},t}^{\max}(aUb) = 0$ for each C-state t , this yields

$$\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathcal{G}}(aUb) \leq \frac{n-1}{n} < 1$$

for each scheduler \mathfrak{S} . Hence, $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\mathfrak{a}U\mathfrak{b}) \leq (n-1)/n < 1$.

The implication “(ii) \implies (i)” is trivial.

The proofs for statements (c) and (d) and the equivalence of $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\min}(\mathfrak{a}U\mathfrak{b}) = 0$ and the existence of a scheduler \mathfrak{S} with $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(\mathfrak{a}U\mathfrak{b}) = 0$ are similar and omitted here. \square

Duality of minimal and maximal long-run probabilities

The following lemma essentially shows that minimizing the long-run probability for $\mathfrak{a}U\mathfrak{b}$ is dual to the task to maximize the long-run probability for AUC . An exception is the case where \mathcal{M} has A-ECs, in which case the minimal lrp-problem can be answered directly.

Lemma A.11 (Min-lrp via max-lrp). *Suppose \mathcal{M} is a strongly connected and let A, B, C as above. If \mathcal{M} has an A-EC then $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\min}(\mathfrak{a}U\mathfrak{b}) = 0$. Otherwise, i.e., if \mathcal{M} has no A-EC, then $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\min}(\mathfrak{a}U\mathfrak{b}) = 1 - \mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(AUC)$.*

Proof. Suppose first that \mathcal{M} has an A-EC \mathcal{E} . For each state t in \mathcal{E} , let $\mathfrak{A}(t) = \{\alpha \in Act : (s, \alpha) \in \mathcal{E}\}$. Consider an MD-scheduler \mathfrak{S} maximizing the probability to reach \mathcal{E} from every state s outside \mathcal{E} and which selects only actions in $\mathfrak{A}(t)$ for every state t in \mathcal{E} . Then, $\Pr_{\mathcal{M},s}^{\mathfrak{S}}(\diamond\Box A) = 1$ for all states s , and therefore $\Pr_{\mathcal{M},s}^{\mathfrak{S}}(AUB) = 0$ (as A and B are disjoint). But then $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(\mathfrak{a}U\mathfrak{b}) = \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(AUB) = 0$ for each state s . This yields $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\min}(\mathfrak{a}U\mathfrak{b}) = 0$.

Suppose now that there are no A-ECs. Then, for each scheduler \mathfrak{S} and each state s :

$$\Pr_{\mathcal{M},s}^{\mathfrak{S}}(\mathfrak{a}U\mathfrak{b}) = \Pr_{\mathcal{M},s}^{\mathfrak{S}}(AUB) = 1 - \Pr_{\mathcal{M},s}^{\mathfrak{S}}(AUC)$$

This yields $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\min}(\mathfrak{a}U\mathfrak{b}) = 1 - \mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(AUC)$. \square

Finite-memory schedulers for $\mathfrak{a}U\mathfrak{b}$

We now show that the maximal long-run probabilities for until properties can be approximated by FM-schedulers.

As before, we suppose that \mathcal{M} is a strongly connected MDP with state space S . Furthermore, we may safely assume that $C \neq \emptyset$ as otherwise all states either belong to A or B , in which case either $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\mathfrak{a}U\mathfrak{b}) = 0$ if $B = \emptyset$ or, if $B \neq \emptyset$, then $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\mathfrak{a}U\mathfrak{b}) = \mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{S}}(\mathfrak{a}U\mathfrak{b}) = 1$ for any MD-scheduler \mathfrak{S} that maximizes the probability for reaching B .

Lemma A.12 (see Lemma IV.6). *For each scheduler \mathfrak{T} for \mathcal{M} , each $\varepsilon > 0$ and each state s of \mathcal{M} , there is a FM-scheduler \mathfrak{S} for \mathcal{M} such that:*

$$\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(\mathfrak{a}U\mathfrak{b}) \geq \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{T}}(\mathfrak{a}U\mathfrak{b}) - \varepsilon$$

Proof. If \mathfrak{S} is a scheduler then briefly we write $p_s^{\mathfrak{S}}$ instead of $\Pr_{\mathcal{M},s}^{\mathfrak{S}}(AUB)$.

By Fatou’s lemma, we have:

$$\begin{aligned} \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(\mathfrak{a}U\mathfrak{b}) &= \mathbb{E}_{\mathcal{M},s}^{\mathfrak{T}} \left(\liminf_{n \rightarrow \infty} \frac{1}{n+1} \sum_{i=0}^n p_{\pi_{[i]}}^{\mathfrak{T} \uparrow \pi_{[0..i]}} \right) \\ &\leq \liminf_{n \rightarrow \infty} \mathbb{E}_{\mathcal{M},s}^{\mathfrak{T}} \left(\frac{1}{n+1} \sum_{i=0}^n p_{\pi_{[i]}}^{\mathfrak{T} \uparrow \pi_{[0..i]}} \right) \end{aligned}$$

So, there exists $k_0 \in \mathbb{N}$ such that for all $k \geq k_0$:

$$\mathbb{E}_{\mathcal{M},s}^{\mathfrak{T}} \left(\frac{1}{k+1} \sum_{i=0}^k p_{\pi_{[i]}}^{\mathfrak{T} \uparrow \pi_{[0..i]}} \right) \geq \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{T}}(\mathfrak{a}U\mathfrak{b}) - \frac{\varepsilon}{2}$$

Let \mathfrak{U}_s be the following FM-scheduler with two modes. If the current state is in A , it starts in the first mode, in which it behaves like an MD-scheduler maximizing the probability of $\mathfrak{a}U\mathfrak{b}$. As soon as a state in $B \cup C$ has been reached, scheduler \mathfrak{U}_s operates in the second mode, in which it memorylessly minimizes the expected number of steps until reaching s . Let $f_{t,s} = \mathbb{E}_{\mathcal{M},t}^{\min}$ (“steps until s ”) denote the expected number of steps this scheduler \mathfrak{U}_s needs to reach s in the second mode starting from state t . We then define $f_s = \max_{t \in S} f_{t,s}$ and $f = \max_{s \in S} f_s$.

We now construct an FM-scheduler \mathfrak{S} satisfying the claim of the lemma. First, choose a natural number k with $k \geq k_0$ and $k+1 > \frac{2f_s}{\varepsilon}$. The behavior of scheduler \mathfrak{S} is as follows. In its first mode, it starts in s and behaves like \mathfrak{T} in the first k steps. Then, it switches to the second mode and behaves like \mathfrak{U}_s until it reaches s (in the second mode of \mathfrak{U}_s). Afterwards, it switches back to the first mode.

As \mathfrak{U}_s maximizes the probability of $\mathfrak{a}U\mathfrak{b}$ whenever it starts in a state in A , we obtain:

$$\frac{1}{k+1} \cdot \sum_{i=0}^k p_{\pi_{[i]}}^{\mathfrak{T} \uparrow \pi_{[0..i]}} \leq \frac{1}{k+1} \cdot \sum_{i=0}^k p_{\pi_{[i]}}^{\mathfrak{S} \uparrow \pi_{[0..i]}}$$

for all paths π . Furthermore, the expected number of steps which \mathfrak{S} takes to follow \mathfrak{T} for $k+1$ steps and to return to s via \mathfrak{U}_s is at most $k+1+f_s$.

Expressing the long-run probability of \mathfrak{S} as a quotient, we obtain:

$$\begin{aligned} \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(\mathfrak{a}U\mathfrak{b}) &\geq \frac{p_{\pi_{[0]}}^{\mathfrak{T} \uparrow \pi_{[0..0]}} + \dots + p_{\pi_{[k]}}^{\mathfrak{T} \uparrow \pi_{[0..k]}}}{k+1+f_s} \\ &\geq \frac{p_{\pi_{[0]}}^{\mathfrak{T} \uparrow \pi_{[0..0]}} + \dots + p_{\pi_{[k]}}^{\mathfrak{T} \uparrow \pi_{[0..k]}}}{(k+1) \cdot (1+\varepsilon/2)} \\ &\geq \frac{p_{\pi_{[0]}}^{\mathfrak{T} \uparrow \pi_{[0..0]}} + \dots + p_{\pi_{[k]}}^{\mathfrak{T} \uparrow \pi_{[0..k]}}}{k+1} \cdot (1-\varepsilon/2) \\ &\geq (\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{T}}(\mathfrak{a}U\mathfrak{b}) - \varepsilon/2) \cdot (1-\varepsilon/2) \end{aligned}$$

by the choice of k . Using the fact that $\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{T}}(aUb)$ is bounded by 1 we obtain:

$$\begin{aligned} \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{S}}(aUb) &\geq (\mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{T}}(aUb) - \varepsilon/2) \cdot (1 - \varepsilon/2) \\ &\geq \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{T}}(aUb) - \varepsilon \end{aligned}$$

This completes the proof of Lemma A.12 and yields the statement of Lemma IV.6. \square

Saturation point

We present the proof of Lemma IV.7. Let us first recall the relevant notations. The value K has been defined in Section IV of the main paper as $\max\{N, \lceil \varepsilon/\delta \rceil\}$ where $N = |A|$ and where ε and δ are as follows.

The value δ has been defined $\delta = \min_{s \in A} \delta_s$ where δ_s denotes the minimal difference between $p_s^{\max} = \Pr_{\mathcal{M},s}^{\max}(aUb)$ and the values $p_{s,\alpha} = \sum_{t \in S} P(s, \alpha, t) \cdot p_t^{\max}$ for $\alpha \in Act(s) \setminus Act^{\max}(s)$, and $\delta_s = \infty$ if $Act(s) = Act^{\max}(s)$. In the special case where $\delta_s = \infty$ for all states $s \in A$, δ is defined to be 1.

For the definition of ε we considered the unfolded MDP \mathcal{M}_N (where again $N = |A|$) and defined ε as the maximum of the values

$$e_{t,s} = \mathbb{E}_{\mathcal{M}_N, t_N}^{\min}(\text{“steps until } s_N\text{”})$$

Here, $s_N = s$ for $s \in B \cup C$ and s_N is any copy (s, i) of state s in \mathcal{M}_N that is reachable from some $(B \cup C)$ -state in \mathcal{M}_N . As $N = |A|$, with \mathcal{M} also \mathcal{M}_N is strongly connected. This ensures the well-definedness (finiteness) of the values $e_{t,s}$ and the existence of an MD-scheduler $\mathfrak{R}_{N,s}$ for \mathcal{M}_N that minimizes the expected number of steps to s_N from every state t_N . This MD-scheduler $\mathfrak{R}_{N,s}$ for \mathcal{M}_N can be viewed as an FMD-scheduler for \mathcal{M} . We then have:

$$e_{t,s} \geq \mathbb{E}_{\mathcal{M},t}^{\mathfrak{R}_{N,s}}(\text{“steps until } s\text{”})$$

for all states $s, t \in S$. Moreover, $\mathfrak{R}_{N,s}$ enjoys the property that $\Pr_{\mathcal{M},t}^{\mathfrak{R}_{N,s} \uparrow \omega}(aUb) = p_t^{\max}$ for each $\mathfrak{R}_{N,s}$ -path ω that has a suffix consisting of N or more A-states where $t = last(\omega)$. In particular, $\mathfrak{R}_{N,s}$ only schedules actions in $Act^{\max}(\cdot)$ when having generated a path that ends with a suffix consisting of N or more A-states.

Lemma A.13 (see Lemma IV.7). *Suppose \mathcal{M} has no $(A \cup B)$ -EC containing at least one B-state. Then, for each FM-scheduler \mathfrak{T} , there is a scheduler $\mathfrak{S} \in FM(K)$ with $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{S}}(aUb) \geq \max_{s \in S} \mathbb{L}\mathbb{P}_{\mathcal{M},s}^{\mathfrak{T}}(aUb)$.*

Proof. Let \mathfrak{T} be an FM-scheduler for \mathcal{M} with modes (memory cells) in the finite set X . Let $\mathcal{C}^{\mathfrak{T}}$ denote the Markov chain induced by \mathfrak{T} . We can think of the states in $\mathcal{C}^{\mathfrak{T}}$ as pairs (s, x) consisting of a state s in \mathcal{M} and a mode $x \in X$. We may assume w.l.o.g. that $\mathcal{C}^{\mathfrak{T}}$ has

a single BSCC, say $\mathcal{B}^{\mathfrak{T}}$. This yields that all states of $\mathcal{C}^{\mathfrak{T}}$ have the same long-run probability for aUb . Let us simply write $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{T}}(aUb)$ for this value.

Given a state $\mathfrak{s} = (s, x)$ in $\mathcal{C}^{\mathfrak{T}}$, we say \mathfrak{s} is an A-state if $s \in A$. The notations B-state and C-state have the analogous meaning. We suppose that all A-states of $\mathcal{C}^{\mathfrak{T}}$ are labeled with a , while the B-states are labeled with b .

If $\mathcal{B}^{\mathfrak{T}}$ consists of $(A \cup C)$ -states then $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{T}}(aUb) = 0$ and the claim is trivial as we can deal with any FM(K)-scheduler.

Suppose now that $\mathcal{B}^{\mathfrak{T}}$ contains at least one B-state. Then, almost all \mathfrak{T} -paths visit infinitely often some B-state and infinitely often some C-state.

We now explain how to modify \mathfrak{T} 's decision for generating a scheduler in FM(K) with the desired property. Our procedure works by induction on the number $k^{\mathfrak{T}}$ of $(B \cup C)$ -states $\mathfrak{s} = (s, x)$ in $\mathcal{B}^{\mathfrak{T}}$ where

$$\Pr_{\mathcal{B}^{\mathfrak{T}}, \mathfrak{s}}(\bigcirc(AU \geq^K D^{\mathfrak{T}})) > 0$$

Here, $D^{\mathfrak{T}}$ denotes the set of A-states $t = (t, y)$ in the BSCC $\mathcal{B}^{\mathfrak{T}}$ where $\mathfrak{T}(t)(\alpha) > 0$ for some action $\alpha \notin Act^{\max}(t)$.

If $k^{\mathfrak{T}} = 0$ then for every path $\omega = \mathfrak{s}_0 \alpha_0 \dots \alpha_{n-1} \mathfrak{s}_n$ in $\mathcal{B}^{\mathfrak{T}}$ of length $n \geq K$ where all but the first state are A-states, the states $\mathfrak{s}_K, \mathfrak{s}_{K+1}, \dots, \mathfrak{s}_n$ do not belong to $D^{\mathfrak{T}}$. That is, \mathfrak{T} schedules only actions in Act^{\max} for these states. But then the probability for aUb in $\mathcal{B}^{\mathfrak{T}}$ from each of the states $\mathfrak{s}_i = (s_i, x_i)$ with $i \geq K$ equals $p_{s_i}^{\max}$.² This implies that if $k^{\mathfrak{T}} = 0$ then \mathfrak{T} is an FM(K)-scheduler and we can deal with $\mathfrak{S} = \mathfrak{T}$.

Suppose now that $k^{\mathfrak{T}} \geq 1$. We show how to transform \mathfrak{T} into a new FM-scheduler \mathfrak{S} with a single BSCC such that $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{S}}(aUb) \geq \mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{T}}(aUb)$ and $k^{\mathfrak{S}} < k^{\mathfrak{T}}$.

Given states $\mathfrak{s} = (s, x)$ and $\mathfrak{t} = (t, y)$ in $\mathcal{B}^{\mathfrak{T}}$, where \mathfrak{s} is a $(B \cup C)$ -state and \mathfrak{t} an A-state, let $\Gamma_{\mathfrak{s}, \mathfrak{t}}$ denote the set of finite \mathfrak{T} -paths $\omega = \mathfrak{s}_0 \alpha_0 \dots \alpha_{n-1} \mathfrak{s}_n$ such that

- $n \geq K$,
- $\mathfrak{s}_0 = \mathfrak{s}$, $\mathfrak{s}_n = \mathfrak{t}$,
- $\mathfrak{s}_1, \dots, \mathfrak{s}_{n-1}$ are A-states, and
- $\mathfrak{T}(\mathfrak{s}_n)(\alpha) > 0$ for some action $\alpha \notin Act^{\max}(\mathfrak{t})$.

Let $\Pi_{\mathfrak{s}, \mathfrak{t}}$ denote the set of paths $\omega \in \Gamma_{\mathfrak{s}, \mathfrak{t}}$ such that no proper prefix of ω belongs to $\Gamma_{\mathfrak{s}, \mathfrak{t}}$, and let $\Pi_{\mathfrak{s}}$ denote the union of the sets $\Pi_{\mathfrak{s}, \mathfrak{t}}$.

As $k^{\mathfrak{T}}$ is positive, we can pick some $(B \cup C)$ -state $\mathfrak{s} = (s, x)$ in $\mathcal{B}^{\mathfrak{T}}$ where $\Pi_{\mathfrak{s}}$ is nonempty.

The definition of FM-scheduler \mathfrak{S} is as follows. Scheduler \mathfrak{S} operates in two phases. Its first phase starts in

²For general MDPs, schedulers that only select actions in Act^{\max} might not achieve the maximal probability for aUb . This, however, is only possible if the schedulers under consideration realize A-ECs. As $\mathcal{B}^{\mathfrak{T}}$ is a BSCC with at least one B-state, this case does not apply to scheduler \mathfrak{T} .

$\mathfrak{s} = (s, \chi)$ and uses additional memory cells to keep track of the number of consecutive A-states that have been traversed since the last visit of \mathfrak{s} . More precisely, if \mathfrak{G} has generated the path $\mathfrak{s}_0 \mathfrak{s}_1 \dots \mathfrak{s}_i$ where $i \leq K$, $\mathfrak{s}_0 = \mathfrak{s}$ and $\mathfrak{s}_1, \dots, \mathfrak{s}_i$ are A-states then the memory cells encode the counter value i . As long as the counter value is smaller than K or if a BUC-state has been reached along a path where the counter value is always smaller than K , scheduler \mathfrak{G} just behaves like \mathfrak{T} . As soon as the counter value equals K , scheduler \mathfrak{G} switches to the second phase and behaves as scheduler $\mathfrak{R}_{N,s}$. More precisely, if \mathfrak{G} 's current state t in \mathcal{M} belongs to A then \mathfrak{G} mimics the behavior of $\mathfrak{R}_{N,s}$ from state (t, \top) . Thus, by following $\mathfrak{R}_{N,s}$'s decisions, \mathfrak{G} will only choose actions in Act^{\max} until A is left. As soon as state s is reached in \mathfrak{G} 's second phase (this will happen with probability 1 as $\mathfrak{R}_{N,s}$ minimizes the expected number of steps to $s = s_N$ from every state in the strongly connected MDP \mathcal{M}_N), \mathfrak{G} switches back to the first phase and restarts to mimic \mathfrak{T} from state s in mode χ , i.e., from state \mathfrak{s} in $\mathcal{B}^{\mathfrak{T}}$. For all states that are not reachable from \mathfrak{s} in this way, \mathfrak{G} behaves as \mathfrak{T} .

As \mathfrak{T} has a single BSCC, so does \mathfrak{G} , although the BSCC $\mathcal{B}^{\mathfrak{G}}$ induced by \mathfrak{G} can be different from $\mathcal{B}^{\mathfrak{T}}$. As \mathfrak{s} belongs to both $\mathcal{B}^{\mathfrak{T}}$ and $\mathcal{B}^{\mathfrak{G}}$, \mathfrak{s} is visited infinitely often almost surely with finite expected return time under both schedulers \mathfrak{G} and \mathfrak{T} .

Let us first observe that we indeed have $k^{\mathfrak{G}} < k^{\mathfrak{T}}$. This is thanks to the fact that (1) $\mathfrak{R}_{N,s}$ maximizes the probability for aUb whenever N or more consecutive A-states have been visited, (2) $K \geq N$ and (3) the reference state s is not an A-state. Thus, for each (BUC)-state u visited by $\mathfrak{R}_{N,s}$ in the return (second) phase of \mathfrak{G} we have:

$$\Pr_{\mathcal{B}^{\mathfrak{G}},u}(\bigcirc(AU \geq^K D^{\mathfrak{G}})) = 0$$

Hence, whenever u is (BUC)-state in $\mathcal{B}^{\mathfrak{G}}$ where $\Pr_{\mathcal{B}^{\mathfrak{G}},u}(\bigcirc(AU \geq^K D^{\mathfrak{G}}))$ is positive then

- $u \neq \mathfrak{s}$,
- u also belongs to $\mathcal{B}^{\mathfrak{T}}$ and
- the \mathfrak{G} -paths from u satisfying $\bigcirc(AU \geq^K D^{\mathfrak{G}})$ are also \mathfrak{T} -paths and satisfy $\bigcirc(AU \geq^K D^{\mathfrak{T}})$.

The last item yields $\Pr_{\mathcal{B}^{\mathfrak{T}},u}(\bigcirc(AU \geq^K D^{\mathfrak{T}})) > 0$. This completes the proof that $k^{\mathfrak{G}}$ is smaller than $k^{\mathfrak{T}}$.

We now show that $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{G}}(aUb) \geq \mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{T}}(aUb)$. To simplify the calculations, we present the proof for the case where $\Pi_{\mathfrak{s}}$ is a singleton, say $\Pi_{\mathfrak{s}} = \{\omega\}$.

Furthermore, let n be the length of ω (then $n \geq K$) and $t = (t, \chi) = \text{last}(\omega)$. Again, to simplify the calculations, let us suppose that there is a single action $\alpha \in Act(t) \setminus Act^{\max}(t)$ that \mathfrak{T} schedules for t with

positive probability p .³ So,

$$p = \mathfrak{T}(t)(\alpha) > 0$$

(Note that $p = 1$ if \mathfrak{T} is a deterministic FM-scheduler.)

The long run probabilities of the two schedulers \mathfrak{G} and \mathfrak{T} can be expressed as follows.

Given a state u in $\mathcal{B}^{\mathfrak{T}}$, let $ens_{u,\mathfrak{s}}^{\mathfrak{T}}$ be the expected number of steps from u to \mathfrak{s} under \mathfrak{T} (via paths from u to \mathfrak{s} of length at least 1 where all intermediate states are different from \mathfrak{s}). Let $eap_{u,\mathfrak{s}}^{\mathfrak{T}}$ denote the expected accumulated probability for aUb that \mathfrak{T} accumulates during this period. So, $ens_{\mathfrak{s},\mathfrak{s}}^{\mathfrak{T}}$ can be understood as the expected return time from and to \mathfrak{s} under \mathfrak{T} . Then:

$$\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{T}}(aUb) = \frac{eap_{\mathfrak{s},\mathfrak{s}}^{\mathfrak{T}}}{ens_{\mathfrak{s},\mathfrak{s}}^{\mathfrak{T}}} \quad (\dagger)$$

For the scheduler \mathfrak{G} we express $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{G}}(aUb)$ as the fraction of the expected accumulated probability along return paths and the expected return time from \mathfrak{s} to \mathfrak{s} as well:

$$\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{G}}(aUb) = \frac{eap_{\mathfrak{s},\mathfrak{s}}^{\mathfrak{G}}}{ens_{\mathfrak{s},\mathfrak{s}}^{\mathfrak{G}}}$$

To provide an upper bound for $ens_{\mathfrak{s},\mathfrak{s}}^{\mathfrak{G}}$ and a lower bound for $eap_{\mathfrak{s},\mathfrak{s}}^{\mathfrak{G}}$, we need several auxiliary notations.

Recall that $e_{t,s}$ is the expected number of steps that $\mathfrak{R}_{N,s}$ needs from t_N to s_N . Hence, $e_{t,s}$ is an upper bound for the expected number of steps $ens_{t,s}^{\mathfrak{G}}$ that \mathfrak{G} needs from state $t = \text{last}(\omega)$ to the reference state \mathfrak{s} . The value e has been defined as the maximum of the values $e_{t,s}$. Hence, we obtain:

$$e \geq e_{t,s} \geq ens_{t,s}^{\mathfrak{G}}$$

Let $ens_{t,\alpha,\mathfrak{s}}^{\mathfrak{T}}$ denote the the expected number of steps that \mathfrak{T} needs from t to \mathfrak{s} , under the assumption that action α is scheduled in t (which happens with probability p). So, if $(u_1, z_1), \dots, (u_\ell, z_\ell)$ denote the α -successors of t in $\mathcal{B}^{\mathfrak{T}}$ then:

$$ens_{t,\alpha,\mathfrak{s}}^{\mathfrak{T}} = 1 + \sum_{i=1}^{\ell} P(t, \alpha, u_i) \cdot ens_{(u_i, z_i), \mathfrak{s}}^{\mathfrak{T}}$$

Similarly, we define:

$$eap_{t,\alpha,\mathfrak{s}}^{\mathfrak{T}} = \Pr_t^{\mathfrak{T}}(aUb) + \sum_{i=1}^{\ell} P(t, \alpha, u_i) \cdot eap_{(u_i, z_i), \mathfrak{s}}^{\mathfrak{T}}$$

For $0 \leq i \leq n$, let $\wp_{[0..i]}$ denote the probability under \mathfrak{T} for generating the path fragment $\omega_{[0..i]}$ from state s in mode χ . So, $\wp_{[0..n]}$ is the probability under \mathfrak{T} for generating the full path ω from \mathfrak{s} .

³At the end of the proof, we briefly explain how to treat the general case where $\Pi_{\mathfrak{s}}$ is a (prefix-free) countable set of paths, for which \mathfrak{T} can schedule multiple actions not in Act^{\max} with positive probability.

For the expected number of steps $ens_{s,s}^{\mathfrak{S}}$ that \mathfrak{S} needs from s to s along paths of length at least 1, we get:

$$ens_{s,s}^{\mathfrak{S}} \leqslant ens_{s,s}^{\mathfrak{T}} + \wp_{[0\dots n]} \cdot p \cdot (e - ens_{t,\alpha,s}^{\mathfrak{T}}) \quad (*)$$

Proof of ():* Let us first observe that $ens_{s,s}^{\mathfrak{S}}$ can be written as an infinite sum:

$$ens_{s,s}^{\mathfrak{S}} = \sum_{\zeta \in \Xi} \Pr^{\mathfrak{S}}(\zeta) \cdot |\zeta|$$

where Ξ denotes the set of all finite paths ζ that start in s and end in s , and none of the intermediate states equals s .⁴

Just to simplify the notations, let us suppose that $\mathfrak{X}_{N,s}$ schedules an action $\beta \in Act^{\max}(t)$ for (t, T) such that $\mathfrak{T}(t)(\beta) = 0$. (This is not a proper restriction as we may extend the action alphabet of \mathcal{M} by a fresh action name for $\mathfrak{X}_{N,s}$'s decision for (t, T) .)

Let now Ξ_0 denote the set of paths ζ in Ξ such that either ω is not a prefix of ζ or ζ has the form $\omega \gamma \zeta'$ where γ is an action different from α and β . Then, all \mathfrak{S} -paths $\zeta \in \Xi_0$ are also \mathfrak{T} -paths with $\Pr^{\mathfrak{T}}(\zeta) = \Pr^{\mathfrak{S}}(\zeta)$, and vice versa, each \mathfrak{T} -path $\zeta \in \Xi_0$ is also a \mathfrak{S} -path.

Let now $\Xi^{\mathfrak{S}}$ denote the set of \mathfrak{S} -paths in $\Xi \setminus \Xi_0$. Then, all paths $\zeta \in \Xi_0$ have the form $\omega \beta \zeta'$. Let $\Xi^{\mathfrak{T}}$ denote the analogously defined set for scheduler \mathfrak{T} . That is, $\Xi^{\mathfrak{T}}$ consists of all finite \mathfrak{T} -paths ζ from s to s where no intermediate state equals s and where ζ has the form $\omega \alpha \zeta'$. Then, we get:

$$ens_{s,s}^{\mathfrak{S}} - ens_{s,s}^{\mathfrak{T}} = \sum_{\zeta \in \Xi^{\mathfrak{S}}} \Pr^{\mathfrak{S}}(\zeta) \cdot |\zeta| - \sum_{\zeta \in \Xi^{\mathfrak{T}}} \Pr^{\mathfrak{T}}(\zeta) \cdot |\zeta|$$

We now have (recall $t = last(\omega)$, $\Pr^{\mathfrak{S}}(\omega) = \Pr^{\mathfrak{T}}(\omega) = \wp_{[0\dots n]}$ and $\mathfrak{T}(t)(\alpha) = p$):

$$\sum_{\zeta \in \Xi^{\mathfrak{S}}} \Pr^{\mathfrak{S}}(\zeta) = \Pr^{\mathfrak{S}}(\omega) \cdot \mathfrak{T}(t)(\alpha) = \wp_{[0\dots n]} \cdot p$$

Each path in $\Xi^{\mathfrak{S}}$ has the form $\omega \beta \zeta'$ where $t \beta \zeta'$ can be viewed as a $\mathfrak{X}_{N,s}$ -path. Thus, the expected number of steps from t to s under \mathfrak{S} is bounded by the expected number of steps from t_N to s_N under $\mathfrak{X}_{N,s}$, which again is bounded by e . Using $n = |\omega|$, we obtain:

$$\sum_{\zeta \in \Xi^{\mathfrak{S}}} \Pr^{\mathfrak{S}}(\zeta) \cdot |\zeta| \leqslant \wp_{[0\dots n]} \cdot p \cdot (n + e)$$

Each path $\zeta \in \Xi^{\mathfrak{T}}$ has a prefix of the form $\omega \alpha(u_i, z_i)$ for some $i \in \{1, \dots, \ell\}$. Hence:

$$\sum_{\zeta \in \Xi^{\mathfrak{T}}} \Pr^{\mathfrak{T}}(\zeta) \cdot |\zeta| = \Pr^{\mathfrak{T}}(\omega) \cdot p \cdot (n + ens_{t,\alpha,s}^{\mathfrak{T}})$$

As $\Pr^{\mathfrak{T}}(\omega) = \Pr^{\mathfrak{S}}(\omega) = \wp_{[0\dots n]}$ we conclude:

$$ens_{s,s}^{\mathfrak{S}} - ens_{s,s}^{\mathfrak{T}} \geqslant \wp_{[0\dots n]} \cdot p \cdot (e - ens_{t,\alpha,s}^{\mathfrak{T}})$$

⁴Here, $\Pr^{\mathfrak{S}}(\zeta)$ stands for the probability for \mathfrak{S} to generate ζ from s . So, formally, $\Pr^{\mathfrak{S}}(\zeta)$ equals the probability measure of the cylinder set spanned by ζ under \mathfrak{S} .

This complete the proof of (*).

The next goal is to provide a lower bound for the expected accumulated probability $eap_{s,s}^{\mathfrak{S}}$. The claim is:

$$eap_{s,s}^{\mathfrak{S}} \geqslant eap_{s,s}^{\mathfrak{T}} + \wp_{[0\dots n]} \cdot p \cdot (e - eap_{t,\alpha,s}^{\mathfrak{T}}) \quad (**)$$

*Proof of (**):* Let Ξ be as in the proof of (*) and let $ppref(\cdot)$ denote the set of proper prefixes of the paths in (\cdot) . Furthermore, we write $\Pr_{last(\zeta)}^{\mathfrak{S}}(aUb)$ for the probability for aUb in the Markov chain induced by \mathfrak{S} from state $last(\zeta)$.

Then, $eap_{s,s}^{\mathfrak{S}}$ can be written as an infinite sum:

$$\begin{aligned} eap_{s,s}^{\mathfrak{S}} &= \sum_{\pi \in \Xi} \Pr^{\mathfrak{S}}(\pi) \cdot \sum_{\zeta \in ppref(\pi)} \Pr_{last(\zeta)}^{\mathfrak{S}}(aUb) \\ &= \sum_{\zeta \in ppref(\Xi)} \Pr_{last(\zeta)}^{\mathfrak{S}}(aUb) \cdot \sum_{\substack{\pi \in \Xi \text{ s.t.} \\ \zeta \in ppref(\pi)}} \Pr^{\mathfrak{S}}(\pi) \\ &= \sum_{\zeta \in ppref(\Xi)} \Pr^{\mathfrak{S}}(\zeta) \cdot \Pr_{last(\zeta)}^{\mathfrak{S}}(aUb) \end{aligned}$$

We deal here with the proper prefixes of the paths in Ξ to avoid that the reference state s , which is the starting and end state of the paths in Ξ , is considered twice in the sums for the accumulated weight (probability values) of the paths $\pi \in \Xi$.

The analogous formula holds for \mathfrak{T} . Let Ξ_0 , $\Xi^{\mathfrak{S}}$ and $\Xi^{\mathfrak{T}}$ be defined as in the proof of (*). For the paths ζ that are prefixes of some path in Ξ_0 , but not a prefix of ω , we have $\Pr^{\mathfrak{S}}(\zeta) = \Pr^{\mathfrak{T}}(\zeta)$ and $\Pr_{last(\zeta)}^{\mathfrak{S}}(aUb) = \Pr_{last(\zeta)}^{\mathfrak{T}}(aUb)$.

Thus, the difference between $eap_{s,s}^{\mathfrak{S}}$ and $eap_{s,s}^{\mathfrak{T}}$ is:

$$\begin{aligned} &\sum_{\zeta \in ppref(\Xi^{\mathfrak{S}})} \Pr^{\mathfrak{S}}(\zeta) \cdot \Pr_{last(\zeta)}^{\mathfrak{S}}(aUb) \\ &- \sum_{\zeta \in ppref(\Xi^{\mathfrak{T}})} \Pr^{\mathfrak{T}}(\zeta) \cdot \Pr_{last(\zeta)}^{\mathfrak{T}}(aUb) \end{aligned}$$

The prefixes $\omega_{[0\dots i]}$ of ω are contained in both sets $ppref(\Xi^{\mathfrak{S}})$ and $ppref(\Xi^{\mathfrak{T}})$. These have the same probability under \mathfrak{S} and \mathfrak{T} (namely $\wp_{[0\dots i]}$) and for $i = 1, \dots, n$:

$$\begin{aligned} &\Pr_{\omega_{[i]}}^{\mathfrak{S}}(aUb) - \Pr_{\omega_{[i]}}^{\mathfrak{T}}(aUb) \\ &\geqslant \wp_{[i\dots n]} \cdot p \cdot (p_t^{\max} - p_{t,\alpha}) \end{aligned}$$

Here, we write $\wp_{[i\dots n]}$ for the probability under \mathfrak{T} or \mathfrak{S} for generating the path fragment $\omega_{[i\dots n]}$ from state $\omega_{[i]}$, and as before $p_{t,\alpha} = \sum_u P(t, \alpha, u) \cdot p_u^{\max}$.

With $ppref(\omega)$ denoting the set of all prefixes of ω (including ω) and using that $first(\omega) = s$ is a $(B \cup C)$ -

state, we get:

$$\begin{aligned}
& \sum_{\zeta \in \text{pref}(\omega)} \Pr^{\mathfrak{S}}(\zeta) \cdot \Pr_{last(\zeta)}^{\mathfrak{S}}(aUb) \\
& - \sum_{\zeta \in \text{pref}(\omega)} \Pr^{\mathfrak{T}}(\zeta) \cdot \Pr_{last(\zeta)}^{\mathfrak{T}}(aUb) \\
& = \sum_{i=1}^n \wp_{[0\dots i]} \cdot \left(\Pr_{\omega[i]}^{\mathfrak{S}}(aUb) - \Pr_{\omega[i]}^{\mathfrak{T}}(aUb) \right) \\
& \geq \sum_{i=1}^n \wp_{[0\dots i]} \cdot \wp_{[i\dots n]} \cdot p \cdot (p_t^{\max} - p_{t,\alpha}) \\
& \geq n \cdot \wp_{[0\dots n]} \cdot p \cdot \delta
\end{aligned}$$

where we used $\wp_{[0\dots n]} = \wp_{[0\dots i]} \cdot \wp_{[i\dots n]}$ and

$$p_t^{\max} - p_{t,\alpha} \geq \delta$$

(by definition of δ and as $\alpha \notin Act^{\max}(t)$).

Furthermore, the \mathfrak{T} -paths $\zeta \in \text{ppref}(\Xi^{\mathfrak{T}})$ that do not belong to $\text{pref}(\omega)$ are extensions of ω of the form $\omega\alpha\zeta'$. Their cumulative eap -value is:

$$\sum_{\substack{\zeta \in \text{ppref}(\Xi^{\mathfrak{T}}) \\ \zeta \notin \text{pref}(\omega)}} \Pr^{\mathfrak{T}}(\zeta) \cdot \Pr_{last(\zeta)}^{\mathfrak{T}} = \Pr^{\mathfrak{T}}(\omega) \cdot p \cdot eap_{t,\alpha,s}^{\mathfrak{T}}$$

Using the fact that $\Pr^{\mathfrak{T}}(\omega)$ equals $\wp_{[0\dots n]}$ and that

$$\sum_{\zeta \in \text{ppref}(\Xi^{\mathfrak{S}})} \Pr^{\mathfrak{S}}(\zeta) \cdot \Pr_{last(\zeta)}^{\mathfrak{S}} \geq \sum_{\zeta \in \text{pref}(\omega)} \Pr^{\mathfrak{S}}(\zeta) \cdot \Pr_{last(\zeta)}^{\mathfrak{S}}$$

we obtain:

$$\begin{aligned}
& eap_{s,s}^{\mathfrak{S}} - eap_{s,s}^{\mathfrak{T}} \\
& \geq n \cdot \wp_{[0\dots n]} \cdot p \cdot \delta - \wp_{[0\dots n]} \cdot p \cdot eap_{t,\alpha,s}^{\mathfrak{T}} \\
& = \wp_{[0\dots n]} \cdot p \cdot (n \cdot \delta - eap_{t,\alpha,s}^{\mathfrak{T}})
\end{aligned}$$

We have $n \geq K$ (as ω is a path of Π_s and by definition of Π_s). Moreover, $K \cdot \delta \geq e$ (by the choice of K). We conclude:

$$n \cdot \delta \geq e$$

and therefore:

$$eap_{s,s}^{\mathfrak{S}} - eap_{s,s}^{\mathfrak{T}} \geq \wp_{[0\dots n]} \cdot p \cdot (e - eap_{t,\alpha,s}^{\mathfrak{T}})$$

This completes the proof of (**).

With $q = \wp_{[0\dots n]} \cdot p$, we obtain by (*) and (**):

$$\begin{aligned}
ens_{s,s}^{\mathfrak{S}} & \leq ens_{s,s}^{\mathfrak{T}} + q \cdot (e - ens_{t,\alpha,s}^{\mathfrak{T}}) \\
eap_{s,s}^{\mathfrak{S}} & \geq eap_{s,s}^{\mathfrak{T}} + q \cdot (e - eap_{t,\alpha,s}^{\mathfrak{T}})
\end{aligned}$$

and therefore:

$$\mathbb{LP}_{\mathcal{M}}^{\mathfrak{S}}(aUb) \geq \frac{eap_{s,s}^{\mathfrak{T}} + q \cdot (e - eap_{t,\alpha,s}^{\mathfrak{T}})}{ens_{s,s}^{\mathfrak{T}} + q \cdot (e - ens_{t,\alpha,s}^{\mathfrak{T}})} \quad (\ddagger)$$

We now use (\ddagger) and (\ddagger) to show that $\mathbb{LP}_{\mathcal{M}}^{\mathfrak{S}}(aUb) \geq \mathbb{LP}_{\mathcal{M}}^{\mathfrak{T}}(aUb)$.

Obviously, the expected number of steps is an upper bound for the expected accumulated probability for aUb . In particular:

$$eap_{t,\alpha,s}^{\mathfrak{T}} \leq ens_{t,\alpha,s}^{\mathfrak{T}}$$

Hence, if $e = ens_{t,\alpha,s}^{\mathfrak{T}}$ then $e - eap_{t,\alpha,s}^{\mathfrak{T}} \geq 0$ and therefore:

$$\mathbb{LP}_{\mathcal{M}}^{\mathfrak{S}}(aUb) \geq \frac{eap_{s,s}^{\mathfrak{T}}}{ens_{s,s}^{\mathfrak{T}}} = \mathbb{LP}_{\mathcal{M}}^{\mathfrak{T}}(aUb)$$

Suppose now that $e \neq ens_{s,\alpha,s}^{\mathfrak{T}}$.

If $e > ens_{s,\alpha,s}^{\mathfrak{T}}$ then

$$\frac{e - eap_{s,\alpha,s}^{\mathfrak{T}}}{e - ens_{s,\alpha,s}^{\mathfrak{T}}} \geq \frac{e - ens_{s,\alpha,s}^{\mathfrak{T}}}{e - ens_{s,\alpha,s}^{\mathfrak{T}}} = 1 \geq \frac{eap_{s,s}^{\mathfrak{T}}}{ens_{s,s}^{\mathfrak{T}}}$$

We now use the fact that if x, y, z, w are non-negative real numbers with $w, y > 0$ then:

$$\frac{x+z}{y+w} \geq \frac{x}{y} \quad \text{iff} \quad \frac{z}{w} \geq \frac{x}{y}$$

This yields:

$$\mathbb{LP}_{\mathcal{M}}^{\mathfrak{S}}(aUb) \geq \frac{eap_{s,s}^{\mathfrak{T}}}{ens_{s,s}^{\mathfrak{T}}} = \mathbb{LP}_{\mathcal{M}}^{\mathfrak{T}}(aUb)$$

It remains to consider the case $e < ens_{s,\alpha,s}^{\mathfrak{T}}$. Here, we use the fact for that all non-negative rational numbers x, y, z, w with $y > w > 0$ then:

$$\frac{x-z}{y-w} \geq \frac{x}{y} \quad \text{iff} \quad \frac{z}{w} \leq \frac{x}{y}$$

In particular, if $x \leq y$ and $0 < z < y$ then $(x-z)/(y-z) \leq x/y$. Hence:

$$\frac{eap_{s,\alpha,s}^{\mathfrak{T}} - e}{ens_{s,\alpha,s}^{\mathfrak{T}} - e} \leq \frac{eap_{s,s}^{\mathfrak{T}}}{ens_{s,s}^{\mathfrak{T}}}$$

which again yields:

$$\begin{aligned}
\mathbb{LP}_{\mathcal{M}}^{\mathfrak{S}}(aUb) & \geq \frac{eap_{s,s}^{\mathfrak{T}} - q \cdot (eap_{t,\alpha,s}^{\mathfrak{T}} - e)}{ens_{s,s}^{\mathfrak{T}} - q \cdot (ens_{t,\alpha,s}^{\mathfrak{T}} - e)} \\
& \geq \frac{eap_{s,s}^{\mathfrak{T}}}{ens_{s,s}^{\mathfrak{T}}} = \mathbb{LP}_{\mathcal{M}}^{\mathfrak{T}}(aUb)
\end{aligned}$$

For the general case where Π_s is a (countable and prefix-free) set of paths ω , for which \mathfrak{T} schedules several actions not in Act^{\max} , the argument is fairly the same. The essential difference is that we have to consider all state-action pairs (t, α) in $\mathcal{B}^{\mathfrak{T}}$ such that state $t = (t, y)$ is reachable from s via a path in Π_s and $\alpha \in Act(t) \setminus Act^{\max}(t)$ such that \mathfrak{T} schedules action α for t with positive probability. So, the lower bound for $\mathbb{LP}_{\mathcal{M},s}^{\mathfrak{S}}(aUb)$ in (\ddagger) then has the form:

$$\mathbb{LP}_{\mathcal{M}}^{\mathfrak{S}}(aUb) \geq \frac{eap_{s,s}^{\mathfrak{T}} + \sum_{t,\alpha} q_{t,\alpha} \cdot (e - eap_{t,\alpha,s}^{\mathfrak{T}})}{ens_{s,s}^{\mathfrak{T}} + \sum_{t,\alpha} q_{t,\alpha} \cdot (e - ens_{t,\alpha,s}^{\mathfrak{T}})}$$

where $q_{t,\alpha}$ is the product of the probability in $\mathcal{B}^{\mathfrak{T}}$ of reaching t from s along a path in Π_s and the probability for taking action α in state t of $\mathcal{B}^{\mathfrak{T}}$. \square

Max-lrp via max-MP

We now turn to the soundness of the constructed MDP \mathcal{K} that encodes a counter for the number of consecutive A-states that have been visited in a suffix of the current path.

Let us observe that \mathcal{K} might be not strongly connected, even if \mathcal{M} is strongly connected. The reason is that not all states in \mathcal{K} might be reachable from the states $s_{\mathcal{K}}$. If, however, \mathcal{M} is strongly connected and we restrict \mathcal{K} to the states reachable from some states $s_{\mathcal{K}}$ then \mathcal{K} is strongly connected. This justifies to drop the starting state when talking maximal long-run probabilities or the maximal expected mean payoff in \mathcal{K} .

We start with the observation that the switch from \mathcal{M} to \mathcal{K} preserves the maximal long-run probabilities for aUb .

Lemma A.14. $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(aUb) = \mathbb{L}\mathbb{P}_{\mathcal{K}}^{\max}(aUb)$

Proof. obvious from Lemma IV.7. \square

The remaining task is to show that $\mathbb{L}\mathbb{P}_{\mathcal{K}}^{\max}(aUb)$ is the maximal expected mean payoff in \mathcal{K} . For this, we switch from \mathcal{K} to the following MDP \mathcal{N} which uses with state- rather than transition-based weights. The definition of the weighted MDP \mathcal{N} is as follows. The state space of \mathcal{N} is:

$$S_{\mathcal{N}} = C \times \{0\} \cup (A \cup B) \times \{1, \dots, K, T\}$$

The new MDP \mathcal{N} has the same action set as \mathcal{M} and the following transition probabilities:

- If $s \in C$ then $P_{\mathcal{N}}((s, 0), \alpha, (s', k)) = P(s, \alpha, s')$ if either $s' \in C \wedge k = 0$ or $s' \in A \cup B \wedge k = 1$.
- If $(s, i) \in A \times \{1, \dots, K-1\}$ then:
 - $P_{\mathcal{N}}((s, i), \alpha, (s', 0)) = P(s, \alpha, s')$ if $s' \in C$
 - $P_{\mathcal{N}}((s, i), \alpha, (s', i+1)) = P(s, \alpha, s')$ if $s' \in A \cup B$
- If $s \in A$, $k \in \{K, T\}$ and $\alpha \in Act^{\max}(s)$ then:
 - $P_{\mathcal{N}}((s, k), \alpha, (s', 0)) = P(s, \alpha, s')$ if $s' \in C$
 - $P_{\mathcal{N}}((s, k), \alpha, (s', T)) = P(s, \alpha, s')$ if $s' \in A \cup B$
- If $s \in B$ and $k \in \{1, \dots, K, T\}$ then:
 - $P_{\mathcal{N}}((s, k), \alpha, (s', 0)) = P(s, \alpha, s')$ if $s' \in C$
 - $P_{\mathcal{N}}((s, k), \alpha, (s', 1)) = P(s, \alpha, s')$ if $s' \in A \cup B$

And $P_{\mathcal{N}}(\cdot) = 0$ in all remaining cases. The weight function of \mathcal{N} assigns rational values to states:

- $wgt_{\mathcal{N}}(s, k) = k$ if $s \in B$ and $k \in \{1, \dots, K\}$
- $wgt_{\mathcal{N}}(s, K) = p_s^{\max} \cdot K$ if $s \in A$
- $wgt_{\mathcal{N}}(s, T) = p_s^{\max}$ if $s \in A \cup B$

Note that for $s \in B$ we have $wgt_{\mathcal{N}}(s, T) = 1$.

If s is a state of \mathcal{M} then $s_{\mathcal{N}}$ denotes the corresponding state in \mathcal{N} , namely $s_{\mathcal{N}} = (s, 0)$ for $s \in C$ and $s_{\mathcal{N}} = (s, 1)$ if $s \in A \cup B$.

As for \mathcal{K} , with the above definition \mathcal{N} might be not strongly connected. If, however, we restrict \mathcal{N} to the fragment that is reachable from the state $s_{\mathcal{N}}$ then \mathcal{N} is strongly connected.

Lemma A.15. $\mathbb{L}\mathbb{P}_{\mathcal{K}}^{\max}(aUb) = \mathbb{L}\mathbb{P}_{\mathcal{N}}^{\max}(aUb)$ and $\mathbb{E}_{\mathcal{K}}^{\max}(\text{MIP}) = \mathbb{E}_{\mathcal{N}}^{\max}(\text{MIP})$

Proof. obvious as \mathcal{N} can be viewed as a variant of \mathcal{K} with state-based (rather than transition-based) weights. \square

Hence, the remaining is to prove that $\mathbb{L}\mathbb{P}_{\mathcal{N}}^{\max}(aUb)$ coincides with $\mathbb{E}_{\mathcal{N}}^{\max}(\text{MIP})$.

Lemma A.16. For each FM-scheduler \mathfrak{S} for \mathcal{N} :

$$\mathbb{E}_{\mathcal{N}, s_{\mathcal{N}}}^{\mathfrak{S}}(\text{MIP}) \geq \mathbb{L}\mathbb{P}_{\mathcal{N}, s_{\mathcal{N}}}^{\mathfrak{S}}(aUb)$$

Furthermore, if \mathfrak{S} leaves $A \times \{T\}$ with probability 1 whenever it reaches this set of states, i.e., if the Markov chain induced by \mathfrak{S} has no BSCC consisting of states in $A \times \{T\}$, then equality holds.

Proof. Let \mathfrak{S} be an FM-scheduler for \mathcal{N} and let $\mathcal{C} = \mathcal{C}^{\mathfrak{S}}$ be the induced finite Markov chain with $S_{\mathcal{C}} = S_{\mathcal{N}} \times X$ where X denotes the set of modes of \mathfrak{S} . Thus, the states of \mathcal{C} are triples (s, k, x) with $s \in S$, $k \in \{0, 1, \dots, K, T\}$ and $x \in X$ where $k = 0$ if $s \in C$ and $k \in \{1, \dots, K, T\}$ if $s \in A \cup B$.

It suffices to prove that the expected mean payoff of each BSCC \mathcal{B} of \mathcal{C} is greater or equal than the long-run probability for aUb in \mathcal{B} . In what follows, we fix a BSCC \mathcal{B} of \mathcal{C} and use the following notations.

- Symbol \mathfrak{s} is used to denote states in \mathcal{B} .
- $\theta_{\mathfrak{s}}$ denotes the long-run frequency (steady-state probability) of state \mathfrak{s} in \mathcal{B} .
- $p_{\mathfrak{s}}$ denotes the probability for aUb in \mathcal{C} (resp. \mathcal{B}) from state \mathfrak{s} .
- $w_{\mathfrak{s}}$ denotes the weight of state \mathfrak{s} in \mathcal{C} , i.e., if $\mathfrak{s} = (s, k, x)$ then $w_{\mathfrak{s}} = wgt_{\mathcal{N}}(s, k)$.

For sets of A- and B-states inside the given BSCC \mathcal{B} we shall use the following notations:

$$\mathfrak{A} = \mathcal{B} \cap (A \times \{1, \dots, K\} \times X)$$

$$\mathfrak{B} = \mathcal{B} \cap (B \times \{1, \dots, K\} \times X)$$

We write \mathfrak{A}_T for \mathfrak{A} extended by the set of states $(s, T, x) \in \mathcal{B}$ with $s \in A$. Similarly, \mathfrak{B}_T is \mathfrak{B} extended by the states $(s, T, x) \in \mathcal{B}$ with $s \in B$.

Treating $\theta = (\theta_{\mathfrak{s}})_{\mathfrak{s} \in S_{\mathcal{C}}}$ as a row vector and $p = (p_{\mathfrak{s}})_{\mathfrak{s} \in S_{\mathcal{C}}}$ and $w = (w_{\mathfrak{s}})_{\mathfrak{s} \in S_{\mathcal{C}}}$ as column vectors:

$$\mathbb{E}_{\mathcal{N}, s_{\mathcal{N}}}^{\mathfrak{S}}(\text{MIP}) - \mathbb{L}\mathbb{P}_{\mathcal{N}, s_{\mathcal{N}}}^{\mathfrak{S}}(aUb) = \langle \theta, w - p \rangle$$

where $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner product. Using the fact that $w_{\mathfrak{s}} = p_{\mathfrak{s}} = 1$ for each state $\mathfrak{s} = (s, k, x)$ in

\mathcal{C} with $s \in B$ and $k = \top$ and $w_s = p_s = 0$ for each state $s = (s, 1, x)$ in \mathcal{C} with $s \in C$ we obtain:

$$\begin{aligned} \langle \theta, w-p \rangle &= \sum_{s \in S_{\mathcal{C}}} \theta_s \cdot (w_s - p_s) \\ &= \sum_{s \in \mathfrak{A} \cup \mathfrak{B}_{\top}} \theta_s \cdot (w_s - p_s) \\ &\geq \sum_{s \in \mathfrak{A} \cup \mathfrak{B}} \theta_s \cdot (w_s - p_s) \end{aligned}$$

If \mathfrak{S} leaves $A \times \{\top\}$ with probability 1 after entering, then equality holds, as then $w_s = p_s = p_s^{\max}$ for each state $s = (s, k, x)$ in \mathcal{B} with $s \in A$ and $k = \top$.

We define the matrix

$$\mathfrak{P} \in [0, 1]^{\mathfrak{A} \cup \mathfrak{B}}$$

as the transition probability matrix of \mathcal{C} restricted to the states in $\mathfrak{A} \cup \mathfrak{B}$ and with all outgoing transitions from states in B and $A \times \{K\}$ removed, i.e., the rows with index (s, k, x) where $s \in B$ or $k = K$ are set to 0. Let θ' , p' and w' be the vectors θ , p and w projected to $\mathfrak{A} \cup \mathfrak{B}$.

Furthermore, let $\theta'[i]$ be the vector obtained from θ' by setting all entries not indexed by (s, i, x) , $s \in S$, $x \in X$, to 0 and define $p'[i]$ analogously. Let now i, j be integers with $K \geq j \geq i \geq 1$. Then:

$$\theta'[j] = \theta'[i] \cdot \mathfrak{P}^{j-i} \quad \text{and} \quad \theta' = \sum_{i=0}^{K-1} \theta'[1] \cdot \mathfrak{P}^i.$$

Intuitively, the left equation states that we can compute the steady state probabilities for states with index i using only the steady state probabilities for states with index j and the $j-i$ steps transition probability matrix \mathfrak{P}^{j-i} . This follows from the fact that states indexed by j can only be reached from states indexed by i in $j-i$ steps.

We now regard the vector $d \in [0, 1]^{\mathfrak{A} \cup \mathfrak{B}}$ with $d_{(s, k, x)} = 1$ if $s \in B$, $d_{(s, k, x)} = p_s^{\max}$, and $d_{(s, k, x)} = 0$ in all other cases. Then:

$$p'[i] \leq \sum_{j=0}^{K-i} \mathfrak{P}^j \cdot d$$

where the inequality is understood componentwise. Again, equality holds if \mathfrak{S} always leaves $A \times \{\top\}$ almost surely.

So, we get:

$$\begin{aligned} &\sum_{s \in \mathfrak{A} \cup \mathfrak{B}} \theta_s \cdot p_s \\ &= \sum_{k=1}^K \langle \theta'[k], p'[k] \rangle \leq \sum_{k=1}^K \langle \theta'[k], \sum_{j=0}^{K-k} \mathfrak{P}^j \cdot d \rangle \\ &= \sum_{k=1}^K \sum_{j=0}^{K-k} \theta'[k] \cdot \mathfrak{P}^j \cdot d \\ &= \sum_{k=1}^K \sum_{j=0}^{K-k} \theta'[1] \cdot \mathfrak{P}^{j+k-1} \cdot d \\ &= \sum_{k=1}^K \sum_{i=K-k}^K \theta'[i] \cdot d = \sum_{k=1}^K k \cdot q'[k] \cdot d \end{aligned}$$

But now, we see that the weight structure in \mathcal{N} was just defined such that the last line sums up to $\langle \theta', w' \rangle$. Thus:

$$\sum_{s \in \mathfrak{A} \cup \mathfrak{B}} \theta_s \cdot p_s = \sum_{s \in \mathfrak{A} \cup \mathfrak{B}} \theta_s \cdot w_s$$

We conclude that $\langle \theta, w-p \rangle \geq 0$ and $\langle \theta, w-p \rangle = 0$ if \mathfrak{S} leaves $A \times \{\top\}$ almost surely.

This yields $\mathbb{E}_{\mathcal{B}}^{\mathfrak{S}}(\text{MIP}) \geq \text{LP}_{\mathcal{B}}^{\mathfrak{S}}(aUb)$ in the general case and $\mathbb{E}_{\mathcal{B}}^{\mathfrak{S}}(\text{MIP}) = \text{LP}_{\mathcal{B}}^{\mathfrak{S}}(aUb)$ if \mathfrak{S} leaves $A \times \{\top\}$ almost surely. \square

Lemma A.17. $\mathbb{E}_{\mathcal{N}}^{\max}(\text{MIP}) = \text{LP}_{\mathcal{N}}^{\max}(aUb)$

Proof. Thanks to Lemma A.16, it suffices to construct a scheduler \mathfrak{T} for \mathcal{N} such that $\text{LP}_{\mathcal{N}, s_{\mathcal{N}}}^{\mathfrak{T}}(aUb) = \mathbb{E}_{\mathcal{N}}^{\max}(\text{MIP})$ for all states $s_{\mathcal{N}}$.

Let \mathfrak{S} be an MD-scheduler for \mathcal{N} that maximizes the expected mean payoff from every state in \mathcal{N} . We may assume w.l.o.g. that the induced Markov chain (restricted to the states that are reachable from the states $s_{\mathcal{N}}$) has a single BSCC, say \mathcal{B} .

Case 1: \mathcal{B} is not contained in $A \times \{\top\}$. Then \mathfrak{S} always leaves $A \times \{\top\}$ with probability 1. The second part of Lemma A.16 yields:

$$\mathbb{E}_{\mathcal{N}}^{\max}(\text{MIP}) = \mathbb{E}_{\mathcal{N}, s_{\mathcal{N}}}^{\mathfrak{S}}(\text{MIP}) = \text{LP}_{\mathcal{M}, s_{\mathcal{N}}}^{\mathfrak{S}}(aUb)$$

Hence, we can deal with $\mathfrak{S} = \mathfrak{T}$.

Case 2: \mathcal{B} is contained in $A \times \{\top\}$. In this case, we can construct an infinite-memory scheduler \mathfrak{T} with $\text{LP}_{\mathcal{N}, s_{\mathcal{N}}}^{\mathfrak{T}}(aUb) = \mathbb{E}_{\mathcal{N}, s_{\mathcal{N}}}^{\mathfrak{S}}(\text{MIP})$. The idea is that scheduler \mathfrak{T} stays inside \mathcal{B} for larger and larger number of steps and in between it leaves \mathcal{B} to maximize the probability for aUb . In this way, the weight of states in \mathcal{B} is equal to the probability of aUb under \mathfrak{T} and the long-run probability then equals this probability. \square

Corollary A.18 (See Lemma IV.8). $\text{LP}_{\mathcal{N}}^{\max}(aUb) = \mathbb{E}_{\mathcal{N}}^{\max}(\text{MIP})$.

Proof. The claim follows by combining the results that have been established so far:

$$\begin{aligned}
\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\alpha U b) &= \mathbb{L}\mathbb{P}_{\mathcal{K}}^{\max}(\alpha U b) && \text{(Lemma A.14)} \\
&= \mathbb{L}\mathbb{P}_{\mathcal{N}}^{\max}(\alpha U b) && \text{(Lemma A.15)} \\
&= \mathbb{E}_{\mathcal{N}}^{\max}(\text{MIP}) && \text{(Lemma A.17)} \\
&= \mathbb{E}_{\mathcal{K}}^{\max}(\text{MIP}) && \text{(Lemma A.15)}
\end{aligned}$$

□

NP-hardness

We now turn to the threshold problem for long-run probabilities and until-properties that takes as input a strongly connected \mathcal{M} , atomic propositions α , b and a rational value ϑ and asks whether $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\alpha U b) \geq \vartheta$. By the results that have been established so far, this problem belongs to EXPTIME. We now provide an NP lower bound.

Theorem A.19. *The threshold problem “is $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\alpha U b) \geq \vartheta$?” is NP-hard.*

Proof. We prove the statement by a polynomial reduction from the intersection problem for unary DFA, i.e., DFA over a one-letter alphabet. This problem is known to be NP-complete [5].

So, we are given a finite number of unary DFA, say $\mathcal{D}_1, \dots, \mathcal{D}_k$ over the alphabet $\Sigma = \{0\}$, where $\mathcal{D}_i = (Q_i, \Sigma, \delta_i, q_{0,i}, F_i)$. We simply write $\delta_i(q)$ rather than $\delta_i(q, 0)$. We may suppose the transition functions δ_i are total and that $Q_i \cap Q_j = \emptyset$ if $i \neq j$. W.l.o.g. we further assume that $|Q_i| \geq 2$ for all $i \leq k$.

We are going to construct an MDP \mathcal{M} over $\text{AP} = \{\alpha, b\}$ and a rational value ϑ such that $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\max}(\alpha U b) \geq \vartheta$ if and only if $\mathcal{L}(\mathcal{D}_1) \cap \dots \cap \mathcal{L}(\mathcal{D}_k)$ is nonempty. Obviously, the latter is equivalent to the statement that there exists some $n \in \mathbb{N}$ with $n < \ell$ and $0^n \in \mathcal{L}(\mathcal{D}_1) \cap \dots \cap \mathcal{L}(\mathcal{D}_k)$ where $\ell = |Q_1| \cdot \dots \cdot |Q_k|$.

Let $\mathcal{A} = (Q, \Sigma, \delta, Q_0, F)$ denote the NFA resulting from the union of $\mathcal{D}_1, \dots, \mathcal{D}_k$. That is, $Q = Q_1 \cup \dots \cup Q_k$, $Q_0 = \{q_{0,1}, \dots, q_{0,k}\}$, $F = F_1 \cup \dots \cup F_k$ and $\delta(q) = \delta_i(q)$ if $q \in Q_i$. That is, besides the nondeterministic choice of the initial state, \mathcal{A} behaves deterministically. The idea is now to treat \mathcal{A} as a substructure of an MDP \mathcal{M} where all states of \mathcal{M} are labeled by α . The substructure \mathcal{A} can be entered in \mathcal{M} via an action that assigns probabilities $1/k$ to each of the k initial states, possibly using a “pumping option” to increase the number of α -states before simulating \mathcal{M} . In each state $q \in Q$ an action α is enabled that mimics \mathcal{A} ’s transition from q with probability $(r-1)/r$ and moves back to the state where \mathcal{A} can be entered (in which neither α nor b holds) with probability $1/r$. For the final states $q \in F$ there is also an β that leads with

probability 1 to a b -state form which \mathcal{M} can return to the state where \mathcal{M} can be entered. Other states and transitions of \mathcal{M} are needed for technical purposes.

The idea of this construction is as follows. To maximize the long-run probability for $\alpha U b$ in \mathcal{M} , a scheduler guesses a natural number n with $n \leq \ell$ and $0^n \in \mathcal{L}(\mathcal{D}_1) \cap \dots \cap \mathcal{L}(\mathcal{D}_k)$. It then uses the pumping option to enter \mathcal{M} after having generated a sequence of $\ell-n$ α -states. It then attempts to make n steps inside \mathcal{A} using action α and leaves \mathcal{A} using the β transition.

Formally, the state space of \mathcal{M} is

$$S = Q \cup \{\alpha, b, c, \text{init}\}$$

where the states in $Q \cup \{\alpha\}$ are labeled by α and b is labeled by b . The action set is $\text{Act} = \{\alpha, \beta, \text{enter}, \text{pump}, \tau\}$. The transition probabilities are as follows (where r and r' are rational numbers > 1 defined later):

- In $s \in \{\text{init}, \alpha\}$, actions *enter* and *pump* are enabled with the transition probabilities:

$$\begin{aligned}
P(s, \text{enter}, q_{0,i}) &= \frac{r-1}{k \cdot r}, \quad i = 1, \dots, k, \\
P(s, \text{enter}, \text{init}) &= \frac{1}{r}, \\
P(s, \text{pump}, \alpha) &= \frac{r-1}{r}, \quad P(s, \text{pump}, \text{init}) = \frac{1}{r}.
\end{aligned}$$

- In each state $q \in Q$, action α is enabled with:

$$P(q, \alpha, \delta(q)) = \frac{r-1}{r}, \quad P(q, \alpha, \text{init}) = \frac{1}{r}$$

For the final states $q \in F$, additionally action β is enabled with $P(q, \beta, b) = 1$.

- In $s' \in \{b, c\}$, action τ is enabled with:

$$P(s', \tau, c) = \frac{r'-1}{r'}, \quad P(s', \tau, \text{init}) = \frac{1}{r'}$$

Scheduler \mathfrak{S}_n and its long-run probability.: Before defining the values r and r' , let us suppose n is an integer with $n < \ell$ and $0^n \in \mathcal{L}(\mathcal{D}_1) \cap \dots \cap \mathcal{L}(\mathcal{D}_k)$. Regard the FMD-scheduler \mathfrak{S}_n which behaves as follows. Let $m = \ell - n - 1$. From state *init*, \mathfrak{S}_n attempts to reach α and stay there via the pumping option *pump* until the generated path ends in m consecutive visits to α . Afterwards the scheduler \mathfrak{S}_n enters \mathcal{A} via *enter*. Let $q_{0,i}$ be the state reached after that sequence. \mathfrak{S}_n then attempts to follow the unique accepting run for 0^n in \mathcal{D}_i (via action α). If successful it reaches a final state $q \in F_i$. It then takes action β and returns to *init* via the τ -transitions in states b and c . Of course, each of the attempts might fail, in which case \mathfrak{S}_n returns to *init* and behaves in the same way.

The long-run probability $\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathfrak{S}_n}(\alpha U b)$ under \mathfrak{S}_n can be computed as the quotient of expected accumulated probability and expected number of steps between two visits to *init*: Let $\rho = \frac{r-1}{r}$. For $1 \leq k \leq \ell$, the scheduler succeeds in generating a path starting with k consecutive α -states with probability ρ^k . The probability to satisfy $\alpha U b$ afterwards is $\rho^{\ell-k}$. In addition b is

reached with probability ρ^ℓ and here the probability of \mathbf{aUb} is 1. So, the expected accumulated probability is

$$\left(\sum_{k=1}^{\ell} \rho^k \cdot \rho^{\ell-k} \right) + \rho^\ell = \rho^\ell (\ell + 1).$$

For the expected number of steps, we get the following: The probability to return to *init* in exactly $1 \leq k \leq \ell$ steps is given by $(1-\rho)\rho^{k-1}$. With probability ρ^ℓ , the state \mathbf{b} is reached in which case the expected return time is $\ell + 1 + r'$. All in all, the expected number of steps is

$$\begin{aligned} & \sum_{k=1}^{\ell} k(1-\rho)\rho^{k-1} + (\ell + r' + 1)\rho^\ell \\ &= \frac{1-\rho^\ell}{1-\rho} - \ell \cdot \rho^\ell + (\ell + r' + 1)\rho^\ell \\ &= r + (r' + 1 - r)\rho^\ell \end{aligned}$$

as $\frac{1}{1-\rho} = r$. We let $r' := r - 1$ and obtain

$$\mathbb{L}\mathbb{P}_{\mathcal{M}}^{\mathcal{S}_n}(\mathbf{aUb}) = \frac{\rho^\ell(\ell + 1)}{r} = (1-\rho)\rho^\ell(\ell + 1).$$

Definition of r (and $\rho = (r-1)/r$): For $j \in \mathbb{N}$, define $\wp_j := (1-\rho)\rho^j(j+1)$. The goal is to choose r in such a way that $\wp_\ell > \wp_j$ for all $j \in \mathbb{N} \setminus \{\ell\}$. In particular, this means that we have to choose r such that $\rho^j(j+1)$ reaches its maximum for $j = \ell$. This is achieved by $r = \ell + 1.5$ as can be seen as follows. Let $j \in \mathbb{N}$ and $\bowtie \in \{<, =, >\}$. Then:

$$\begin{aligned} \rho^j(j+1) \bowtie \rho^{j+1}(j+2) & \text{ iff } j+1 \bowtie \frac{r-1}{r}(j+2) \\ & \text{ iff } \frac{j+1}{j+2} \bowtie \frac{r-1}{r} \\ & \text{ iff } j+2 \bowtie r \end{aligned}$$

In particular, if $r = \ell + 1.5$ then $\ell + 2 > r$ and therefore:

$$\wp_\ell > \wp_{\ell+1} > \wp_{\ell+2} > \dots$$

On the other hand, $(\ell-1)+2 = \ell+1 < r$ and therefore:

$$\wp_\ell > \wp_{\ell-1} > \wp_{\ell-2} > \dots$$

Definition of the threshold value ϑ : The idea is to use the observation that the maximal long-run probability that can be achieved when the intersection language is empty is bounded by (maybe, needs to be checked):

$$\vartheta_\ell \stackrel{\text{def}}{=} \wp_\ell - \frac{\mu}{k}$$

where μ is the minimal difference between the value \wp_ℓ and one of the values \wp_j for $j \neq \ell$.

Note that when entering \mathcal{A} after having visited state \mathbf{a} m -times, then at least one of the DFA $\mathcal{D}_1, \dots, \mathcal{D}_k$ does not accept the word $0^{\ell-m}$. And taking α resp. *pump* fewer or more times than ℓ (which means firing the β -transition to state \mathbf{b} after having visited j \mathbf{a} -states for some j different from ℓ) yields a value $\wp_j < \wp_\ell$. Same

for combinations of such options depending on which of the states $q_{0,i}$ is entered.

As the values \wp_j are strictly increasing for $j < \ell$ and strictly decreasing for $j > \ell$ we have:

$$\mu = \min \{ \wp_\ell - \wp_{\ell-1}, \wp_\ell - \wp_{\ell+1} \}$$

Using the fact that

$$\rho^\ell = \left(1 - \frac{1}{r}\right)^\ell \geq 1 - \frac{\ell}{r}$$

and using $r = \ell + 1.5$ and $(1-\rho) = 1/r$ we obtain:

$$\begin{aligned} \wp_\ell - \wp_{\ell-1} &= (1-\rho)(\rho^\ell(\ell+1) - \rho^{\ell-1}\ell) \\ &\geq \frac{1}{r} \left(1 - \frac{\ell}{r}\right) (1 + \ell - \ell/\rho) \\ &= \frac{3}{2r^2} \cdot \left(1 - \frac{\ell}{r-1}\right) = \frac{3}{4r^2(r-1)} \geq \frac{3}{4r^3}. \end{aligned}$$

and

$$\begin{aligned} \wp_\ell - \wp_{\ell+1} &= (1-\rho)(\rho^\ell(\ell+1) - \rho^{\ell+1}(\ell+2)) \\ &\geq \frac{1}{r} \left(1 - \frac{\ell}{r}\right) (1 + \ell(1-\rho) - 2\rho) \\ &= \frac{3}{2r^2} \cdot \frac{1}{2r} = \frac{3}{4r^3}. \end{aligned}$$

This yields $\mu \geq \frac{3}{4r^3}$. In addition, $k \leq \log(\ell)$ as we assume that all \mathcal{D}_i have at least two states. To get a safe approximation, we observe that for $\ell \geq 1000$, we have $\frac{100}{\ell^4} < \frac{\mu}{k}$.

The crux is now to find a rational threshold ϑ with $\vartheta_\ell < \vartheta \leq \wp_\ell$ that is computable in polynomial time. We cannot use \wp_ℓ as its logarithmic length is exponential in the input size (the sum of the sizes of $\mathcal{D}_1, \dots, \mathcal{D}_k$). Instead, we will compute an approximation \wp' of $\wp_\ell \cdot \ell^4$ up to an absolute error < 50 in polynomial time. Then for $\ell \geq 1000$, we can choose the threshold ϑ to be $\frac{\wp' - 50}{\ell^4}$ as

$$\wp_\ell \geq \frac{\wp' - 50}{\ell^4} \geq \wp_\ell - \frac{100}{\ell^4} > \wp_\ell - \mu/k = \vartheta_\ell.$$

Approximation of $\wp_\ell \cdot \ell^4$: We define the real function

$$f(z) := \frac{(1/z + 1/2)^{1/z}}{(1/z + 3/2)^{1/z+1}} (1/z + 1)$$

for $z \in [-1/2, 1/2] \setminus \{0\}$ and $f(0) := \frac{1}{e}$ where e is Euler's number. The idea is that $f(1/\ell) = \wp_\ell$ and that we can provide good approximations of $f(z)$ for z close to 0 using Taylor's theorem: By standard methods from calculus, we see that the function f is 5 times continuously differentiable on $[-1/2, 1/2]$. Calculating the derivatives, we obtain the following approximation:

$$f(z) = \frac{1}{e} + \frac{z}{2e} - \frac{z^2}{3e} + \frac{z^3}{4e} - \frac{313z^4}{1440e} + \mathcal{O}(z^5)$$

for $z \rightarrow 0$. So there are reals $c_0, z_0 > 0$ such that

$$\left| f(z) - \left(\frac{1}{e} + \frac{z}{2e} - \frac{z^2}{3e} + \frac{z^3}{4e} - \frac{313z^4}{1440e} \right) \right| \leq c_0 z^5$$

for $|z| < z_0$. Using $\wp_\ell = f(1/\ell)$ we obtain

$$\left| \wp_\ell \ell^4 - \frac{1}{e} \left(\ell^4 + \frac{1}{2}\ell^3 - \frac{1}{3}\ell^2 + \frac{1}{4}\ell - \frac{313}{1440} \right) \right| \leq c_0 \ell^{-1}$$

for $\ell > 1/z_0$.

Let $L = \max\{1/z_0, c_0, 1000\}$. To obtain the desired approximation of $\wp_\ell \ell^4$ for $\ell > L$, it is enough to approximate

$$\frac{1}{e} \left(\ell^4 + \frac{1}{2}\ell^3 - \frac{1}{3}\ell^2 + \frac{1}{4}\ell - \frac{313}{1440} \right)$$

up to an absolute error of 49. Hence, approximating $\frac{1}{e}$ up to an absolute error of $\frac{1}{\ell^4}$ is sufficient. This can be done in polynomial time.

So, we can compute a threshold value for $\ell > L$ which completes the reduction of the intersection problem for unary DFA to the threshold problem for maximal long-run probabilities in MDPs. As L is defined in terms of the function f , i.e. independent of all variables, and as there are only finitely many instances with $\ell \leq L$, this finishes the proof. \square