

The Value Problem for Multiple-Environment MDPs with Parity Objective

Krishnendu Chatterjee 

IST Austria

Laurent Doyen 

CNRS & LMF, ENS Paris-Saclay, France

Jean-François Raskin 

Université Libre de Bruxelles, Belgium

Ocan Sankur 

Université de Rennes, CNRS, Inria, France & Mitsubishi Electric R&D Centre Europe, France

Abstract

We consider multiple-environment Markov decision processes (MEMDP), which consist of a finite set of MDPs over the same state space, representing different scenarios of transition structure and probability. The value of a strategy is the probability to satisfy the objective, here a parity objective, in the worst-case scenario, and the value of an MEMDP is the supremum of the values achievable by a strategy.

We show that deciding whether the value is 1 is a PSPACE-complete problem, and even in P when the number of environments is fixed, along with new insights to the almost-sure winning problem, which is to decide if there exists a strategy with value 1. Pure strategies are sufficient for these problems, whereas randomization is necessary in general when the value is smaller than 1. We present an algorithm to approximate the value, running in double exponential space. Our results are in contrast to the related model of partially-observable MDPs where all these problems are known to be undecidable.

2012 ACM Subject Classification Theory of computation → Logic and verification; Theory of computation → Probabilistic computation

Keywords and phrases Markov decision processes, imperfect information, randomized strategies, limit-sure winning

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.144

Funding *Krishnendu Chatterjee*: ERC CoG 863818 (ForM-SMArt) and Austrian Science Fund (FWF) 10.55776/COE12;

Jean-François Raskin: PDR Weave project FORM-LEARN-POMDP funded by FNRS and DFG, and the support of the Fondation ULB;

Ocan Sankur: ANR BisoUS (ANR-22-CE48-0012) and ANR EpiRL (ANR-22-CE23-0029).

1 Introduction

We consider Markov decision processes (MDP), a well-established state-transition model for decision making in a stochastic environment. The decisions involve choosing an action from a finite set, which together with the current state determine a probability distribution over the successor state. The question of constructing a strategy that maximizes the probability to satisfy a logical specification is a classical synthesis problem with a wide range of applications [20, 12, 3, 21].

The stochastic transitions in MDPs capture the uncertainty in the effect of an action. Another form of uncertainty arises when the states are (partially) hidden to the decision-maker, as in the classical model of partially-observable MDPs (POMDP) [16, 19]. Recently,



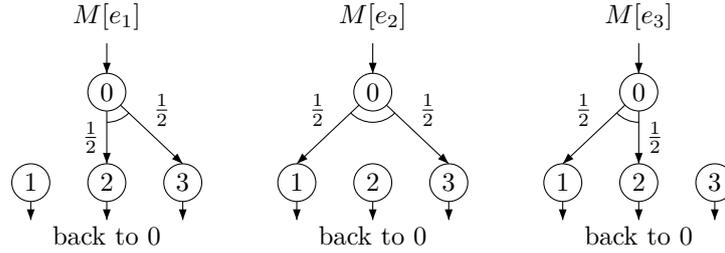
© Krishnendu Chatterjee and Laurent Doyen and Jean-François Raskin and Ocan Sankur; licensed under Creative Commons License CC-BY 4.0

52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025).

Editors: Keren Censor-Hillel, Fabrizio Grandoni, Joel Ouaknine, and Gabriele Puppis; Article No. 144; pp. 144:1–144:17



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Multiple-environment MDP for the missing card (over 3-card deck). Each $M[e_i]$ represents the behavior of the MEMDP under environment e_i where card i has been removed. The environment can be identified almost-surely (with probability 1).

an alternative model of MDPs with partial information has attracted attention, the multiple-environment MDPs (MEMDP) [22], which consists of a finite set of MDPs over the same state space. Each MDP represents a possible environment, but the decision-maker does not know in which environment they are operating. The synthesis problem is then to construct a single strategy that can be executed in all environments to ensure the objective be satisfied independently of the environment. This model is natural in applications where the structure of the transitions and their probability are uncertain such as in robust planning or population models with individual variability [4, 6, 1, 24, 23].

In contrast to what previous work suggest, the two models of POMDP and MEMDP are (syntactically) incomparable: the choice of the environment in MEMDP is adversarial, which cannot be expressed in a POMDP, and the partial observability of POMDP can occur throughout the execution, whereas the uncertainty in MEMDP is only initial. In particular, MEMDP are *not* a subclass of POMDP since pure strategies are sufficient in POMDPs [18, 7] while randomization is necessary in general in MEMDPs [22, Lemma 3].

The synthesis problem has been considered for traditional ω -regular objectives, defined as parity [22] or Rabin [23] condition, in three variants: the almost-sure problem is to decide whether there exists a strategy that is winning with probability 1 in all environments, the limit-sure problem is to decide whether, for every $\varepsilon > 0$, there exists a strategy that is winning with probability at least $1 - \varepsilon$ in all environments, and the gap problem, which is an approximate version of the quantitative problem to decide, given a threshold $0 < \lambda \leq 1$, whether there exists a strategy that is winning with probability at least λ in all environments. The limit-sure problem is also called the value-1 problem, where the value of an MEMDP is defined as the supremum of the values achievable by a strategy. The value is 1 if and only if the answer to the limit-sure problem is **Yes**.

A classical example to illustrate the difference between almost-sure and limit-sure winning is to consider an environment consisting of 51 cards, obtained by removing one card from a standard 52-card deck (see Figure 1). The decision-maker has two possible actions: the action *sample* reveals the top card of the deck and then shuffles the cards (including the top card, which remains in the deck); the action *guess*(x), where $1 \leq x \leq 52$ is a card, stops the game with a win if x is the missing card, and a lose otherwise. If no guess is ever made, the game is also losing. An almost-sure winning strategy is to sample until each of the 51 cards has been revealed at least once, then to make a correct guess. It is easy to see that the strategy wins with probability 1, even if there exist scenarios (though with probability 0) where some of the 51 cards are never revealed and no correct guess is made. Hence the MEMDP is almost-sure winning, and we say that it is not sure winning because a losing scenario exists

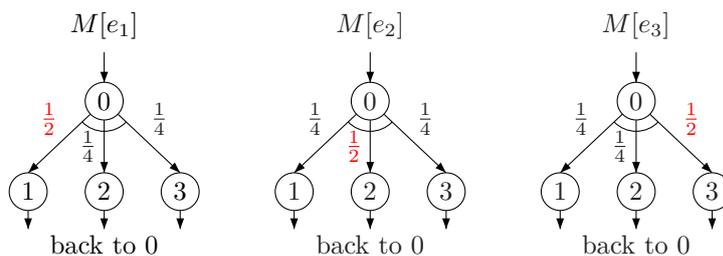


Figure 2 Multiple-environment MDP for the duplicate card (over 3-card deck). Each $M[e_i]$ represents the behavior of the MEMDP under environment e_i where card i has been duplicated. The environment can be identified limit-surely (with probability arbitrarily close to 1).

in every strategy. Consider now an environment consisting of 53 cards, obtained by adding one duplicate card c to the standard deck, and the same action set and rules of the game, except that a correct guess is now the duplicate card $x = c$ (see Figure 2). The strategy that samples for a long time and then makes a guess based on the most frequent card wins with probability close to 1 – and closer to 1 as the sampling time is longer – but not equal to 1, since no matter how long is the sampling phase there is always a nonzero probability that the duplicate card does not have the highest frequency at the time of the guess. In this case, the MEMDP is limit-sure winning, but not almost-sure winning. Intuitively, the solution of almost-sure winning relies on the analysis of *revealing* transitions, which give a sure information allowing to exclude some environment (seeing card c is a guarantee that we are not in the environment where c is missing); the solution of limit-sure winning involves *learning* by sampling, which also allows to exclude some environment, but possibly with a nonzero probability to be mistaken.

For MEMDPs with two environments, it is known that the almost-sure and limit-sure problem for parity objectives are solvable in polynomial time [22, Theorem 33, Theorem 40], while the gap problem is decidable in 2-fold exponential space [22, Theorem 30] and is NP-hard, even for acyclic MEMDPs with two environments [22, Theorem 26]. With an arbitrary number of environments, the almost-sure problem becomes PSPACE-complete [23, Theorem 41], even for reachability objectives [24, Lemma 11]. For comparison, in the close model of POMDP, the decidability frontier lies between limit-sure winning and almost-sure winning: with reachability objectives, the almost-sure problem is decidable (and EXPTIME-complete [2]), whereas the limit-sure problem is undecidable [13]. The gap problem is also undecidable [17].

In this paper, we consider the limit-sure problem and the gap problem for parity objectives in MEMDPs with an arbitrary number of environments. Our main result is to show that (a) the limit-sure problem is PSPACE-complete and can be solved in polynomial time for a fixed number of environments, and (b) the gap problem can be solved in double exponential space. Correspondingly, our algorithms significantly extend the solutions that are known for two environments, relying on a non-trivial recursive (inductive) analysis.

The PSPACE upper bound is obtained by a characterization of limit-sure winning for a subclass of MEMDPs, in terms of almost-sure winning conditions (Lemma 12). A pre-processing phase transforms general MEMDPs into the subclass. We present a PSPACE algorithm to compute the pre-processing and verify the characterization. Since our algorithm relies on almost-sure winning, we also give a new characterization of almost-sure winning for parity objectives in MEMDPs (Lemma 3), which gives a conceptually simple alternative

algorithm to the known solution [23]. The PSPACE lower bound straightforwardly follows from the same reduction as for almost-sure winning [23, Theorem 7]. A corollary of our characterizations is a refined strategy complexity: pure (non-randomized) strategies are sufficient for both limit-sure and almost-sure winning, which was known only for acyclic MEMDPs and almost-sure reachability objectives [24, Lemma 12], and exponential memory is sufficient. In the last part of the paper, we present an algorithm running in double exponential space for solving the gap problem, by computing an approximation of the value of the MEMDP. To win with probability at least λ in all environments, randomized strategies are more powerful [22, Lemma 3], and thus need to be considered for solving the gap problem.

In conclusion, the model of MEMDP is a valuable alternative to POMDPs, from a theoretical perspective since the limit-sure problem and gap problem are undecidable for POMDPs whereas our results establish decidability for MEMDPs, and from a practical perspective since many applications of POMDPs can be expressed by MEMDPs, as was observed previously [1, 24].

Omitted proofs and material are provided in the extended version [8].

2 Definitions

A *probability distribution* on a finite set Q is a function $d : Q \rightarrow [0, 1]$ such that $\sum_{q \in Q} d(q) = 1$. The support of d is $\text{Supp}(d) = \{q \in Q \mid d(q) > 0\}$. A Dirac distribution assigns probability 1 to some $q \in Q$. We denote by $\mathcal{D}(Q)$ the set of all probability distributions on Q .

2.1 Markov Decision Processes

A *Markov decision process (MDP)* over a finite set A of actions is a tuple $M = \langle Q, (A_q)_{q \in Q}, \delta \rangle$ consisting of a finite set Q of *states*, a nonempty set $A_q \subseteq A$ of actions for each state $q \in Q$, and a partial probabilistic transition function $\delta : Q \times A \rightarrow \mathcal{D}(Q)$. We say that (q, a, q') is a transition if $\delta(q, a)(q') > 0$. A state $q \in Q$ is a *sink* if $\delta(q, a)(q) = 1$ for all $a \in A_q$.

A *run* of M from an initial state $q_0 \in Q$ is an infinite sequence $\pi = q_0 a_0 q_1 a_1 \dots$ of interleaved states and actions such that $a_i \in A_{q_i}$ and $\delta(q_i, a_i)(q_{i+1}) > 0$ for all $i \geq 0$. Finite prefixes $\rho = q_0 a_0 \dots q_n$ of runs ending in a state are called *histories* and we denote by $\text{last}(\rho) = q_n$ the last state of ρ . We denote by $\text{Hist}^\omega(M)$ (resp., $\text{Hist}(M)$) the set of all runs (resp., histories) of M , and by $\text{Inf}(\pi)$ the set of states that occur infinitely often along the run π .

A *sub-MDP* of M is an MDP $M' = \langle Q', (A'_q)_{q \in Q'}, \delta' \rangle$ such that $Q' \subseteq Q$ and $\text{Supp}(\delta(q, a)) \subseteq Q'$ for all states $q \in Q'$ and actions $a \in A'_q$ (recall the requirement that $A'_q \neq \emptyset$). Consider a set $Q' \subseteq Q$ such that for all $q \in Q'$, there exists $a \in A_q$ with $\text{Supp}(\delta(q, a)) \subseteq Q'$. We define the *sub-MDP of M induced by Q'* , denoted by $M|_{Q'}$, as the sub-MDP $M' = \langle Q', (A'_q)_{q \in Q'}, \delta' \rangle$ where $A'_q = \{a \in A_q \mid \text{Supp}(\delta(q, a)) \subseteq Q'\}$ for all $q \in Q'$.

End-components An *end-component* of $M = \langle Q, (A_q)_{q \in Q}, \delta \rangle$ is a pair $(Q', (A'_q)_{q \in Q'})$ such that $(Q', (A'_q)_{q \in Q'}, \delta')$ is a sub-MDP of M , where δ' denotes the restriction of δ to $\{(q, a) \mid q \in Q', a \in A'_q\}$, and where the graph $\langle Q', E' \rangle$ with $E' = \{(q, q') \in Q' \times Q' \mid \exists a \in A'_q : \delta(q, a)(q') > 0\}$ is strongly connected [10, 3]. We often identify an end-component as the set $Q' \cup \{(q, a) \mid q \in Q', a \in A_q\}$ of states and state-action pairs, and we say that it is *supported* by the set Q' of states. The (componentwise) union of two end-components with nonempty intersection is an end-component, thus one can define *maximal* end-components. We denote by $\text{MEC}(M)$ the set of maximal end-components of M , which is computable in polynomial time [10], and by $\text{EC}(M)$ the set of all end-components of M .

Histories and Strategies A *strategy* is a function $\sigma : \text{Hist}(M) \rightarrow \mathcal{D}(A)$ such that $\text{Supp}(\sigma(\rho)) \subseteq A_q$ for all histories $\rho \in \text{Hist}(M)$ ending in $\text{last}(\rho) = q$. A strategy is *pure* if all histories are mapped to Dirac distributions. A strategy σ is *memoryless* if $\sigma(\rho) = \sigma(\rho')$ for all histories ρ, ρ' such that $\text{last}(\rho) = \text{last}(\rho')$. We sometimes view memoryless strategies as functions $\sigma : Q \rightarrow \mathcal{D}(A)$. A strategy σ uses *finite memory* (of size k) if there exists a right congruence \approx over $\text{Hist}(M)$ (i.e., such that if $\rho \approx \rho'$, then $\rho \cdot a \cdot q \approx \rho' \cdot a \cdot q$ for all $\rho, \rho' \in \text{Hist}(M)$ and $(a, q) \in A \times Q$) of finite index k such that $\sigma(\rho) = \sigma(\rho')$ for all histories $\rho \approx \rho'$ with $\text{last}(\rho) = \text{last}(\rho')$.

Objectives An objective φ is a Borel set of runs. We denote by $\mathbb{P}_q^\sigma(M, \varphi)$ the standard probability measure on the sigma-algebra over the set of (infinite) runs of M with initial state q , generated by the cylinder sets spanned by the histories [3]. Given a history $\rho = q_0 a_0 q_1 \dots q_k$, the cylinder set $\text{Cyl}(\rho) = \rho(AQ)^\omega$ has probability $\mathbb{P}_q^\sigma(M, \text{Cyl}(\rho)) = \prod_{i=0}^{k-1} \sigma(q_0 a_0 q_1 \dots q_i)(a_i) \cdot \delta(q_i, a_i)(q_{i+1})$ if $q_0 = q$, and probability 0 otherwise. We say that a run ρ is compatible with strategy σ if $\mathbb{P}_q^\sigma(M, \text{Cyl}(\rho)) > 0$.

We consider the following standard objectives for an MDP M :

- safety objective: given a set $T \subseteq Q$ of states, let $\text{Safe}(T) = \{q_0 a_0 q_1 a_1 \dots \in \text{Hist}^\omega(M) \mid \forall i \geq 0 : q_i \in T\}$;
- reachability objective: given a set $T \subseteq Q$ of states, let $\text{Reach}(T) = \{q_0 a_0 q_1 a_1 \dots \in \text{Hist}^\omega(M) \mid \exists i \geq 0 : q_i \in T\}$;
- parity objective: given a priority function $p : Q \rightarrow \mathbb{N}$, let $\text{Parity}(p) = \{\pi \in \text{Hist}^\omega(M) \mid \min\{p(q) \mid q \in \text{Inf}(\pi)\} \text{ is even}\}$.

It is standard to cast safety and reachability objectives as special cases of parity objectives, using sink states. Given an objective φ , we denote by $\neg\varphi = \text{Hist}^\omega(M) \setminus \varphi$ the complement of φ . We say that a run $\pi \in \text{Hist}^\omega(M)$ *satisfies* φ if $\pi \in \varphi$, and that it *violates* φ otherwise.

It is known that under arbitrary strategies, with probability 1 the set $\text{Inf}(\pi)$ of states occurring infinitely often along a run π is the support of an end-component [9, 10].

► **Lemma 1** ([9, 10]). *Given an MDP M , for all states $q \in Q$ and all strategies σ , we have $\mathbb{P}_q^\sigma(M, \{\pi \mid \text{Inf}(\pi) \text{ is the support of an end-component}\}) = 1$.*

An end-component $D \in \text{EC}(M)$ is *positive* under strategy σ from q if $\mathbb{P}_q^\sigma(M, \{\pi \mid \text{Inf}(\pi) = D\}) > 0$. By Lemma 1, we have $\sum_{D \in \text{EC}(M)} \mathbb{P}_q^\sigma(M, \{\pi \mid \text{Inf}(\pi) = D\}) = 1$.

Value and qualitative satisfaction A strategy σ is winning for objective φ from q with probability (at least) α if $\mathbb{P}_q^\sigma(M, \varphi) \geq \alpha$. We denote by $\text{Val}_q^*(M, \varphi) = \sup_\sigma \mathbb{P}_q^\sigma(M, \varphi)$ the *value* of objective φ from state q . A strategy σ is *optimal* if $\mathbb{P}_q^\sigma(M, \varphi) = \text{Val}_q^*(M, \varphi)$.

We consider the following classical qualitative modes of winning. Given an objective φ , a state q is:

- *almost-sure winning* if there exists a strategy σ such that is winning with probability 1, that is $\mathbb{P}_q^\sigma(M, \varphi) = 1$.
- *limit-sure winning* if $\text{Val}_q^*(M, \varphi) = 1$, or equivalently for all $\varepsilon > 0$ there exists a strategy σ such that $\mathbb{P}_q^\sigma(M, \varphi) \geq 1 - \varepsilon$.

We denote by $\text{AS}(M, \varphi)$ and $\text{LS}(M, \varphi)$ the set of almost-sure and limit-sure winning states, respectively. In MDPs, it is known that $\text{AS}(M, \varphi) = \text{LS}(M, \varphi)$ and pure memoryless optimal strategies exist for parity objectives φ [20, 9].

We recall that the value of a parity objective $\varphi = \text{Parity}(p)$ from every state of an end-component D is the same, and is either 0 or 1, which does not depend on the precise value of the (non-zero) transition probabilities, but only on the supports $\text{Supp}(\delta(q, a))$ of the transition function at the state-action pairs (q, a) in D [10]. When the value 1, there exists a

pure memoryless strategy σ such that $\mathbb{P}_q^\sigma(M, \varphi) = 1$ for all states $q \in D$. If such a strategy exists, then D is said to be φ -winning, and otherwise φ -losing.

2.2 Multiple-Environment MDP

A *multiple-environment MDP* (MEMDP) over a finite set E of environments is a tuple $M = \langle Q, (A_q)_{q \in Q}, (\delta_e)_{e \in E} \rangle$, where $M[e] = \langle Q, (A_q)_{q \in Q}, \delta_e \rangle$ is an MDP that models the behaviour of the system in the environment $e \in E$. The state space is identical in all $M[e]$ ($e \in E$), only the transition probabilities may differ. We sometimes refer to the environments of M as the MDPs $\{M[e] \mid e \in E\}$ rather than the set E itself. For $E' \subset E$, let $M[E']$ be the MEMDP M over set E' of environments. We denote by $M[-e]$ the MEMDP M over environments $E \setminus \{e\}$, and by $\cup_{e \in E} M[e]$ the MDP $\langle Q, (A_q)_{q \in Q}, \delta_\cup \rangle$ such that $\delta_\cup(q, a)$ is the uniform distribution over $\cup_{e \in E} \text{Supp}(\delta_e(q, a))$ for all $q \in Q$ and $a \in A$.

A transition $t = (q, a, q')$ is *revealing* in M if $K_t = \{e \in E \mid q' \in \text{Supp}(\delta_e(q, a))\}$ is a strict subset of E ($K_t \subsetneq E$). We say that K_t , which is the set of environments where the transition $t = (q, a, q')$ is possible, is the *knowledge* after observing transition t . An MEMDP is in *revealed form* if for all revealing transitions $t = (q, a, q')$, the state q' is a sink in all environments, that is $\text{Supp}(\delta_e(q', a)) = \{q'\}$ for all environments $e \in E$ and all actions $a \in A_{q'}$. By extension, we call knowledge after a history ρ the set of environments in which all transitions of ρ are possible.

Decision Problems We are interested in synthesizing a *single* strategy σ with guarantees in *all* environments, without knowing in which environment σ is executing. We consider reachability, safety, and parity objectives.

A state q is *almost-sure winning* in M for objective φ if there exists a strategy σ such that in all environments $e \in E$, we have $\mathbb{P}_q^\sigma(M[e], \varphi) = 1$, and we call such a strategy σ almost-sure winning. A state q is *limit-sure winning* in M for objective φ if for all $\varepsilon > 0$, there exists a strategy σ such that in all environments $e \in E$ we have $\mathbb{P}_q^\sigma(M[e], \varphi) \geq 1 - \varepsilon$, and we say that such a strategy σ is $(1 - \varepsilon)$ -winning.

We denote by $\text{AS}(M, \varphi)$ (resp., $\text{LS}(M, \varphi)$) the set of all almost-sure (resp., limit-sure) winning states in M for objective φ . We consider the *membership problem for almost-sure* (resp., *limit-sure*) *winning*, which asks whether a given state q is almost-sure (resp., limit-sure) winning in M for objective φ . We refer to these membership problems as *qualitative* problems.

We are also interested in the *quantitative* problems. Given MEMDP M , a parity objective φ , and probability threshold $\alpha \geq 0$, we are interested in the existence of a strategy σ satisfying $\mathbb{P}_q^\sigma(M[e], \varphi) \geq \alpha$ for all $e \in E$. We present an approximation algorithm for the quantitative problem, solving the *gap problem* consisting, given MEMDP M , state q , parity objective φ , and thresholds $0 < \alpha < 1$ and $\varepsilon > 0$, in answering

- Yes if there exists a strategy σ such that for all $e \in E$, we have $\mathbb{P}_q^\sigma(M[e], \varphi) \geq \alpha$,
- No if for all strategies σ , there exists $e \in E$ with $\mathbb{P}_q^\sigma(M[e], \varphi) < \alpha - \varepsilon$,
- and arbitrarily otherwise.

The gap problem is an instance of promise problems which guarantee a correct answer in two disjoint sets of inputs, namely positive and negative instances – which do not necessarily cover all inputs, while giving no guarantees in the rest of the input [11, 14].

Results We solve the membership problem for limit-sure winning with parity objectives φ (i.e., deciding whether a given state q is limit-sure winning, that is $q \in \text{LS}(M, \varphi)$), providing a PSPACE algorithm with a matching complexity lower bound, and showing that the problem is solvable in polynomial time when the number of environments is fixed. Our solution relies on the solution of almost-sure winning, which is known to be PSPACE-complete for

reachability [24] and Rabin objectives [23]. We revisit the solution of almost-sure winning and give a simple characterization for safety objectives (which is also PSPACE-complete), that can easily be extended to parity objectives. A corollary of our characterization is that pure (non-randomized) strategies are sufficient for both limit-sure and almost-sure winning, which was known only for acyclic MEMDPs and reachability objectives [24, Lemma 12].

For the gap problem, we present an double exponential-space procedure to approximate the value α that can be achieved in all environments, up to an arbitrary precision ε .

3 Almost-Sure Winning

It is known that the membership problem for almost-sure winning in MEMDPs is PSPACE-complete with reachability objectives [24] as well as with Rabin objectives [23], an expressively equivalent of the parity objectives. We revisit the membership problem for almost-sure winning with parity objectives, as it will be instrumental to the solution of limit-sure winning. We present a conceptually simple characterization of the winning region for almost-sure winning, from which we derive a PSPACE algorithm, thus matching the known complexity for almost-sure Rabin objectives. A corollary of our characterization is that pure (non-randomized) strategies are sufficient for both limit-sure and almost-sure winning, which was known only for acyclic MEMDPs and reachability objectives [24, Lemma 12].

► **Theorem 2** ([24],[23]). *The membership problem for almost-sure winning in MEMDPs with a reachability, safety, or Rabin objective is PSPACE-complete.*

To solve the membership problem for a parity objective φ , we first convert M into an MEMDP M' in revealed form with state space $Q \uplus \{q_{\text{win}}, q_{\text{lose}}\}$ and each revealing transition $t = (q, a, q')$ in M is redirected in M' to q_{win} if $q' \in \text{AS}(M[K_t], \varphi)$ is almost-sure winning when the set of environments is the knowledge K_t after observing transition t , and to q_{lose} otherwise. In order to decide if $q' \in \text{AS}(M[K_t], \varphi)$, we need to solve the membership problem for an MEMDP with strictly fewer environments than in M as $K_t \subsetneq E$, which will lead to a recursive algorithm. The base case of the solution is MEMDPs with one environment, which is equivalent to plain MDPs.

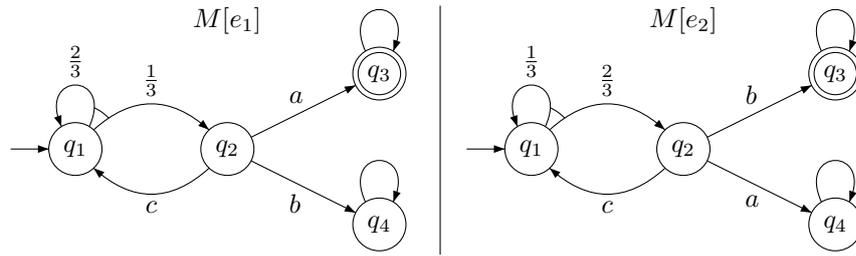
It is easy to see that $\text{AS}(M, \varphi) \cup \{q_{\text{win}}\} = \text{AS}(M', \varphi \cup \text{Reach}(q_{\text{win}}))$ for all prefix-independent objectives φ , and we can transform the objective $\varphi \cup \text{Reach}(q_{\text{win}})$ into an objective of the same type as φ (for example, if φ is a parity objective then assigning the smallest even priority to q_{win} turns the objective $\varphi \cup \text{Reach}(q_{\text{win}})$ into a pure parity objective).

Hence, the main difficulty is to solve the membership problem for MEMDP in revealed form.

By definition, the almost-sure winning region $W = \text{AS}(M, \text{Parity}(p))$ for a parity objective in an MEMDP M is such that there exists a strategy σ that is almost-sure winning for the parity objective from every state $q \in W$ in every MDP $M[e]$ (where e is an environment of M). In contrast, we show the following characterization (note the order of the quantifiers).

► **Lemma 3.** *Given an MEMDP M in revealed form with state space Q , if $W \subseteq Q$ is such that in every environment e , from every state $q \in W$, there exists a strategy σ_e that is almost-sure winning for the parity objective $\text{Parity}(p)$ in $M|_W[e]$ from q , then $W \subseteq \text{AS}(M, \text{Parity}(p))$. Moreover, for all $q \in W$, there exists a pure $(|Q| \cdot |E|)$ -memory strategy ensuring $\text{Parity}(p)$ from q in M .*

The characterization in the first part of Lemma 3 holds simply because parity objectives are prefix-independent (runs that differ by a finite prefix are either both winning or both losing), and thus the characterization holds for all prefix-independent objectives.



■ **Figure 3** An end-component $\{q_1, q_2\}$ with different transition probabilities in environments e_1 and e_2 .

Following the characterization in Lemma 3, we can construct an algorithm that computes the revealed form (using recursive calls to determine the winning region of MEMDPs with smaller sets of environments), and then computes the winning region by a fixpoint iteration computable in polynomial time for a fixed set of environments. The depth of the recursive calls is bounded by the number of environments, giving a PSPACE algorithm. Correspondingly, we show that pure strategies with exponential memory $O(|Q| \cdot |E| \cdot 2^{|E|})$ are sufficient. It is easy to show that exponential memory is necessary using the construction establishing PSPACE lower bound of the problem [23, Theorem 7].

► **Theorem 4.** *The membership problem for almost-sure parity in MEMDPs is PSPACE-complete. Pure exponential-memory strategies are sufficient for almost-sure winning in MEMDPs with parity (thus also reachability and safety) objectives. When the number of environments is fixed, the problem is solvable in polynomial time.*

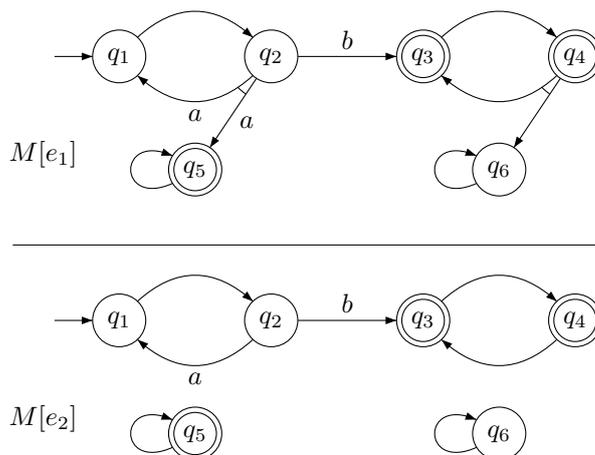
4 Limit-Sure Winning

We refer to the examples of the duplicate card and the missing card in Section 1 to illustrate the difference between limit-sure and almost-sure winning. We present in Section 4.1 two other scenarios where limit-sure winning and almost-sure winning do not coincide, which will be useful to illustrate the key ideas in the algorithmic solution.

4.1 Examples

In the example of Figure 3, the set $D = \{q_1, q_2\}$ is an end-component in both environments e_1 and e_2 (the actions are shown in the figures only when relevant, that is in q_2). However, the transition probabilities from q_1 are different in the two environments e_1 and e_2 , and intuitively we can learn (with high probability) in which environment we are by playing c for a long enough (but finite) time and collecting the frequency of the visits to q_1 and q_2 . Then, in order to reach the target q_3 , if there are more q_1 's than q_2 's in the history we play a in q_2 , otherwise b . The intuition is that the histories with more q_1 's than q_2 's have a high probability (more than $1 - \varepsilon$) in $M[e_1]$ and a small probability (less than ε) in $M[e_2]$, where ε can be made arbitrarily small (however not 0) by playing c for sufficiently long. Hence q_1 is limit-sure winning, but not almost-sure winning.

In the second scenario (Figure 4), the transition probabilities do not matter. The objective is to visit some state in $\{q_3, q_4, q_5\}$ infinitely often (those states have priority 0, the other states have priority 1). The state q_1 is limit-sure winning, but not almost-sure winning. To



■ **Figure 4** The set $\{q_1, q_2\}$ is an end-component in e_2 , not in e_1 .

win with probability $1 - \varepsilon$, a strategy can play a (in q_2) for a sufficiently long time, then switch to playing b (unless q_5 was reached before that). The crux is that playing a does not harm, as it does not leave the limit-sure winning region, but ensures in at least one environment (namely, e_1) that the objective is satisfied with probability 1 (by reaching q_5). This allows to “discard” the environment e_1 if q_5 was not reached, and to switch to a strategy that is winning with probability at least $1 - \varepsilon$ in e_2 , namely by playing b . With an arbitrary number of environments, the difficulty is to determine in which order the environments can be “discarded”.

Note that the transition (q_2, a, q_1) is not revealing, since it is present in both environments. However, after crossing this transition a large number of times, we can still learn that the environment is e_2 (and be mistaken with arbitrarily small probability). In contrast, the transition (q_2, a, q_5) is revealing and the environment is e_1 with certainty upon crossing that transition.

To solve the membership problem for limit-sure parity, we first convert M into a revealed-form MEMDP M' , similar to the case of almost-sure winning, with the obvious difference that revealing transitions $t = (q, a, q')$ of $M[e]$ are redirected in $M'[e]$ to q_{win} if $q' \in \text{LS}(M[K_t], \varphi)$ is limit-sure winning when the set of environments is the knowledge K_t after observing transition t . Thus, we aim for a recursive algorithm, where the base case is limit-sure winning in MEMDPs with one environment, which are equivalent to plain MDPs, for which limit-sure and almost-sure parity coincide. Note that the examples of Figure 3 and Figure 4 are in revealed form.

4.2 Common End-Components and Learning

A *common end-component (CEC)* of an MEMDP $M = \langle Q, (A_q)_{q \in Q}, (\delta_e)_{e \in E} \rangle$ is a pair (Q', A') that is an end-component in $M[e]$ for all environments $e \in E$. A CEC D is *trivial* if it contains a single state. D is said *winning* for a parity condition $\text{Parity}(p)$, if for all $e \in E$, there is a strategy in $M[e]$ which, when started inside D , ensures $\text{Parity}(p)$ with probability 1. Notice that since D is a common end-component, such a strategy ensures $\text{Parity}(p)$ with probability 1 in $M[e]$ iff it does in $M[e']$.

We note that the common end-components of an MEMDP are the end-components of the

144:10 The Value Problem for Multiple-Environment MDPs with Parity Objective

MDP $\cup_{e \in E} M[e]$ assuming M is in revealed form, and thus can be computed using standard algorithm for end-components [10].

► **Lemma 5.** *Consider an MEMDP M in revealed form. The common end-components of M are exactly the end-components of $\cup_{e \in E} M[e]$.*

A CEC may have different transition probabilities in different environments. We call a CEC *distinguishing* if it contains a transition (q, a, q') (called a distinguishing transition) such that $\delta_e(q, a)(q') \neq \delta_f(q, a)(q')$ for some environments $e, f \in E$. Given a distinguishing transition (q, a, q') and environment e , define $K_1 = \{f \in E \mid \delta_f(q, a)(q') = \delta_e(q, a)(q')\}$ and $K_2 = E \setminus K_1$. We say that (K_1, K_2) is a *distinguishing partition* of D that is *induced* by the distinguishing transition (q, a, q') and environment e .

Distinguishing transitions can be used to learn the partition (K_1, K_2) , that is to guess (correctly with high probability) whether the current environment is in K_1 or K_2 , as in the example of Figure 3, where the set $D = \{q_1, q_2\}$ is a distinguishing end-component with distinguishing transition (q_1, \cdot, q_2) and partition $(\{e_1\}, \{e_2\})$. A distinguishing CEC may have several distinguishing transitions and induced partitions.

We formalize how a strategy can distinguish between K_1 and K_2 with high probability inside a distinguishing CEC. First let us recall Hoeffding's inequality.

► **Theorem 6 (Hoeffding's Inequality [15]).** *Let X_1, X_2, \dots, X_n be a sequence of independent and identical Bernoulli variables with $\mathbb{P}[X_i] = p$, and write $S_n = X_1 + \dots + X_n$. For all $t > 0$, $\mathbb{P}[S_n - \mathbb{E}[S_n] \geq t] \leq e^{-2t^2/n}$, and $\mathbb{P}[\mathbb{E}[S_n] - S_n \geq t] \leq e^{-2t^2/n}$.*

Given a distinguishing CEC with distinguishing partition (K_1, K_2) induced by a transition (q, a, q') , a strategy can sample the distribution $\delta_e(q, a)$ by repeating the following two phases: first, use a pure memoryless strategy to almost-surely visit q , then play action a ; by repeating this long enough (the precise bound depends on a given ε and is derived from Theorem 6) while storing the frequency of visits to q' in the second phase, we can learn and guess in which block K_i belongs the environment, with sufficiently small probability of mistake to ensure winning with probability $1 - \varepsilon$.

► **Lemma 7.** *Given an MEMDP M containing a distinguishing common end-component D with partition (K_1, K_2) induced by a distinguishing transition, and parity objective φ , for all states q_0 in D , all pairs of strategies σ_1, σ_2 , and all $\varepsilon > 0$, there exists a strategy σ such that:*

$$\begin{aligned} \mathbb{P}_{q_0}^{\sigma} (M[e], \varphi) &\geq (1 - \varepsilon) \mathbb{P}_{q_0}^{\sigma_1} (M[e], \varphi) \text{ for all } e \in K_1, \\ \mathbb{P}_{q_0}^{\sigma} (M[e], \varphi) &\geq (1 - \varepsilon) \mathbb{P}_{q_0}^{\sigma_2} (M[e], \varphi) \text{ for all } e \in K_2. \end{aligned}$$

Moreover, the strategy σ is pure if both σ_i are pure; and if each strategy σ_i uses a memory of size m_i , then σ uses finite memory of size $m_1 + m_2 + \lceil 8 \frac{\log(1/\varepsilon)}{\eta^2} \rceil$ where $\eta = \min(\{|\delta_e(q, a)(q') - \delta_f(q, a)(q')| \mid e, f \in E, q, q' \in Q, a \in A\} \setminus \{0\})$.

It follows that the membership problem for limit-sure winning can be decomposed into subproblems where the set of environments is one of the blocks K_i in the partition.

► **Lemma 8.** *Given an MEMDP M containing a distinguishing common end-component D with a partition (K_1, K_2) induced by a distinguishing transition, and a parity objective φ the following equivalence holds: $D \subseteq \text{LS}(M, \varphi)$ if and only if $D \subseteq \text{LS}(M[K_1], \varphi)$ and $D \subseteq \text{LS}(M[K_2], \varphi)$.*

4.3 Characterization and Algorithm

Here, we assume that MEMDPs are in revealed form with sink states q_{win} and q_{lose} .

We show that the winning region $W = \text{LS}(M, \varphi)$ for limit-sure parity is a closed set: from every state $q \in W$, there exists an action a ensuring in all environments that all successors of q are in W . We call such actions *limit-sure safe* for q . We show in Lemma 9 that a limit-sure safe action always exists in limit-sure winning states. Note that playing actions that are *not* limit-sure safe may be useful for limit-sure winning, as in the example of Figure 3 where action a is limit-sure safe, but action b is not (from q_2).

By definition of limit-sure winning, if a state q is not limit-sure winning, there exists $\varepsilon_q > 0$ such that for all strategies σ , there exists an environment $e \in E$ such that $\mathbb{P}_q^\sigma(M[e], \varphi) < 1 - \varepsilon_q$. We denote by $\varepsilon_0 = \min\{\varepsilon_q \mid q \in Q \setminus \text{LS}(M, \varphi)\}$ a uniform bound.

► **Lemma 9.** *Given an MEMDP M (in revealed form) over environments E , a parity objective φ , and a state q , if $q \in \text{LS}(M, \varphi)$ is limit-sure winning, then there exists an action a such that for all environments $e \in E$, all successors of q are limit-sure winning, i.e. $\text{Supp}(\delta_e(q, a)) \subseteq \text{LS}(M, \varphi)$.*

Given an MEMDP M , consider the limit-sure winning region $W = \text{LS}(M, \varphi)$ for $\varphi = \text{Parity}(p)$. For the purpose of the analysis, consider the (memoryless) randomized strategy σ_{LS} that plays uniformly at random all limit-sure safe actions in every state $q \in W$, which is well-defined by Lemma 9.

Consider an arbitrary environment e , and an end-component D in $M[e]$ that is positive under σ_{LS} (recall Lemma 1 and the definition afterward). There are three possibilities:

1. D is not a common end-component (as in the example of Figure 4, for $D = \{q_1, q_2\}$ in $M[e_2]$), that is, D is not an end-component in some environment e' (in the example $e' = e_1$), then we can learn (and be mistaken with arbitrarily small probability) that we are not in e' , reducing the problem to an MEMDP with fewer environments (namely, $M[\neg e']$);
2. D is a common end-component and is distinguishing (as in the example of Figure 3, for $D = \{q_1, q_2\}$), then we can also learn a distinguishing partition (K_1, K_2) and reduce the problem to MEMDPs with fewer environments (namely, $M[K_1]$ and $M[K_2]$);
3. D is a common end-component and is non-distinguishing, then we show in Lemma 10 below that D is almost-sure winning ($D \subseteq \text{AS}(M, \varphi)$), obviously in all environments.

► **Lemma 10.** *Given an MEMDP M over environments E (in revealed form), a parity objective φ , and a state q , if $q \in \text{LS}(M, \varphi)$, then all non-distinguishing common end-components D that are positive under strategy σ_{LS} from q in $M[e]$ (for some $e \in E$) are almost-sure winning for φ (that is $D \subseteq \text{AS}(M, \varphi)$).*

Our approach to compute the limit-sure winning states is to first identify the distinguishing CECs that are limit-sure winning. We can compute the maximal CECs using Lemma 5, and note that a maximal CEC containing a distinguishing CEC is itself distinguishing, so it is sufficient to consider maximal CECs. By Lemma 8, we can decide if a given distinguishing CEC is limit-sure winning using a recursive procedure on MEMDPs with fewer environments. We show in Lemma 11 below that we can replace the limit-sure CECs by a sink state q_{win} .

► **Lemma 11.** *Given an MEMDP M with parity objective φ and a set $T \subseteq \text{LS}(M, \varphi)$ of limit-sure winning states, we have $\text{LS}(M, \varphi) = \text{LS}(M, \varphi \cup \text{Reach}(T))$.*

Algorithm 1 LS_Parity(M, p)

Input : $M = \langle Q, (A_q)_{q \in Q}, (\delta_e)_{e \in E} \rangle$ an MEMDP, $p : Q \rightarrow \mathbb{N}$ a priority function.
Output: The winning region $\text{LS}(M, \text{Parity}(p))$ for limit-sure parity.

begin

```

1  | if  $|E| = 1$  then return  $\text{AS}(M, \text{Parity}(p))$                                 /* pre-processing */
2  |   put  $M$  in revealed form (defined in Section 3)
3  |    $\text{MEC} \leftarrow$  maximal end-components of the MDP  $\cup_{e \in E} M'[e]$ 
4  |   for  $D \in \text{MEC}$  do
5  |     | if  $D$  is distinguishing in  $M$  then
6  |       |   Let  $(K_1, K_2)$  be a distinguishing partition in  $D$ 
7  |       |   if  $D \subseteq \text{LS\_Parity}(M[K_1], p) \cap \text{LS\_Parity}(M[K_2], p)$  then
8  |       |     |  $\perp$  replace  $D$  by sink  $q_{\text{win}}$  in  $M$  with  $p(q_{\text{win}}) = 0$ 
9  |       |
10 |       |   /*  $M$  is in revealed form and Lemma 12 applies */
11 |   for  $e \in E$  do
12 |     |  $\perp T_e = \text{LS\_Parity}(M[\neg e], p)$ 
13 |    $Q \leftarrow \text{AS}\left(M, \text{Reach}\left(\bigcup_{e \in E} \text{AS}(M[e], \text{Parity}(p) \cap \text{Safe}(T_e))\right)\right)$ 
14 |   return  $Q \setminus \{q_{\text{win}}\}$ 

```

end

We can now assume that MEMDPs contain no limit-sure winning distinguishing CEC, and present a characterization for the remaining possibility, illustrated by the scenario of Figure 4, where playing the action a (in q_2 , forever) ensures, in some environment (namely, e_1), almost-sure satisfaction of the parity objective while remaining inside the limit-sure winning region in all other environments.

► **Lemma 12.** *Consider an MEMDP M (in revealed form) over environments E with $|E| \geq 2$, that contains no limit-sure winning distinguishing common end-component, and a parity objective φ . Writing $T_e = \text{LS}(M[\neg e], \varphi)$, we have the following: $\text{LS}(M, \varphi) = \text{AS}\left(M, \text{Reach}\left(\bigcup_{e \in E} \text{AS}(M[e], \varphi \cap \text{Safe}(T_e))\right)\right)$.*

Algorithm Overview Given a MEMDP $M = (Q, (A_q)_{q \in Q}, (\delta_e)_{e \in E})$, the algorithm proceeds by recursion on the size of the environment set E (Algorithm 1). The base case is that of a singleton set E where $\text{LS}(M, \varphi) = \text{AS}(M, \varphi)$ and this can be computed in polynomial time.

Assume $|E| \geq 2$. We first convert M into an MEMDP M' in revealed form with state space $Q \uplus \{q_{\text{win}}, q_{\text{lose}}\}$ and each revealing transition $t = (q, a, q')$ in M is redirected in M' to q_{win} if $q' \in \text{LS}(M[K_t], \varphi)$ is limit-sure winning when the set of environments is the knowledge K_t after observing transition t , and to q_{lose} otherwise. Notice that each query $q' \in \text{LS}(M[K_t], \varphi)$ uses a set K_t that is strictly smaller than E .

We now assume that M is in revealed form and we compute the maximal end-components of the MDP $\cup_{e \in E} M'[e]$; these are maximal common end-components of M by Lemma 5. For each distinguishing maximal CEC D , we determine whether it is limit-sure winning using the condition of Lemma 8, namely that $D \subseteq \text{LS}(M[K_i], \varphi)$ (for $i = 1, 2$) where (K_1, K_2) is a partition of E induced by a distinguishing transition of D , which is computed by a

recursive calls to the algorithm. We replace D by q_{win} if it is limit-sure winning, which yields an MEMDP without limit-sure winning distinguishing CECs, and we can apply Lemma 12: for each environment $e \in E$, we compute $T_e = \text{LS}(M[\neg e], \varphi)$ which is done by $|E|$ separate recursive calls, and we compute the sets $\text{AS}(M[e], \varphi \cap \text{Safe}(T_e))$ using standard MDP algorithms (we restrict the state space to T_e and compute the almost-sure winning states for φ). We then solve the almost-sure reachability problem in M for the target set $\bigcup_{e \in E} \text{AS}(M[e], \varphi \cap \text{Safe}(T_e))$.

Thus each recursive step takes polynomial time (besides the recursive calls), and because each recursive call decreases the size of E , the depth of the recursion is bounded by $|E|$. It follows that the procedure runs in polynomial space. The PSPACE lower bound follows from the same reduction as for almost-sure winning [23, Theorem 7], since the MEMDP constructed in the reduction is acyclic, thus almost-sure and limit-sure winning coincide.

Note that Lemma 12 constructs a pure strategy that achieves the objective with probability at least $1 - \varepsilon$ from the limit-sure winning states, and that the strategies constructed in Lemmas 8 and 11 to witness limit-sure winning are also pure (in Lemma 8, the construction assumes that pure strategies are sufficient for fewer environments, which allows a proof by induction since pure strategies are sufficient in MDPs, *i.e.* in a single environment).

► **Theorem 13.** *The membership problem for limit-sure parity objectives in MEMDPs is PSPACE-complete and pure exponential-memory strategies are sufficient, *i.e.*, if a state q is limit-sure winning, then for all $\varepsilon > 0$ there exists a pure exponential-memory strategy that ensures the objective is satisfied with probability at least $1 - \varepsilon$ from q . When the number of environments is fixed, the problem is solvable in polynomial time.*

The time complexity of Algorithm 1 is established as follows. Let us consider a single recursive call. The maximal end-components of $\bigcup_{e \in E} M'[e]$ can be computed in $O(|Q| \cdot |\delta|)$ where $|\delta|$ denotes the number of transitions. Then, determining whether each MEC is distinguishing, and replacing them with sink states can be done in time $O(|\delta| \cdot |E|)$ since one needs to go over each transition and check whether their probability differs in two environments. The last step requires solving almost-sure parity and safety for MDPs defined for each $e \in E$, which can be done in time $O(|E| \cdot |Q| \cdot |\delta|)$ (similarly as in the discussion following Theorem 4). The most costly operation is almost-sure reachability for the MEMDP M , which by Theorem 4 takes $O(|Q|^4 \cdot |E| \cdot |A| \cdot 2^{|E|})$. There are $2^{|E|}$ recursive calls (the algorithm can be run once for each subset of E using memoization), so overall we get $O(|Q|^4 \cdot |E| \cdot |A| \cdot 2^{2|E|})$.

5 The Gap Problem

The goal of this section is to give a procedure to solve the gap problem for parity objectives. For this, we show that an arbitrary strategy in M can be imitated by a finite-memory one (with a computable bound on the memory size) while achieving the same probability of winning up to ε in all environments. Once this is established, we show how to guess a finite-memory strategy of appropriate size, and check that it is a solution to the gap problem.

To establish the memory bound for such an ε -approximation, we need a few intermediate lemmas. First, we define a transformation on MEMDPs consisting in collapsing non-distinguishing maximal CECs (MCECs) of the MEMDP M ; the resulting MEMDP is denoted $\text{purge}(M)$. We show that M and $\text{purge}(M)$ have the same probabilities of satisfaction of the considered parity objective under all environments.

Intuitively, removing non-distinguishing MCECs ensures that in $\text{purge}(M)$, under all strategies, with high probability, within a fixed number of steps, either a maximal CEC

is reached (which is either distinguishing, or non-distinguishing but trivial – recall that a trivial CEC contains a single absorbing state) or enough samples are gathered to improve the knowledge about the current environment. This observation will help us constructing the finite-memory strategy inductively since in each case the knowledge can be improved correctly with high probability: in trivial MCECs, the strategy is extended arbitrarily; inside distinguishing MCECs, the strategy can be extended so that it stays inside the MCEC while sampling distinguishing transitions with any desired precision as in Lemma 7; last, if no MCECs are reached but enough samples are gathered along the way, we prove that the knowledge can also be improved with high probability. The final strategy is obtained by combining finite-memory strategies constructed inductively for smaller sets of environments.

Purge: Removing Non-Distinguishing MCECs. We present key properties of the transformation $\text{purge}(M)$ in Lemma 14 below, showing that strategies of M can be imitated in $\text{purge}(M)$, and vice versa. Hence the value of parity objectives in M and $\text{purge}(M)$ coincide.

► **Lemma 14.** *Consider an MEMDP $M = \langle Q, A, (\delta_e)_{e \in E} \rangle$, and objective $\varphi = \text{Parity}(p)$. There exists an MEMDP $\text{purge}(M)$, with a map $f : Q \rightarrow Q'$ relating states of M and those of $\text{purge}(M) = \langle Q', A', (\delta'_e)_{e \in E} \rangle$, with the following properties:*

- *The only non-distinguishing MCECs of $\text{purge}(M)$ are the trivial q_{win} and q_{lose} .*
- *For all states $q \in Q$ and strategies σ for M , there exists σ' with $\mathbb{P}_q^\sigma(M, \varphi) \leq \mathbb{P}_{f(q)}^{\sigma'}(\text{purge}(M), \varphi)$.*
- *For all states $q' \in Q'$ and strategies σ' for $\text{purge}(M)$, and all $q \in f^{-1}(q')$, there exists a strategy σ with $\mathbb{P}_q^\sigma(M, \varphi) = \mathbb{P}_{q'}^{\sigma'}(\text{purge}(M), \varphi)$. Furthermore, if σ' is a m -memory strategy, then σ can be chosen to be a $(m + |Q| \cdot |A|)$ -memory strategy.*

Note that while $\text{purge}(M)$ is computable in polynomial time, our results only need that $\text{purge}(M)$ can be constructed to establish the existence of finite-memory strategies (Lemma 16).

Learning While Playing. We now show that after collapsing non-distinguishing MCECs, within n steps (for n large enough), with high probability, we either reach a MCEC (which is either distinguishing or trivial) or collect a large number of samples of distinguishing transitions whose empirical average is close to their mean. Intuitively, this means that either the knowledge can be improved after n steps using the collected samples while bounding the probability of error, or a MCEC is reached.

For a history h , let $|h|_{q,a}$ denote the number of occurrences of the state-action pair (q, a) , and $|h|_{q,a,q'}$ the number of times these are followed by q' , where $q' \in \text{Supp}(\delta(q, a))$. For a distinguishing transition $t = (q, a, q')$, we say that a history h is a *bad* (t, η) -classification in MDP $M[e]$ if $\left| \frac{|h|_{q,a,q'}}{|h|_{q,a}} - \delta_e(q, a)(q') \right| \geq \eta/2$, that is the measured and theoretical frequency of t are too far apart. It is a *good* (t, η) -classification otherwise. Intuitively, over long histories, good classifications have high probability.

Given MEMDP M , let T_M be a set obtained by selecting one distinguishing transition (q, a, q') for each state-action pair (q, a) for which such transitions exist (i.e., such that $\delta_e(q, a) \neq \delta_f(q, a)$ for some environments e, f). Let $|h|_{T_M} = \sum_{(q,a,q') \in T_M} |h|_{q,a}$, and fix η such that $\eta < \frac{1}{2} \min(\{|\delta_e(q, a)(q') - \delta_f(q, a)(q')| \mid e, f \in E, q, q' \in Q, a \in A\} \setminus \{0\})$.

Let us define the set of *good histories with n_0 samples*, denoted by Good_{n_0} , as the set of histories h satisfying:

1. $|h|_{T_M} \geq n_0$,
2. for all $t = (q, a, \cdot) \in T_M$ satisfying $|h|_{q,a} = \max_{(q',a',\cdot) \in T_M} |h|_{q',a'}$, the history h is a good (t, η) -classification.

► **Lemma 15.** *Consider an MEMDP M whose only non-distinguishing MCECs are trivial, and fix $\varepsilon > 0$, Let $n_0 = \lceil \frac{2(|Q||A|)^3}{\varepsilon\eta^2} \rceil$, and $n \geq 2p^{-2|Q|} \max(\log(\frac{4}{\varepsilon}), n_0)$ where p is the smallest*

nonzero probability that appears in M . Then, from any starting state, and under any strategy, with probability at least $1 - \varepsilon$, within n steps, the history either visits a MCEC (distinguishing or trivial), or belongs to Good_{n_0} .

Lemma 15 shows that long histories that do not enter MCECs have high probability of being good classifications, which means that the frequency of the distinguishing transitions is close to their actual probability. Technically, this result is nontrivial to prove. In fact, successive visits of distinguishing transitions are not independent trials so classical reasoning based, for instance, on Hoeffding's inequality cannot be applied. It turns out that although these trials can be dependent, their covariances are 0. We exploit this fact to prove Lemma 15 by carefully modeling the successive trials, and relying on Chebyshev's inequality.

Constructing Approximate Finite-Memory Strategies. We are now ready to construct a finite-memory strategy that approximates an arbitrary strategy σ . We first construct a finite-memory strategy for $\text{purge}(M)$ and then transfer it to M using Lemma 14. The finite-memory strategy imitates σ for n steps, where n is defined in Lemma 15. By this lemma, because all nontrivial MCECs of $\text{purge}(M)$ are distinguishing, when we play for n steps, with high probability, we either visit a trivial MCEC (which is either winning for all environments or losing for all environments), or reach a distinguishing MCEC, or observe enough samples of distinguishing transitions. The strategy is extended arbitrarily in trivial MCECs. Inside distinguishing MCECs, it gathers samples of distinguishing transitions as in Lemma 7, which improves the knowledge (with an arbitrarily small probability of error). The knowledge is also correctly improved with high probability if enough samples are gathered outside of MCECs. In both cases, the strategy switches to a finite-memory strategy for the improved knowledge constructed recursively for smaller sets of environments.

Lemma 16 formalizes this reasoning and gives a bound N on the memory of the resulting strategy. In the memory bound, the term $\lceil 8 \frac{\log(8/\varepsilon)}{\eta^2} \rceil$ comes from the application of Lemma 7 for distinguishing MCECs for each subset of E ; and the term $(2|Q|)^{n(|E|+1)}$ corresponds to the recursive analysis, since the strategy is defined inductively for each subset of E . Notice that n is exponential, and that N is doubly-exponential in the size of the input.

► **Lemma 16.** *Consider an MEMDP $M = \langle Q, A, (\delta_e)_{e \in E} \rangle$, state $q \in Q$, parity objective φ . For all strategies σ , and $\varepsilon > 0$, there exists a strategy σ' using at most $N = (2|Q|)^{n(|E|+1)} |A| \lceil 8 \frac{\log(8/\varepsilon)}{\eta^2} \rceil$ memory where $n = \left\lceil 2p^{-2|Q|} \max\left(\frac{8(|Q||A|)^3}{\varepsilon\eta^2}, \log(16/\varepsilon)\right) \right\rceil$, with p the smallest nonzero probability in M , and that satisfies $\mathbb{P}_q^{\sigma'}(M, \varphi) \geq \mathbb{P}_q^\sigma(M, \varphi) - \varepsilon$.*

Approximation Algorithm. We now sketch a procedure for solving the gap problem with threshold α for parity objectives in MEMDPs. Informally, when the memory size is fixed, one can define a system of polynomial equations of linear size that describe the reachability probabilities for a target set in MDPs [3]; and this can be extended easily to MEMDPs. Here, given any $\varepsilon > 0$, we consider the bound N of Lemma 16, and solve the polynomial equation system whose solutions correspond to finite-memory strategies with N memory states. Solving such equations can be done in polynomial space [5] in the size of the equations, which means a double exponential-space procedure here because N is double exponential.

More precisely, the equation system is written so that any solution corresponds to a strategy achieving the probability threshold $\alpha - \varepsilon$. If some strategy achieves the threshold α , by Lemma 16, we know that there is one with a memory size N , and the equation system has a solution. Conversely, if no strategy achieves the threshold $\alpha - \varepsilon$, then, in particular, no finite-memory strategy achieves this threshold and the equation system has no solution.

The case of parity objectives is solved by enumerating the supports of the strategies that are being guessed since the supports determine bottom strongly connected components, and the problem can then be reduced to reachability.

► **Theorem 17.** *The gap problem can be solved in double exponential space for MEMDPs with parity objectives.*

Note that the double exponential complexity remains even if $|E|$ is fixed. Whether the complexity can be improved when $|E|$ is fixed, and the decidability of the value problem (and not just the ε -gap problem) are interesting directions for future research.

References

- 1 T. S. Badings, T. D. Simão, M. Suilen, and N. Jansen. Decision-making under uncertainty: beyond probabilities. *Int. J. Softw. Tools Technol. Transf.*, 25(3):375–391, 2023.
- 2 C. Baier, M. Größer, and N. Bertrand. Probabilistic ω -automata. *J. ACM*, 59(1):1, 2012.
- 3 C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
- 4 P. Buchholz and D. Scheftelowitsch. Computation of weighted sums of rewards for concurrent MDPs. *Math. Methods Oper. Res.*, 89(1):1–42, 2019.
- 5 J. Canny. Some algebraic and geometric computations in PSPACE. In *Proc. of STOC: Symposium on Theory of Computing*, page 460–467. ACM, 1988.
- 6 K. Chatterjee, M. Chmelík, D. Karkhanis, P. Novotný, and A. Royer. Multiple-environment Markov decision processes: Efficient analysis and applications. In *Proc. of ICAPS: Automated Planning and Scheduling*, pages 48–56. AAAI Press, 2020.
- 7 K. Chatterjee, L. Doyen, H. Gimbert, and T. A. Henzinger. Randomness for free. *Information and Computation*, 245:3–16, 2017.
- 8 K. Chatterjee, L. Doyen, J.-F. Raskin, and O. Sankur. The value problem for multiple-environment MDPs with parity objective. *CoRR*, abs/2504.15960, 2025.
- 9 C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *J. ACM*, 42(4):857–907, July 1995.
- 10 L. de Alfaro. *Formal verification of probabilistic systems*. Ph.d. thesis, Stanford University, 1997.
- 11 S. Even, A. L. Selman, and Y. Yacobi. The complexity of promise problems with applications to public-key cryptography. *Information and Control*, 61(2):159 – 173, 1984.
- 12 E. A. Feinberg and A. Shwartz, editors. *Handbook of Markov Decision Processes - Methods and Applications*. Kluwer, 2002.
- 13 H. Gimbert and Y. Oualhadj. Probabilistic automata on finite words: Decidable and undecidable problems. In *Proc. of ICALP (2)*, LNCS 6199, pages 527–538. Springer, 2010.
- 14 O. Goldreich. On promise problems (a survey in memory of Shimon Even [1935-2004]). *Manuscript*, 2005.
- 15 W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- 16 K. J. Åström. Optimal control of Markov processes with incomplete state information I. *Journal of Mathematical Analysis and Applications*, 10:174–205, 1965.
- 17 O. Madani, S. Hanks, and A. Condon. On the undecidability of probabilistic planning and related stochastic optimization problems. *Artif. Intell.*, 147(1-2):5–34, 2003.
- 18 D. A. Martin. The determinacy of Blackwell games. *The Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
- 19 C. H. Papadimitriou and J. N. Tsitsiklis. The complexity of Markov decision processes. *Math. Oper. Res.*, 12(3):441–450, 1987.
- 20 M. L. Puterman. *Markov decision processes*. John Wiley and Sons, 1994.
- 21 N. M. van Dijk R. J. Boucherie. *Markov decision processes in practice*. Springer, 2017.

- 22 J.-F. Raskin and O. Sankur. Multiple-environment Markov decision processes. In *Proc. of FSTTCS: Foundation of Software Technology and Theoretical Computer Science*, volume 29 of *LIPICs*, pages 531–543. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2014.
- 23 M. Suilen, M. van der Vegt, and S. Junges. A PSPACE algorithm for almost-sure Rabin objectives in multi-environment MDPs. In *Proc. of CONCUR: Concurrency Theory*, volume 311 of *LIPICs*, pages 40:1–40:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.
- 24 M. van der Vegt, N. Jansen, and S. Junges. Robust almost-sure reachability in multi-environment MDPs. In *Proc. of TACAS: Tools and Algorithms for the Construction and Analysis of Systems*, LNCS 13993, pages 508–526. Springer, 2023.