The Complexity of Synchronizing Markov Decision Processes[☆],[☆☆]

Laurent Doyen[a],[*], Thierry Massart[b], Mahsa Shirmohammadi[a],[b]

[a]*LSV, ENS Cachan & CNRS, France*
[b]*Université Libre de Bruxelles, Belgium*

Abstract

We consider Markov decision processes (MDP) as generators of sequences of probability distributions over states. A probability distribution is p-synchronizing if the probability mass is at least p in a single state, or in a given set of states. We consider four temporal synchronizing modes: a sequence of probability distributions is always p-synchronizing, eventually p-synchronizing, weakly p-synchronizing, or strongly p-synchronizing if, respectively, all, some, infinitely many, or all but finitely many distributions in the sequence are p-synchronizing.

For each synchronizing mode, an MDP can be (*i*) *sure* winning if there is a strategy that produces a 1-synchronizing sequence; (*ii*) *almost-sure* winning if there is a strategy that produces a sequence that is, for all $\varepsilon > 0$, a (1-ε)-synchronizing sequence; (*iii*) *limit-sure* winning if for all $\varepsilon > 0$, there is a strategy that produces a (1-ε)-synchronizing sequence.

We provide fundamental results on the expressiveness, decidability, and complexity of synchronizing properties for MDPs. For each synchronizing mode, we consider the problem of deciding whether an MDP is sure, almost-sure, or limit-sure winning, and we establish matching upper and lower complexity bounds of the problems: for all winning modes, we show that the problems are PSPACE-complete for eventually and weakly synchronizing, and PTIME-complete for always and strongly synchronizing. We establish the memory requirement for winning strategies, and we show that all winning modes coincide for always synchronizing, and that the almost-sure and limit-sure winning modes coincide for weakly and strongly synchronizing.

1. Introduction

Markov decision processes (MDP) are finite-state stochastic models of dynamic systems studied in many applications such as planning [50], randomized algorithms [3, 56], communication protocols [32], and in many problems related to reactive system design and verification [31, 6, 29]. MDPs exhibit both stochastic and nondeterministic behavior, as in the control problem for reactive systems: nondeterminism represents the possible choice of actions of the controller, and stochasticity represents the uncertainties about the system response (see Figure 1).

The controller synthesis problem is to compute the largest probability with which a control strategy can ensure that the system satisfies a given specification, and to construct an optimal strategy [12, 31]. The qualitative variant of the problem is to decide if the system can satisfy the specification with probability 1. Fundamental well-studied specifications are state-based and describe correct behaviors as infinite sequences of states of the MDP, including safety and liveness properties such as reachability, Büchi, and co-Büchi conditions, which require the system to visit a set of target states once, infinitely often, and ultimately always, respectively [38, 24].

In contrast to this traditional approach, we consider a distribution-based semantics where the specification describes correct behaviors of MDPs as infinite sequences of probability distributions $d_i : Q \to [0,1]$ over the finite state space Q of the system, where $d_i(q)$ is the probability that the MDP is in state $q \in Q$ after i execution steps. The distribution-based semantics is adequate in large-population models, such as systems biology [39], robot planning [8], distributed systems [35], etc. where the system consists of several copies of the same process (molecules, robots, sensors, etc.), and the relevant information along the execution of the system is the number of processes in each state, or the relative frequency (i.e., the probability) of each state. In the context of several identical processes, the same control strategy is used in every process, but the internal state of each process need not be the same along the execution, since probabilistic transitions may have a different outcome in each process. Therefore, the global execution of the system (consisting of all the processes) is better described by the sequence of probability distributions over states along the execution. However, the control strategy is local to each process and can select control actions depending on the full history of the process execution, which corresponds to general perfect-information strategies that we consider in this work.

Previously, the special case of blind strategies has been considered, which in each step select the same control action at all states, and thus only depend on the number of execution steps of the system. In automata theory, a blind strategy corresponds simply to an input word. In MDPs with blind strategies, also known as probabilistic automata [58, 55], several basic problems are undecidable such as deciding if there exists a blind strategy that ensures a coBüchi condition with probability 1 [7], or deciding if a reachability condition can be ensured with probability arbitrarily close to 1 [36].
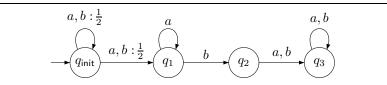
Figure 1: An MDP with four states and set of actions $\{a, b\}$. All transitions are deterministic except from q_{init} where on all actions, the successors are q_{init} and q_1 with probability $\frac{1}{2}$. The incoming arrow in q_{init} denotes an initial Dirac distribution that assigns probability 1 to q_{init}.

The main contribution of this article is to establish the decidability and optimal complexity of deciding *synchronizing* properties for the distribution-based semantics of MDPs under general strategies. Synchronizing properties require that the sequence of probability distributions accumulate all the probability mass in a single state, or in a given set of states. They generalize synchronizing properties of finite automata [63, 26]. Formally, for $0 \leq p \leq 1$ let a probability distribution $d : Q \to [0, 1]$ be p-synchronized if it assigns probability at least p to some state. A sequence $\bar{d} = d_0 d_1 \ldots$ of probability distributions is

(a) *always p-synchronizing* if d_i is p-synchronized for all i;

(b) *eventually p-synchronizing* if d_i is p-synchronized for some i;

(c) *weakly p-synchronizing* if d_i is p-synchronized for infinitely many i's;

(d) *strongly p-synchronizing* if d_i is p-synchronized for all but finitely many i's.

We present a consistent and comprehensive theory of the qualitative synchronizing properties, corresponding to the case where either $p = 1$, or p tends to 1, which are analogous to the traditional safety, reachability, Büchi, and coBüchi conditions [23].

Applications. A typical application scenario of synchronizing properties is the design of a control program for a group of mobile robots running in a stochastic environment [52]. The possible behaviors of the robots and the stochastic response of the environment (such as obstacle encounters) are represented by an MDP, and a synchronizing strategy corresponds to a control program that can be embedded in every robot to ensure that they meet (or synchronize) all the time, eventually once, infinitely often, or eventually forever.

Synchronization properties are central in large-population models in biology, such as yeast, where experimental synchronization methods have been developed to get a population of yeast in the same cell cycle stage [11, 34]. A simple abstraction of large populations of identical finite-state stochastic agents is to consider a continuum of agents, described by the relative fraction of agents in each possible state, i.e. by a distribution. For example, consider a model of cells where at each time instant half of the cells get activated, and once activated we can block them for a while, or release them to reach a fluorescent state. In the MDP of Figure 1, the state q_1 corresponds to activation, action a is

blocking, and action b is releasing. The fluorescent state is q_2. Probability mass arbitrarily close to 1 can be accumulated in q_2, thus for all $\varepsilon > 0$ we can generate an eventually $(1 - \varepsilon)$-synchronizing sequence in q_2, but not an eventually 1-synchronizing sequence. If it was possible to reset the cell state after fluorescence, such as in the MDP of Figure 8, then we can obtain a sequence of distribution that is weakly $(1 - \varepsilon)$-synchronizing, for all $\varepsilon > 0$.

We consider the following qualitative winning modes, summarized in Table 1: (i) *sure* winning, if there is a strategy that generates an {always, eventually, weakly, strongly} 1-synchronizing sequence; (ii) *almost-sure* winning, if there is a strategy that generates a sequence that is, for all $\varepsilon > 0$, {always, eventually, weakly, strongly} $(1 - \varepsilon)$-synchronizing; (iii) *limit-sure* winning, if for all $\varepsilon > 0$, there is a strategy that generates an {always, eventually, weakly, strongly} $(1 - \varepsilon)$-synchronizing sequence.

Contribution. The contributions of this article are summarized as follows:

- *Expressiveness.* We show that the three winning modes form a strict hierarchy for eventually synchronizing: there are limit-sure winning MDPs that are not almost-sure winning, and there are almost-sure winning MDPs that are not sure winning. This is in contrast with the traditional state-based reachability objectives for which the notions of almost-sure and limit-sure winning coincide in MDPs. In this context, a more unexpected and difficult result is that the almost-sure and limit-sure modes coincide for weakly and strongly synchronizing. Thus those two synchronizing modes are more robust than eventually synchronizing, although we show that almost-sure weakly synchronizing strategies can be constructed from the analysis of eventually synchronizing (in limit-sure winning mode). Finally, for always synchronizing the three winning modes coincide, and we show that they coincide with a traditional safety objective.

- *Complexity.* For each synchronizing and winning mode, we consider the problem of deciding if a given initial distribution is winning. The complexity results are shown in Table 2 (p. 13). We establish the decidability and optimal complexity bounds for all winning modes. Under general strategies, the decision problems have much lower complexity than with blind strategies. We show that all decision problems are decidable, in polynomial time for always and strongly synchronizing, and PSPACE-complete for eventually and weakly synchronizing. This is also in contrast with almost-sure winning in the traditional semantics of MDPs, which is solvable in polynomial time for both safety and reachability [22].

- *Memory bounds.* We complete the picture by proving optimal memory bounds for winning strategies, summarized in Table 3 (p. 14). Memoryless strategies are sufficient for always synchronizing (like for safety objectives). We show that linear memory is sufficient for strongly synchronizing, and we identify a variant of strongly synchronizing for which

4

memoryless strategies are sufficient. For eventually and weakly synchronizing, exponential memory is sufficient and may be necessary for sure winning strategies, and in general infinite memory is necessary for almost-sure winning.

Some results in this article rely on insights about games and alternating automata that are of independent interest. Firstly, the sure-winning problem for eventually synchronizing is equivalent to a two-player game with a synchronized reachability objective, where the goal for the first player is to ensure that a target state is reached after a number of steps that is independent of the strategy of the opponent (and thus this number can be fixed in advance by the first player). This condition is stronger than plain reachability, and while the winner in two-player reachability games can be decided in polynomial time, deciding the winner for synchronized reachability is PSPACE-complete. This result is obtained by turning the synchronized reachability game into a one-letter alternating automaton for which the emptiness problem (i.e., deciding if there exists a word accepted by the automaton) is PSPACE-complete [41, 44]. Secondly, our PSPACE lower bound for the limit-sure winning problem in eventually synchronizing uses a PSPACE-completeness result that we establish for the *universal finiteness problem*, which is to decide, given a one-letter alternating automata, whether from every state the accepted language is finite.

Related Works. The traditional state-based semantics of MDPs has been studied extensively [57, 21, 31] and plays a central role in recent developments of system verification and controller synthesis, including expressiveness and complexity analysis of various classes of properties [33], using techniques such as symbolic algorithms for Büchi objectives [18], game-based abstraction techniques [45], and multi-objective analysis for assume-guarantee model-checking [29].

On the other hand, the distribution-based semantics has received a greater

Table 1: Winning modes for always, eventually, weakly, and strongly synchronizing objectives (where $\mathcal{M}_n^\alpha(T)$ denotes the probability that under strategy α, after n steps the MDP $\mathcal{M}$ is in a state of T).

	Always			Eventually		
Sure	$\exists \alpha$	$\forall n$	$\mathcal{M}_n^\alpha(T) = 1$	$\exists \alpha$	$\exists n$	$\mathcal{M}_n^\alpha(T) = 1$
Almost-sure	$\exists \alpha$	$\inf_n$	$\mathcal{M}_n^\alpha(T) = 1$	$\exists \alpha$	$\sup_n$	$\mathcal{M}_n^\alpha(T) = 1$
Limit-sure	$\sup_\alpha$	$\inf_n$	$\mathcal{M}_n^\alpha(T) = 1$	$\sup_\alpha$	$\sup_n$	$\mathcal{M}_n^\alpha(T) = 1$
	Weakly			Strongly		
Sure	$\exists \alpha$	$\forall N \, \exists n \geq N$	$\mathcal{M}_n^\alpha(T) = 1$	$\exists \alpha$	$\exists N \, \forall n \geq N$	$\mathcal{M}_n^\alpha(T) = 1$
Almost-sure	$\exists \alpha$	$\limsup_{n\to\infty}$	$\mathcal{M}_n^\alpha(T) = 1$	$\exists \alpha$	$\liminf_{n\to\infty}$	$\mathcal{M}_n^\alpha(T) = 1$
Limit-sure	$\sup_\alpha$	$\limsup_{n\to\infty}$	$\mathcal{M}_n^\alpha(T) = 1$	$\sup_\alpha$	$\liminf_{n\to\infty}$	$\mathcal{M}_n^\alpha(T) = 1$

interest only recently, as it is shown that relevant key properties of MDPs can only be expressed in a distribution-based logical framework [9, 47] and that a new useful notion of probabilistic bisimulation can be obtained in the distribution-based semantics [40]. Several recent works have investigated this new approach showing that the verification of quantitative properties of the distribution-based semantics is undecidable [47, 30], and decidability can be obtained for special subclasses of systems [15], or through approximations [1]. In this context, a challenging goal is to identify useful decidable properties for the distribution-based semantics.

Synchronization problems were first considered for deterministic finite automata (DFA) where a *synchronizing word* is a finite sequence of control actions that can be executed from any state of an automaton and leads to the same state (see [63] for a survey of results and applications). While the existence of a synchronizing word can be decided in NLOGSPACE for DFA, extensive research effort is devoted to establishing a tight bound on the length of the shortest synchronizing word [10, 53, 60], which is conjectured to be $(n-1)^2$ for automata with n states [14]. Various extensions of the notion of synchronizing word have been proposed for non-deterministic and probabilistic automata [13, 42, 46, 59], leading to results of PSPACE-completeness [51], or even undecidability [46].

For probabilistic systems, it is natural to consider infinite input words (i.e., blind strategies) in order to study synchronization at the limit. In particular, almost-sure weakly and strongly synchronizing with blind strategies has been studied [27] and the main result is that the problem of deciding the existence of a blind almost-sure winning strategy is undecidable for weakly synchronizing, and PSPACE-complete for strongly synchronizing [26, 28]. In contrast in this article, for general strategies, we establish the PSPACE-completeness and PTIME-completeness of deciding almost-sure weakly and strongly synchronizing respectively.

Synchronization has been studied recently in a variety of classical computation models that extend finite automata, such as timed automata [25], weighted automata [43, 25], visibly pushdown automata [20], and register automata [4]. Automata with partial observability have been considered to model systems that disclose information along their execution, which can help the synchronizing strategy [49, 48]. An elegant extension of the computation tree logic (CTL) has been proposed to express synchronizing properties [17].

2. Markov Decision Processes and Synchronizing Properties

A *probability distribution* over a finite set S is a function $d : S \rightarrow [0,1]$ such that $\sum_{s \in S} d(s) = 1$. The *support* of d is the set $\mathsf{Supp}(d) = \{s \in S \mid d(s) > 0\}$. We denote by $\mathcal{D}(S)$ the set of all probability distributions over S. Given a set $T \subseteq S$, let

$$d(T) = \sum_{s \in T} d(s) \quad \text{and} \quad \|d\|_T = max_{s \in T} d(s).$$

For $T \neq \varnothing$, the *uniform distribution* on T assigns probability $\frac{1}{|T|}$ to every

state in T. Given $s \in S$, we denote by ξ^s the *Dirac distribution* on s that assigns probability 1 to s.

A *Markov decision process* (MDP) is a tuple $\mathcal{M} = \langle Q, \mathsf{A}, \delta \rangle$ where Q is a finite set of states, A is a finite set of actions, and $\delta : Q \times \mathsf{A} \to \mathcal{D}(Q)$ is a probabilistic transition function. A state q is *absorbing* if $\delta(q, a)$ is the Dirac distribution on q for all actions $a \in \mathsf{A}$. Given state $q \in Q$ and action $a \in \mathsf{A}$, the successor state of q under action a is q' with probability $\delta(q, a)(q')$. Denote by $\mathsf{post}(q, a)$ the set $\mathsf{Supp}(\delta(q, a))$, and given $T \subseteq Q$ let

$$\mathsf{Pre}(T) = \{q \in Q \mid \exists a \in \mathsf{A} : \mathsf{post}(q, a) \subseteq T\}$$

be the set of states from which there is an action to ensure that the successor state is in T. For $k > 0$, let $\mathsf{Pre}^k(T) = \mathsf{Pre}(\mathsf{Pre}^{k-1}(T))$ with $\mathsf{Pre}^0(T) = T$.

Note that the sequence $\mathsf{Pre}^k(T)$ of iterated predecessors is ultimately periodic, precisely there exist $k < k' < 2^{|Q|}$ such that $\mathsf{Pre}^k(T) = \mathsf{Pre}^{k'}(T)$.

A *path* in $\mathcal{M}$ is an infinite sequence $\pi = q_0 a_0 q_1 a_1 \ldots$ such that $q_{i+1} \in \mathsf{post}(q_i, a_i)$ for all $i \geq 0$. A finite prefix $\rho = q_0 a_0 q_1 a_1 \ldots q_n$ of a path (or simply a finite path) has length $|\rho| = n$ and last state $\mathsf{Last}(\rho) = q_n$. We denote by $\mathsf{Path}(\mathcal{M})$ and $\mathsf{Pref}(\mathcal{M})$ the set of all paths and finite paths in $\mathcal{M}$ respectively.

For the decision problems considered in this article, only the support of the probability distributions in the transition function is relevant (i.e., the exact value of the positive probabilities does not matter); therefore, we can assume that MDPs are encoded as A-labelled transition systems (Q, R) with $R \subseteq Q \times \mathsf{A} \times Q$ such that $(q, a, q') \in R$ is a transition if $q' \in \mathsf{post}(q, a)$.

Strategies. A *randomized strategy* for $\mathcal{M}$ (or simply a strategy) is a function

$$\alpha : \mathsf{Pref}(\mathcal{M}) \to \mathcal{D}(\mathsf{A})$$

that, given a finite path ρ, returns a probability distribution $\alpha(\rho)$ over the action set, used to select a successor state q' of ρ with probability $\sum_{a \in \mathsf{A}} \alpha(\rho)(a) \cdot \delta(q, a)(q')$ where $q = \mathsf{Last}(\rho)$.

A strategy α is *pure* if for all $\rho \in \mathsf{Pref}(\mathcal{M})$, there exists an action $a \in \mathsf{A}$ such that $\alpha(\rho)(a) = 1$; and *memoryless* if $\alpha(\rho) = \alpha(\rho')$ for all ρ, ρ' such that $\mathsf{Last}(\rho) = \mathsf{Last}(\rho')$. We view pure strategies as functions $\alpha : \mathsf{Pref}(\mathcal{M}) \to \mathsf{A}$, and memoryless strategies as functions $\alpha : Q \to \mathcal{D}(\mathsf{A})$.

Finally, a strategy α uses *finite-memory* if it can be represented by a finite-state transducer $T = \langle \mathsf{Mem}, m_0, \alpha_u, \alpha_n \rangle$ where Mem is a finite set of modes (the memory of the strategy), $m_0 \in \mathsf{Mem}$ is the initial mode, $\alpha_u : \mathsf{Mem} \times (\mathsf{A} \times Q) \to \mathsf{Mem}$ is an update function that, given the current memory, last action, and state updates the memory, and $\alpha_n : \mathsf{Mem} \times Q \to \mathcal{D}(\mathsf{A})$ is a next-move function that selects the probability distribution $\alpha_n(m, q)$ over actions when the current mode is m and the current state of $\mathcal{M}$ is q. For pure strategies, we assume that $\alpha_n : \mathsf{Mem} \times Q \to \mathsf{A}$. The *memory size* of the strategy is the number $|\mathsf{Mem}|$ of modes. For a finite-memory strategy α, let $\mathcal{M}(\alpha)$ be the Markov chain obtained as the product of $\mathcal{M}$ with the transducer defining α.

2.1. State-based semantics

In the traditional state-based semantics, given an initial distribution $d_0 \in \mathcal{D}(Q)$ and a strategy α in an MDP $\mathcal{M}$, a *path-outcome* is a path $\pi = q_0 a_0 q_1 a_1 \ldots$ in $\mathcal{M}$ such that $q_0 \in \mathsf{Supp}(d_0)$ and $a_i \in \mathsf{Supp}(\alpha(q_0 a_0 \ldots q_i))$ for all $i \geq 0$. The probability of a finite prefix $\rho = q_0 a_0 q_1 a_1 \ldots q_n$ of π is

$$d_0(q_0) \cdot \prod_{j=0}^{n-1} \alpha(q_0 a_0 \ldots q_j)(a_j) \cdot \delta(q_j, a_j)(q_{j+1}).$$

We denote by $Outcome(d_0, \alpha)$ the set of all path-outcomes from d_0 under strategy α. An *event* $\Omega \subseteq \mathsf{Path}(\mathcal{M})$ is a measurable set of paths, and given an initial distribution d_0 and a strategy α, the probability $Pr^\alpha(\Omega)$ of Ω is uniquely defined [61]. We consider the following classical winning modes. Given an initial distribution d_0 and an event Ω, we say that $\mathcal{M}$ is: *sure winning* if there exists a strategy α such that $Outcome(d_0, \alpha) \subseteq \Omega$; *almost-sure winning* if there exists a strategy α such that $Pr^\alpha(\Omega) = 1$; *limit-sure winning* if $\sup_\alpha Pr^\alpha(\Omega) = 1$, that is the event Ω can be realized with probability arbitrarily close to 1. Given a set $T \subseteq Q$ of target states, and $k \in \mathbb{N}$, we define the following events:

- $\Box T = \{q_0 a_0 q_1 \cdots \in \mathsf{Path}(\mathcal{M}) \mid \forall i : q_i \in T\}$ the safety event of always staying in T;

- $\Diamond T = \{q_0 a_0 q_1 \cdots \in \mathsf{Path}(\mathcal{M}) \mid \exists i : q_i \in T\}$ the event of reaching T;

- $\Diamond^k T = \{q_0 a_0 q_1 \cdots \in \mathsf{Path}(\mathcal{M}) \mid q_k \in T\}$ the event of reaching T after exactly k steps;

- $\Diamond^{\leq k} T = \bigcup_{j \leq k} \Diamond^j T$ the event of reaching T within at most k steps.

For example, if $Pr^\alpha(\Diamond T) = 1$ then almost-surely a state in T is reached under strategy α.

It is known for reachability objectives $\Diamond T$, that an MDP is almost-sure winning if and only if it is limit-sure winning, and the set of initial distributions for which an MDP is sure (resp., almost-sure or limit-sure) winning can be computed in polynomial time [22].

2.2. Distribution-based semantics

In contrast to the state-based semantics, we consider a symbolic outcome of MDPs viewed as generators of sequences of probability distributions over states [47]. Given an initial distribution $d_0 \in \mathcal{D}(Q)$ and a strategy α in $\mathcal{M}$, the *symbolic outcome* of $\mathcal{M}$ from d_0 is the sequence $(\mathcal{M}_n^\alpha)_{n \in \mathbb{N}}$ of probability distributions defined by $\mathcal{M}_k^\alpha(q) = Pr^\alpha(\Diamond^k \{q\})$ for all $k \geq 0$ and $q \in Q$. Hence, $\mathcal{M}_k^\alpha$ is the probability distribution over states after k steps under strategy α. Note that $\mathcal{M}_0^\alpha = d_0$ and the symbolic outcome is a deterministic sequence of distributions: each distribution $\mathcal{M}_k^\alpha$ has a unique (deterministic) successor.

Informally, synchronizing objectives require that the probability of some state (or some group of states) tends to 1 in the sequence $(\mathcal{M}_n^\alpha)_{n \in \mathbb{N}}$, either

always, once, infinitely often, or always after some point. Given a set $T \subseteq Q$, consider the functions

$$sum_T : \mathcal{D}(Q) \to [0,1] \;\; \text{defined by} \;\; sum_T(d) = \sum_{q \in T} d(q), \;\; \text{and}$$
$$max_T : \mathcal{D}(Q) \to [0,1] \;\; \text{defined by} \;\; max_T(d) = \max_{q \in T} d(q).$$

For $f \in \{sum_T, max_T\}$ and $p \in [0,1]$, we say that a probability distribution d is p-synchronized according to f if $f(d) \geq p$, and that a sequence $\bar{d} = d_0 d_1 \ldots$ of probability distributions is:

(a) *always p-synchronizing* if d_i is p-synchronized for all $i \geq 0$;

(b) *event* (or *eventually*) *p-synchronizing* if d_i is p-synchronized for some $i \geq 0$;

(c) *weakly p-synchronizing* if d_i is p-synchronized for infinitely many i's;

(d) *strongly p-synchronizing* if d_i is p-synchronized for all but finitely many i's.

For $p = 1$, these definitions are analogous to the traditional safety, reachability, Büchi, and coBüchi conditions [23]. In this article, we consider the following winning modes where either $p = 1$, or p tends to 1 (we do not consider $p < 1$, see the discussion in Section 6). Given an initial distribution d_0 and a function $f \in \{sum_T, max_T\}$, we say that for the objective of {always, eventually, weakly, strongly} synchronizing from d_0, the MDP $\mathcal{M}$ is:

- *sure winning* if there exists a strategy α such that the symbolic outcome of α from d_0 is {always, eventually, weakly, strongly} 1-synchronizing according to f;

- *almost-sure winning* if there exists a strategy α such that for all $\varepsilon > 0$ the symbolic outcome of α from d_0 is {always, eventually, weakly, strongly} $(1 - \varepsilon)$-synchronizing according to f;

- *limit-sure winning* if for all $\varepsilon > 0$, there exists a strategy α such that the symbolic outcome of α from d_0 is {always, eventually, weakly, strongly} $(1 - \varepsilon)$-synchronizing according to f;

Note that the winning modes for synchronizing objectives differ from the traditional winning modes in MDPs: synchronizing objectives specify sequences of distributions, in a deterministic transition system with infinite state space (the states are the probability distributions). Since the transitions are deterministic and the probabilities are embedded in the state space, the behavior of the system is non-stochastic and the specification is simply a set of sequences (of distributions). In contrast, the traditional almost-sure and limit-sure winning modes of MDPs specify probability measures over sequences of states (called paths) in a probabilistic system with finite state space. Since the probabilities influence the transitions, the behavior of the system is stochastic and the specification is a set of probability measures over paths. For instance almost-sure reachability requires that the probability measure of *all* paths that visit a

9

target state is 1, while almost-sure eventually synchronizing requires that the *single* symbolic outcome belongs to the set of sequences of distributions that are $(1 - \varepsilon)$-synchronizing for all $\varepsilon > 0$.

We often write $\|d\|_T$ instead of $max_T(d)$ (and we omit the subscript when $T = Q$) and $d(T)$ instead of $sum_T(d)$, as in Table 1 where the definitions of the various winning modes and synchronizing objectives for $f = sum_T$ are summarized.

2.3. Membership problem

For $f \in \{sum_T, max_T\}$ and $\lambda \in \{always, event, weakly, strongly\}$, the *winning region* $\langle\!\langle 1 \rangle\!\rangle^{\lambda}_{sure}(f)$ is the set of initial distributions such that $\mathcal{M}$ is sure winning for λ-synchronizing (we assume that $\mathcal{M}$ is clear from the context). We define analogously the sets $\langle\!\langle 1 \rangle\!\rangle^{\lambda}_{almost}(f)$ and $\langle\!\langle 1 \rangle\!\rangle^{\lambda}_{limit}(f)$ of almost-sure and limit-sure winning distributions.

By an abuse of notation, if a Dirac distribution ξ^q belongs to $\langle\!\langle 1 \rangle\!\rangle^{\lambda}_{\mu}(f)$, we often write $q \in \langle\!\langle 1 \rangle\!\rangle^{\lambda}_{\mu}(f)$ instead of $\xi^q \in \langle\!\langle 1 \rangle\!\rangle^{\lambda}_{\mu}(f)$. For a singleton $T = \{q\}$ we have $sum_T = max_T$, and we simply write $\langle\!\langle 1 \rangle\!\rangle^{\lambda}_{\mu}(q)$ (where $\mu \in \{sure, almost, limit\}$). We are interested in the algorithmic complexity of the *membership problem*, which is to decide, given a probability distribution d_0 and a function f, whether $d_0 \in \langle\!\langle 1 \rangle\!\rangle^{\lambda}_{\mu}(f)$.

We show that the winning region is identical for always synchronizing in the three winning modes (Lemma 2), whereas for eventually synchronizing, the winning regions of the three winning modes are in general different (Lemma 3).

Remark 1. *First, note that it follows from the definitions that for all $f \in \{sum_T, max_T\}$, for all $\lambda \in \{always, event, weakly, strongly\}$, and for all $\mu \in \{sure, almost, limit\}$:*

- $\langle\!\langle 1 \rangle\!\rangle^{always}_{\mu}(f) \subseteq \langle\!\langle 1 \rangle\!\rangle^{strongly}_{\mu}(f) \subseteq \langle\!\langle 1 \rangle\!\rangle^{weakly}_{\mu}(f) \subseteq \langle\!\langle 1 \rangle\!\rangle^{event}_{\mu}(f)$, *and*

- $\langle\!\langle 1 \rangle\!\rangle^{\lambda}_{sure}(f) \subseteq \langle\!\langle 1 \rangle\!\rangle^{\lambda}_{almost}(f) \subseteq \langle\!\langle 1 \rangle\!\rangle^{\lambda}_{limit}(f)$.

Lemma 2. *Let T be a set of states. For all functions $f \in \{max_T, sum_T\}$, we have $\langle\!\langle 1 \rangle\!\rangle^{always}_{sure}(f) = \langle\!\langle 1 \rangle\!\rangle^{always}_{almost}(f) = \langle\!\langle 1 \rangle\!\rangle^{always}_{limit}(f)$.*

Proof. By Remark 1, we obtain a cyclic chain of inclusions (thus an overall equality) if we show that

$$\langle\!\langle 1 \rangle\!\rangle^{always}_{limit}(f) \subseteq \langle\!\langle 1 \rangle\!\rangle^{always}_{sure}(f),$$

that is for all distributions d_0, if $\mathcal{M}$ is limit-sure always synchronizing from d_0, then $\mathcal{M}$ is sure always synchronizing from d_0. For $f = max_T$, consider ε smaller than the smallest positive probability in the initial distribution d_0 and in the transitions of the MDP $\mathcal{M} = \langle Q, \mathsf{A}, \delta \rangle$. If $\mathcal{M}$ is limit-sure always synchronizing, then by definition there exists an always $(1 - \varepsilon)$-synchronizing strategy α, and it is easy to show by induction on k that the distributions $\mathcal{M}^{\alpha}_k$ are Dirac for all $k \geq 0$. In particular d_0 is Dirac, and let $q_{\mathsf{init}} \in T$ be such that

10

$d_0(q_{\mathsf{init}}) = 1$. It follows that there is an infinite path from q_{init} in the graph $\langle T, E \rangle$ where $(q, q') \in E$ if there exists an action $a \in \mathsf{A}$ such that $\delta(q, a)(q') = 1$. The existence of this path entails that there is a loop reachable from q_{init} in the graph $\langle T, E \rangle$, and this naturally defines a sure-winning always synchronizing strategy in $\mathcal{M}$. A similar argument for $f = sum_T$ shows that for sufficiently small ε, an always $(1 - \varepsilon)$-synchronizing strategy α must produce a sequence of distributions with support contained in T, until some support repeats in the sequence. This naturally induces an always 1-synchronizing strategy.   □

The results established in this article will entail that the almost-sure and limit-sure modes coincide for weakly and strongly synchronizing (see Theorem 7, Corollary 26, and Corollary 28). The other winning regions are distinct, as shown in the following lemma.

Lemma 3. *There exists an MDP $\mathcal{M}$ and states q_1, q_2 such that:*

(i) $\langle\!\langle 1 \rangle\!\rangle_{sure}^{\lambda}(q_1) \subsetneq \langle\!\langle 1 \rangle\!\rangle_{almost}^{\lambda}(q_1)$ *for all* $\lambda \in \{event, weakly, strongly\}$*, and*

(ii) $\langle\!\langle 1 \rangle\!\rangle_{almost}^{event}(q_2) \subsetneq \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(q_2)$.

Proof. Consider the MDP $\mathcal{M}$ with states $q_{\mathsf{init}}, q_1, q_2, q_3$ and actions a, b as shown in Figure 1. All transitions are deterministic except from q_{init} where on all actions, the successors are q_{init} and q_1 with probability $\frac{1}{2}$.

To establish (i), it is sufficient to prove that

$$q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{strongly}(q_1) \quad \text{and} \quad q_{\mathsf{init}} \notin \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(q_1),$$

because, by Remark 1, it implies that

$$q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{\lambda}(q_1) \quad \text{and} \quad q_{\mathsf{init}} \notin \langle\!\langle 1 \rangle\!\rangle_{sure}^{\lambda}(q_1)$$

for all $\lambda \in \{event, weakly, strongly\}$, establishing all strict inclusions at once. To prove that $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{strongly}(q_1)$, consider the pure strategy that always plays a. The outcome is such that the probability to be in q_1 after k steps is $1 - \frac{1}{2^k}$, showing that $\mathcal{M}$ is almost-sure winning for the strongly synchronizing objective in q_1 (from q_{init}). On the other hand, $q_{\mathsf{init}} \notin \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(q_1)$ because for all strategies α, the probability in q_{init} remains always positive, and thus in q_1 we have $\mathcal{M}_n^{\alpha}(q_1) < 1$ for all $n \geq 0$, showing that $\mathcal{M}$ is not sure winning for the eventually synchronizing objective in q_1 (from q_{init}).

To establish (ii), we first show that $\mathcal{M}$ is limit-sure winning for the eventually synchronizing objective in q_2 (from q_{init}): for $k \geq 0$ consider a strategy that plays a for k steps, and then plays b. Then the probability to be in q_2 after $k + 1$ steps is $1 - \frac{1}{2^k}$, showing that this strategy is eventually $(1 - \frac{1}{2^k})$-synchronizing in q_2.

Second, we show that almost-sure eventually synchronizing is impossible because, to get probability $1 - \varepsilon$ in q_2, the probability mass needs to accumulate for more and more steps in q_1 as ε gets smaller, which cannot be achieved by a single strategy. Formally, for all strategies, since the probability in q_{init} remains
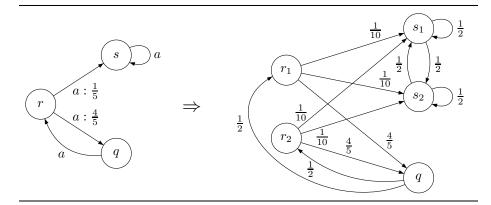
Figure 2: State duplication ensures that the probability mass can never be accumulated in a single state except in q (we omit action a for readability).

always positive, the probability in q_2 is always smaller than 1. Moreover, if the probability p in q_2 is positive after n steps ($p > 0$), then after any number $m > n$ of steps, the probability in q_2 is bounded by $1 - p < 1$. It follows that the probability in q_2 is never equal to 1 and cannot tend to 1 for $m \to \infty$, showing that $\mathcal{M}$ is not almost-sure winning for the eventually synchronizing objective in q_2 (from q_{init}). $\qquad\square$

Finally, for eventually and weakly synchronizing we present in Lemma 4 a reduction of the membership problem with function max_T to the membership problem with function $sum_{T'}$ for a singleton T'. It follows that the complexity results established in this article for eventually and weakly synchronizing with function sum_T also hold with function max_T (this is trivial for the upper bounds, and for the lower bounds it follows from the fact that our hardness results hold for sum_T with singleton T, and thus for max_T as well since in this case $sum_T = max_T$).

Lemma 4. *For eventually and weakly synchronizing, in each winning mode the following problems are polynomial-time equivalent:*

- *the membership problem with a function max_T where T is an arbitrary subset of the state space, and*

- *the membership problem with a function $sum_{T'}$ where T' is a singleton.*

Proof. Let $\mu \in \{sure, almost, limit\}$ and $\lambda \in \{event, weakly\}$. First we have $\langle\!\langle 1 \rangle\!\rangle^\lambda_\mu(max_T) = \bigcup_{q \in T} \langle\!\langle 1 \rangle\!\rangle^\lambda_\mu(q)$, showing that the membership problems for max and max_T are polynomial-time reducible to the corresponding membership problem for $sum_{T'}$ with singleton T'.

The reverse reduction is as follows. Given an MDP $\mathcal{M}$, a state q and an initial distribution d_0, we can construct an MDP $\mathcal{M}'$ and initial distribution d'_0 such that $d_0 \in \langle\!\langle 1 \rangle\!\rangle^\lambda_\mu(q)$ iff $d'_0 \in \langle\!\langle 1 \rangle\!\rangle^\lambda_\mu(max_{Q'})$ where Q' is the state space

Table 2: Computational complexity of the membership problem.

	Always	Eventually	Weakly	Strongly
Sure		PSPACE-C	PSPACE-C	PTIME-C
Almost-sure	PTIME-C	PSPACE-C	PSPACE-C	PTIME-C
Limit-sure		PSPACE-C		

of $\mathcal{M}'$ (thus $max_{Q'}$ is simply the function max). The idea is to construct $\mathcal{M}'$ and d_0' as a copy of $\mathcal{M}$ and d_0 where all states except q are duplicated, and the initial and transition probabilities are equally distributed between the copies (see Figure 2). Therefore, if the probability tends to 1 in some state, it has to be in q (since the probability in any other state of $\mathcal{M}'$ can be at most $1/2$). □

The rest of this article is devoted to the solution of the membership problem. By definition, a sequence of probability distribution is sure always synchronizing according to sum_T if all supports are contained in T, i.e., all states in all path-outcomes are in T, that is all path-outcomes are contained in $\Box T$. Hence, it follows from the proof of Lemma 2 that the winning region for always synchronizing according to sum_T coincides with the set of winning initial distributions for the safety objective $\Box T$ in the traditional semantics, which can be computed in polynomial time [19]. Moreover, always synchronizing according to max_T is equivalent to the existence of an infinite path staying in T in the transition system $\langle Q, R \rangle$ of the MDP restricted to transitions $(q, a, q') \in R$ such that $\delta(q,a)(q') = 1$, which can also be decided in polynomial time. In both cases, pure memoryless strategies are sufficient.

Theorem 1. *The membership problem for always synchronizing can be solved in polynomial time, and pure memoryless strategies are sufficient.*

Remark 5. *For the other synchronizing modes (namely, eventually, weakly, and strongly synchronizing), it is sufficient to consider Dirac initial distributions (i.e., assuming that MDPs have a single initial state) because the answer to the general membership problem for an MDP $\mathcal{M}$ with initial distribution d_0 can be obtained by solving the membership problem for a copy of $\mathcal{M}$ with a new initial state from which the successor distribution on all actions is d_0.*

In the rest of the article, we present algorithms to decide the membership problem and we establish matching upper and lower bounds for the complexity of the problem: we show that eventually and weakly synchronizing are PSPACE-complete, whereas strongly synchronizing is PTIME-complete (like always synchronizing). We also establish optimal memory bounds for the memory needed by strategies to win. Our results will show that pure strategies are sufficient in all modes. The complexity results are summarized in Table 2, and we present the memory requirement for winning strategies in Table 3.

Table 3: Memory requirement.

	Always	Eventually	Weakly	Strongly	
				sum_T	max_T
Sure	memoryless	exponential	exponential	memoryless	linear
Almost-sure		infinite	infinite	memoryless	linear
Limit-sure		unbounded			

2.4. One-Letter Alternating Automata

In this section, we consider one-letter alternating automata (1L-AFA) and show a tight connection with MDP. We present complexity results for 1L-AFA that are useful to establish the PSPACE lower bounds for eventually and weakly synchronizing in MDPs (in Theorem 2, Lemma 10, 17, and 19).

In 1L-AFA, the alphabet is a singleton, and thus only the length of a word is relevant. The transitions of an alternating automaton are described by Boolean formulas over the set of automaton states using only $\land$ and $\lor$ (but no negation). For example, if the formula $(q_2 \land q_3) \lor q_4$ describes the transitions from a state q_1, then the word of length n is accepted from q_1 if the word of length $n-1$ is accepted either from both q_2 and q_3, or from q_4.

One-letter alternating automata. Let $\mathsf{B}^+(Q)$ be the set of positive Boolean formulas over a set Q, i.e. Boolean formulas built from elements in Q using $\land$ and $\lor$ (but no negation). A set $S \subseteq Q$ *satisfies* a formula $\varphi \in \mathsf{B}^+(Q)$ (denoted $S \models \varphi$) if φ is satisfied when replacing in φ the elements in S by true, and the elements in $Q \setminus S$ by false.

A *one-letter alternating finite automaton* is a tuple $\mathcal{A} = \langle Q, \delta_{\mathcal{A}}, \mathcal{F} \rangle$ where Q is a finite set of states, $\delta_{\mathcal{A}} : Q \to \mathsf{B}^+(Q)$ is the transition function, and $\mathcal{F} \subseteq Q$ is the set of accepting states. We assume that the formulas in transition function are in disjunctive normal form. Note that the alphabet of the automaton is omitted, as it consists of a single letter. As in the language of a 1L-AFA, only the length of words is relevant, define for all $n \geq 0$, the set $Acc_{\mathcal{A}}(n, \mathcal{F}) \subseteq Q$ of states from which the word of length n is accepted by $\mathcal{A}$ as follows:

- $Acc_{\mathcal{A}}(0, \mathcal{F}) = \mathcal{F}$;

- $Acc_{\mathcal{A}}(n, \mathcal{F}) = \{q \in Q \mid Acc_{\mathcal{A}}(n-1, \mathcal{F}) \models \delta(q)\}$ for all $n > 0$.

The set $\mathcal{L}(\mathcal{A}_q) = \{n \in \mathbb{N} \mid q \in Acc_{\mathcal{A}}(n, \mathcal{F})\}$ is the *language* accepted by $\mathcal{A}$ from state q (called *initial* state in this context).

Example. Consider the 1L-AFA $\mathcal{A}$ in Figure 3(a) with initial state q_{init}. Transition function is defined by $\delta_{\mathcal{A}}(q_{\mathsf{init}}) = q_1 \land q_2$ and $\delta_{\mathcal{A}}(q_1) = (q_{\mathsf{init}} \land q_2) \lor (q_2 \land q_3)$, etc. The word of length 3 is accepted by $\mathcal{A}$, as witnessed by the execution tree in Figure 4(a): for every node of the tree (let q be its label), the labels of the successor nodes form a set that satisfies the transition function at q. The root
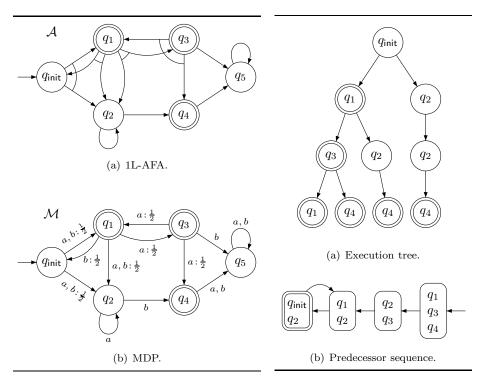
(a) 1L-AFA.

(b) MDP.

Figure 3: 1L-AFA and MDP with the same underlying structure.

(a) Execution tree.

(b) Predecessor sequence.

Figure 4: Execution tree and predecessor sequence (determinization of $\mathcal{A}$).

of the tree is labeled by q_{init} and all leaves are accepting. Note that all branches are of the same length, namely 3, the length of the input word.

For every 1L-AFA with n states, there is an equivalent deterministic automaton with at most 2^n states (that accepts the same language), which can be constructed as follows [16]. It is easier to think that the deterministic automaton accepts the *reverse* image of the words in the language of $\mathcal{A}$, which is the same as the language of $\mathcal{A}$. For all $n \geq 1$, we view $Acc_{\mathcal{A}}(n, \cdot)$ as an operator on 2^Q that, given a set $\mathcal{F} \subseteq Q$ computes the set $Acc_{\mathcal{A}}(n, \mathcal{F})$. Note that $Acc_{\mathcal{A}}(n, \mathcal{F}) = Acc_{\mathcal{A}}(1, Acc_{\mathcal{A}}(n-1, \mathcal{F}))$ for all $n \geq 1$. Call $Acc_{\mathcal{A}}(1, s)$ the *predecessor* of s. The deterministic automaton has state space $\{s_0, \ldots, s_{2^n}\}$ where $s_i = Acc_{\mathcal{A}}(i, \mathcal{F})$, and a deterministic transition from s_i to its predecessor $Acc_{\mathcal{A}}(1, s_i)$. The sequence of predecessors for the 1L-AFA of Figure 3(a) is shown in Figure 4(b). It is easy to see that this sequence is always ultimately periodic (for all $k > 2^n$ there exists $k' \leq 2^n$ such that $s_k = s_{k'}$), and therefore the transitions of the deterministic automaton are well defined. Let $s_0 = \mathcal{F}$ be the initial state, and let all s_i such that $q_{\text{init}} \in s_i$ be the accepting states, then its language is $\mathcal{L}(\mathcal{A}_{q_{\text{init}}})$. Considering the sequence of predecessors in Figure 4(b) it is easy to see that the language of $\mathcal{A}$ is the set $\{n > 1 \mid n \text{ is odd}\}$ of odd numbers greater than 1. Note that the language of $\mathcal{A}$ is nonempty because

15

in the sequence of predecessors there is a state s such that $q_{\mathsf{init}} \in s$, and the language of $\mathcal{A}$ is infinite because there is such a state s in the periodic part of the sequence.

Relation with MDPs. Consider the MDP $\mathcal{M}$ in Figure 3(b), obtained from $\mathcal{A}$ by transforming the transition function, which is a disjunction of conjunctions of states as follows: each conjunction is replaced by a uniform probability distribution over its elements, and the elements of the disjunction are labeled by a letter from the alphabet of $\mathcal{M}$.

The correspondence between 1L-AFA and MDPs is based on the observation that they have the same underlying structure of alternating graph (or AND-OR graph):

- the combinatorial structure of an MDP $\langle Q, \mathsf{A}, \delta \rangle$ can be described as an alternating graph, with *existential* vertices $q \in Q$ with successors (q, a) for all actions $a \in \mathsf{A}$, and *universal* vertices $(q, a) \in Q \times \mathsf{A}$ with successors $q' \in \mathsf{Supp}(\delta(q, a))$;

- the structure of a 1L-AFA $\langle Q, \delta_{\mathcal{A}}, \mathcal{F} \rangle$ is also an alternating graph, where states $q \in Q$ are *existential* vertices with successors the clauses $c_1, \ldots, c_m$ such that $\delta_{\mathcal{A}}(q) = c_1 \vee \cdots \vee c_m$ where each c_i is a conjunctive clause, and the conjunctive clauses c_i are *universal* vertices with successors the states that belong to c_i.

The common structure of 1L-AFA and MDP is illustrated in Figure 3 where we show only the existential vertices. The correspondence is defined formally as follows. Given a 1L-AFA $\mathcal{A} = \langle Q, \delta_{\mathcal{A}}, \mathcal{F} \rangle$, assume without loss of generality that the transition function $\delta_{\mathcal{A}}$ is such that $\delta_{\mathcal{A}}(q) = c_1 \vee \cdots \vee c_m$ has the same number m of conjunctive clauses for all $q \in Q$ (we can duplicate clauses if necessary). From $\mathcal{A}$, construct the MDP $\mathcal{M}_{\mathcal{A}} = \langle Q, \mathsf{A}, \delta_{\mathcal{M}} \rangle$ where $\mathsf{A} = \{a_1, \ldots, a_m\}$ and $\delta_{\mathcal{M}}(q, a_k)$ is the uniform distribution over the states occurring in the k-th clause c_k in $\delta_{\mathcal{A}}(q)$, for all $q \in Q$ and $a_k \in \mathsf{A}$. Then, we have $Acc_{\mathcal{A}}(n, \mathcal{F}) = \mathsf{Pre}_{\mathcal{M}}^n(\mathcal{F})$ for all $n \geq 0$.

Similarly, from an MDP $\mathcal{M}$ and a set T of states, we can construct a 1L-AFA $\mathcal{A} = \langle Q, \delta_{\mathcal{A}}, \mathcal{F} \rangle$ with $\mathcal{F} = T$ such that $Acc_{\mathcal{A}}(n, \mathcal{F}) = \mathsf{Pre}_{\mathcal{M}}^n(T)$ for all $n \geq 0$: let $\delta_{\mathcal{A}}(q) = \bigvee_{a \in \mathsf{A}} \bigwedge_{q' \in \mathsf{post}(q,a)} q'$ for all $q \in Q$. For example, for $\mathsf{A} = \{a, b\}$ if $\delta_{\mathcal{M}}(q_3, a)(q_1) = \delta_{\mathcal{M}}(q_3, a)(q_4) = \frac{1}{2}$ and $\delta_{\mathcal{M}}(q_3, b)(q_5) = 1$, then $\delta_{\mathcal{A}}(q_3) = (q_1 \wedge q_4) \vee q_5$ (see Figure 3).

It follows that, up to the correspondence between 1L-AFA and MDPs established above, $Acc_{\mathcal{A}}(n, T) = \mathsf{Pre}_{\mathcal{M}}^n(T)$. In the sequel we denote the operator $Acc_{\mathcal{A}}(1, \cdot)$ by $\mathsf{Pre}_{\mathcal{A}}(\cdot)$ (note the subscript). Then for all $n \geq 0$ the operator $Acc_{\mathcal{A}}(n, \cdot)$ coincides with $\mathsf{Pre}_{\mathcal{A}}^n(\cdot)$, the n-th iterate of $\mathsf{Pre}_{\mathcal{A}}(\cdot)$.

Decision problems. Questions related to MDPs have a corresponding formulation in terms of alternating automata. We show that such connections exist between synchronizing problems for MDPs and language-theoretic questions for alternating automata.

Several decision problems for 1L-AFA can be solved by computing the sequence $Acc_{\mathcal{A}}(n, \mathcal{F})$ (i.e., $\mathsf{Pre}_{\mathcal{A}}^n(\mathcal{F})$), and analogously we show that synchronizing problems for MDPs can also be solved by computing the sequence $\mathsf{Pre}_{\mathcal{M}}^n(\mathcal{F})$. Therefore, the above relationship between 1L-AFA and MDPs provides a tight connection that we use in Section 3 to transfer complexity results between 1L-AFA and MDPs.

We review classical decision problems for 1L-AFA, namely the emptiness and finiteness problems, and establish the complexity of a new problem, the *universal finiteness problem* which is to decide if from every initial state the language of a given 1L-AFA is finite. These results of independent interest are useful to establish the PSPACE lower bounds for eventually and weakly synchronizing in MDPs.

- The *emptiness problem* for 1L-AFA is to decide, given a 1L-AFA $\mathcal{A}$ and an initial state q, whether $\mathcal{L}(\mathcal{A}_q) = \varnothing$. The emptiness problem can be solved by checking whether $q \in \mathsf{Pre}_{\mathcal{A}}^n(\mathcal{F})$ for some $n \geq 0$. It is known that the emptiness problem is PSPACE-complete, even for transition functions in disjunctive normal form [41, 44].

- The *finiteness problem* is to decide, given a 1L-AFA $\mathcal{A}$ and an initial state q, whether $\mathcal{L}(\mathcal{A}_q)$ is finite. The finiteness problem can be solved in (N)PSPACE by guessing $n, k \leq 2^{|Q|}$ such that $\mathsf{Pre}_{\mathcal{A}}^{n+k}(\mathcal{F}) = \mathsf{Pre}_{\mathcal{A}}^n(\mathcal{F})$ and $q \in \mathsf{Pre}_{\mathcal{A}}^n(\mathcal{F})$. The finiteness problem is PSPACE-complete by a simple reduction from the emptiness problem: from an instance $(\mathcal{A}, q)$ of the emptiness problem, construct $(\mathcal{A}', q')$ where $q' = q$ and $\mathcal{A}' = \langle Q, \delta', \mathcal{F} \rangle$ is a copy of $\mathcal{A} = \langle Q, \delta, \mathcal{F} \rangle$ with a self-loop on q (formally, $\delta'(q) = q \vee \delta(q)$ and $\delta'(r) = \delta(r)$ for all $r \in Q \setminus \{q\}$). It is easy to see that $\mathcal{L}(\mathcal{A}_q) = \varnothing$ iff $\mathcal{L}(\mathcal{A}'_{q'})$ is finite.

- The *universal finiteness problem* is to decide, given a 1L-AFA $\mathcal{A}$, whether $\mathcal{L}(\mathcal{A}_q)$ is finite for all states q. This problem can be solved by checking whether $\mathsf{Pre}_{\mathcal{A}}^n(\mathcal{F}) = \varnothing$ for some $n \leq 2^{|Q|}$, and thus it is in PSPACE. Note that if $\mathsf{Pre}_{\mathcal{A}}^n(\mathcal{F}) = \varnothing$, then $\mathsf{Pre}_{\mathcal{A}}^m(\mathcal{F}) = \varnothing$ for all $m \geq n$.

Given the PSPACE-hardness proofs of the emptiness and finiteness problems, it is not easy to see that the universal finiteness problem is PSPACE-hard.

Lemma 6. *The universal finiteness problem for 1L-AFA is PSPACE-hard.*

Proof. We show the result by a reduction from the emptiness problem for 1L-AFA, which is PSPACE-complete [41, 44]. We first present a basic fact about 1L-AFA, then an overview of the reduction, and a detailed description of the reduction and the correctness argument.

Basic result. The language of a 1L-AFA $\mathcal{A} = \langle Q, \delta, \mathcal{F} \rangle$ from initial state q_0 is non-empty if $q_0 \in \mathsf{Pre}_{\mathcal{A}}^i(\mathcal{F})$ for some $i \geq 0$. Since the sequence $\mathsf{Pre}_{\mathcal{A}}^i(\mathcal{F})$ is ultimately periodic, it is sufficient to compute $\mathsf{Pre}_{\mathcal{A}}^i(\mathcal{F})$ for every $i \leq 2^{|Q|}$ to decide emptiness.
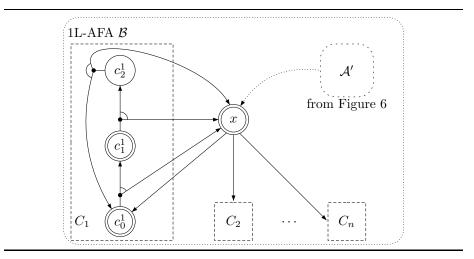
Figure 5: Sketch of reduction to show PSPACE-hardness of the universal finiteness problem for 1l-AFA.

Overview of the reduction. From $\mathcal{A}$, we construct a 1L-AFA $B = \langle Q', \delta', \mathcal{F}' \rangle$ with set $\mathcal{F}'$ of accepting states such that the sequence $\mathsf{Pre}^i_B(\mathcal{F}')$ in B mimics the sequence $\mathsf{Pre}^i_\mathcal{A}(\mathcal{F})$ in $\mathcal{A}$ for $2^{|Q|}$ steps. The automaton B contains the state space of $\mathcal{A}$, i.e. $Q \subseteq Q'$. The goal is to have

$$\mathsf{Pre}^i_B(\mathcal{F}') \cap Q = \mathsf{Pre}^i_\mathcal{A}(\mathcal{F}) \text{ for all } i \leq 2^{|Q|}, \text{ as long as } q_0 \notin \mathsf{Pre}^i_\mathcal{A}(\mathcal{F}).$$

Moreover, if $q_0 \in \mathsf{Pre}^i_\mathcal{A}(\mathcal{F})$ for some $i \geq 0$, then $\mathsf{Pre}^j_B(\mathcal{F}')$ will contain q_0 for all $j \geq i$ (the state q_0 has a self-loop in B), and if $q_0 \notin \mathsf{Pre}\mathcal{A}^i(\mathcal{F})$ for all $i \geq 0$, then B is constructed such that $\mathsf{Pre}^j_B(\mathcal{F}') = \varnothing$ for sufficiently large j (roughly for $j > 2^{|Q|}$). Hence, the language of $\mathcal{A}$ is non-empty if and only if the sequence $\mathsf{Pre}^j_B(\mathcal{F}')$ is not ultimately empty, that is if and only if the language of B is infinite from some state (namely q_0).

Detailed reduction. The key is to let B simulate $\mathcal{A}$ for exponentially many steps, and to ensure that the simulation stops if and only if q_0 is not reached within $2^{|Q|}$ steps. We achieve this by defining B as the gadget in Figure 5 connected to a modified copy $\mathcal{A}'$ of $\mathcal{A}$ with the same state space. The transitions in $\mathcal{A}'$ are defined as follows, where x is the entry state of the gadget (see Figure 6): for all $q \in Q$ let (*i*) $\delta_B(q) = x \wedge \delta_\mathcal{A}(q)$ if $q \neq q_0$, and (*ii*) $\delta_B(q_0) = q_0 \vee (x \wedge \delta_\mathcal{A}(q_0))$. Thus, q_0 has a self-loop, and given a set $S \subseteq Q$ in the automaton $\mathcal{A}$, if $q_0 \notin S$, then $\mathsf{Pre}_\mathcal{A}(S) = \mathsf{Pre}_B(S \cup \{x\})$ that is Pre_B mimics $\mathsf{Pre}_\mathcal{A}$ when x is in the argument (and q_0 has not been reached yet). Note that if $x \notin S$ (and $q_0 \notin S$), then $\mathsf{Pre}_B(S) = \varnothing$, that is unless q_0 has been reached, the simulation of $\mathcal{A}$ by B stops. Since we need that B mimics $\mathcal{A}$ for $2^{|Q|}$ steps, we define the gadget and the set $\mathcal{F}'$ to ensure that $x \in \mathcal{F}'$ and if $x \in \mathsf{Pre}^i_B(\mathcal{F}')$, then $x \in \mathsf{Pre}^{i+1}_B(\mathcal{F}')$ for all $i \leq 2^{|Q|}$.
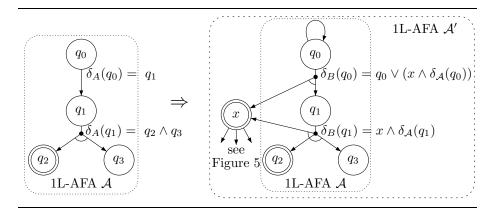
18

Figure 6: Detail of the copy $\mathcal{A}'$ obtained from $\mathcal{A}$ in the reduction of Figure 5.

In the gadget (Figure 5), the state x has nondeterministic transitions

$$\delta_B(x) = c_0^1 \vee c_0^2 \vee \cdots \vee c_0^n$$

to n components with state space $C_i = \{c_0^i, \ldots, c_{p_i-1}^i\}$ where p_i is the $(i+1)$-th prime number, and the transitions[1] $\delta_B(c_j^i) = x \wedge c_{j+1}^i$ $(i = 1, \ldots, n)$ form a loop in each component. We choose n such that $p_n^{\#} = \prod_{i=1}^n p_i > 2^{|Q|}$ (take $n = |Q|$). Note that the number of states in the gadget is $1 + \sum_{i=1}^n p_i \in O(n^2 \log n)$ [5] and thus the construction is polynomial in the size of $\mathcal{A}$.

By construction, for all sets S, we have $x \in \mathsf{Pre}_B(S)$ whenever the first state c_0^i of some component C_i is in S, and if $x \in S$, then $c_j^i \in S$ implies $c_{j-1}^i \in \mathsf{Pre}_B(S)$. Thus, if $x \in S$, the operator $\mathsf{Pre}_B(S)$ 'shifts' backward the states in each component; and, x is in the next iteration (i.e., $x \in Pre_B(S)$) as long as $c_0^i \in S$ for some component C_i.

Now, define the set of accepting states $\mathcal{F}'$ in B in such a way that all states c_0^i disappear simultaneously only after $p_n^{\#}$ iterations. Let $\mathcal{F}' = \mathcal{F} \cup \{x\} \cup \bigcup_{1 \le i \le n}(C_i \setminus \{c_{p_i-1}^i\})$, thus $\mathcal{F}'$ contains all states of the gadget except the last state of each component.

Correctness argument. It is easy to check that, irrespective of the transition relation in $\mathcal{A}$, we have $x \in \mathsf{Pre}_B^i(\mathcal{F}')$ if and only if $0 \le i < p_n^{\#}$. Therefore, if $q_0 \in \mathsf{Pre}_{\mathcal{A}}^i(\mathcal{F})$ for some i, then $q_0 \in \mathsf{Pre}_B^j(\mathcal{F}')$ for all $j \ge i$ by the self-loop on q_0. On the other hand, if $q_0 \notin \mathsf{Pre}_{\mathcal{A}}^i(\mathcal{F})$ for all $i \ge 0$, then since $x \notin \mathsf{Pre}_B^i(\mathcal{F}')$ for all $i > p_n^{\#}$, we have $\mathsf{Pre}_B^i(\mathcal{F}') = \varnothing$ for all $i > p_n^{\#}$. This shows that the language of $\mathcal{A}$ is non-empty if and only if the language of B is infinite from some state (namely q_0), and establishes the correctness of the reduction. □

---

[1] In expression c_j^i, we assume that j is interpreted modulo p_i.

3. Eventually Synchronizing

In this section, we show the PSPACE-completeness of the membership problem for eventually synchronizing objectives and the three winning modes. By Lemma 4 and Remark 5, we consider without loss of generality the membership problem with function *sum* and Dirac initial distributions (i.e., single initial state).

The eventually synchronizing objective is reminiscent of a reachability objective in the distribution-based semantics: it requires that in the sequence of distributions of an MDP $\mathcal{M}$ under strategy α we have $\sup_n \mathcal{M}_n^\alpha(T) = 1$ (and that the sup is reached in the case of sure winning, that is $\mathcal{M}_n^\alpha(T) = 1$ for some $n \geq 0$).

The sure winning mode can be solved by a reachability analysis in the alternating graph underlying the MDP (Section 3.1). We show that the almost-sure winning mode can be solved by a reduction to the limit-sure winning mode (Section 3.2). We solve the limit-sure winning mode by a reduction to a reachability question in a modified MDP of exponential size that ensures the probability mass reaches the target set synchronously (Section 3.3). We present reductions to show PSPACE-hardness of each winning mode, matching our PSPACE upper bounds.

3.1. Sure eventually synchronizing

Given a target set T, the membership problem for sure-winning eventually synchronizing objective in T can be solved by computing the sequence $\mathsf{Pre}^n(T)$ of iterated predecessors, like in 1L-AFA, as shown in the following lemma.

Lemma 7. *Let $\mathcal{M}$ be an MDP and T be a target set. For all states q_{init}, we have $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_T)$ if and only if there exists $n \geq 0$ such that $q_{\mathsf{init}} \in \mathsf{Pre}_{\mathcal{M}}^n(T)$.*

Proof. We prove the following equivalence by induction (on the length i): for all initial states q_{init}, there exists a strategy α sure-winning in i steps from q_{init} (i.e., such that $\mathcal{M}_i^\alpha(T) = 1$) if and only if $q_{\mathsf{init}} \in \mathsf{Pre}^i(T)$. The case $i = 0$ trivially holds since for all strategies α, we have $\mathcal{M}_0^\alpha(T) = 1$ if and only if $q_{\mathsf{init}} \in T$.

Assume that the equivalence holds for all $i < n$. For the induction step, show that $\mathcal{M}$ is sure eventually synchronizing from q_{init} (in n steps) if and only if there exists an action a such that $\mathcal{M}$ is sure eventually synchronizing (in $n-1$ steps) from all states $q' \in \mathsf{post}(q_{\mathsf{init}}, a)$ (equivalently, $\mathsf{post}(q_{\mathsf{init}}, a) \subseteq \mathsf{Pre}^{n-1}(T)$ by the induction hypothesis, that is $q_{\mathsf{init}} \in \mathsf{Pre}^n(T)$ by definition of Pre). First, if all successors q' of q_{init} under some action a are sure eventually synchronizing, then so is q_{init} by playing a followed by a winning strategy from each successor q'. For the other direction, assume towards contradiction that $\mathcal{M}$ is sure eventually synchronizing from q_{init} (in n steps), but for each action a, there is a state $q' \in \mathsf{post}(q_{\mathsf{init}}, a)$ that is not sure eventually synchronizing. Then, from q' there is a positive probability to reach a state not in T after $n-1$ steps, no matter the strategy played. Hence from q_{init}, for all strategies, the probability mass in T cannot be 1 after n steps, in contradiction with the fact that $\mathcal{M}$ is sure
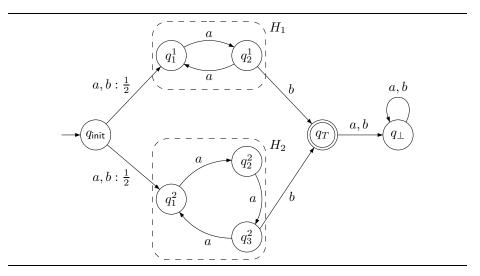
Figure 7: The MDP $\mathcal{M}_2$.

eventually synchronizing from q_{init} in n steps. It follows that the induction step holds, and the proof is complete. □

The following theorem summarizes the results for sure eventually synchronizing.

Theorem 2. *For sure eventually synchronizing in MDPs:*

1. *(Complexity). The membership problem is PSPACE-complete.*
2. *(Memory). Exponential memory is necessary and sufficient for both pure and randomized strategies, and pure strategies are sufficient.*

Proof. By Lemma 7, the membership problem for sure eventually synchronizing is equivalent to the emptiness problem of 1L-AFA, and thus PSPACE-complete. Moreover, if $q_{\mathsf{init}} \in \mathsf{Pre}_{\mathcal{M}}^n(T)$, a finite-memory strategy with n modes that at mode i in a state q plays an action a such that $\mathsf{post}(q,a) \subseteq \mathsf{Pre}^{i-1}(T)$ is sure winning for eventually synchronizing. Note that this strategy is pure.

We present a family of MDPs $\mathcal{M}_n$ ($n \in \mathbb{N}$) over alphabet $\{a,b\}$ that are sure winning for eventually synchronizing, and where the sure winning strategies require exponential memory. The MDP $\mathcal{M}_2$ is shown in Figure 7. The structure of $\mathcal{M}_n$ is an initial uniform probabilistic transition to n components $H_1, \ldots, H_n$ where H_i is a cycle of length p_i the ith prime number. On action a, the next state in the cycle is reached, and on action b the target state q_T is reached, only from the last state in the cycles. From other states, the action b leads to $q_\perp$ (transitions not depicted). A sure winning strategy for eventually synchronizing in $\{q_T\}$ is to play a in the first $p_n^\# = \prod_{i=1}^n p_i$ steps, and then play b. This requires memory of size $p_n^\# > 2^n$ while the size of $\mathcal{M}_n$ is in $O(n^2 \log n)$ [5]. It can be proved by standard pumping arguments that no strategy of size smaller than $p_n^\#$ is sure winning. □
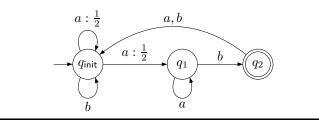
21

Figure 8: An MDP where infinite memory is necessary for almost-sure eventually and almost-sure weakly synchronizing strategies.

3.2. Almost-sure eventually synchronizing

We show an example where infinite memory is necessary to win for almost-sure eventually synchronizing. Consider the MDP in Figure 8 with initial state q_{init}. We construct a strategy that is almost-sure eventually synchronizing in q_2, showing that $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{event}(q_2)$. First, observe that for all $\varepsilon > 0$ we can have probability at least $1 - \varepsilon$ in q_2 after finitely many steps: playing n times a and then b leads to probability $1 - \frac{1}{2^n}$ in q_2 (and $\frac{1}{2^n}$ in q_{init}) . Choosing n sufficiently large (namely, $n > \log_2(\frac{1}{\varepsilon})$) shows that the MDP is limit-sure eventually synchronizing in q_2. Moreover, the remaining probability mass is in q_{init}. By playing a we get again support $\{q_{\mathsf{init}}\}$, thus from any (initial) distribution with support $\{q_{\mathsf{init}}, q_2\}$, the MDP is again limit-sure eventually synchronizing in q_2, and with support in $\{q_{\mathsf{init}}, q_2\}$. Therefore, we can take a smaller value of ε and play a strategy to have probability at least $1 - \varepsilon$ in q_2 in finitely many steps, then reaching back support $\{q_{\mathsf{init}}\}$, and we can repeat this for $\varepsilon \to 0$. This strategy ensures probability mass $1 - \varepsilon$ in q_2 for all $\varepsilon > 0$, hence it is almost-sure eventually synchronizing in q_2. The next result shows that infinite memory is necessary for almost-sure winning in this example.

Lemma 8. *There exists an almost-sure eventually synchronizing MDP for which all almost-sure eventually synchronizing strategies require infinite memory.*

Proof. Consider the MDP $\mathcal{M}$ shown in Figure 8. We argued before the lemma that $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{event}(q_2)$ and we now show that infinite memory is necessary from q_{init} for almost-sure eventually synchronizing in q_2. Note that $\mathcal{M}$ is not sure eventually synchronizing in q_2 since the probability in q_{init} is positive at all times (for all strategies).

Assume towards contradiction that there exists a (possibly randomized) finite-memory strategy α that is almost-sure eventually synchronizing in q_2. Consider the Markov chain $\mathcal{M}(\alpha)$ (the product of the MDP $\mathcal{M}$ with the finite-state transducer defining α). A state (q, m) in $\mathcal{M}(\alpha)$ is called a q-*state*. Since α is almost-sure eventually synchronizing (but is not sure eventually synchronizing) in q_2, there is a q_2-state in the recurrent states of $\mathcal{M}(\alpha)$. Since on all actions q_{init} is a successor of q_2, and q_{init} is a successor of itself, it follows that there is a recurrent q_{init}-state in $\mathcal{M}(\alpha)$, and that all periodic classes of recurrent

22

states in $\mathcal{M}(\alpha)$ contain a q_{init}-state. Hence, in each stationary distribution there is a q_{init}-state with a positive probability, and therefore the probability mass in q_{init} is bounded away from zero. It follows that the probability mass in q_2 is bounded away from 1 thus α is not almost-sure eventually synchronizing in q_2, a contradiction. $\qquad\square$

The membership problem for almost-sure eventually synchronizing can be reduced to other winning modes since an almost-sure eventually synchronizing strategy is either sure eventually synchronizing or almost-sure weakly synchronizing. Nevertheless we give a direct proof that the problem is decidable in PSPACE, using a characterization that will be useful later for almost-sure weakly synchronizing.

It turns out that in general, almost-sure eventually synchronizing strategies can be constructed from a family of limit-sure eventually synchronizing strategies if we can also ensure that the probability mass remains in the winning region (as in the MDP in Figure 8). We present a characterization of the winning region for almost-sure winning based on an extension of the limit-sure eventually synchronizing objective *with exact support*. This objective requires to ensure probability arbitrarily close to 1 in the target set T, and moreover that after the same number of steps the support of the probability distribution is contained in the given set U. Formally, given an MDP $\mathcal{M}$, let $\langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_T, U)$ for $T \subseteq U$ be the set of all initial distributions such that for all $\varepsilon > 0$ there exists a strategy α and $n \in \mathbb{N}$ such that $\mathcal{M}_n^\alpha(T) \geq 1 - \varepsilon$ and $\mathcal{M}_n^\alpha(U) = 1$. We say that α is limit-sure eventually synchronizing in T with support in U (consider the example at the beginning of Section 3.2 with $T = \{q_2\}$ and $U = \{q_{\mathsf{init}}, q_2\}$).

We will present an algorithmic solution to limit-sure eventually synchronizing objectives with exact support in Section 3.3. Our characterization of the winning region for almost-sure winning is as follows.

Lemma 9. *Let $\mathcal{M}$ be an MDP and T be a target set. For all states q_{init}, we have $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{event}(sum_T)$ if and only if there exists a set U of states such that:*

- *$q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_U)$, and*

- *$d_U \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_T, U)$ where d_U is the uniform distribution over U.*

Proof. First, if $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{event}(sum_T)$, then there is a strategy α such that $\sup_{n \in \mathbb{N}} \mathcal{M}_n^\alpha(T) = 1$. Then either $\mathcal{M}_n^\alpha(T) = 1$ for some $n \geq 0$, or $\limsup_{n \to \infty} \mathcal{M}_n^\alpha(T) = 1$. If $\mathcal{M}_n^\alpha(T) = 1$, then q_{init} is sure winning for eventually synchronizing in T, thus $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_T)$ and we can take $U = T$. Otherwise, for all $i > 0$ there exists $n_i \in \mathbb{N}$ such that $\mathcal{M}_{n_i}^\alpha(T) \geq 1 - 2^{-i}$, and moreover $n_{i+1} > n_i$ for all $i > 0$. Let $s_i = \mathsf{Supp}(\mathcal{M}_{n_i}^\alpha)$ be the support of $\mathcal{M}_{n_i}^\alpha$. Since the state space is finite, there is a set U that occurs infinitely often in the sequence $s_0 s_1 \dots$, thus for all $k > 0$ there exists $m_k \in \mathbb{N}$ such that $\mathcal{M}_{m_k}^\alpha(T) \geq 1 - 2^{-k}$ and $\mathcal{M}_{m_k}^\alpha(U) = 1$. It follows that α is sure eventually synchronizing in U from q_{init}, hence $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_U)$. Moreover, $\mathcal{M}$ with

initial distribution $d_1 = \mathcal{M}_{m_1}^\alpha$ is limit-sure eventually synchronizing in T with exact support in U. Since $\mathsf{Supp}(d_1) = U = \mathsf{Supp}(d_U)$, it follows by Corollary 15 that $d_U \in \langle\langle 1 \rangle\rangle_{limit}^{event}(sum_T, U)$.

To establish the converse, note that since $d_U \in \langle\langle 1 \rangle\rangle_{limit}^{event}(sum_T, U)$, it follows from Corollary 15 that from all initial distributions with support in U, for all $\varepsilon > 0$ there exists a strategy α_ε and a position n_ε such that $\mathcal{M}_{n_\varepsilon}^{\alpha_\varepsilon}(T) \geq 1 - \varepsilon$ and $\mathcal{M}_{n_\varepsilon}^{\alpha_\varepsilon}(U) = 1$. We construct an almost-sure limit eventually synchronizing strategy α as follows. Since $q_{\mathsf{init}} \in \langle\langle 1 \rangle\rangle_{sure}^{event}(sum_U)$, play according to a sure eventually synchronizing strategy from q_{init} until all the probability mass is in U. Then for $i = 1, 2, \dots$ and $\varepsilon_i = 2^{-i}$, repeat the following procedure: given the current probability distribution, select the corresponding strategy α_{ε_i} and play according to α_{ε_i} for n_{ε_i} steps, ensuring probability mass at least $1 - 2^{-i}$ in T, and since after that the support of the probability mass is again in U, play according to $\alpha_{\varepsilon_{i+1}}$ for $n_{\varepsilon_{i+1}}$ steps, etc. This strategy α ensures that $\sup_{n \in \mathbb{N}} \mathcal{M}_n^\alpha(T) = 1$ from q_{init}, hence $q_{\mathsf{init}} \in \langle\langle 1 \rangle\rangle_{almost}^{event}(sum_T)$. Note that α is a pure strategy. $\qquad\square$

As we show in Section 3.3 that the membership problem for limit-sure eventually synchronizing with exact support can be solved in PSPACE, it follows from the characterization in Lemma 9 that the membership problem for almost-sure eventually synchronizing is in PSPACE, using the following (N)PSPACE algorithm: guess the set U, and check that $q_{\mathsf{init}} \in \langle\langle 1 \rangle\rangle_{sure}^{event}(sum_U)$, and that $d_U \in \langle\langle 1 \rangle\rangle_{limit}^{event}(sum_T, U)$ where d_U is the uniform distribution over U (this can be done in PSPACE by Theorem 2 and Theorem 4). We present a matching lower bound.

Lemma 10. *The membership problem for $\langle\langle 1 \rangle\rangle_{almost}^{event}(sum_T)$ is PSPACE-hard even if T is a singleton.*

Proof. We show the result by a reduction from the membership problem for sure eventually synchronizing, which is PSPACE-complete by Theorem 2. Given an MDP $\mathcal{M} = \langle Q, \mathsf{A}, \delta \rangle$, an initial state $q_{\mathsf{init}} \in Q$, and a state $\hat{q} \in Q$, we construct an MDP $\mathcal{N} = \langle Q', \mathsf{A}', \delta' \rangle$ with $Q \subseteq Q'$ and a state $\hat{p} \in Q'$ such that $q_{\mathsf{init}} \in \langle\langle 1 \rangle\rangle_{sure}^{event}(\hat{q})$ in $\mathcal{M}$ if and only if $q_{\mathsf{init}} \in \langle\langle 1 \rangle\rangle_{almost}^{event}(\hat{p})$ in $\mathcal{N}$. The MDP $\mathcal{N}$ is a copy of $\mathcal{M}$ with two new states $\hat{p}$ and sink reachable only by a new action $\sharp$ (see Figure 9). Formally, $Q' = Q \cup \{\hat{p}, \mathsf{sink}\}$ and $\mathsf{A}' = \mathsf{A} \cup \{\sharp\}$, and the transition function δ' is defined as follows, for all $q \in Q$ and $a \in \mathsf{A}$: $\delta'(q, a) = \delta(q, a)$, and $\delta'(q, \sharp)(\mathsf{sink}) = 1$ if $q \neq \hat{q}$, and $\delta'(\hat{q}, \sharp)(\hat{p}) = 1$; finally, for all $a \in \mathsf{A}'$, let $\delta'(\hat{p}, a)(\mathsf{sink}) = \delta'(\mathsf{sink}, a)(\mathsf{sink}) = 1$.

The goal is that $\mathcal{N}$ simulates $\mathcal{M}$ until the action $\sharp$ is played in $\hat{q}$ to move the probability mass from $\hat{q}$ to $\hat{p}$, ensuring that if $\mathcal{M}$ is sure-winning for eventually synchronizing in $\hat{q}$, then $\mathcal{N}$ is also sure-winning (and thus almost-sure winning) for eventually synchronizing in $\hat{p}$. Moreover, the only way to be almost-sure eventually synchronizing in $\hat{p}$ is to have probability 1 in $\hat{p}$ at some point, because the state $\hat{p}$ is transient under all strategies, thus the probability mass cannot accumulate and tend to 1 in $\hat{p}$ in the long run. Therefore (from all initial states q_{init}) $\mathcal{M}$ is sure-winning for eventually synchronizing in $\hat{q}$ if and only if $\mathcal{N}$ is almost-sure winning for eventually synchronizing in $\hat{p}$. It follows
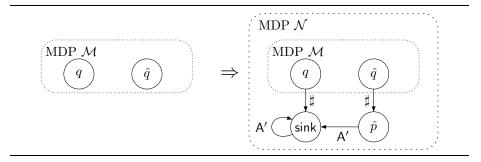
Figure 9: Sketch of the reduction to show PSPACE-hardness of the membership problem for almost-sure eventually synchronizing.

from this reduction that the membership problem for almost-sure eventually synchronizing objective is PSPACE-hard. □

The results of this section are summarized as follows.

Theorem 3. *For almost-sure eventually synchronizing in MDPs:*

1. *(Complexity). The membership problem is PSPACE-complete.*
2. *(Memory). Infinite memory is necessary in general for both pure and randomized strategies, and pure strategies are sufficient.*

3.3. Limit-sure eventually synchronizing

In this section, we present the algorithmic solution for limit-sure eventually synchronizing with exact support, which requires to get probability arbitrarily close to 1 in a target set T while all the probability mass is contained in a given set U. Note that the limit-sure eventually synchronizing objective is a special case where the support is the state space of the MDP. Consider the MDP in Figure 1 which is limit-sure eventually synchronizing in $\{q_2\}$, as shown in Lemma 3. For $i = 0, 1, \ldots$, the sequence $\mathsf{Pre}^i(T)$ of predecessors of $T = \{q_2\}$ is ultimately periodic: $\mathsf{Pre}^0(T) = \{q_2\}$, and $\mathsf{Pre}^i(T) = \{q_1\}$ for all $i \geq 1$. Given $\varepsilon > 0$, a strategy to get probability $1 - \varepsilon$ in q_2 first accumulates probability mass in the *periodic* subsequence of predecessors (here $\{q_1\}$), and when the probability mass is greater than $1 - \varepsilon$ in q_1, the strategy injects the probability mass in q_2 (through the aperiodic prefix of the sequence of predecessors). This is the typical shape of a limit-sure eventually synchronizing strategy. Note that in this scenario, the MDP is also limit-sure eventually synchronizing in every set $\mathsf{Pre}^i(T)$ of the sequence of predecessors. A special case is when it is possible to get probability 1 in the sequence of predecessors after finitely many steps. In this case, the probability mass injected in T is 1 and the MDP is even sure-winning. The algorithm for deciding limit-sure eventually synchronizing relies on the above characterization, generalized in Lemma 11 to limit-sure eventually synchronizing with exact support, saying that limit-sure eventually synchronizing in T with support in U is equivalent to either sure eventually synchronizing

25

in T (and therefore also in U), or limit-sure eventually synchronizing in $\mathsf{Pre}^k(T)$ with support in $\mathsf{Pre}^k(U)$ (for arbitrary k). The intuition of the proof is that if an MDP is limit-sure eventually synchronizing in T with support in U, then either a bounded number of steps is sufficient to get probability $1-\varepsilon$ in T (and then we argue that the MDP is sure eventually synchronizing), or unbounded number of steps is required, which means that k steps before getting probability $1-\varepsilon$ in T, the probability mass in $\mathsf{Pre}^k(T)$ must also be close to 1 (and arbitrarily close to 1 as ε tends to 0).

Lemma 11. *For all $T \subseteq U$ and all $k \geq 0$, we have*

$$\langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_T, U) = \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_T) \cup \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_R, Z)$$

where $R = \mathsf{Pre}^k(T)$ and $Z = \mathsf{Pre}^k(U)$.

Proof. We establish the equality in the lemma by showing inclusions in the two directions. First we show that

$$\langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_T) \cup \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_R, Z) \subseteq \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_T, U).$$

Since $T \subseteq U$, it follows that $\langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_T) \subseteq \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_T, U)$; to show that $\langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_R, Z) \subseteq \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_T, U)$ in an MDP $\mathcal{M}$, let $\varepsilon > 0$ and consider an initial distribution d_0 and a strategy α such that for some $i \geq 0$ we have $\mathcal{M}_i^\alpha(R) \geq 1-\varepsilon$ and $\mathcal{M}_i^\alpha(Z) = 1$. We construct a strategy β that plays like α for the first i steps, and then since $R = \mathsf{Pre}^k(T)$ and $Z = \mathsf{Pre}^k(U)$ plays from states in R according to a sure eventually synchronizing strategy with target T, and from states in $Z \setminus R$ according to a sure eventually synchronizing strategy with target U (such strategies exist by Lemma 7 since $R = \mathsf{Pre}^k(T)$). The strategy β ensures from d_0 that $\mathcal{M}_{i+k}^\beta(T) \geq 1 - \varepsilon$ and $\mathcal{M}_{i+k}^\beta(U) = 1$, showing that $\mathcal{M}$ is limit-sure eventually synchronizing in T with support in U.

Second we show the converse inclusion, namely that

$$\langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_T, U) \subseteq \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_T) \cup \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_R, Z).$$

Consider an initial distribution $d_0 \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_T, U)$ in the MDP $\mathcal{M}$ and for $\varepsilon_i = \frac{1}{i}$ ($i \in \mathbb{N}$) let α_i be a strategy and $n_i \in \mathbb{N}$ such that $\mathcal{M}_{n_i}^{\alpha_i}(T) \geq 1 - \varepsilon_i$ and $\mathcal{M}_{n_i}^{\alpha_i}(U) = 1$. We consider two cases.

(a) If the set $\{n_i \mid i \geq 0\}$ is bounded, then there exists a number n that occurs *infinitely often* in the sequence $(n_i)_{i \in \mathbb{N}}$. It follows that for all $i \geq 0$, there exists a strategy β_i such that $\mathcal{M}_n^{\beta_i}(T) \geq 1 - \varepsilon_i$ and $\mathcal{M}_n^{\beta_i}(U) = 1$. Since n is fixed, we can assume w.l.o.g. that the strategies β_i are pure, and since there is a finite number of pure strategies over paths of length at most n, it follows that there is a strategy β that occurs infinitely often among the strategies β_i and such that for all $\varepsilon > 0$ we have $\mathcal{M}_n^\beta(T) \geq 1 - \varepsilon$, hence $\mathcal{M}_n^\beta(T) = 1$, showing that $\mathcal{M}$ is sure winning for eventually synchronizing in T, that is $d_0 \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_T)$.

(*b*) otherwise, the set $\{n_i \mid i \geq 0\}$ is unbounded and we can assume w.l.o.g. that $n_i \geq k$ for all $i \geq 0$. We claim that the family of strategies α_i ensures limit-sure eventually synchronizing in $R = \mathsf{Pre}^k(T)$ with support in $Z = \mathsf{Pre}^k(U)$. Essentially this is because if the probability in T is close to 1 after n_i steps, then k steps before the probability in $\mathsf{Pre}^k(T)$ must be close to 1 as well. Formally, we show that α_i is such that $\mathcal{M}_{n_i-k}^{\alpha_i}(R) \geq 1 - \frac{\varepsilon_i}{\eta^k}$ and $\mathcal{M}_{n_i-k}^{\alpha_i}(Z) = 1$ where η is the smallest positive probability in the transitions of $\mathcal{M}$. Towards contradiction, assume that $\mathcal{M}_{n_i-k}^{\alpha_i}(R) < 1 - \frac{\varepsilon_i}{\eta^k}$. Then $\mathcal{M}_{n_i-k}^{\alpha_i}(Q \setminus R) > \frac{\varepsilon_i}{\eta^k}$ and from every state $q \in Q \setminus R$, no matter which sequence of actions is played by α_i for the next k steps, there is a path from q to a state outside T (by Lemma 7 since $R = \mathsf{Pre}^k(T)$), thus with probability at least η^k. Hence, the probability in $Q \setminus T$ after n_i steps is greater than $\frac{\varepsilon_i}{\eta^k} \cdot \eta^k$, and therefore $\mathcal{M}_{n_i}^{\alpha_i}(T) < 1 - \varepsilon_i$, in contradiction with the definition of α_i. This shows that $\mathcal{M}_{n_i-k}^{\alpha_i}(R) \geq 1 - \frac{\varepsilon_i}{\eta^k}$, and an argument analogous to the proof of Lemma 7 shows that $\mathcal{M}_{n_i-k}^{\alpha_i}(Z) = 1$. It follows that $d_0 \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_R, Z)$ and the proof is complete. $\qquad\square$

Thanks to Lemma 11, since sure-winning is already solved in Section 3.1, it suffices to solve the limit-sure eventually synchronizing problem for target $R = \mathsf{Pre}^k(T)$ and support $Z = \mathsf{Pre}^k(U)$ with arbitrary k, instead of T and U. We choose k such that both $\mathsf{Pre}^k(T)$ and $\mathsf{Pre}^k(U)$ lie in the periodic part of the sequence of pairs of predecessors $(\mathsf{Pre}^i(T), \mathsf{Pre}^i(U))$.

Note that $\mathsf{Pre}^i(T) \subseteq \mathsf{Pre}^i(U) \subseteq Q$ for all $i \geq 0$, and there are at most $3^{|Q|}$ different pairs (A, B) with $A \subseteq B \subseteq Q$ (each state $q \in Q$ belongs either to A, or to $B \setminus A$, or to $Q \setminus B$). Hence, we can assume that $k \leq 3^{|Q|}$.

For such value of k the limit-sure problem is conceptually simpler: once some probability is injected in $R = \mathsf{Pre}^k(T)$, it can loop through the sequence of predecessors and visit R infinitely often (every r steps, where $r \leq 3^{|Q|}$ is the period of the sequence of pairs of predecessors). It follows that if a strategy ensures with probability 1 that the set R can be reached by finite paths whose lengths are congruent modulo r, then the whole probability mass can indeed synchronously accumulate in R in the limit.

Therefore, limit-sure eventually synchronizing in R reduces to standard limit-sure reachability (in the state-based semantics) with target set R and the additional requirement that the numbers of steps at which the target set is reached be congruent modulo r. In the case of limit-sure eventually synchronizing with support in Z, we also need to ensure that no mass of probability leaves the sequence $\mathsf{Pre}^i(Z)$. In a state $q \in \mathsf{Pre}^i(Z)$, we say that an action $a \in \mathsf{A}$ is *Z-safe* at position i if[2] $\mathsf{post}(q, a) \subseteq \mathsf{Pre}^{i-1}(Z)$. In states $q \notin \mathsf{Pre}^i(Z)$ there is no Z-safe action at position i.

To encode the above requirements, we construct an MDP $\mathcal{M}_Z \times [r]$ that allows only Z-safe actions to be played (and then mimics the original MDP),

---

[2] Since $\mathsf{Pre}^r(Z) = Z$ and $\mathsf{Pre}^r(R) = R$, we assume a modular arithmetic for exponents of Pre, that is $\mathsf{Pre}^x(\cdot)$ is defined as $\mathsf{Pre}^{x \bmod r}(\cdot)$. For example $\mathsf{Pre}^{-1}(Z)$ is $\mathsf{Pre}^{r-1}(Z)$.

and tracks the position (modulo r) in the sequence of predecessors, thus simply decrementing the position on each transition since all successors of a state $q \in \mathsf{Pre}^i(Z)$ on a safe action are in $\mathsf{Pre}^{i-1}(Z)$.

Formally, if $\mathcal{M} = \langle Q, \mathsf{A}, \delta \rangle$ then $\mathcal{M}_Z \times [r] = \langle Q', \mathsf{A}, \delta' \rangle$ where

- $Q' = Q \times \{r-1, \ldots, 1, 0\} \cup \{\mathsf{sink}\}$; a state $\langle q, i \rangle$ consisting of a state q of $\mathcal{M}$ and a *position* i in the predecessor sequence corresponds to the promise that $q \in \mathsf{Pre}^i(Z)$;

- δ' is defined as follows for all $\langle q, i \rangle \in Q'$ and $a \in \mathsf{A}$ (assuming an arithmetic modulo r on positions): if a is a Z-safe action in q at position i, then

$$\delta'(\langle q, i \rangle, a)(\langle q', i-1 \rangle) = \delta(q, a)(q'),$$

otherwise $\delta'(\langle q, i \rangle, a)(\mathsf{sink}) = 1$ (and sink is absorbing).

Note that the size of the MDP $\mathcal{M}_Z \times [r]$ is exponential in the size of $\mathcal{M}$ (since r is at most $3^{|Q|}$).

Lemma 12. *Let $\mathcal{M}$ be an MDP and $R \subseteq Z$ be two sets of states such that $\mathsf{Pre}^r(R) = R$ and $\mathsf{Pre}^r(Z) = Z$ where $r > 0$. Then a state q_{init} is limit-sure eventually synchronizing in R with support in Z ($q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_R, Z)$) if and only if there exists $0 \le t < r$ such that $\langle q_{\mathsf{init}}, t \rangle$ is limit-sure winning for the reachability objective $\Diamond(R \times \{0\})$ in the MDP $\mathcal{M}_Z \times [r]$.*

Proof. For the first direction of the lemma, assume that q_{init} is limit-sure eventually synchronizing in R with support in Z, and for $\varepsilon > 0$ let β be a strategy such that $\mathcal{M}_k^\beta(Z) = 1$ and $\mathcal{M}_k^\beta(R) \ge 1 - \varepsilon$ for some number k of steps. Let $0 \le t \le r$ such that $t = k \mod r$. Let $R_0 = R \times \{0\}$. We show that from initial state (q_{init}, t) the strategy α in $\mathcal{M}_Z \times [r]$ that mimics (copies) the strategy β is limit-sure winning for the reachability objective $\Diamond R_0$: it follows from Lemma 7 that α plays only Z-safe actions, and since $Pr^\alpha(\Diamond R_0) \ge Pr^\alpha(\Diamond^k R_0) = \mathcal{M}_k^\beta(R) \ge 1 - \varepsilon$, the result follows.

For the converse direction, assuming that there exists $0 \le t < r$ such that $\langle q_{\mathsf{init}}, t \rangle$ is limit-sure winning for the reachability objective $\Diamond R_0$ in $\mathcal{M}_Z \times [r]$, we show that q_{init} is limit-sure synchronizing in target set R with exact support in Z. Since the winning region of limit-sure and almost-sure reachability coincide for MDPs [24], there exists a (pure) strategy α in $\mathcal{M}_Z \times [r]$ with initial state $\langle q_{\mathsf{init}}, t \rangle$ such that $\mathrm{Pr}^\alpha(\Diamond R_0) = 1$.

Given $\varepsilon > 0$, we construct from α a pure strategy β in $\mathcal{M}$ that is $(1 - \varepsilon)$-synchronizing in R with support in Z. Given a finite path $\rho = q_0 a_0 q_1 a_1 \ldots q_n$ in $\mathcal{M}$ (with $q_0 = q_{\mathsf{init}}$), there is a corresponding path $\rho' = \langle q_0, k_0 \rangle a_0 \langle q_1, k_1 \rangle a_1 \ldots \langle q_n, k_n \rangle$ in $\mathcal{M}_Z \times [r]$ where $k_0 = t$ and $k_{i+1} = k_i - 1$ for all $i \ge 0$. Since the sequence $k_0, k_1, \ldots$ is uniquely determined from ρ, there is a clear bijection between the paths in $\mathcal{M}$ starting in q_{init} and the paths in $\mathcal{M}_Z \times [r]$ starting in $\langle q_{\mathsf{init}}, t \rangle$. In the sequel, we freely omit to apply and mention this bijection. Define the strategy β as follows: if $q_n \in \mathsf{Pre}^{k_n}(R)$, then there exists an action a such that $\mathsf{post}(q_n, a) \subseteq \mathsf{Pre}^{k_n-1}(R)$ and we define $\beta(\rho) = a$,

otherwise let $\beta(\rho) = \alpha(\rho')$. Thus β mimics α (thus playing only Z-safe actions) unless a state q is reached at step n such that $q \in \mathsf{Pre}^{t-n}(R)$, and then β switches to always playing actions that are R-safe (and thus also Z-safe since $R \subseteq Z$). We now prove that β is limit-sure eventually synchronizing in target set R with support in Z. First since β plays only Z-safe actions, it follows for all k such that $t - k = 0$ (modulo r), all states reached from q_{init} with positive probability after k steps are in Z. Hence, $\mathcal{M}_k^\beta(Z) = 1$ for all such k. Second, we show that given $\varepsilon > 0$ there exists k such that $t - k = 0$ and $\mathcal{M}_k^\beta(R) \geq 1 - \varepsilon$, thus also $\mathcal{M}_k^\beta(Z) = 1$ and β is limit-sure eventually synchronizing in target set R with support in Z. To show this, recall that $\Pr^\alpha(\Diamond R_0) = 1$, and therefore $\Pr^\alpha(\Diamond^{\leq k} R_0) \geq 1 - \varepsilon$ for all sufficiently large k. Without loss of generality, consider such a k satisfying $t - k = 0$ (modulo r). For $i = 1, \ldots, r - 1$, let $R_i = \mathsf{Pre}^i(R) \times \{i\}$. Then trivially $\Pr^\alpha(\Diamond^{\leq k} \bigcup_{i=0}^r R_i) \geq 1 - \varepsilon$ and since β agrees with α on all finite paths that do not (yet) visit $\bigcup_{i=0}^r R_i$, given a path ρ that visits $\bigcup_{i=0}^r R_i$ (for the first time), only R-safe actions will be played by β and thus all continuations of ρ in the outcome of β will visit R after k steps (in total). It follows that $\Pr^\beta(\Diamond^k R_0) \geq 1 - \varepsilon$, that is $\mathcal{M}_k^\beta(R) \geq 1 - \varepsilon$. Note that we used the same pure strategy β for all $\varepsilon > 0$ and thus β is also almost-sure eventually synchronizing in R. □

From the proof of Lemma 12 (last sentence), it follows that if the MDP $\mathcal{M}$ is limit-sure eventually synchronizing in R with support in Z, then $\mathcal{M}$ is also almost-sure eventually synchronizing in R. Since almost-sure synchronization implies limit-sure synchronization by definition, the two notions coincide in this case.

Corollary 13. *Given $R \subseteq Z$ two sets of states in an MDP such that $\mathsf{Pre}^r(R) = R$ and $\mathsf{Pre}^r(Z) = Z$ where $r > 0$, we have $\langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_R, Z) = \langle\!\langle 1 \rangle\!\rangle_{almost}^{event}(sum_R)$.*

Since deciding limit-sure reachability is PTIME-complete, it follows from Lemma 12 that limit-sure eventually synchronizing (with exact support) can be decided in EXPTIME.

We show in Lemma 14 that the problem can be solved in PSPACE by exploiting the special structure of the exponential MDP used in Lemma 12. We conclude this section by Lemma 17 showing that limit-sure eventually synchronizing with exact support is PSPACE-complete (even in the special case where the support is the whole state space).

Lemma 14. *The membership problem for limit-sure eventually synchronizing with exact support is in PSPACE.*

Proof. We present a (nondeterministic) PSPACE algorithm to decide, given an MDP $\mathcal{M} = \langle Q, \mathsf{A}, \delta \rangle$, a state q_{init}, and two sets $T \subseteq U$, whether q_{init} is limit-sure eventually synchronizing in T with support in U.

First, the algorithm computes numbers $k \geq 0$ and $r > 0$ such that for $R = \mathsf{Pre}^k(T)$ and $Z = \mathsf{Pre}^k(U)$ we have $\mathsf{Pre}^r(R) = R$ and $\mathsf{Pre}^r(Z) = Z$. As

discussed before, this can be done by guessing $k, r \leq 3^{|Q|}$. By Lemma 11, we have

$$\langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_T, U) = \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_R, Z) \cup \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_T),$$

and since sure eventually synchronizing in T can be decided in PSPACE (by Theorem 2), it suffices to decide limit-sure eventually synchronizing in R with support in Z in PSPACE. According to Lemma 12, it is therefore sufficient to show that deciding limit-sure winning for the (standard) reachability objective $\Diamond(R \times \{0\})$ in the MDP $\mathcal{M}_Z \times [r]$ can be done in polynomial space. As we cannot afford to construct the exponential-size MDP $\mathcal{M}_Z \times [r]$, the algorithm relies on the following characterization of the limit-sure winning set for reachability objectives in MDPs. It is known that the winning region for limit-sure and almost-sure reachability coincide [24], and pure memoryless strategies are sufficient. Therefore, we can see that the almost-sure winning set W for the reachability objective $\Diamond(R \times \{0\})$ satisfies the following property: there exists a memoryless strategy $\alpha : W \to \mathsf{A}$ such that:

(1) W is closed, that is $\mathsf{post}(q, \alpha(q)) \subseteq W$ for all $q \in W$, and

(2) in the graph of the Markov chain $M(\alpha)$, for every state $q \in W$, there is a path (of length at most $|W|$) from q to some state in $R \times \{0\}$.

This property ensures that from every state in W, the target set $R \times \{0\}$ is reached within $|W|$ steps with positive (and bounded) probability, and since W is closed it ensures that $R \times \{0\}$ is reached with probability 1 in the long run. Thus any set W satisfying the above property is almost-sure winning.

Our algorithm will guess and explore on the fly a set W to ensure that it satisfies this property, and contains the state $\langle q_{\mathsf{init}}, t \rangle$ for some $t < r$. As we cannot afford to explicitly guess W (remember that W could be of exponential size), we decompose W into *slices* $W_0, W_1, \dots$ such that $W_i \subseteq Q$ and $W_i \times \{-i \mod r\} = W \cap (Q \times \{-i \mod r\})$. We start by guessing W_0, and we use the property that in $\mathcal{M}_Z \times [r]$, from a state (q, j) under all Z-safe actions, all successors are of the form $(\cdot, j-1)$. It follows that the successors of the states in $W_i \times \{-i\}$ should lie in the slice $W_{i+1} \times \{-i-1\}$, and we can guess on the fly the next slice $W_{i+1} \subseteq Q$ by guessing for each state q in a slice W_i an action a_q such that $\bigcup_{q \in W_i} \mathsf{post}(q, a_q) \subseteq W_{i+1}$. Moreover, we need to check the existence of a path from every state in W to $R \times \{0\}$. As W is closed, it is sufficient to check that there is a path from every state in $W_0 \times \{0\}$ to $R \times \{0\}$. To do this we guess along with the slices $W_0, W_1, \dots$ a sequence of sets $P_0, P_1, \dots$ where $P_i \subseteq W_i$ contains the states of slice W_i that belong to the guessed paths. Formally, $P_0 = W_0$, and for all $i \geq 0$, the set P_{i+1} is such that $\mathsf{post}(q, a_q) \cap P_{i+1} \neq \varnothing$ for all $q \in P_i'$ (where $P_i' = P_i \setminus R$ if i is a multiple of r, and $P_i' = P_i$ otherwise), that is P_{i+1} contains a successor of every state in P_i that is not already in the target R (at position 0 modulo r).

We need polynomial space to store the first slice W_0, the current slice W_i and the set P_i, and the value of i (in binary). As $\mathcal{M}_Z \times [r]$ has $|Q| \cdot r$ states, the algorithm runs for $|Q| \cdot r$ iterations and then checks that:

(1) $W_{|Q|\cdot r} \subseteq W_0$ to ensure that $W = \bigcup_{i \leq |Q|\cdot r} W_i \times \{i \mod r\}$ is closed,

(2) $P_{|Q|\cdot r} = \varnothing$ showing that from every state in $W_0 \times \{0\}$ there is a path to $R \times \{0\}$ (and thus also from all states in W), and

(3) the state q_{init} occurs in some slice W_i.

The correctness of the algorithm follows from the characterization of the almost-sure winning set for reachability in MDPs: if some state $\langle q_{\mathsf{init}}, t \rangle$ is limit-sure winning, then the algorithm accepts by guessing (slice by slice) the almost-sure winning set W and the paths from $W_0 \times \{0\}$ to $R \times \{0\}$ (at position 0 modulo r), and otherwise any set (and paths) correctly guessed by the algorithm would not contain q_{init} in any slice. $\qquad \square$

It follows from the proof of Lemma 12 that all winning modes for eventually synchronizing are independent of the numerical value of the positive transition probabilities.

Corollary 15. *Let $\mu \in \{sure, almost, limit\}$ and $T \subseteq U$ be two sets. For two distributions d, d' with $\mathsf{Supp}(d) = \mathsf{Supp}(d')$, we have $d \in \langle\!\langle 1 \rangle\!\rangle_\mu^{event}(sum_T, U)$ if and only if $d' \in \langle\!\langle 1 \rangle\!\rangle_\mu^{event}(sum_T, U)$.*

Remark 16. *Corollary 15 ensures that knowing the support of the initial distribution is sufficient to establish that it is eventually synchronizing. However, this corollary should be used carefully in the case of limit-sure eventually synchronizing: given a support $S \subseteq Q$, if for all $\varepsilon > 0$ there exists a distribution d_ε with support S that is eventually $(1 - \varepsilon)$-synchronizing, this does not imply that the distributions with support S are limit-sure eventually synchronizing.*

For example, consider an MDP with set of states $Q = \{q_1, q_2\}$, self-loops on both q_1 and q_2, and target set $T = \{q_1\}$ (with function sum_T). For $\varepsilon > 0$, the initial distribution d defined by $d(q_1) = 1 - \varepsilon$ and $d(q_2) = \varepsilon$ has support $S = Q$ and ensures probability $1 - \varepsilon$ in T. Thus for all $\varepsilon > 0$, we have an initial distribution that satisfies the requirement, but the uniform distribution over S is obviously not limit-sure eventually synchronizing in T.

To establish the PSPACE-hardness for limit-sure eventually synchronizing in MDPs, we use a reduction from the universal finiteness problem for 1L-AFAs.

Lemma 17. *The membership problem for $\langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_T)$ is PSPACE-hard even if T is a singleton.*

Proof. We show the result by a reduction from the universal finiteness problem for one-letter alternating automata (1L-AFA), which is PSPACE-complete (by Lemma 6). It is easy to see that this problem remains PSPACE-complete even if the set T of accepting states of the 1L-AFA is a singleton, and given the tight relation between 1L-AFA and MDP (see Section 2.4), it follows from the definition of the universal finiteness problem that deciding, in an MDP $\mathcal{M}$, whether the sequence $\mathsf{Pre}_{\mathcal{M}}^n(T) \neq \varnothing$ for all $n \geq 0$ is PSPACE-complete.
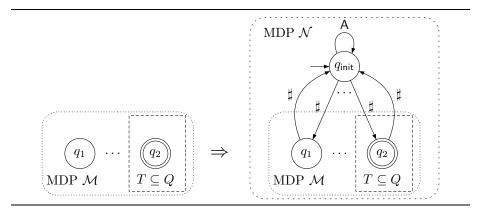
Figure 10: Sketch of the reduction to show PSPACE-hardness of the membership problem for limit-sure eventually and almost-sure weakly synchronizing.

The reduction is as follows (see also Figure 10). Given an MDP $\mathcal{M} = \langle Q, \mathsf{A}, \delta \rangle$ and a singleton $T \subseteq Q$, we construct an MDP $\mathcal{N} = \langle Q', \mathsf{A}', \delta' \rangle$ with state space $Q' = Q \uplus \{q_{\mathsf{init}}\}$ such that $\mathsf{Pre}^n_{\mathcal{M}}(T) \neq \varnothing$ for all $n \geq 0$ if and only if q_{init} is limit-sure eventually synchronizing in T. The MDP $\mathcal{N}$ is essentially a copy of $\mathcal{M}$ with alphabet $\mathsf{A} \uplus \{\sharp\}$ and the transition function on action $\sharp$ is the uniform distribution on Q from q_{init}, and the Dirac distribution on q_{init} from the other states $q \in Q$. There are self-loops on q_{init} for all other actions $a \in \mathsf{A}$. Formally, the transition function δ' is defined as follows, for all $q \in Q$:

- $\delta'(q, a) = \delta(q, a)$ for all $a \in \mathsf{A}$ (copy of $\mathcal{M}$), and $\delta'(q, \sharp)(q_{\mathsf{init}}) = 1$;

- $\delta'(q_{\mathsf{init}}, a)(q_{\mathsf{init}}) = 1$ for all $a \in \mathsf{A}$, and $\delta'(q_{\mathsf{init}}, \sharp)(q) = \frac{1}{|Q|}$.

We establish the correctness of the reduction as follows. For the first direction, assume that $\mathsf{Pre}^n_{\mathcal{M}}(T) \neq \varnothing$ for all $n \geq 0$. Then since $\mathcal{N}$ embeds a copy of $\mathcal{M}$ it follows that $\mathsf{Pre}^n_{\mathcal{N}}(T) \neq \varnothing$ for all $n \geq 0$ and there exist numbers $k_0, r \leq 2^{|Q|}$ such that $\mathsf{Pre}^{k_0+r}_{\mathcal{N}}(T) = \mathsf{Pre}^{k_0}_{\mathcal{N}}(T) \neq \varnothing$. Using Lemma 11 with $k = k_0$ and $R = \mathsf{Pre}^{k_0}_{\mathcal{N}}(T)$ (and $U = Z = Q'$ is the trivial support), it is sufficient to prove that $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle^{event}_{limit}(R)$ to get $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle^{event}_{limit}(T)$ (in $\mathcal{N}$). We show the stronger statement that q_{init} is actually almost-sure eventually synchronizing in R with the pure strategy α defined as follows, for all play prefixes ρ (let $m = |\rho| \bmod r$):

- if $\mathsf{Last}(\rho) = q_{\mathsf{init}}$, then $\alpha(\rho) = \sharp$;

- if $\mathsf{Last}(\rho) = q \in Q$, then

  - if $q \in \mathsf{Pre}^{r-m}_{\mathcal{N}}(R)$, then $\alpha(\rho)$ plays a R-safe action at position $r - m$;
  - otherwise, $\alpha(\rho) = \sharp$.

The strategy α ensures that the probability mass that is not (yet) in the sequence of predecessors $\mathsf{Pre}_{\mathcal{N}}^n(R)$ goes to q_{init}, where by playing $\sharp$ at least a fraction $\frac{1}{|Q|}$ of it would reach the sequence of predecessors (at a synchronized position). It follows that after $2i$ steps, the probability mass in q_{init} is $(1 - \frac{1}{|Q|})^i$ and the probability mass in the sequence of predecessors is $1 - (1 - \frac{1}{|Q|})^i$. For $i \to \infty$, the probability in the sequence of predecessors tends to 1 and since $\mathsf{Pre}_{\mathcal{N}}^n(R) = R$ for all positions n that are a multiple of r, we get $\sup_n \mathcal{M}_n^\alpha(R) = 1$ and $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{event}(R)$.

For the converse direction, assume that $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(T)$ is limit-sure eventually synchronizing in T. By Lemma 11, either (1) q_{init} is limit-sure eventually synchronizing in $\mathsf{Pre}_{\mathcal{N}}^n(T)$ for all $n \geq 0$, and then it follows that $\mathsf{Pre}_{\mathcal{N}}^n(T) \neq \varnothing$ for all $n \geq 0$, or (2) q_{init} is sure eventually synchronizing in T, and then since only the action $\sharp$ leaves the state q_{init} (and $\mathsf{post}(q_{\mathsf{init}}, \sharp) = Q$), the characterization of sure eventually synchronizing in Lemma 7 shows that $Q \subseteq \mathsf{Pre}_{\mathcal{N}}^k(T)$ for some $k \geq 0$, and since $Q \subseteq \mathsf{Pre}_{\mathcal{N}}(Q)$ and $\mathsf{Pre}_{\mathcal{N}}(\cdot)$ is a monotone operator, it follows that $Q \subseteq \mathsf{Pre}_{\mathcal{N}}^n(T)$ for all $n \geq k$ and thus $\mathsf{Pre}_{\mathcal{N}}^n(T) \neq \varnothing$ for all $n \geq 0$. We conclude the proof by noting that $\mathsf{Pre}_{\mathcal{M}}^n(T) = \mathsf{Pre}_{\mathcal{N}}^n(T) \cap Q$ and therefore $\mathsf{Pre}_{\mathcal{M}}^n(T) \neq \varnothing$ for all $n \geq 0$. $\qquad\square$

The example in the proof of Lemma 8 can be used to show that the memory needed by a family of strategies to win limit-sure eventually synchronizing objective (in target $T = \{q_2\}$) is unbounded. Observe that given $\varepsilon > 0$, the required memory to accumulate $1 - \varepsilon$ in T is finite, but the memory size increases and cannot be bounded as ε tends to 0. The following theorem summarizes the results for limit-sure eventually synchronizing.

Theorem 4. *For limit-sure eventually synchronizing (with or without exact support) in MDPs:*

1. *(Complexity). The membership problem is PSPACE-complete.*
2. *(Memory). Unbounded memory is required for both pure and randomized strategies, and pure strategies are sufficient.*

4. Weakly Synchronizing

We establish the complexity and memory requirement for weakly synchronizing objectives. We show that the membership problem is PSPACE-complete for sure and almost-sure winning, that exponential memory is necessary and sufficient for sure winning while infinite memory is necessary for almost-sure winning, and we show that limit-sure and almost-sure winning coincide. By Lemma 4, the complexity results established in this section for function sum_T hold for function $\max_T$ as well.

The weakly synchronizing objective is reminiscent of a Büchi objective in the distribution-based semantics: it requires that in the sequence of distributions of an MDP $\mathcal{M}$ under strategy α we have $\limsup_{n \to \infty} \mathcal{M}_n^\alpha(T) = 1$ (and that $\mathcal{M}_n^\alpha(T) = 1$ for infinitely many n in the case of sure winning).

The sure winning mode can be solved by a technique similar to the search for a lasso in Büchi automata [62] (Section 4.1). We show that the almost-sure winning mode can be solved by a reduction analogous to the case of eventually synchronizing (Section 4.2). For the limit-sure winning mode, we show that it coincides with the almost-sure winning mode. The proof of this result is technical and requires a careful characterization of the limit-sure winning mode. We present examples to provide intuitive illustration of the proof (Section 4.3).

4.1. Sure weakly synchronizing

The PSPACE upper bound of the membership problem for sure weakly synchronizing is obtained by the following characterization.

Lemma 18. *Let $\mathcal{M}$ be an MDP and T be a target set. For all states q_{init}, we have*

$q_{\text{init}} \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{weakly}(sum_T)$ if and only if there exists a set $S \subseteq T$ such that $q_{\text{init}} \in \text{Pre}^m(S)$ for some $m \geq 0$ and $S \subseteq \text{Pre}^n(S)$ for some $n \geq 1$.

Proof. First, if $q_{\text{init}} \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{weakly}(sum_T)$, then let α be a sure winning weakly synchronizing strategy. Then there are infinitely many positions n such that $\mathcal{M}_n^\alpha(T) = 1$, and since the state space is finite, there is a set S of states such that for infinitely many positions n we have $\text{Supp}(\mathcal{M}_n^\alpha) = S$ and $\mathcal{M}_n^\alpha(T) = 1$, and thus $S \subseteq T$. By the characterization of sure eventually synchronizing in Lemma 7, it follows that $q_{\text{init}} \in \text{Pre}^m(S)$ for some $m \geq 0$, and by considering two positions $n_1 < n_2$ where $\text{Supp}(\mathcal{M}_{n_1}^\alpha) = \text{Supp}(\mathcal{M}_{n_2}^\alpha) = S$, it follows that $S \subseteq \text{Pre}^n(S)$ for $n = n_2 - n_1 \geq 1$.

The reverse direction is straightforward by considering a strategy α that ensures $\mathcal{M}_m^\alpha(S) = 1$ for some $m \geq 0$, and then ensures that the probability mass from all states in S remains in S after every multiple of n steps where $n > 0$ is such that $S \subseteq \text{Pre}^n(S)$, showing that α is a sure winning weakly synchronizing strategy in S (and thus in T) from q_{init}, thus $q_{\text{init}} \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{weakly}(sum_T)$. Note that α is a pure strategy. □

The PSPACE upper bound follows from the characterization in Lemma 18. A (N)PSPACE algorithm is to guess the set $S \subseteq T$, and the numbers m, n (with $m, n \leq 2^{|Q|}$ since the sequence $\text{Pre}^n(S)$ of predecessors is ultimately periodic), and check that $q_{\text{init}} \in \text{Pre}^m(S)$ and $S \subseteq \text{Pre}^n(S)$. We present a matching PSPACE lower bound in the following lemma.

Lemma 19. *The membership problem for $\langle\!\langle 1 \rangle\!\rangle_{sure}^{weakly}(sum_T)$ is PSPACE-hard even if T is a singleton.*

Proof. We show the result by a reduction from the membership problem for $\langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_T)$ with a singleton T, which is PSPACE-complete (Theorem 2). From an MDP $\mathcal{M} = \langle Q, \mathsf{A}, \delta \rangle$ with initial state q_{init} and target state $\hat{q}$, we construct another MDP $\mathcal{N} = \langle Q', \mathsf{A}', \delta' \rangle$ and a target state $\hat{p}$ such that $q_{\text{init}} \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(\hat{q})$ in $\mathcal{M}$ if and only if $q_{\text{init}} \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{weakly}(\hat{p})$ in $\mathcal{N}$.
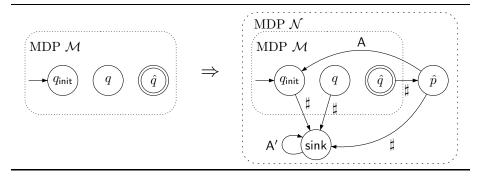
Figure 11: The reduction sketch to show PSPACE-hardness of the membership problem for sure weakly synchronizing in MDPs.

The MDP $\mathcal{N}$ is a copy of $\mathcal{M}$ with two new states $\hat{p}$ and sink that are reachable only by a new action $\sharp$ (see Figure 11). Formally, $Q' = Q \cup \{\hat{p}, \text{sink}\}$ and $A' = A \cup \{\sharp\}$. The transition function δ' is defined as follows: $\delta'(q,a) = \delta(q,a)$ for all states $q \in Q$ and $a \in A$, $\delta(q, \sharp)(\text{sink}) = 1$ for all $q \in Q' \setminus \{\hat{q}\}$ and $\delta(\hat{q}, \sharp)(\hat{p}) = 1$. The state sink is absorbing and from state $\hat{p}$ all other transitions lead to the initial state, i.e. $\delta(\text{sink}, a)(\text{sink}) = 1$ and $\delta(\hat{p}, a)(q_{\text{init}}) = 1$ for all $a \in A$.

We establish the correctness of the reduction as follows. First, if $q_{\text{init}} \in \langle\langle 1 \rangle\rangle_{sure}^{event}(\hat{q})$ in $\mathcal{M}$, then let α be a sure winning strategy in $\mathcal{M}$ for eventually synchronizing in $\{\hat{q}\}$. A sure winning strategy in $\mathcal{N}$ for weakly synchronizing in $\{\hat{p}\}$ is to play according to α until the whole probability mass is in $\hat{q}$, then play $\sharp$ followed by some $a \in A$ to visit $\hat{p}$ and get back to the initial state q_{init}, and then repeat the same strategy from q_{init}. Hence, $q_{\text{init}} \in \langle\langle 1 \rangle\rangle_{sure}^{weakly}(\hat{p})$ in $\mathcal{N}$.

Second, if $q_{\text{init}} \in \langle\langle 1 \rangle\rangle_{sure}^{weakly}(\hat{p})$ in $\mathcal{N}$, then consider a strategy α such that $\mathcal{N}_n^\alpha(\hat{p}) = 1$ for some $n \geq 0$. By construction of $\mathcal{N}$, it follows that $\mathcal{N}_{n-1}^\alpha(\hat{q}) = 1$, that is all path-outcomes of α of length $n-1$ reach $\hat{q}$, and α plays $\sharp$ in the next step. If α never plays $\sharp$ before position $n-1$, then α is a valid strategy in $\mathcal{M}$ up to step $n-1$ and it shows that $q_{\text{init}} \in \langle\langle 1 \rangle\rangle_{sure}^{event}(\hat{q})$ is sure winning in $\mathcal{M}$ for eventually synchronizing in $\{\hat{q}\}$. Otherwise let m be the largest number such that there is a finite path-outcome ρ of α of length $m < n-1$ with $\sharp \in \text{Supp}(\alpha(\rho))$. Thus between position m and $n-1$, the strategy α does not play $\sharp$. Note that the action $\sharp$ can be played by α only in the state $\hat{q}$, and thus $\text{Last}(\rho) = \hat{q}$. Hence two steps later, in the path-outcome ρ' of length $m+2$ that extends ρ, we have $\text{Last}(\rho') = q_{\text{init}}$. Since the action $\sharp$ is not played by α until position $n-1$, after position $m+2$ in ρ' the strategy α corresponds to a valid strategy from $\text{Last}(\rho')$ in $\mathcal{M}$ that brings all the probability mass of $\text{Last}(\rho') = q_{\text{init}}$ to $\hat{q}$, witnessing that $q_{\text{init}} \in \langle\langle 1 \rangle\rangle_{sure}^{event}(\hat{q})$. □

The proof of Lemma 18 suggests an exponential-memory pure strategy for sure weakly synchronizing that in $q \in \text{Pre}^n(S)$ plays an action a such that $\text{post}(q,a) \subseteq \text{Pre}^{n-1}(S)$, which can be realized with exponential memory since $n \leq 2^{|Q|}$. It can be shown that exponential memory is necessary in general,

using an argument similar to the proof of exponential memory lower bound for sure eventually synchronizing (Theorem 2), and by modifying the MDPs $\mathcal{M}_n$ (illustrated in Figure 7) as follows: let the transitions from state q_T go to q_{init} (instead of the absorbing state $q_\perp$).

Theorem 5. *For sure weakly synchronizing in MDPs:*

1. *(Complexity). The membership problem is PSPACE-complete.*
2. *(Memory). Exponential memory is necessary and sufficient for both pure and randomized strategies, and pure strategies are sufficient.*

4.2. Almost-sure weakly synchronizing

   We present a characterization of almost-sure weakly synchronizing that gives a PSPACE upper bound for the membership problem. Our characterization, similar to Lemma 9 for almost-sure eventually synchronizing, uses the limit-sure eventually synchronizing objectives with exact support introduced in Section 3.2. We show that an MDP is almost-sure weakly synchronizing in target T if (and only if), for some set U, there is a sure eventually synchronizing strategy in target U, and from the probability distributions with support U there is a limit-sure winning strategy for eventually synchronizing in $\mathsf{Pre}(T)$ with support in $\mathsf{Pre}(U)$. This ensures that from the initial state we can have the whole probability mass in U, and from U have probability $1 - \varepsilon$ in $\mathsf{Pre}(T)$ (and in T in the next step), while the whole probability mass is back in $\mathsf{Pre}(U)$ (and in U in the next step), allowing to repeat the strategy for $\varepsilon \to 0$, thus ensuring infinitely often probability at least $1 - \varepsilon$ in T (for all $\varepsilon > 0$).

Lemma 20. *Let $\mathcal{M}$ be an MDP and T be a target set. For all states q_{init}, we have $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{weakly}(sum_T)$ if and only if there exists a set U of states such that*

- *$q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_U)$, and*

- *$d_U \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_{\mathsf{Pre}(T)}, \mathsf{Pre}(U))$ where d_U is the uniform distribution over U.*

   Proof. First, if $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{weakly}(sum_T)$, then there exists a strategy α such that for all $i \geq 0$ there exists $n_i \in \mathbb{N}$ such that $\mathcal{M}_{n_i}^\alpha(T) \geq 1 - 2^{-i}$, and moreover $n_{i+1} > n_i$ for all $i \geq 0$. Let $s_i = \mathsf{Supp}(\mathcal{M}_{n_i}^\alpha)$ be the support of $\mathcal{M}_{n_i}^\alpha$. Since the state space is finite, there is a set U that occurs infinitely often in the sequence $s_0 s_1 \ldots$, thus for all $k > 0$ there exists $m_k \in \mathbb{N}$ such that $\mathcal{M}_{m_k}^\alpha(T) \geq 1 - 2^{-k}$ and $\mathcal{M}_{m_k}^\alpha(U) = 1$. It follows that α is sure eventually synchronizing in U from q_{init}, i.e. $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_U)$. Moreover, we can assume that $m_{k+1} > m_k$ for all $k > 0$ and thus $\mathcal{M}$ is also limit-sure eventually synchronizing in $\mathsf{Pre}(T)$ with exact support in $\mathsf{Pre}(U)$ from the initial distribution[3] $d_1 = \mathcal{M}_{m_1}^\alpha$. Since

---

[3]Note that the initial distribution $d_1 = \mathcal{M}_{m_1}^\alpha$ can be fixed before the other quantifications in the statement that we want to prove, namely: $\exists d_1 \in \mathcal{D}(U) \cdot \forall \varepsilon > 0 \cdot \exists \alpha \cdot \exists m_k : \mathcal{M}_{m_k}^\alpha(T) \geq$

$\mathsf{Supp}(d_1) = U = \mathsf{Supp}(d_U)$ and since only the support of the initial probability distributions is relevant for the limit-sure eventually synchronizing objective (Corollary 15), it follows that $d_U \in \langle\!\langle 1 \rangle\!\rangle^{event}_{limit}(sum_{\mathsf{Pre}(T)}, \mathsf{Pre}(U))$.

To establish the converse, note that since $d_U \in \langle\!\langle 1 \rangle\!\rangle^{event}_{limit}(sum_{\mathsf{Pre}(T)}, \mathsf{Pre}(U))$, it follows from Corollary 15 that from all initial distributions with support in U, for all $\varepsilon > 0$ there exists a strategy α_ε and a position n_ε such that $\mathcal{M}^{\alpha_\varepsilon}_{n_\varepsilon}(T) \geq 1 - \varepsilon$ and $\mathcal{M}^{\alpha_\varepsilon}_{n_\varepsilon}(U) = 1$. We construct an almost-sure weakly synchronizing strategy α as follows:

- Since $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle^{event}_{sure}(sum_U)$, play according to a sure eventually synchronizing strategy from q_{init} until all the probability mass is in U.

- Then for $i = 1, 2, \ldots$ and $\varepsilon_i = 2^{-i}$, repeat the following procedure:

  - given the current probability distribution, play according to α_{ε_i} for n_{ε_i} steps (ensuring probability mass at least $1 - 2^{-i}$ in $\mathsf{Pre}(T)$ and support of the probability mass in $\mathsf{Pre}(U)$);

  - then from states in $\mathsf{Pre}(T)$, play an action to ensure reaching T in the next step, and from states in $\mathsf{Pre}(U)$ ensure reaching U.

  - continue playing according to $\alpha_{\varepsilon_{i+1}}$ for $n_{\varepsilon_{i+1}}$ steps, etc.

Since $n_{\varepsilon_i} + 1 > 0$ for all $i \geq 0$, this strategy ensures that $\limsup_{n \to \infty} \mathcal{M}^{\alpha}_n(T) = 1$ from q_{init}, hence $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle^{weak}_{almost}(sum_T)$. Note that this strategy is pure. □

Since the membership problems for sure eventually synchronizing and for limit-sure eventually synchronizing with exact support are PSPACE-complete (Theorem 2 and Theorem 4), the membership problem for almost-sure weakly synchronizing is in PSPACE by guessing the set U, and checking that $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle^{event}_{sure}(sum_U)$, and that $d_U \in \langle\!\langle 1 \rangle\!\rangle^{event}_{limit}(sum_{\mathsf{Pre}(T)}, \mathsf{Pre}(U))$. We establish a matching PSPACE lower bound.

Lemma 21. *The membership problem for $\langle\!\langle 1 \rangle\!\rangle^{weakly}_{almost}(sum_T)$ is PSPACE-hard even if T is a singleton.*

Proof. We use the same reduction and construction as in the PSPACE-hardness proof of Lemma 17 where from an MDP $\mathcal{M}$ and a singleton T, we constructed N and q_{init}. Referring to that construction, we show that $\mathsf{Pre}^n_{\mathcal{M}}(T) \neq \varnothing$ for all $n \geq 0$ if and only if $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle^{weakly}_{almost}(T)$.

First, if $\mathsf{Pre}^n_{\mathcal{M}}(T) \neq \varnothing$ for all $n \geq 0$, then by Lemma 20 we need to show that (i) $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle^{event}_{sure}(sum_Q)$, and (ii) $d_Q \in \langle\!\langle 1 \rangle\!\rangle^{event}_{limit}(sum_{\mathsf{Pre}(T)}, \mathsf{Pre}(Q))$ where d_Q is the uniform distribution over Q. To show (i), we can play $\sharp$ from q_{init} to get

---

$1 - 2^{-k}$ where we compute $\mathcal{M}^\alpha$ with initial distribution d_1. This is because we fixed the strategy α in the first step of the proof, and this is why we need that q_{init} is almost-sure weakly synchronizing. Otherwise, if q_{init} is only limit-sure weakly synchronizing, we would get a possibly different initial distribution d_1 for each $\varepsilon > 0$ (induced by a possibly different strategy α for each ε) which would be problematic (see the example in Remark 16, p.31).
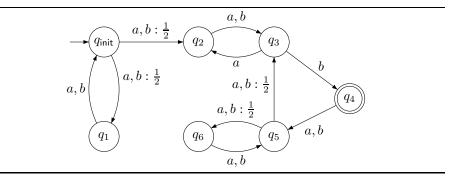
Figure 12: An example to show $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{weakly}(q_4)$ implies $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{weakly}(q_4)$.

the probability mass synchronized in Q. To show (ii), since playing $\sharp$ from d_Q ensures to reach q_{init}, it suffices to prove that $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_T, Q)$, which is done in the proof of Lemma 17.

For the converse direction, if q_{init} is almost-sure weakly synchronizing in T, then q_{init} is also limit-sure eventually synchronizing in T, and we can directly use that argument in the proof of Lemma 17 to show that $\mathsf{Pre}_{\mathcal{M}}^n(T) \neq \varnothing$ for all $n \geq 0$.

It follows from this reduction that the membership problem for almost-sure weakly synchronization is PSPACE-hard. □

It is easy to show that winning strategies require infinite memory for almost-sure weakly synchronizing in the same example that we used in the proof of Lemma 8 to show that infinite memory may be necessary for almost-sure eventually synchronizing (Figure 8),

Theorem 6. *For almost-sure weakly synchronizing in MDPs:*

1. *(Complexity). The membership problem is PSPACE-complete.*
2. *(Memory). Infinite memory is necessary in general for both pure and randomized strategies, and pure strategies are sufficient.*

4.3. Limit-sure weakly synchronizing

We show that the winning regions for almost-sure and limit-sure weakly synchronizing coincide. The result is not intuitively obvious (recall that it does not hold for eventually synchronizing, by Lemma 3(ii)) and requires a careful analysis of the structure of limit-sure winning strategies to show that they always imply the existence of an almost-sure winning strategy. The construction of an almost-sure winning strategy from a family of limit-sure winning strategies is illustrated in the following example.

Consider the MDP $\mathcal{M}_{12}$ in Figure 12 with initial state q_{init} and target set $T = \{q_4\}$. Note that there is a relevant strategic choice only in q_3, where we can either loop through q_2, or go to the target q_4. First we argue that $\mathcal{M}_{12}$ is limit-sure weakly synchronizing, then we explain why limit-sure weakly synchronizing

implies that we can construct an almost-sure weakly synchronizing strategy in this example, using the same line of arguments as in our proof of the general result (that limit-sure winning implies almost-sure winning) presented further as Lemma 22, Lemma 23, and Theorem 7.

4.3.1. The MDP $\mathcal{M}_{12}$ is limit-sure weakly synchronizing

To show that $\mathcal{M}_{12}$ (Figure 12) is limit-sure weakly synchronizing, we rely on the following claims:

- q_{init} is limit-sure eventually synchronizing with target $T = \{q_4\}$;

- q_4, which can be viewed as a uniform distribution over T, is also limit-sure eventually synchronizing with target $T = \{q_4\}$ (even after at least one step).

The above claims hold since for arbitrarily small $\varepsilon > 0$, from both q_{init} and q_4, we can inject probability mass $1 - \varepsilon$ in q_3 (by playing a long enough in q_3), and then switching to playing b in q_3 gets probability $1 - \varepsilon$ in T.

Now, these two claims are sufficient to show that q_{init} is limit-sure weakly synchronizing in $T = \{q_4\}$, and to define a family α_ε of limit-sure winning strategies as follows: given $\varepsilon > 0$, let α_ε play from q_{init} a strategy to ensure probability at least $p_1 = 1 - \frac{\varepsilon}{2}$ in q_4 (in finitely many steps), and then play according to a strategy that ensures from q_4 probability $p_2 = p_1 - \frac{\varepsilon}{4}$ in q_4 (in finitely many, and at least one step), and repeat this process using strategies that ensure, if the probability mass in q_4 is at least p_i, that (in at least one step) the probability in q_4 is at least $p_{i+1} = p_i - \frac{\varepsilon}{2^{i+1}}$. It follows that $p_i = 1 - \frac{\varepsilon}{2} - \frac{\varepsilon}{4} - \cdots - \frac{\varepsilon}{2^i} > 1 - \varepsilon$ for all $i \geq 1$, and thus $\limsup_{i \to \infty} p_i \geq 1 - \varepsilon$, thus α_ε is weakly $(1-\varepsilon)$-synchronizing. Therefore q_{init} is limit-sure weakly synchronizing for target $\{q_4\}$.

Illustration of Lemma 23. We show in Lemma 23 that in general the above two claims hold in a limit-sure weakly synchronizing MDP (and it is easy to generalize the argument we used for $\mathcal{M}_{12}$ to show that the converse implication of Lemma 23 holds as well, although we do not need to prove this for our purpose). Hence, Lemma 23 shows that limit-sure weakly synchronizing strategies can always be decomposed as a repetition of eventually $(1 - \varepsilon)$-synchronizing strategies, played for finitely many steps (and with decreasing ε).

4.3.2. The MDP $\mathcal{M}_{12}$ is almost-sure weakly synchronizing

The following claims are central to show that $\mathcal{M}_{12}$ (Figure 12) is almost-sure weakly synchronizing (note the slight difference with the claims in Section 4.3.1):

- q_{init} is limit-sure eventually synchronizing with target $\{q_3\}$;

- q_4, which can be viewed as a uniform distribution over T, is also limit-sure eventually synchronizing with target $\{q_3\}$.

The above claims hold by the exact same argument as in Section 4.3.1 (and follow directly from the fact that $\mathcal{M}_{12}$ is limit-sure weakly synchronizing). Intuitively, an almost-sure weakly synchronizing strategy in $\mathcal{M}_{12}$ repeats the following phases (informally):

1. accumulate probability mass (arbitrarily close to 1, say $1 - \varepsilon_0$) in q_3;
2. transfer the probability mass from q_3 to q_4;
3. given the current distribution, decrease ε_0 by half and repeat from (1.).

Such a strategy would ensure, for all $\varepsilon > 0$, probability mass at least $1 - \varepsilon$ in q_4 infinitely often, and thus it is almost-sure weakly synchronizing. To show that such a strategy exists and is well defined, we need to show that at every iteration from the distribution at the beginning of step (1.) we can indeed accumulate probability mass in q_3. This is true in the first iteration, as we start from q_{init}. After one iteration, the distribution has support $S = \{q_1, q_2, q_4\} = \{q_1, q_2\} \cup \{q_4\}$ where the distributions over $\{q_1, q_2\}$ are limit-sure eventually synchronizing to $\{q_3\}$ (by analogous argument as the first claim above), and q_4 is also limit-sure eventually synchronizing to $\{q_3\}$ (by the second claim above). In the next iterations, from the distribution at step 1 the situation is similar (as in fact all states are limit-sure weakly synchronizing with target $\{q_3\}$).

Illustration of Theorem 7 (claim 1 of the proof). The state q_3 plays a crucial role here because $\{q_3\} = \mathsf{Pre}(T)$ and $\{q_3\} = \mathsf{Pre}^2(\{q_3\})$, thus $R = \{q_3\}$ occurs infinitely often in the sequence $\mathsf{Pre}^i(T)$ (for $i \geq 0$), which is ultimately periodic with period $r = 2$. It follows from the general result established in Claim 1 of the proof of Theorem 7 that limit-sure weakly synchronizing with target $T = \{q_4\}$ implies limit-sure eventually (and even almost-sure weakly) synchronizing with target $\{q_3\}$.

Intuitively, from the fact that the distributions over both $\{q_1, q_2\}$ and $\{q_4\}$ are limit-sure eventually synchronizing to $\{q_3\}$, it may not be obvious that the distributions over $\{q_1, q_2, q_4\}$ are limit-sure eventually synchronizing to $\{q_3\}$. For instance in the example of $\mathcal{M}_{12}$ (Figure 12), (almost all) the probability mass in $T = \{q_4\}$ can move to q_3 in an even number of steps, while from $\{q_1, q_2\}$ an odd number of steps is required, resulting in a *shift* of the probability mass.

Illustration of Theorem 7 (claim 2 of the proof). Although, the simplest strategy accumulates probability mass in q_3 after *even* number of steps from $\{q_4\}$, by repeating the same strategy two times from q_4 (injecting large probability mass in q_3, moving to q_4, and injecting in q_3 again), we can accumulate probability mass in q_3 after *odd* number of steps from $\{q_4\}$, thus in synchronization with the probability mass accumulated in q_3 from $\{q_1, q_2\}$. However, by doing that, we also hit several other states and the remaining (small) probability mass is distributed over support $\{q_1, q_2, q_3, q_4, q_5, q_6\}$ when the next iteration starts. By a similar argument, we can again construct a strategy to implement the phases described above, and this can be done for all iterations and for $\varepsilon \to 0$. Indeed, the result of Claim 2 in the proof of Theorem 7 shows that by repeating strategies with shifting, we can eventually synchronize all the shifts.
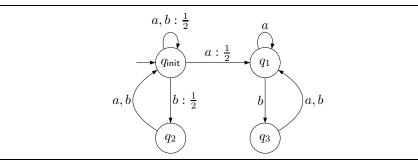
Figure 13: The state q_2 is vanishing for target set $T = \{q_2, q_3\}$ and strategies $(\alpha)_{i \in \mathbb{N}}$ where α_i repeats playing i times a, then playing b forever.

Vanishing states and Lemma 22. In the example of $\mathcal{M}_{12}$ (Figure 12), the target T is a singleton, which makes the result easier to prove than for an arbitrary set T. In particular, the second claim at the beginning of this section (that q_4 is limit-sure eventually synchronizing with target $\{q_3\}$) follows from the fact that q_4 is limit-sure weakly synchronizing to itself: it is easy to argue that if the probability is infinitely often arbitrarily close to 1 in $T = \{q_4\}$, then (starting with probability 1) from q_4 there must be a way to inject (almost all) the probability mass back to q_4 (via q_3). However, if T is not a singleton, the same argument is more difficult because when the probability mass is $1 - \varepsilon$ in T, it may still be that some state q in T holds only a tiny (less than ε) probability mass, which makes it more difficult to argue that we must be able to inject (almost all) the probability mass from q back to T (because if the tiny probability in q could not be injected at all in T, there would be no contradiction to the fact that probability $1 - \varepsilon$ is in T infinitely often).

Therefore, given an arbitrary target set T, we need to get rid of the states in T that do not contribute a significant (i.e., bounded away from 0) probability mass in the limit, that we call the vanishing states. We show that the vanishing states can be removed from T without changing the winning region for limit-sure winning. When the target set has no vanishing state, we can construct an almost-sure winning strategy similarly to the case of a singleton target set.

Given an MDP $\mathcal{M}$ with initial state $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{weakly}(sum_T)$ that is limit-sure winning for the weakly synchronizing objective in target set T, let $(\alpha_i)_{i \in \mathbb{N}}$ be a family of limit-sure winning strategies such that $\limsup_{n \to \infty} \mathcal{M}_n^{\alpha_i}(T) \geq 1 - \varepsilon_i$ where $\lim_{i \to \infty} \varepsilon_i = 0$. Hence, by definition of $\limsup$, for all $i \geq 0$ there exists a strictly increasing sequence $k_{i,0} < k_{i,1} < \cdots$ of positions such that $\mathcal{M}_{k_{i,j}}^{\alpha_i}(T) \geq 1 - 2\varepsilon_i$ for all $j \geq 0$. A state $q \in T$ is *vanishing* if $\liminf_{i \to \infty} \liminf_{j \to \infty} \mathcal{M}_{k_{i,j}}^{\alpha_i}(q) = 0$ for some family of limit-sure weakly synchronizing strategies $(\alpha_i)_{i \in \mathbb{N}}$. Intuitively, the contribution of a vanishing state q to the probability in T tends to 0 and therefore $\mathcal{M}$ is also limit-sure winning for the weakly synchronizing objective in target set $T \setminus \{q\}$.

Example. Consider the MDP in Figure 13 where all transitions are deterministic except from the initial state q_{init}. The state q_{init} has two successors on all actions:

$$\delta(q_{\mathsf{init}}, a)(q_{\mathsf{init}}) = \delta(q_{\mathsf{init}}, a)(q_1) = \frac{1}{2} \quad \text{and} \quad \delta(q_{\mathsf{init}}, b)(q_{\mathsf{init}}) = \delta(q_{\mathsf{init}}, b)(q_2) = \frac{1}{2}.$$

Let $T = \{q_2, q_3\}$ be the target set and for all $i \in \mathbb{N}$, let α_i be the strategy that repeats forever the following template in every state: playing i times a and then playing b. The family of strategies $(\alpha_i)_{i \in \mathbb{N}}$ is a witness to show that $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{weakly}(sum_T)$ where the state q_2 is a vanishing state. The contribution of q_2 in accumulating the probability mass in $\{q_2, q_3\}$ tends to 0 when $i \to \infty$. As a result, $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{weakly}(q_3)$ too.

4.3.3. Proof that limit-sure weakly and almost-sure weakly coincide

We present the formal proof of the main result (Theorem 7) along with the intermediate lemmas that we illustrated in Section 4.3.1 and Section 4.3.2.

Lemma 22. *If an MDP $\mathcal{M}$ is limit-sure weakly synchronizing in target set T, then there exists a set $T' \subseteq T$ such that $\mathcal{M}$ is limit-sure weakly synchronizing in T' without vanishing states.*

Proof. If there is no vanishing state for $(\alpha_i)_{i \in \mathbb{N}}$, then take $T' = T$ and the proof is complete. Otherwise, let $(\alpha_i)_{i \in \mathbb{N}}$ be a family of limit-sure winning strategies such that $\limsup_{n \to \infty} \mathcal{M}_n^{\alpha_i}(T) \geq 1 - \varepsilon_i$ where $\lim_{i \to \infty} \varepsilon_i = 0$ and let q be a vanishing state for $(\alpha_i)_{i \in \mathbb{N}}$. We show that $(\alpha_i)_{i \in \mathbb{N}}$ is limit-sure weakly synchronizing in $T \setminus \{q\}$. For every $i \geq 0$ let $k_{i,0} < k_{i,1} < \cdots$ be a strictly increasing sequence such that (a) $\mathcal{M}_{k_{i,j}}^{\alpha_i}(T) \geq 1 - 2\varepsilon_i$ for all $i, j \geq 0$, and (b) $\liminf_{i \to \infty} \liminf_{j \to \infty} \mathcal{M}_{k_{i,j}}^{\alpha_i}(q) = 0$.

It follows from (b) that for all $\varepsilon > 0$ and all $x > 0$ there exists $i > x$ such that for all $y > 0$ there exists $j > y$ such that $\mathcal{M}_{k_{i,j}}^{\alpha_i}(q) < \varepsilon$, and thus

$$\mathcal{M}_{k_{i,j}}^{\alpha_i}(T \setminus \{q\}) \geq 1 - 2\varepsilon_i - \varepsilon$$

by (a). Since this holds for infinitely many i's, we can choose i such that $\varepsilon_i < \varepsilon$ and we have

$$\limsup_{j \to \infty} \mathcal{M}_{k_{i,j}}^{\alpha_i}(T \setminus \{q\}) \geq 1 - 3\varepsilon$$

and thus

$$\limsup_{n \to \infty} \mathcal{M}_n^{\alpha_i}(T \setminus \{q\}) \geq 1 - 3\varepsilon$$

since the sequence $(k_{i,j})_{j \in \mathbb{N}}$ is strictly increasing. This shows that $(\alpha_i)_{i \in \mathbb{N}}$ is limit-sure weakly synchronizing in $T \setminus \{q\}$.

By repeating this argument as long as there is a vanishing state (thus at most $|T| - 1$ times), we can construct the desired set $T' \subseteq T$ without vanishing state. □

For a limit-sure weakly synchronizing MDP in target set T (without vanishing states), we show that from a probability distribution with support T, a

probability mass arbitrarily close to 1 can be injected synchronously back in T (in at least one step), that is $d_T \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_{\mathsf{Pre}(T)})$. The same holds from the initial state q_{init} of the MDP. This property is the key to construct an almost-sure weakly synchronizing strategy.

Lemma 23. *If an MDP $\mathcal{M}$ with initial state q_{init} is limit-sure weakly synchronizing in a target set T without vanishing states, then we have $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_{\mathsf{Pre}(T)})$ and $d_T \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_{\mathsf{Pre}(T)})$ where d_T is the uniform distribution over T.*

Proof. As $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{weakly}(sum_T)$ and $\langle\!\langle 1 \rangle\!\rangle_{limit}^{weakly}(sum_T) \subseteq \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_T)$ (Remark 1), we have $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_T)$ and thus it suffices to prove that $d_T \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_{\mathsf{Pre}(T)})$. This is because then from q_{init}, probability arbitrarily close to 1 can be injected in $\mathsf{Pre}(T)$ through a distribution with support in T (since by Corollary 15 only the support of the initial probability distribution is important for limit-sure eventually synchronizing).

Let $(\alpha_i)_{i \in \mathbb{N}}$ be a family of limit-sure winning strategies such that

$$\limsup_{n \to \infty} \mathcal{M}_n^{\alpha_i}(T) \geq 1 - \varepsilon_i \text{ where } \lim_{i \to \infty} \varepsilon_i = 0,$$

and such that there is no vanishing state. For every $i \geq 0$ let $k_{i,0} < k_{i,1} < \cdots$ be a strictly increasing sequence such that $\mathcal{M}_{k_{i,j}}^{\alpha_i}(T) \geq 1 - 2\varepsilon_i$ for all $i, j \geq 0$, and let

$$B = \min_{q \in T} \liminf_{i \to \infty} \liminf_{j \to \infty} \mathcal{M}_{k_{i,j}}^{\alpha_i}(q) = \liminf_{i \to \infty} \liminf_{j \to \infty} \min_{q \in T} \mathcal{M}_{k_{i,j}}^{\alpha_i}(q).$$

Note that $B > 0$ since there is no vanishing state. It follows that there exists $x > 0$ such that for all $i > x$ there exists $y_i > 0$ such that for all $j > y_i$ and all $q \in T$ we have $\mathcal{M}_{k_{i,j}}^{\alpha_i}(q) \geq \frac{B}{2}$.

Given $\nu > 0$, let $i > x$ such that $\varepsilon_i < \frac{\nu B}{4}$, and for $j > y_i$, consider the positions $n_1 = k_{i,j}$ and $n_2 = k_{i,j+1}$. We have $n_1 < n_2$ and $\mathcal{M}_{n_1}^{\alpha_i}(T) \geq 1 - 2\varepsilon_i$ and $\mathcal{M}_{n_2}^{\alpha_i}(T) \geq 1 - 2\varepsilon_i$, and $\mathcal{M}_{n_1}^{\alpha_i}(q) \geq \frac{B}{2}$ for all $q \in T$. Consider the strategy β that plays like α_i plays from position n_1 and thus transforms the distribution $\mathcal{M}_{n_1}^{\alpha_i}$ into $\mathcal{M}_{n_2}^{\alpha_i}$. For all states $q \in T$, from the Dirac distribution on q under strategy β, the probability to reach $Q \setminus T$ in $n_2 - n_1$ steps is thus at most $\frac{\mathcal{M}_{n_2}^{\alpha_i}(Q \setminus T)}{\mathcal{M}_{n_1}^{\alpha_i}(q)} \leq \frac{2\varepsilon_i}{B/2} < \nu$.

Therefore, from an arbitrary probability distribution with support T we have $\mathcal{M}_{n_2-n_1}^{\beta}(T) > 1 - \nu$, showing that d_T is limit-sure eventually synchronizing in T and thus in $\mathsf{Pre}(T)$ since $n_2 - n_1 > 0$ (it is easy to show that if the mass of probability in T is at least $1 - \nu$, then the mass of probability in $\mathsf{Pre}(T)$ one step before is at least $1 - \frac{\nu}{\eta}$ where η is the smallest positive probability in $\mathcal{M}$). $\square$

To show that limit-sure and almost-sure winning coincide for weakly synchronizing objectives, from a family of limit-sure winning strategies we construct an almost-sure winning strategy that uses the eventually synchronizing strategies

of Lemma 23. The construction consists in using successively strategies that ensure probability mass $1 - \varepsilon_i$ in the target T, for a decreasing sequence $\varepsilon_i \to 0$. Such strategies exist by Lemma 23, both from the initial state and from the set T. However, the mass of probability that can be guaranteed to be synchronized in T by the successive strategies is always smaller than 1, and therefore we need to argue that the remaining mass of probability (of total size ε_i) scattered in the state space can also get synchronized in T, despite the variable shifts with the main mass of probability.

Two main key arguments are needed to establish the correctness of the construction: (1) eventually synchronizing implies that a finite number of steps is sufficient to obtain a probability mass of $1 - \varepsilon_i$ in T, and thus the construction of the strategy is well defined, and (2) by the finiteness of the period r (such that $R = \mathsf{Pre}^r(R)$ where $R = \mathsf{Pre}^k(T)$ for some k) from every state, we can accumulate shifts such that their sum is $0 \mod r$, and thus the probability mass from every state contributes (synchronously) to the probability accumulated in the target.

Theorem 7. $\langle\!\langle 1 \rangle\!\rangle_{limit}^{weakly}(sum_T) = \langle\!\langle 1 \rangle\!\rangle_{almost}^{weakly}(sum_T)$ *for all MDPs and target sets T.*

Proof. Since $\langle\!\langle 1 \rangle\!\rangle_{almost}^{weakly}(sum_T) \subseteq \langle\!\langle 1 \rangle\!\rangle_{limit}^{weakly}(sum_T)$ holds by the definition, it is sufficient to prove that $\langle\!\langle 1 \rangle\!\rangle_{limit}^{weakly}(sum_T) \subseteq \langle\!\langle 1 \rangle\!\rangle_{almost}^{weakly}(sum_T)$ and by Lemma 22 it is sufficient to prove that if $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{weakly}(sum_T)$ is limit-sure weakly synchronizing in T without vanishing state, then q_{init} is almost-sure weakly synchronizing in T. If T has vanishing states, then consider $T' \subseteq T$ as in Lemma 22 and it will follows that q_{init} is almost-sure weakly synchronizing in T' and thus also in T. We proceed with the proof that $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{weakly}(sum_T)$ implies $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{weakly}(sum_T)$.

For $i = 1, 2, \ldots$ consider the sequence of predecessors $\mathsf{Pre}^i(T)$, which is ultimately periodic: let $1 \le k, r \le 2^{|Q|}$ such that $\mathsf{Pre}^k(T) = \mathsf{Pre}^{k+r}(T)$, and let $R = \mathsf{Pre}^k(T)$. Thus $R = \mathsf{Pre}^{k+r}(T) = \mathsf{Pre}^r(R)$.

Claim 1. We have $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{event}(sum_R)$ and $d_T \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{event}(sum_R)$.

Proof of Claim 1. By Lemma 23, since there is no vanishing state in T we have $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_{\mathsf{Pre}(T)})$ and $d_T \in \langle\!\langle 1 \rangle\!\rangle_{limit}^{event}(sum_{\mathsf{Pre}(T)})$. The characterization of the winning region for limit-sure eventually synchronizing given by Lemma 11, and the fact that almost-sure and limit-sure coincide for eventually synchronizing in the set R (Corollary 13) give the following:

- either (1) $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_{\mathsf{Pre}(T)})$ or (2) $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{event}(sum_R)$, and

- either (a) $d_T \in \langle\!\langle 1 \rangle\!\rangle_{sure}^{event}(sum_{\mathsf{Pre}(T)})$ or (b) $d_T \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{event}(sum_R)$.

We show that (a) implies (b), hence that (b) holds. Then we show that (1) implies (2), hence that (2) holds. We conclude that both (2) and (b), which establishes Claim 1.
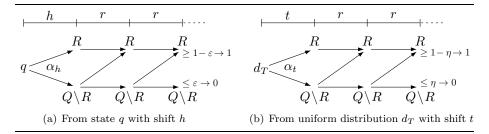
Figure 14: Sketch of the outcome of almost-sure eventually synchronizing strategies (with shifts).

To show that (a) implies (b): by the characterization of sure eventually synchronizing (Lemma 7), if (a) holds, then $T \subseteq \mathsf{Pre}^i(T)$ for some $i \geq 1$, and thus $T \subseteq \mathsf{Pre}^{n \cdot i}(T)$ for all $n \geq 0$ by monotonicity of $\mathsf{Pre}^i(\cdot)$. This entails for $n \cdot i \geq k$ that $T \subseteq \mathsf{Pre}^m(R)$ where $m = (n \cdot i - k) \mod r$ and thus d_T is sure (and almost-sure) winning for the eventually synchronizing objective in target R (by Lemma 7), hence (b) holds.

To show that (1) implies (2): if (1) holds, then we can play a sure-winning strategy from q_{init} to ensure in finitely many steps probability 1 in $\mathsf{Pre}(T)$ and in the next step probability 1 in T, and by (b) play an almost-sure winning strategy for eventually synchronizing in R. Hence, $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{event}(sum_R)$, i.e. (2) holds. The proof of Claim 1 is done.

We now show that there exists an almost-sure winning strategy for the weakly synchronizing objective in target T. Recall that $\mathsf{Pre}^r(R) = R$ and thus once some probability mass p is in R, it is possible to ensure that the probability mass in R after r steps is at least p, and thus that (with period r) the probability in R does not decrease. By the result of Lemma 12, almost-sure winning for eventually synchronizing in R implies that there exists a strategy α such that the probability in R tends to 1 at periodic positions: for some $0 \leq h < r$ the strategy α is *almost-sure eventually synchronizing in R with shift h*, that is $\forall \varepsilon > 0 \cdot \exists N \cdot \forall n \geq N : n \equiv h \mod r \implies \mathcal{M}_n^\alpha(R) \geq 1 - \varepsilon$. We also say that the initial distribution $d_0 = \mathcal{M}_0^\alpha$ is almost-sure eventually synchronizing in R with shift h. Almost-sure eventually synchronizing strategies with shift are illustrated in Figure 14.

Claim 2.

$(\star)$ If $\mathcal{M}_0^\alpha$ is almost-sure eventually synchronizing in R with some shift h, then $\mathcal{M}_i^\alpha$ is almost-sure eventually synchronizing in R with shift $h - i \mod r$.

$(\star\star)$ Let t such that d_T is almost-sure eventually synchronizing in R with shift t. If a distribution is almost-sure eventually synchronizing in R with some shift h, then it is also almost-sure eventually synchronizing in R with shift $h + k + t \mod r$ (where we chose k such that $R = \mathsf{Pre}^k(T)$).
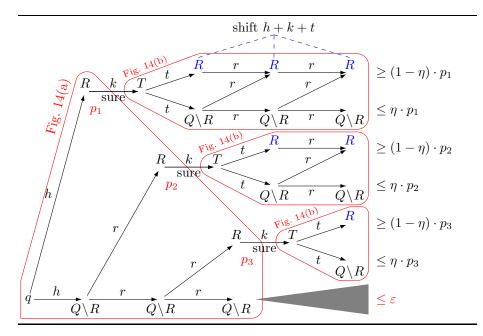
Figure 15: Proof of Claim 2($\star\star$) for Theorem 7.

Proof of Claim 2. The result ($\star$) immediately follows from the definition of shift, and we prove ($\star\star$) as follows. We show that almost-sure eventually synchronizing in R with shift h implies almost-sure eventually synchronizing in R with shift $h + k + t \mod r$. The argument is illustrated in Figure 15. Intuitively, the probability mass that is in R with shift h can be injected in T in k steps, and then from T we can play an almost-sure eventually synchronizing strategy in target R with shift t, thus a total shift of $h + k + t \mod r$. Precisely, an almost-sure winning strategy α is constructed as follows (Figure 15):

- given a finite prefix of play ρ, if there is no state $q \in R$ that occurs in ρ at a position $n \equiv h \mod r$, then play in ρ according to the almost-sure winning strategy α_h for eventually synchronizing in R with shift h;

- otherwise,

   - if there is no $q \in T$ that occurs in ρ at a position $n \equiv h + k \mod r$, then we play according to a sure winning strategy α_{sure} for eventually synchronizing in T,

   - and otherwise we play according to an almost-sure winning strategy α_t from T for eventually synchronizing in R with shift t.

To show that α is almost-sure eventually synchronizing in R with shift $h+k+t$, note that α_h ensures with probability 1 that R is reached at positions $n \equiv h$ mod r (see Figure 15). Consider positions $h, h+r, h+2r, \ldots$ and the probability
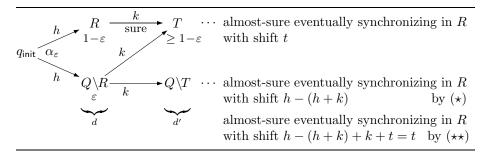
$$
\begin{array}{l}
q_{\text{init}} \xrightarrow{\ h\ } R \xrightarrow[\text{sure}]{\ k\ } T \quad \cdots \quad \text{almost-sure eventually synchronizing in } R \\
\qquad \quad 1-\varepsilon \qquad\qquad \geq 1-\varepsilon \qquad \text{with shift } t
\end{array}
$$

$$
q_{\text{init}} \xrightarrow{\alpha_\varepsilon}
$$

almost-sure eventually synchronizing in R
with shift $h-(h+k)$ by $(\star)$

almost-sure eventually synchronizing in R
with shift $h-(h+k)+k+t=t$ by $(\star\star)$

Figure 16: Construction of an almost-sure weakly synchronizing strategy.

mass p_i in R at position $h+ir$. Then for all $\varepsilon > 0$, by considering sufficiently long sequence of positions, we have $\sum_i p_i \geq 1-\varepsilon$ (Figure 14(a)). Since α_{sure} is sure eventually synchronizing in T, we also have probability mass at least p_i in T at position $h+k+ir$. From the states in T the strategy α_t ensures with probability 1 that R is reached at positions $h+k+t \mod r$, thus for all $\eta > 0$, by considering sufficiently long sequence of positions (and Figure 14(b)), we have probability mass at least $\sum_i p_i \cdot (1-\eta) \geq (1-\varepsilon) \cdot (1-\eta)$ at in R at some position $h+k+t+ir$, thus with shift $h+k+t$ (see also Figure 15). This concludes the proof of Claim 2.

Construction of an almost-sure winning strategy. We construct strategies α_ε for $\varepsilon > 0$ that ensure, from a distribution that is almost-sure eventually synchronizing in R (with some shift h), that after finitely many steps, a distribution d' is reached such that $d'(T) \geq 1-\varepsilon$ and d' is almost-sure eventually synchronizing in R (with some shift h'). Since q_{init} is almost-sure eventually synchronizing in R (with some shift h), it follows that the strategy α_{as} that plays successively the strategies (each for finitely many steps) $\alpha_{\frac{1}{2}}, \alpha_{\frac{1}{4}}, \alpha_{\frac{1}{8}}, \ldots$ is almost-sure winning from q_{init} for the weakly synchronizing objective in target T.

We define the strategies α_ε as follows (the construction is illustrated in Figure 16). Given an initial distribution that is almost-sure eventually synchronizing in R with a shift h and given $\varepsilon > 0$, let α_ε be the strategy that plays according to the almost-sure winning strategy α_h for eventually synchronizing in R with shift h for a number of steps $n \equiv h \mod r$ until a distribution d is reached such that $d(R) \geq 1-\varepsilon$, and then from d it plays according to a sure winning strategy α_{sure} for eventually synchronizing in T from the states in R (for k steps), and keeps playing according to α_h from the states in $Q \setminus R$ (for k steps). The distribution d' reached from d after k steps is such that $d'(T) \geq 1-\varepsilon$ and we claim that it is almost-sure eventually synchronizing in R with shift t. This holds by definition of α_t from the states in $\mathsf{Supp}(d') \cap T$, and by $(\star)$ the states in $\mathsf{Supp}(d') \setminus T$ are almost-sure eventually synchronizing in R with shift $h-(h+k) \mod r$, and by $(\star\star)$ with shift $h-(h+k)+k+t=t$.

It follows that the strategy α_{as} is well-defined and ensures, for all $\varepsilon > 0$, that the probability mass in T is infinitely often at least $1-\varepsilon$, thus is almost-sure weakly synchronizing in T. This concludes the proof of Theorem 7. $\qquad\square$
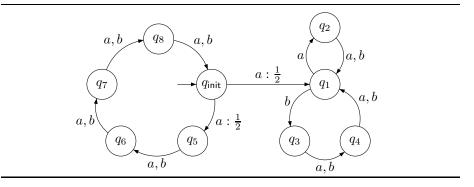
Figure 17: An example to show $q_{\text{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{strong}(max_Q)$ reduces to synchronized reachability of a state in a simple deterministic cycle.

5. Strongly Synchronizing

The strongly synchronizing objective is reminiscent of a coBüchi objective in the distribution-based semantics: with function sum_T it requires that in the sequence of distributions of an MDP $\mathcal{M}$ under strategy α we have $\liminf_{n\to\infty} \mathcal{M}_n^\alpha(T) = 1$ (and that $\mathcal{M}_n^\alpha(T) = 1$ from some point on in the case of sure winning).

We show that the membership problem for strongly synchronizing objectives can be solved in polynomial time, for all winning modes, and both with function max_T (Section 5.1) and function sum_T (Section 5.2). We show that linear-size memory is necessary in general for max_T, and memoryless strategies are sufficient for sum_T. It follows from our results that the limit-sure and almost-sure winning modes coincide for strongly synchronizing.

5.1. Strongly synchronizing with function max

First, note that for strongly synchronizing the membership problem with function max_T reduces to the membership problem with function max_Q where Q is the entire state space, by a construction similar to the proof of Lemma 4: states in $Q \setminus T$ are duplicated, ensuring that only states in T are used to accumulate probability.

The strongly synchronizing objective with function max_Q requires that from some point on, almost all the probability mass is at every step in a single state. Intuitively, the sequence of states that contain almost all the probability corresponds to a sequence of deterministic transitions in the MDP, and thus eventually to a cycle of deterministic transitions.

Consider the MDP in Figure 17 with initial state q_{init}: all transitions are deterministic except from q_{init} where on both actions a and b, the successors are q_1 and q_5 with probability $\frac{1}{2}$. The strategic choice is only relevant in q_1 where $\delta(q_1, a)(q_2) = 1$ and $\delta(q_1, b)(q_3) = 1$. We present a strategy such that the sequence of states that contain almost all the probability is the cycle $q_1 q_2 q_1 q_3 q_4 q_1$ of deterministic transitions.

48

The state q_{init} is almost-sure strongly synchronizing (according to function *max*) with the strategy α defined as follows, for all paths ρ such that $\mathsf{Last}(\rho) = q_1$:

- if the number of occurrences of q_1 in ρ is odd (i.e., the length of ρ is 1 modulo 5), then play action a;

- if the number of occurrences of q_1 in ρ is even (i.e., the length of ρ is 3 modulo 5), then play action b.

The strategy α ensures the probability mass injected from q_{init} in q_1 after every other 5 steps loops in the cycle $q_1 q_2 q_1 q_3 q_4 q_1$ (with length 5). Hence, the probability mass from q_{init} is always injected in q_1 synchronously (i.e., when the probability mass in the cycle is also in q_1).

It follows that after $5i$ steps, the probability mass in q_{init} is $\frac{1}{2^i}$ and the probability mass in q_1 is $1 - \frac{1}{2^i}$. Considering $i \to \infty$, we then get $\liminf_{n \to \infty} \|\mathcal{M}_n^\alpha\| = 1$ and $q_{\mathsf{init}} \in \langle\!\langle 1 \rangle\!\rangle_{almost}^{strong}(max)$. Note that only the states in the cycle $q_1 q_2 q_1 q_3 q_4 q_1$ (of deterministic transitions) are used to accumulate the probability mass tending to 1.

Cycles consisting of deterministic transitions are keys to decide strongly synchronizing. A *deterministic cycle* of length $\ell \geq 1$ in an MDP $\mathcal{M}$ is a finite sequence $\hat{q}_0 \hat{q}_1 \ldots \hat{q}_\ell$ of states such that $\hat{q}_0 = \hat{q}_\ell$ and for all $0 \leq i < \ell$, there exists an action a_i such that $\delta(q_i, a_i)(q_{i+1}) = 1$. The cycle is *simple* if $\hat{q}_i \neq \hat{q}_j$ for all $1 \leq i < j \leq \ell$.

We show that sure (resp., almost-sure and limit-sure) strongly synchronizing is equivalent to sure (resp., almost-sure and limit-sure) reachability to a state in a *simple* deterministic cycle, with the requirement that the state can be reached in a synchronized way (i.e., by finite paths whose lengths are congruent modulo the length ℓ of the cycle).

In the MDP of Figure 17, we can construct an almost-sure strongly synchronizing strategy β that accumulates the probability mass only in the simple cycle $q_1 q_3 q_4 q_1$. The strategy β is defined as follows, for all paths ρ such that $\mathsf{Last}(\rho) = q_1$:

- if the length of ρ is 0 modulo 3, then play action b;

- if the length of ρ is 1 or 2 modulo 3, then play action a.

Note that if the length of ρ is a multiple of 3 and the action b is played, then on the next visit to q_1 the length of the path is also a multiple of 3, and the action b is played again. Hence, once a probability mass follows the cycle $q_1 q_3 q_4 q_1$, it will follow this cycle forever. Whenever probability mass is injected in q_1 (from q_{init}) on a path ρ of length 1 or 2 modulo 3, the action a is played to visit the other cycle $q_1 q_2 q_1$ until getting back to q_1 with a path whose length is a multiple of 3. The probability mass is then injected (synchronously) into the cycle $q_1 q_3 q_4 q_1$ where eventually the probability mass tends to 1, thus the strategy β is almost-sure strongly synchronizing and it ensures with probability 1 that q_1 is reached with by paths whose length is a multiple of 3.

We show in Lemma 24 that simple deterministic cycles are always sufficient for strongly synchronizing in MDPs, and that strongly synchronizing reduces to a synchronized reachability problem of reaching a state q_1 of a simple deterministic cycle by paths of length that is a multiple of the length ℓ of the cycle. To check synchronized reachability,

we keep track of a modulo-ℓ counter along the path. Define the MDP $\mathcal{M} \times [\ell] = \langle Q', \mathsf{A}, \delta' \rangle$ where $Q' = Q \times \{0, 1, \ldots, \ell - 1\}$ and $\delta'(\langle q, i \rangle, a)(\langle q', i - 1 \rangle) = \delta(q, a)(q')$ (where $i - 1$ is $\ell - 1$ for $i = 0$) for all states $q, q' \in Q$, actions $a \in \mathsf{A}$, and $0 \le i \le \ell - 1$. Note that given a finite path $\rho = q_0 a_0 q_1 a_1 \ldots q_n$ in $\mathcal{M}$, there is a corresponding path $\rho' = \langle q_0, k_0 \rangle a_0 \langle q_1, k_1 \rangle a_1 \ldots \langle q_n, k_n \rangle$ in $\mathcal{M} \times [\ell]$ where $k_i = -i \bmod \ell$. Since the sequence $k_0 k_1 \ldots$ is uniquely defined, there is a clear bijection between the paths in $\mathcal{M}$ (starting from q_0) and the paths in $\mathcal{M} \times [\ell]$ (starting from $\langle q_0, 0 \rangle$) that we often omit to apply and mention in the sequel.

Lemma 24. *Let η be the smallest positive probability in the transitions of $\mathcal{M}$, and let $\frac{1}{1+\eta} < p \le 1$. There exists a strategy α such that $\liminf_{n \to \infty} \|\mathcal{M}_n^\alpha\| \ge p$ from an initial state q_{init} if and only if there exist a simple deterministic cycle $\hat{q}_0 \hat{q}_1 \ldots \hat{q}_\ell$ in $\mathcal{M}$ and a strategy β in $\mathcal{M} \times [\ell]$ such that $\mathrm{Pr}^\beta(\Diamond\{\langle \hat{q}_0, 0 \rangle\}) \ge p$ from $\langle q_{\mathsf{init}}, 0 \rangle$.*

Proof. For the first direction of the lemma, assume that there exists a strategy α such that $\liminf_{n \to \infty} \|\mathcal{M}_n^\alpha\| \ge p$ from q_{init}. Thus for all $\varepsilon > 0$ (in particular, we consider $\varepsilon < p - \frac{1}{1+\eta}$), there exists $k \in \mathbb{N}$ such that for all $n \ge k$ we have $\|\mathcal{M}_n^\alpha\| \ge p - \varepsilon$, and let $\hat{p}_n$ be a state such that $\mathcal{M}_n^\alpha(\hat{p}_n) \ge p - \varepsilon$. We claim that for all $n \ge k$, there exists an action $a \in \mathsf{A}$ such that $\mathsf{post}(\hat{p}_n, a) = \{\hat{p}_{n+1}\}$ i.e., there is a deterministic transition from $\hat{p}_n$ to $\hat{p}_{n+1}$. Assume towards contradiction that for some $n \ge k$, for all $a \in \mathsf{A}$ there exists $q_a \ne \hat{p}_{n+1}$ such that $q_a \in \mathsf{post}(\hat{p}_n, a)$. Then no matter the actions played by α at step n, we have $\mathcal{M}_{n+1}^\alpha(\{q_a \mid a \in \mathsf{A}\}) \ge \mathcal{M}_n^\alpha(\hat{p}_n) \cdot \eta \ge (p - \varepsilon) \cdot \eta$, and since $\hat{p}_{n+1} \ne q_a$ for all $a \in \mathsf{A}$, it follows that

$$\mathcal{M}_{n+1}^\alpha(\hat{p}_{n+1}) \le 1 - \mathcal{M}_{n+1}^\alpha(\{q_a \mid a \in \mathsf{A}\}) \le 1 - (p - \varepsilon) \cdot \eta \le 1 - \frac{\eta}{1 + \eta} < p - \varepsilon,$$

in contradiction with the fact that $\hat{p}_{n+1}$ is a state such that $\mathcal{M}_{n+1}^\alpha(\hat{p}_{n+1}) \ge p - \varepsilon$. This concludes the argument showing that for all $n \ge k$, there exists an action $a \in \mathsf{A}$ such that $\mathsf{post}(\hat{p}_n, a) = \{\hat{p}_{n+1}\}$.

Now in the sequence $\hat{p}_k \hat{p}_{k+1} \ldots$, we can extract a simple (and deterministic) cycle $\mathcal{C} = \hat{p}_i \hat{p}_{i+1} \ldots \hat{p}_{i+\ell}$ since the state space is finite. Let $\hat{q}_0 = \hat{p}_{i+j}$ where $j \le \ell$ is such that $i + j \bmod \ell = 0$. Then $\hat{q}_0$ is on a simple deterministic cycle, and is reachable after a multiple of ℓ steps with probability at least $p - \varepsilon$ by a strategy β in $\mathcal{M} \times [\ell]$ that copies the strategy α. Hence, we have $\mathrm{Pr}^\beta(\Diamond\{\langle \hat{q}_0, 0 \rangle\}) \ge p - \varepsilon$ from $\langle q_{\mathsf{init}}, 0 \rangle$. Since for every $\varepsilon > 0$, we can find such a cycle and state $\hat{q}_0$, and since the state space is finite (as well as the number of simple cycles), it follows that there is a cycle $\mathcal{C}$ and state $\hat{q}_0$ in $\mathcal{C}$ such that for all $\varepsilon > 0$ we have $\mathrm{Pr}^\beta(\Diamond\{\langle \hat{q}_0, 0 \rangle\}) \ge p - \varepsilon$, and thus $\mathrm{Pr}^\beta(\Diamond\{\langle \hat{q}_0, 0 \rangle\}) \ge p$.

For the second direction of the lemma, assume that there exist a simple deterministic cycle $\hat{q}_0 \hat{q}_1 \ldots \hat{q}_\ell$ and a strategy β in $\mathcal{M} \times [\ell]$ that ensures the

target set $\{\langle \hat{q}_0, 0 \rangle\}$ is reached with probability at least p from $\langle q_{\mathsf{init}}, 0 \rangle$. Since randomization is not necessary for reachability objectives in MDPs, we can assume that β is a pure strategy. We show that there exists a strategy α such that $\liminf_{n \to \infty} \|\mathcal{M}_n^\alpha\| \geq p$ from q_{init}. From β, we construct a pure strategy α in $\mathcal{M}$. Given $\rho = q_0 a_0 q_1 a_1 \dots q_n$, we define $\alpha(\rho)$ as follows: if $q_n = \hat{q}_{n \ mod \ \ell}$, then there exists an action a such that $\mathsf{post}(q_n, a) = \{\hat{q}_{n+1 \ mod \ \ell}\}$ and we define $\alpha(\rho) = a$, otherwise let $\alpha(\rho) = \beta(\rho)$. Thus α mimics β until a state $\hat{q}_k$ of the cycle is reached at step n such that $k = n \ mod \ \ell$, and then α switches to always playing actions that keeps $\mathcal{M}$ in the simple deterministic cycle $\hat{q}_0 \hat{q}_1 \dots \hat{q}_\ell$. Note that α is a pure strategy.

We claim that given $\varepsilon > 0$ there exists k such that for all $n \geq k$, we have $\|\mathcal{M}_n^\alpha\| \geq p - \varepsilon$, which entails that $\liminf_{n \to \infty} \|\mathcal{M}_n^\alpha\| \geq p$ from q_{init} and concludes the proof. To show the claim, since $\mathrm{Pr}^\beta(\Diamond \{\langle \hat{q}_0, 0 \rangle\}) \geq p$, consider k such that $\mathrm{Pr}^\beta(\Diamond^{\leq k} \{\langle \hat{q}_0, 0 \rangle\}) \geq p - \varepsilon$, and for $i = 1, 2, \dots, \ell$, let $R_i = \{\langle \hat{q}_i, \ell - i \rangle\}$. Note that $R_\ell = \{\langle \hat{q}_0, 0 \rangle\}$. Then trivially $\mathrm{Pr}^\beta(\Diamond^{\leq k} \bigcup_{i=1}^\ell R_i) \geq p - \varepsilon$ and since α agrees with β on all finite paths that do not (yet) visit $\bigcup_{i=1}^\ell R_i$, given a path ρ that visits $\bigcup_{i=1}^\ell R_i$ (for the first time), only actions that keep $\mathcal{M}$ in the simple cycle $\hat{q}_0 \hat{q}_1 \dots \hat{q}_\ell$ are played by α and thus all continuations of ρ in the outcome of α will visit $\hat{q}_0$ after a multiple of ℓ steps, say $j \cdot \ell$ steps (in total). Since next, α will always play actions that keeps $\mathcal{M}$ looping through the cycle $\hat{q}_0 \hat{q}_1 \dots \hat{q}_\ell$, we have $\mathcal{M}_{j \cdot \ell + i}^\alpha(\hat{q}_i) \geq p - \varepsilon$ for all $0 \leq i < \ell$, and thus $\|\mathcal{M}_n^\alpha\| \geq p - \varepsilon$ for all $n \geq j \cdot \ell$. $\square$

It follows directly from Lemma 24 with $p = 1$ that almost-sure strongly synchronizing is equivalent to almost-sure reachability to a deterministic cycle in $\mathcal{M} \times [\ell]$. The same equivalence holds for the sure and limit-sure winning modes.

Lemma 25. *A state q_{init} is sure (resp., almost-sure or limit-sure) winning for the strongly synchronizing objective (according to max_Q) in $\mathcal{M}$ if and only if there exists a simple deterministic cycle $\hat{q}_0 \hat{q}_1 \dots \hat{q}_\ell$ such that $\langle q_{\mathsf{init}}, 0 \rangle$ is sure (resp., almost-sure or limit-sure) winning for the reachability objective $\Diamond \{\langle \hat{q}_0, 0 \rangle\}$ in $\mathcal{M} \times [\ell]$.*

Proof. We consider the three winning modes:

(1) sure winning mode. The proof is similar to the proof of Lemma 24. For the first direction, given a strategy α and k such that for all $n \geq k$ we have $\|\mathcal{M}_n^\alpha\| = 1$ from the initial state q_{init}, we can construct a sequence $\hat{p}_k \hat{p}_{k+1} \dots$ of states where there is deterministic transition from $\hat{p}_n$ to $\hat{p}_{n+1}$ for all $n \geq k$ (let $\hat{p}_n$ be the state such that $\mathcal{M}_n^\alpha(\hat{p}_n) = 1$). This sequence contains a simple deterministic cycle and a state $\hat{q}_0$ in this cycle occurs in the sequence at a position $\hat{p}_{j \cdot \ell}$ that is a multiple of the length ℓ of the cycle. Hence, the strategy α played in $\mathcal{M} \times [\ell]$ ensures to reach $\langle \hat{q}_0, 0 \rangle$ surely from $\langle q_{\mathsf{init}}, 0 \rangle$.

For the second direction, if a strategy β ensures to reach a state $\langle \hat{q}_0, 0 \rangle$ in $\mathcal{M} \times [\ell]$ where $\hat{q}_0$ belongs to a simple deterministic cycle of length ℓ, then a strategy α that mimics β until $\langle \hat{q}_0, 0 \rangle$ is reached, and then switches to playing actions to follow the simple cycle, ensures sure strongly synchronizing with function max_Q in $\mathcal{M}$. Note that α is a pure strategy.

(2) almost-sure winning mode. This case follows from Lemma 24 with $p = 1$.

(3) limit-sure winning mode: For the first direction, if q_{init} is limit-sure winning for the strongly synchronizing objective, then for all $\varepsilon > 0$, there exists a strategy α such that $\liminf_{n \to \infty} \|\mathcal{M}_n^{\alpha_i}\| \geq 1 - \varepsilon$. By Lemma 24, for a decreasing sequence $\varepsilon_i \to 0$ such that $\varepsilon_i < 1 - \frac{1}{1+\eta}$ there exist a simple deterministic cycle $\mathcal{C}_i$ of length ℓ_i, a state $\hat{q}_0^i$ in $\mathcal{C}_i$, and a strategy β_i in $\mathcal{M} \times [\ell_i]$ such that $\mathrm{Pr}^{\beta_i}(\diamond\{\langle \hat{q}_0, 0 \rangle\}) \geq 1 - \varepsilon_i$ from $\langle q_{\mathsf{init}}, 0 \rangle$. Since there is a finite number of simple deterministic cycles in $\mathcal{M}$, some simple cycle $\mathcal{C} = \hat{q}_0 \hat{q}_1 \ldots \hat{q}_\ell$ and state $\hat{q}_0$ occurs infinitely often in the sequence of $(\mathcal{C}_i, \hat{q}_0^i)$, and thus $\langle \hat{q}_{\mathsf{init}}, 0 \rangle$ is limit-sure winning for the reachability objective $\diamond\{\langle \hat{q}_0, 0 \rangle\})$ in $\mathcal{M} \times [\ell]$.

For the second direction, since limit-sure winning implies almost-sure winning for reachability objectives in MDPs, it follows from case **(2)** that q_{init} is almost-sure (and thus also limit-sure) winning for the strongly synchronizing objective in $\mathcal{M}$. $\square$

Since the winning regions of almost-sure and limit-sure winning coincide for reachability objectives in MDPs [24], the next corollary follows from Lemma 25.

Corollary 26. $\langle\!\langle 1 \rangle\!\rangle_{limit}^{strongly}(max_T) = \langle\!\langle 1 \rangle\!\rangle_{almost}^{strongly}(max_T)$ *for all target sets* T.

If there exists a cycle $\mathcal{C}$ satisfying the condition in Lemma 25, then all cycles reachable from $\mathcal{C}$ in the graph G of deterministic transitions also satisfies the condition. Hence, it is sufficient to check the condition for an arbitrary simple cycle in each strongly connected component (SCC) of G. As shown in the next theorem, it follows that strongly synchronizing can be decided in polynomial time and the length of the cycle gives a linear bound on the memory needed to win.

Theorem 8. *For the three winning modes of strongly synchronizing according to* max_T:

1. *(Complexity). The membership problem is PTIME-complete.*
2. *(Memory). Linear memory is necessary and sufficient for both pure and randomized strategies, and pure strategies are sufficient.*

Proof. First, we prove the PTIME upper bound. Given an MDP $\mathcal{M} = \langle Q, \mathsf{A}, \delta \rangle$ and a state q_{init}, we say that a simple deterministic cycle $\mathcal{C} = \hat{q}_0 \hat{q}_1 \ldots \hat{q}_\ell$ is sure (resp., almost-sure, and limit-sure) winning from q_{init} if $\langle q_{\mathsf{init}}, 0 \rangle$ is sure (resp., almost-sure, and limit-sure) winning for the reachability objective $\diamond\{\langle \hat{q}_0, 0 \rangle\}$ in $\mathcal{M} \times [\ell]$.

We claim that if $\mathcal{C}$ is sure (resp., almost-sure, and limit-sure) winning from q_{init}, then so are all simple cycles $\mathcal{C}'$ reachable from $\mathcal{C}$ in the graph of deterministic transitions induced by $\mathcal{M}$. Given a strategy to reach a state $\hat{q}_0$ of $\mathcal{C}$ surely (resp., with probability p), we can use the path of deterministic transitions from $\mathcal{C}$ to $\mathcal{C}'$ to obtain a strategy to reach a state $\hat{q}_0'$ of $\mathcal{C}'$ surely (resp., with probability p): since $\hat{q}_0$ is reached after a multiple of ℓ steps ($\{\langle \hat{q}_0, 0 \rangle\}$ is
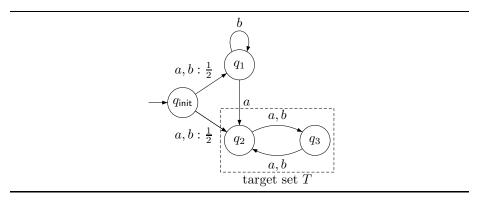
Figure 18: An MDP where all strategies to win sure strongly synchronizing with function $max_{\{q_2,q_3\}}$ require memory.

reached in $\mathcal{M} \times [\ell]$), we can let the probability mass loop through the cycle $\mathcal{C}$, and transfer it to $\mathcal{C}'$ after a number of steps that is also a multiple of ℓ', and then let it loop in $\mathcal{C}'$, ensuring that $\langle \hat{q}'_0, 0 \rangle$ is reached surely (resp., with probability p) in $\mathcal{M} \times [\ell']$. This establishes the claim for the three winning modes.

Using this claim and Lemma 25, it suffices to decide sure (resp., almost-sure, and limit-sure) winning for one simple cycle in each bottom SCC (reachable from q_{init}) of the graph of deterministic transitions. Since SCC decomposition for graphs, as well as sure, almost-sure, and limit-sure reachability for MDPs can be computed in polynomial time, and the number of bottom SCCs is at most the size $|Q|$ of the graph, the PTIME upper bound for the membership problem follows.

For PTIME-hardness, the proof is by a reduction from the monotone Boolean circuit value problem, which is PTIME-complete [37]. This problem is to compute the output value of a given Boolean circuit consisting of AND-gates, OR-gates, and fixed Boolean input values. From a circuit, we construct an MDP $\mathcal{M}$ with actions L and R, where the states correspond to the gates and input values of the circuit, and with three new absorbing states q_1, q_2, and sync. The successors of an AND-gate $n_1 \wedge n_2$ are n_1 and n_2 with probability $\frac{1}{2}$ on all actions, the successors of an OR-gate $n_1 \vee n_2$ are n_1 on action L, and n_2 on action R. On all actions, a node defining input value 1 has unique successor sync, and a node defining input value 0 has successors q_1 and q_2 with probability $\frac{1}{2}$. Let q_{init} be the state corresponding to the output node. Then $\mathcal{M}$ is sure (resp., almost-sure, limit-sure) strongly synchronizing (in sync) from q_{init} if and only if the value of the circuit is 1, which establishes PTIME-hardness of strongly synchronizing in the three winning modes.

Finally, the result on memory requirement is established as follows. Since memoryless strategies are sufficient for reachability objectives in MDPs, it follows from the proof of Lemma 24 and Lemma 25 that the (memoryless) winning strategies in $\mathcal{M} \times [\ell]$ can be transferred to winning strategies with memory $\{0, 1, \ldots, \ell - 1\}$ in $\mathcal{M}$. Since $\ell \leq |Q|$, linear-size memory is sufficient to win
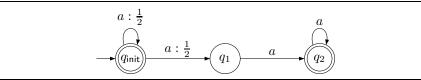
Figure 19: An MDP such that q_{init} is sure-winning for coBüchi objective in $T = \{q_{\mathsf{init}}, q_2\}$ but not for strongly synchronizing according to sum_T.

strongly synchronizing objectives. We present a family of MDPs $\mathcal{M}_n$ ($n \in \mathbb{N}$) that are sure winning for strongly synchronizing (according to max_Q), and where the sure winning strategies require linear memory. The MDP $\mathcal{M}_2$ is shown in Figure 18, and the MDP $\mathcal{M}_n$ is obtained from $\mathcal{M}_2$ by replacing the cycle $q_2 q_3$ of deterministic transitions by a simple cycle of length n. Note that only in q_1 there is a relevant strategic choice. Since both q_1 and q_2 contain probability mass after one step, we need to wait in q_1 (by playing b) until the probability mass in q_2 comes back to q_2 through the cycle. It is easy to show that to ensure strongly synchronizing, we need to play $n-1$ times b in q_1 before playing a, and this requires linear memory. $\square$

5.2. Strongly synchronizing with function sum

The strongly synchronizing objective with function sum_T requires that eventually all the probability mass remains in T. We show that this is equivalent to a traditional reachability objective with target defined by the set S of sure winning initial distributions for the safety objective $\Box T$.

It follows that almost-sure (and limit-sure) winning for strongly synchronizing is equivalent to almost-sure (or equivalently limit-sure) winning for the coBüchi objective $\Diamond\Box T = \{q_0 a_0 q_1 \cdots \in \mathsf{Path}(\mathcal{M}) \mid \exists j \cdot \forall i > j : q_i \in T\}$ in the state-based semantics. However, sure strongly synchronizing is not equivalent to sure winning for the coBüchi objective, as shown by the MDP in Figure 19 which is:

- sure winning for the coBüchi objective $\Diamond\Box\{q_{\mathsf{init}}, q_2\}$ from q_{init} (because in *all* possible infinite paths from q_{init}, there is a point from which only states in $\{q_{\mathsf{init}}, q_2\}$ are visited), but

- not sure winning for the reachability objective $\Diamond S$ where $S = \{q_2\}$ is the winning region for the safety objective $\Box\{q_{\mathsf{init}}, q_2\}$, thus not sure strongly synchronizing (the probability mass assigned to q_1 is always positive after the first step).

Note that this MDP is almost-sure strongly synchronizing in target $T = \{q_{\mathsf{init}}, q_2\}$ from q_{init}, and almost-sure winning for the coBüchi objective $\Diamond\Box T$, as well as almost-sure winning for the reachability objective $\Diamond S$.

Lemma 27. *Given a target set T, an MDP $\mathcal{M}$ is sure (resp., almost-sure or limit-sure) winning for the strongly synchronizing objective according to sum_T if*

and only if $\mathcal{M}$ is sure (resp., almost-sure or limit-sure) winning for the reachability objective $\Diamond S$ where S is the sure winning region for the safety objective $\Box T$.

Proof. First, assume that a state q_{init} of $\mathcal{M}$ is sure (resp., almost-sure or limit-sure) winning for the strongly synchronizing objective according to sum_T, and show that q_{init} is sure (resp., almost-sure or limit-sure) winning for the reachability objective $\Diamond S$.

(*i*) *Limit-sure winning.* For all $\varepsilon > 0$, let $\varepsilon' = \frac{\varepsilon}{|Q|} \cdot \eta^{|Q|}$ where η is the smallest positive probability in the transitions of $\mathcal{M}$. By the assumption, from q_{init} there exists a strategy α and $N \in \mathbb{N}$ such that for all $n \geq N$, we have $\mathcal{M}_n^\alpha(T) \geq 1 - \varepsilon'$. We claim that at step N, all non-safe states have probability at most $\frac{\varepsilon}{|Q|}$, that is $\mathcal{M}_N^\alpha(q) \leq \frac{\varepsilon}{|Q|}$ for all $q \in Q \setminus S$. Towards contradiction, assume that $\mathcal{M}_N^\alpha(q) > \frac{\varepsilon}{|Q|}$ for some non-safe state $q \in Q \setminus S$. Since $q \notin S$ is not safe, there is a path of length $\ell \leq |Q|$ from q to a state in $Q \setminus T$, thus with probability at least $\eta^{|Q|}$. It follows that after $N+\ell$ steps we have $\mathcal{M}_{N+\ell}^\alpha(Q \setminus T) > \frac{\varepsilon}{|Q|} \cdot \eta^{|Q|} = \varepsilon'$, in contradiction with the fact $\mathcal{M}_n^\alpha(T) \geq 1 - \varepsilon'$ for all $n \geq N$. Now, since all non-safe states have probability at most $\frac{\varepsilon}{|Q|}$ at step N, it follows that $\mathcal{M}_N^\alpha(Q \setminus S) \leq \frac{\varepsilon}{|Q|} \cdot |Q| = \varepsilon$ and thus $\mathrm{Pr}^\alpha(\Diamond S) \geq 1 - \varepsilon$. Therefore, $\mathcal{M}$ is limit-sure winning for the reachability objective $\Diamond S$ from q_{init}.

(*ii*) *Almost-sure winning.* Since almost-sure strongly synchronizing implies limit-sure strongly synchronizing, it follows from (*i*) that $\mathcal{M}$ is limit-sure (and thus also almost-sure) winning for the reachability objective $\Diamond S$, as limit-sure and almost-sure reachability coincide for MDPs [24].

(*iii*) *Sure winning.* From q_{init} there exists a strategy α and $N \in \mathbb{N}$ such that for all $n \geq N$, we have $\mathcal{M}_n^\alpha(T) = 1$. Hence, α is sure winning for the reachability objective $\Diamond \mathsf{Supp}(\mathcal{M}_N^\alpha)$, and from all states in $\mathsf{Supp}(\mathcal{M}_N^\alpha)$ the strategy α ensures that only states in T are visited. It follows that $\mathsf{Supp}(\mathcal{M}_N^\alpha) \subseteq S$ is sure winning for the safety objective $\Box T$, and thus α is sure winning for the reachability objective $\Diamond S$ from q_{init}.

For the converse direction of the lemma, assume that a state q_{init} is sure (resp., almost-sure or limit-sure) winning for the reachability objective $\Diamond S$. We construct a winning strategy for strongly synchronizing in T as follows: play according to a sure (resp., almost-sure or limit-sure) winning strategy for the reachability objective $\Diamond S$, and whenever a state of S is reached, then switch to a winning strategy for the safety objective $\Box T$. The constructed strategy is sure (resp., almost-sure or limit-sure) winning for strongly synchronizing according to sum_T because for sure winning, after finitely many steps all paths from q_{init} end up in $S \subseteq T$ and stay in S forever, and for almost-sure (or equivalently limit-sure) winning, for all $\varepsilon > 0$, after sufficiently many steps, the set S is reached with probability at least $1 - \varepsilon$, showing that the outcome is strongly $(1 - \varepsilon)$-synchronizing in $S \subseteq T$, thus the strategy is almost-sure (and also limit-sure) strongly synchronizing. $\qquad\Box$

Corollary 28. $\langle\!\langle 1 \rangle\!\rangle_{limit}^{strongly}(sum_T) = \langle\!\langle 1 \rangle\!\rangle_{almost}^{strongly}(sum_T)$ *for all target sets T.*

The following result follows from Lemma 27 and the fact that the winning

region for sure safety, sure reachability, and almost-sure reachability can be computed in polynomial time for MDPs [24]. Moreover, memoryless strategies are sufficient for these objectives.

Theorem 9. *For the three winning modes of strongly synchronizing according to sum_T in MDPs:*

1. *(Complexity). The membership problem is PTIME-complete.*
2. *(Memory). Pure memoryless strategies are sufficient.*

6. Conclusion

We studied synchronizing properties for Markov decision processes and presented comprehensive expressiveness and decidability results, identifying the expressively equivalent winning modes (Lemma 2, Theorem 7, Corollary 28), and showing, in all winning modes, PSPACE-completeness for eventually and weakly synchronizing, and PTIME-completeness for always and strongly synchronizing (Table 2). We showed that pure strategies are sufficient for all synchronizing objectives and winning modes, and the memory requirements are given in Table 3.

The p-synchronizing objectives we considered are qualitative in the sense that they are defined for $p = 1$ (sure-winning) or for $p \rightarrow 1$ (almost-sure and limit-sure winning). A natural generalization is to consider the same objectives with $p < 1$. However, the quantitative problem, which is to decide, given a rational number $p < 1$ whether an MDP is eventually p-synchronizing (in a given target state) is at least as hard as the Skolem problem (which is to decide whether a linear recurrence sequence over the integers has a zero) whose decidability is a long-standing open question [54]. The proof is by a reduction that can even be carried out for the special case of Markov chains [2, Theorem 3]. A variant of the problem where it is asked whether there exists $p' > p$ such that the given MDP is eventually p'-synchronizing is also Skolem-hard [2, Corollary 4]. An interesting direction for future research is to consider approximation problems such as deciding, given p and $\varepsilon > 0$, whether an MDP is eventually p'-synchronizing for some $p' \in [p - \varepsilon, p + \varepsilon]$. In another direction, the qualitative problem can be generalized to multiple synchronizing objectives (e.g., conjunctions of objectives, in the flavor of limit-sure winning with exact support), and to Boolean combinations of synchronizing objectives, which is completely open.

As we mention in the paragraph on related work (Section 1), synchronizing properties have been considered in several other models of computation, such as weighted automata, register automata, timed systems, and partial-observation systems. An intriguing question is to consider two-player stochastic games and to determine if (or which) synchronizing objectives are decidable. In two-player stochastic games, some states are controlled by an adversary and the synchronizing objectives need to be achieved no matter the choice of the adversary at

their state. The presence of an adversary makes the problem significantly more challenging, as it incurs an alternation of quantifiers over the strategies.

Finally, given the previous works, it is also interesting to extend the results of this article to continuous-time Markov decision processes, pushdown Markov decision processes, and the subclass of one-counter Markov decision processes.

References

[1] Agrawal, M., Akshay, S., Genest, B., Thiagarajan, P. S., 2012. Approximate verification of the symbolic dynamics of Markov chains. In: Proc. of LICS: Logic in Computer Science. IEEE, pp. 55–64.

[2] Akshay, S., Antonopoulos, T., Ouaknine, J., Worrell, J., 2015. Reachability problems for Markov chains. Inf. Process. Lett. 115 (2), 155–158.

[3] Aspnes, J., Herlihy, M., 1990. Fast randomized consensus using shared memory. J. Algorithm 11 (3), 441–461.

[4] Babari, P., Quaas, K., Shirmohammadi, M., 2016. Synchronizing data words for register automata. In: Proc. of MFCS: Mathematical Foundations of Computer Science. Vol. 58 of LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 15:1–15:15.

[5] Bach, E., Shallit, J., 1996. Algorithmic Number Theory, Vol. 1: Efficient Algorithms. MIT Press.

[6] Baier, C., Bertrand, N., Schnoebelen, P., 2006. On computing fixpoints in well-structured regular model checking, with applications to lossy channel systems. In: Proc. of LPAR: Logic for Programming, Artificial Intelligence, and Reasoning. LNCS 4246. Springer, pp. 347–361.

[7] Baier, C., Größer, M., Bertrand, N., 2012. Probabilistic ω-automata. J. ACM 59 (1), 1–52.

[8] Baldoni, R., Bonnet, F., Milani, A., Raynal, M., 2008. On the solvability of anonymous partial grids exploration by mobile robots. In: Proc. of OPODIS: Principles of Distributed Systems. LNCS 5401. Springer, pp. 428–445.

[9] Beauquier, D., Rabinovich, A. M., Slissenko, A., 2002. A logic of probability with decidable model-checking. In: Proc. of CSL: Computer Science Logic. LNCS 2471. Springer, pp. 306–321.

[10] Berlinkov, M. V., 2016. On the probability of being synchronizable. In: Proc. of CALDAM: Algorithms and Discrete Applied Mathematics. LNCS 9602. Springer, pp. 73–84.

[11] Bertrand, N., Dewaskar, M., Genest, B., Gimbert, H., 2017. Controlling a population. In: Proc. of CONCUR: Concurrency Theory. Vol. 85 of LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 12:1–12:16.

[12] Bianco, A., de Alfaro, L., 1995. Model checking of probabalistic and nondeterministic systems. In: Proc. of FSTTCS: Foundations of Software Technology and Theoretical Computer Science. LNCS 1026. Springer, pp. 499–513.

[13] Burkhard, H.-D., 1976. Zum Längenproblem homogener Experimente an determinierten und nicht-deterministischen Automaten. Elektronische Informationsverarbeitung und Kybernetik 12 (6), 301–306.

[14] Cerný, J., 1964. Poznámka k. homogénnym experimentom s konecnymi automatmi. In: Matematicko-fyzikálny časopis. Vol. 14(3). pp. 208–216.

[15] Chadha, R., Korthikanti, V. A., Viswanathan, M., Agha, G., Kwon, Y., 2011. Model checking MDPs with a unique compact invariant set of distributions. In: Proc. of QEST: Quantitative Evaluation of Systems. IEEE Computer Society, pp. 121–130.

[16] Chandra, A. K., Kozen, D., Stockmeyer, L. J., 1981. Alternation. J. ACM 28 (1), 114–133.

[17] Chatterjee, K., Doyen, L., 2016. Computation tree logic for synchronization properties. In: Proc. of ICALP: Automata, Languages, and Programming. Vol. 55 of LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 98:1–98:14.

[18] Chatterjee, K., Henzinger, M., Joglekar, M., Shah, N., 2011. Symbolic algorithms for qualitative analysis of Markov decision processes with Büchi objectives. In: Proc. of CAV: Computer Aided Verification. LNCS 6806. Springer, pp. 260–276.

[19] Chatterjee, K., Henzinger, T. A., 2012. A survey of stochastic ω-regular games. J. Comput. Syst. Sci. 78 (2), 394–413.

[20] Chistikov, D., Martyugin, P., Shirmohammadi, M., 2016. Synchronizing automata over nested words. In: Proc. of FoSSaCS: Foundations of Software Science and Computation Structures. LNCS 9634. Springer, pp. 252–268.

[21] Courcoubetis, C., Yannakakis, M., 1995. The complexity of probabilistic verification. J. ACM 42 (4), 857–907.

[22] de Alfaro, L., 1997. Formal verification of probabilistic systems. Ph.D. thesis, Stanford University.

[23] de Alfaro, L., Henzinger, T. A., 2000. Concurrent omega-regular games. In: Proc. of LICS: Logic in Computer Science. IEEE, pp. 141–154.

[24] de Alfaro, L., Henzinger, T. A., Kupferman, O., 2007. Concurrent reachability games. Theor. Comput. Sci. 386 (3), 188–217.

[25] Doyen, L., Juhl, L., Larsen, K. G., Markey, N., Shirmohammadi, M., 2014. Synchronizing words for weighted and timed automata. In: Proc. of FSTTCS: Foundation of Software Technology and Theoretical Computer Science. Vol. 29 of LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 121–132.

[26] Doyen, L., Massart, T., Shirmohammadi, M., 2011. Infinite synchronizing words for probabilistic automata. In: Proc. of MFCS: Mathematical Foundations of Computer Science. LNCS 6907. Springer, pp. 278–289.

[27] Doyen, L., Massart, T., Shirmohammadi, M., 2011. Synchronizing objectives for Markov decision processes. In: Proc. of iWIGP: Interactions, Games and Protocols. EPTCS 50. pp. 61–75.

[28] Doyen, L., Massart, T., Shirmohammadi, M., 2012. Infinite synchronizing words for probabilistic automata (Erratum). CoRR abs/1206.0995.

[29] Etessami, K., Kwiatkowska, M. Z., Vardi, M. Y., Yannakakis, M., 2008. Multi-objective model checking of Markov decision processes. Logical Methods in Computer Science 4 (4).

[30] Fijalkow, N., Kiefer, S., Shirmohammadi, M., 2016. Trace refinement in labelled markov decision processes. In: Proc. of FoSSaCS: Foundations of Software Science and Computation Structures. LNCS 9634. Springer, pp. 303–318.

[31] Filar, J., Vrieze, K., 1997. Competitive Markov Decision Processes. Springer.

[32] Fokkink, W., Pang, J., 2006. Variations on Itai-Rodeh leader election for anonymous rings and their analysis in PRISM. Journal of Universal Computer Science 12 (8), 981–1006.

[33] Forejt, V., Kwiatkowska, M., Norman, G., Parker, D., 2011. Automated verification techniques for probabilistic systems. In: Proc. of SFM: Formal Methods for Eternal Networked Software Systems. LNCS 6659. Springer, pp. 53–113.

[34] Futcher, B., 1999. Cell cycle synchronization. Methods in Cell Science 21 (2), 79–86.

[35] Gast, N., Gaujal, B., Le Boudec, J.-Y., 2012. Mean field for Markov decision processes: From discrete to continuous optimization. IEEE Trans. Automat. Contr. 57 (9), 2266–2280.

[36] Gimbert, H., Oualhadj, Y., 2010. Probabilistic automata on finite words: Decidable and undecidable problems. In: Proc. of ICALP (2): Automata, Languages and Programming. LNCS 6199. Springer, pp. 527–538.

[37] Goldschlager, L. M., 1977. The monotone and planar circuit value problems are log space complete for P. SIGACT News 9 (2), 25–29.

[38] Grädel, E., Thomas, W., Wilke, T. (Eds.), 2002. Automata, Logics, and Infinite Games: A Guide to Current Research. LNCS 2500. Springer.

[39] Henzinger, T. A., Mateescu, M., Wolf, V., 2009. Sliding window abstraction for infinite Markov chains. In: Proc. of CAV. LNCS 5643. Springer, pp. 337–352.

[40] Hermanns, H., Krcál, J., Kretínský, J., 2014. Probabilistic bisimulation: Naturally on distributions. In: Proc. of CONCUR: Concurrency Theory. LNCS 8704. Springer, pp. 249–265.

[41] Holzer, M., 1995. On emptiness and counting for alternating finite automata. In: Developments in Language Theory. pp. 88–97.

[42] Imreh, B., Steinby, M., 1999. Directable nondeterministic automata. Acta Cybern. 14 (1), 105–115.

[43] Iván, S., 2014. Synchronizing weighted automata. In: Proc. of AFL: Automata and Formal Languages. Vol. 151 of EPTCS. pp. 301–313.

[44] Jancar, P., Sawa, Z., 2007. A note on emptiness for alternating finite automata with a one-letter alphabet. Inf. Process. Lett. 104 (5), 164–167.

[45] Kattenbelt, M., Kwiatkowska, M. Z., Norman, G., Parker, D., 2010. A game-based abstraction-refinement framework for Markov decision processes. Formal Methods in System Design 36 (3), 246–280.

[46] Kfoury, D. J., 1970. Synchronizing sequences for probabilistic automata. Studies in Applied Mathematics 29, 101–103.

[47] Korthikanti, V. A., Viswanathan, M., Agha, G., Kwon, Y., 2010. Reasoning about MDPs as transformers of probability distributions. In: Proc. of QEST: Quantitative Evaluation of Systems. IEEE Computer Society, pp. 199–208.

[48] Kretínský, J., Larsen, K. G., Laursen, S., Srba, J., 2015. Polynomial time decidability of weighted synchronization under partial observability. In: Proc. of CONCUR: Concurrency Theory. Vol. 42 of LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 142–154.

[49] Larsen, K. G., Laursen, S., Srba, J., 2014. Synchronizing strategies under partial observability. In: Proc. of CONCUR: Concurrency Theory. LNCS 8704. Springer, pp. 188–202.

[50] Madani, O., Hanks, S., Condon, A., 2003. On the undecidability of probabilistic planning and related stochastic optimization problems. Artif. Intell. 147 (1-2), 5–34.

[51] Martyugin, P., 2014. Computational complexity of certain problems related to carefully synchronizing words for partial automata and directing words for nondeterministic automata. Theory Comput. Syst. 54 (2), 293–304.

[52] Mukovskiy, A., Slotine, J. J., Giese, M. A., 2010. Design of the dynamic stability properties of the collective behavior of articulated bipeds. In: Proc. of Humanoids: Conference on Humanoid Robots. IEEE, pp. 66–73.

[53] Nicaud, C., 2016. Fast synchronization of random automata. In: Proc. of RANDOM: Workshop on Randomization and Computation. Vol. 60 of LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 43:1–43:12.

[54] Ouaknine, J., Worrell, J., 2014. Positivity problems for low-order linear recurrence sequences. In: Proc. of SODA: Symposium on Discrete Algorithms. SIAM, pp. 366–379.

[55] Paz, A., 1971. Introduction to probabilistic automata. Academic Press, Inc. Orlando, FL, USA.

[56] Pogosyants, A., Segala, R., Lynch, N. A., 2000. Verification of the randomized consensus algorithm of Aspnes and Herlihy: a case study. Distributed Computing 13 (3), 155–186.

[57] Puterman, M. L., 1994. Markov Decision Processes. John Wiley and Sons.

[58] Rabin, M. O., 1963. Probabilistic automata. Information and Control 6, 230–245.

[59] Shirmohammadi, M., 2014. Qualitative analysis of probabilistic synchronizing systems. Ph.D. thesis, Université Libre de Bruxelles.

[60] Szykula, M., 2018. Improving the upper bound on the length of the shortest reset word. In: Proc. of STACS: Symposium on Theoretical Aspects of Computer Science. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 56:1–56:13.

[61] Vardi, M. Y., 1985. Automatic verification of probabilistic concurrent finite-state programs. In: Proc. of FOCS: Foundations of Computer Science. IEEE Computer Society, pp. 327–338.

[62] Vardi, M. Y., 2007. Automata-theoretic model checking revisited. In: Proc. of VMCAI: Verification, Model Checking, and Abstract Interpretation. LNCS 4349. Springer, pp. 137–150.

[63] Volkov, M. V., 2008. Synchronizing automata and the Cerny conjecture. In: Proc. of LATA: Language and Automata Theory and Applications. LNCS 5196. Springer, pp. 11–27.