

Threshold Constraints with Guarantees for Parity Objectives in Markov Decision Processes^{*†}

Raphaël Berthon¹, Mickael Randour², and Jean-François Raskin³

1 ENS Rennes, Rennes, France
raphael.berthon@ens-rennes.fr

2 Computer Science Department, ULB – Université libre de Bruxelles, Brussels, Belgium
mickael.randour@gmail.com

3 Computer Science Department, ULB – Université libre de Bruxelles, Brussels, Belgium
jraskin@ulb.ac.be

Abstract

The *beyond worst-case synthesis problem* was introduced recently by Bruyère et al. [10]: it aims at building system controllers that provide strict worst-case performance guarantees against an antagonistic environment while ensuring higher expected performance against a stochastic model of the environment. Our work extends the framework of [10] and follow-up papers, which focused on quantitative objectives, by addressing the case of ω -regular conditions encoded as parity objectives, a natural way to represent functional requirements of systems.

We build strategies that satisfy a main parity objective on all plays, while ensuring a secondary one with sufficient probability. This setting raises new challenges in comparison to quantitative objectives, as one cannot easily mix different strategies without endangering the functional properties of the system. We establish that, for all variants of this problem, deciding the existence of a strategy lies in $\text{NP} \cap \text{coNP}$, the same complexity class as classical parity games. Hence, our framework provides additional modeling power while staying in the same complexity class.

1998 ACM Subject Classification G.3 Probability and Statistics

Keywords and phrases Markov decision processes, parity objectives, beyond worst-case synthesis

Digital Object Identifier 10.4230/LIPIcs.ICALP.2017.121

1 Introduction


Beyond worst-case synthesis. *Two-player zero-sum games* [18, 21] and *Markov decision processes* (MDPs) [17, 4] are popular frameworks for decision making in adversarial and uncertain environments respectively. In the former, a system controller (player 1) and its environment (player 2) compete antagonistically, and synthesis aims at strategies that ensure a specified behavior *against all possible strategies of the environment*. In the latter, the system is faced with a given stochastic model of its environment, and the focus is on satisfying a given level of expected performance, or a *specified behavior with a sufficient probability*.

The *beyond worst-case synthesis* framework [10] unites both views: we look for strategies that provide both *strict worst-case guarantees* and a *good level of performance against the*

* Full version is available on arXiv [5], <http://arxiv.org/abs/1702.05472>.

† Work partially supported by the ERC Starting grant 279499 (inVEST) and the ARC project “Non-Zero Sum Game Graphs: Applications to Reactive Synthesis and Beyond” (Fédération Wallonie-Bruxelles). J.-F. Raskin is Professeur Francqui de Recherche, M. Randour is an F.R.S.-FNRS postdoctoral researcher.





■ **Figure 1** An MDP where player 1 can ensure p_1 surely and p_2 almost-surely.

stochastic model. Such requirements are natural in practical situations (e.g., see [9, 23] for applications to the shortest path problem). The original paper [10] dealt with mean-payoff and shortest path objectives. Follow-up work include, e.g., multi-dimensional extensions [14], optimization of the expected mean-payoff under hard Boolean constraints [1] or under energy constraints [7], or integration of beyond worst-case concepts in the tool UPPAAL [15].

Parity objectives. We study the beyond worst-case problem for ω -regular conditions encoded as *parity objectives*. Parity games have been under close scrutiny for a long time due to their importance (e.g., they subsume modal μ -calculus model checking [16]) and their intriguing complexity: they belong to the class of problems in $\text{NP} \cap \text{coNP}$ [19] and despite many efforts (see [11] for pointers), whether they belong to P is still an open question.

In the aforementioned papers dealing with beyond worst-case problems, the focus was on *quantitative* objectives (e.g., mean-payoff). While it is usually the case that *qualitative* objectives, such as parity, are easier to deal with than quantitative ones, this is not true in the setting considered in this paper. Indeed, in the context of quantitative objectives, it is conceivable to alternate between two strategies along a play, such that one – efficient – strategy balances the performance loss due to playing the other – less efficient – strategy for a limited stretch of play infinitely often. In the context of qualitative objectives, this is no more possible in general, as one strategy may induce behaviors (such as invalidating the parity condition infinitely often) that can never be counteracted by the other one. Hence, in comparison, we need to define more elaborate analysis techniques to detect when satisfying *both* the worst-case and the probabilistic constraints with a single strategy is actually possible.

Example. Consider the MDP of Figure 1. Circle states are owned by player 1 (system) and square states are owned by player 2 (environment). In the stochastic model of the environment, square states are probabilistic, and, when not specified, we consider the uniform distribution over their successors. Each state is labelled with a name and two integers x, y representing priorities defined by two functions, p_1 and p_2 . An infinite path in the graph is winning for player 1 and parity objective p_i , $i \in \{1, 2\}$, if the *maximal* priority seen infinitely often along the path for function p_i is *even*. We claim that player 1 has a strategy λ to ensure that (i) all plays consistent with λ satisfy p_1 (i.e., p_1 is surely satisfied) and (ii) the probability measure induced by λ on this MDP ensures that p_2 is satisfied with probability one (i.e., almost-surely).

One such λ is as follows. It plays an infinite sequence of rounds of n_i steps, $i \in \mathbb{N}$. In round i , in state a , the strategy chooses b for n_i steps, such that the probability to reach c during round i is larger than $1 - 2^{-i}$ (this is possible as at each step c is reached from b with probability $\frac{1}{2}$). If during round i , c is not reached (which can happen with a small probability) then λ goes to d once. Then the next round $i+1$ is started. This infinite-memory strategy ensures both (i) and (ii). Indeed, it can be shown that the probability that λ plays d infinitely often is zero. Also, during each round, the maximal priority for p_1 is guaranteed to be even because if c is not visited, d is systematically played.

Finally, we can prove that player 1 needs infinite memory to ensure p_1 surely and p_2 almost-surely, and also, that this is the best that player 1 can do here: he has no strategy to enforce surely both p_1 and p_2 at the same time.

Outline and contributions. We consider MDPs with two parity objectives (i.e., using different priority functions). We study the problem of deciding the existence of a strategy that ensures the first parity objective *surely* (i.e., on all plays) while yielding a *probability at least equal to (resp. greater than) a given rational threshold* to satisfy the second parity objective. In Section 2, we formally define the framework and recall important results from the literature. In Section 3, as an intermediate step, we solve the problem of ensuring the first parity objective surely while visiting a target set of states with sufficient probability: this tool will help us several times later. We prove that the corresponding decision problem is in $\text{NP} \cap \text{coNP}$ and at least as hard as parity games, and that finite-memory strategies are sufficient. In Section 4, we solve the problem for the two parity objectives, where the second one must hold *almost-surely* (i.e., with probability one). Our main tools are the novel notion of *ultra-good end-components*, as well as the reachability problem solved in Section 3. We generalize our approach to arbitrary probability thresholds in Section 5, in which we introduce the notion of *very-good end-components*. In both the almost-sure and the arbitrary threshold cases, we prove that the decision problem belongs to $\text{NP} \cap \text{coNP}$ and is at least as hard as parity games. In contrast to the reachability case, we prove that infinite memory is in general necessary. Full proofs are presented in the extended version of this paper [5].

Additional related work. The beyond worst-case synthesis framework illustrates the usefulness of non-zero-sum games for reactive synthesis [8, 22]. Other types of multi-objective specifications in stochastic models have been considered: e.g., percentile queries generalize the classical threshold probability problem to several dimensions [24]. In [3], Baier et al. study the quantitative analysis of MDPs under weak and strong fairness constraints. They provide algorithms for computing the probability for ω -regular properties in worst and best-case scenarios, when considering strategies that in addition satisfy weak or strong fairness constraints almost-surely. In contrast, we are able to consider similar objectives but for strategies that satisfy weak or strong fairness constraints surely, i.e., with certainty and not only with probability one. In [1], Almagor et al. consider the optimization of the expected mean-payoff under hard Boolean constraints in weighted MDPs. Our concept of ultra-good end-component builds upon their notion of super-good one. A reduction to mean-payoff parity games [12] is part of the identification process of both types of end-components.

2 Preliminaries

Directed graphs. A *directed graph* is a pair $G = (S, E)$ with S a set of vertices, called *states*, and $E \subseteq S \times S$ a set of directed edges. We focus here on *finite* graphs (i.e., $|S| < \infty$). Given a state $s \in S$, we denote by $\text{Succ}(s) = \{s' \in S \mid (s, s') \in E\}$ the set of successors of s by edges in E . We assume that graphs are non-blocking, i.e., for all $s \in S$, $\text{Succ}(s) \neq \emptyset$.

A *play* in G from an initial state $s \in S$ is an infinite sequence of states $\pi = s_0 s_1 s_2 \dots$ such that $s_0 = s$ and $(s_i, s_{i+1}) \in E$ for all $i \geq 0$. The *prefix* up to the $(n+1)$ -th state of π is the finite sequence $\pi(0, n) = s_0 s_1 \dots s_n$. We resp. denote the first and last states of a prefix $\rho = s_0 s_1 \dots s_n$ by $\text{First}(\rho) = s_0$ and $\text{Last}(\rho) = s_n$. For a play π , we naturally extend the notation to $\text{First}(\pi)$. Finally, for $i \in \mathbb{N}$, $\pi(i) = s_i$, and for $j > i$, $\pi(i, j) = s_i \dots s_j$. The set of plays of G is $\text{Plays}(G)$ and the set of prefixes is $\text{Pref}(G)$. For a set of plays Π , we

denote by $\text{Pref}(\Pi)$ the set of prefixes of these plays. Given two prefixes $\rho = s_0 \dots s_m$ and $\rho' = s'_0 \dots s'_n$ in $\text{Pref}(G)$, we denote their *concatenation* as $\rho \cdot \rho' = s_0 \dots s_m s'_0 \dots s'_n$. This is not necessarily a valid prefix of G . The same holds for a prefix concatenated with a play.

Probability distributions. Given a countable set A , a (rational) *probability distribution* on A is a function $p: A \rightarrow [0, 1] \cap \mathbb{Q}$ such that $\sum_{a \in A} p(a) = 1$. We write $\mathcal{D}(A)$ the set of probability distributions on A . The *support* of $p \in \mathcal{D}(A)$ is $\text{Supp}(p) = \{a \in A \mid p(a) > 0\}$.

Markov decision processes. An MDP is a tuple $\mathcal{M} = (G, S_1, S_2, \delta)$ where (i) $G = (S, E)$ is a directed graph; (ii) (S_1, S_2) is a partition of S into states of player 1 (denoted by \mathcal{P}_1 and representing the system) and states of player 2 (denoted by \mathcal{P}_2 and representing the *stochastic* environment); (iii) $\delta: S_2 \rightarrow \mathcal{D}(S)$ is the transition function that, given a stochastic state $s \in S_2$, defines the probability distribution $\delta(s)$ over the successors of s , such that for all $s \in S_2$, $\text{Supp}(\delta(s)) = \text{Succ}(s)$. An MDP where for all $s \in S_1$, $|\text{Succ}(s)| = 1$ is a fully-stochastic process called a *Markov chain* (MC). A prefix $\rho \in \text{Pref}(\mathcal{M})$ belongs to \mathcal{P}_i , $i \in \{1, 2\}$, if $\text{Last}(\rho) \in S_i$. The set of prefixes that belong to \mathcal{P}_i is denoted by $\text{Pref}_i(\mathcal{M})$.

Strategies. A *strategy* for \mathcal{P}_1 is a function $\lambda: \text{Pref}_1(\mathcal{M}) \rightarrow \mathcal{D}(S)$, such that for all $\rho \in \text{Pref}_1(\mathcal{M})$, we have $\text{Supp}(\lambda(\rho)) \subseteq \text{Succ}(\text{Last}(\rho))$. The set of all strategies in \mathcal{M} is denoted by Λ . *Pure* strategies have their support equal to a singleton for all prefixes. We mention that a strategy is *randomized* to stress on the need for randomness in general.

A strategy λ for \mathcal{P}_1 can be encoded by a stochastic state machine with outputs, called *stochastic Moore machine*, \mathbb{M} . A strategy λ is *finite-memory* if \mathbb{M} is finite, and *memoryless* if it has only one state. That is, it does not depend on the history but only on the current state of the MDP: in this case, we have that $\lambda: S_1 \rightarrow \mathcal{D}(S)$. Finally, if the same strategy can be used regardless of the initial state, we say that a *uniform* strategy exists.

A play π is *consistent* with a strategy λ if for all $n \geq 0$ such that $\pi(n) \in S_1$, we have that $\pi(n+1) \in \text{Supp}(\lambda(\pi(0, n)))$. It is defined similarly for prefixes. We write $\text{Out}^{\mathcal{M}}(\lambda) \subseteq \text{Plays}(G)$ the set of plays consistent with λ . We use $\text{Out}_s^{\mathcal{M}}(\lambda)$ when fixing an initial state s .

Markov chain induced by a strategy. An MDP $\mathcal{M} = (G = (S, E), S_1, S_2, \delta)$ and a strategy λ for \mathcal{P}_1 determine an MC $\mathcal{C} = (G', \delta')$. Given $s \in S$ an initial state and $\mathcal{A} \subseteq \text{Plays}(G)$ a measurable set, we denote by $\mathbb{P}_{\mathcal{M}, s}^{\lambda}[\mathcal{A}]$ the *probability* of event \mathcal{A} when \mathcal{M} is executed with initial state s and strategy λ .

Objectives. Given an MDP $\mathcal{M} = (G, S_1, S_2, \delta)$, an *objective* is a set of plays $\mathcal{A} \subseteq \text{Plays}(G)$. We consider two classical objectives from the literature. Both define measurable events. To define them, we introduce the following notation: given a play $\pi \in \text{Plays}(G)$, let $\text{inf}(\pi) = \{s \in S \mid \forall i \geq 0, \exists j \geq i, \pi(j) = s\}$ be the set of states seen infinitely often along π .

Reachability. Given a target $T \subseteq S$, this objective asks for plays that visit T : $\text{Reach}(T) = \{\pi \in \text{Plays}(G) \mid \exists n \geq 0, \pi(n) \in T\}$. We later use the LTL notation $\diamond T$ for event $\text{Reach}(T)$.

Parity. Let $p: S \rightarrow \{1, 2, \dots, d\}$ be a *priority function* that maps each state to an integer priority, where $d \leq |S| + 1$ (w.l.o.g.). The parity objective asks that, among the priorities seen infinitely often, the *maximal* one be even: $\text{Parity}(p) = \{\pi \in \text{Plays}(G) \mid \max_{s \in \text{inf}(\pi)} p(s) \text{ is even}\}$. We later simply use p to denote the event $\text{Parity}(p)$.

End-components and sub-MDPs. Let $\mathcal{M} = (G = (S, E), S_1, S_2, \delta)$ be an MDP. An *end-component* (EC) of \mathcal{M} is a set $C \subseteq S$ such that (i) $\forall s \in C \cap S_2, \text{Succ}(s) \subseteq C$ and $\forall s \in C \cap S_1,$

$\text{Succ}(s) \cap C \neq \emptyset$; and (ii) C is *strongly connected*, i.e., for any two states $s, s' \in C$, there exists a path from s to s' that stays in C . It is well-known that inside an EC C , \mathcal{P}_1 can force the visit of any state $s \in C$ with probability 1 (that is, when \mathcal{P}_2 is seen as stochastic and obeys the strategy δ), see e.g., [4]. The union of two ECs with non-empty intersection is an EC. An EC C is thus *maximal* if, for every EC C' , $C' \subseteq C \vee C' \cap C = \emptyset$.

Given an EC $C \subseteq S$ of \mathcal{M} , we write $\mathcal{M}_{|C}$ the *sub-MDP* defined by $\mathcal{M}_{|C} = (G' = (C, E \cap C \times C), S'_1 = S_1 \cap C, S'_2 = S_2 \cap C, \delta')$, where $\delta': S'_2 \rightarrow \mathcal{D}(C)$ is the restriction of δ to the domain C . Note that $\mathcal{M}_{|C}$ is a well-defined MDP: it has no deadlock since C is strongly connected and in all stochastic states s , $\text{Supp}(\delta'(s)) \subseteq C$ (as C was an EC in \mathcal{M}).

Technical lemma. We recall a classical result about MDPs that will be useful later on.

► **Lemma 1** (Optimal reachability [4]). *Given an MDP $\mathcal{M} = (G = (S, E), S_1, S_2, \delta)$ and a target set $T \subseteq S$, we can compute for each state $s \in S$ the maximal probability $v_s^* = \sup_{\lambda \in \Lambda} \mathbb{P}_{\mathcal{M}, s}^\lambda[\diamond T]$ to reach T , in polynomial time. There is an optimal uniform pure memoryless strategy λ^* that enforces v_s^* from all $s \in S$. Now, fix $s \in S$ and $c \in \mathbb{Q}$ such that $c < v_s^*$. Then there exists $k \in \mathbb{N}$ such that by playing λ^* from s for k steps, we reach T with probability larger than c .*

Events and probabilistic operators. Consider an MDP $\mathcal{M} = (G = (S, E), S_1, S_2, \delta)$. Recall that we have defined two types of measurable events (specific subsets of $\text{Plays}(G)$) with respective notations $\diamond T$ for $T \subseteq S$ (reachability), and p for $p: S \rightarrow \{1, \dots, d\}$ a priority function (parity). We define three operators to reason about the probabilities of these events: \mathbf{S} , $\mathbf{P}_{\sim c}$, and \mathbf{AS} . Given an event \mathcal{A} and a state s , they are used as follows:

- \mathcal{A} is *sure* from s , denoted $s \models \mathbf{S}(\mathcal{A})$, if there exists a strategy λ of \mathcal{P}_1 such that $\text{Out}_s^{\mathcal{M}}(\lambda) \subseteq \mathcal{A}$. Here probabilities are ignored and we consider \mathcal{P}_2 as antagonistic.
- \mathcal{A} holds with probability at least equal to (resp. greater than) $c \in \mathbb{Q}$ from s , denoted $s \models \mathbf{P}_{\geq c}(\mathcal{A})$ (resp. $s \models \mathbf{P}_{> c}(\mathcal{A})$) if there exists λ such that $\mathbb{P}_{\mathcal{M}, s}^\lambda[\mathcal{A}] \geq c$ (resp. $> c$).
- \mathcal{A} is *almost-sure* from s , denoted $s \models \mathbf{AS}(\mathcal{A})$, if there exists λ such that $\mathbb{P}_{\mathcal{M}, s}^\lambda[\mathcal{A}] = 1$.

For any operator \mathbf{O} , we say that such a λ is a *witness strategy* for $s \models \mathbf{O}(\mathcal{A})$ and we write $s, \lambda \models \mathbf{O}(\mathcal{A})$ to denote it. We will also consider *combinations* of the type $s \models \mathbf{O}_1(\mathcal{A}_1) \wedge \mathbf{O}_2(\mathcal{A}_2)$ for two operators and events: in this case, we require that the same strategy be a witness for both conjuncts, i.e., that there exists λ such that $s, \lambda \models \mathbf{O}_1(\mathcal{A}_1)$ and $s, \lambda \models \mathbf{O}_2(\mathcal{A}_2)$. Finally, we will sometimes use different MDPs, in which case we add the considered MDP \mathcal{M} as a subscript on \models , e.g., $s \models_{\mathcal{M}} \mathbf{O}(\mathcal{A})$. We drop this subscript when the context is clear.

Beyond worst-case problems. Let $\mathcal{M} = (G = (S, E), S_1, S_2, \delta)$ be an MDP, $s \in S$ be an initial state, and p_1, p_2 be two priority functions. We provide algorithms to decide the existence of a witness strategy — and synthesize it — for the following formulae: (i) $s \models \mathbf{S}(p_1) \wedge \mathbf{AS}(p_2)$, and (ii) $s \models \mathbf{S}(p_1) \wedge \mathbf{P}_{\sim c}(p_2)$ for $\sim \in \{>, \geq\}$ and $c \in \mathbb{Q} \cap [0, 1)$.

3 Reachability under parity constraints

We study two variants, given $s \in S$, $T \subseteq S$, and $p: S \rightarrow \{1, \dots, d\}$: (i) $s \models \mathbf{S}(p) \wedge \mathbf{AS}(\diamond T)$, and (ii) $s \models \mathbf{S}(p) \wedge \mathbf{P}_{\sim c}(\diamond T)$ for $\sim \in \{>, \geq\}$ and $c \in \mathbb{Q} \cap [0, 1)$.

Almost-sure reachability. This case can be solved by reduction to a slight variant studied in [2, Lemma 3] (extended version of [1]). The approach of [2, Lemma 3] relies on a reduction

to a Büchi-parity game: sufficiency of finite memory follows from this reduction. The lower complexity bound is trivial: it suffices to fix $T = S$ to obtain a classical parity game [19].

► **Theorem 2.** *Given an MDP $\mathcal{M} = (G = (S, E), S_1, S_2, \delta)$, a state $s_0 \in S$, a priority function $p: S \rightarrow \{1, \dots, d\}$, and a target set of states $T \subseteq S$, it can be decided in $\text{NP} \cap \text{coNP}$ if $s_0 \models \mathbf{S}(p) \wedge \mathbf{AS}(\diamond T)$. If the answer is YES, then there exists a finite-memory witness strategy. This decision problem is at least as hard as solving parity games.*

Reachability with threshold probability. We first study strategies that maximize the probability of reaching a target $T \subseteq S$ in an MDP \mathcal{M} . By Lemma 1, we have an optimal uniform pure memoryless strategy λ^* that enforces v_s^* from all $s \in S$. We define the set $E^{-\text{opt}} = \{(s, s') \in E \mid s \in S_1 \wedge v_s^* > v_{s'}^*\}$ that contains all edges that are *non-optimal choices* for \mathcal{P}_1 in the sense that they result in a strict decrease of the probability to reach T . We show that playing, for a finite number of steps, edges that are optimal (i.e., in $E^{\text{opt}} = E \setminus E^{-\text{opt}}$), and then switching to an optimal strategy, like λ^* , produces an optimal strategy too.

► **Lemma 3.** *Let λ^* be an optimal uniform pure memoryless strategy in \mathcal{M} to reach T , from all states in S . If λ is a strategy that plays only edges in E^{opt} for m steps, for $m \in \mathbb{N}$, and then switches to λ^* , then λ is also optimal to reach T from all states in S .*

We now turn to the problem $s_0 \models \mathbf{S}(p) \wedge \mathbf{P}_{\sim c}(\diamond T)$ and establish the following result.


► **Theorem 4.** *Given an MDP $\mathcal{M} = (G = (S, E), S_1, S_2, \delta)$, a state $s_0 \in S$, a priority function $p: S \rightarrow \{1, \dots, d\}$, a target set of states $T \subseteq S$, and a probability threshold $c \in [0, 1) \cap \mathbb{Q}$, it can be decided in $\text{NP} \cap \text{coNP}$ if $s_0 \models \mathbf{S}(p) \wedge \mathbf{P}_{\sim c}(\diamond T)$ for $\sim \in \{>, \geq\}$. If the answer is YES, then there exists a finite-memory witness strategy. This decision problem is at least as hard as solving parity games.*

Proof Sketch. First, we restrict \mathcal{M} to the \subseteq -maximal sub-MDP \mathcal{M}^w in which \mathcal{P}_1 can ensure $\mathbf{S}(p_1)$ from all states, by solving a classical parity game, which is in $\text{NP} \cap \text{coNP}$ [19]. Indeed, if $s_0 \not\models \mathbf{S}(p)$, the answer is NO. In \mathcal{M}^w , \mathcal{P}_1 has a uniform pure memoryless strategy λ^p that ensures $\mathbf{S}(p)$ from every state.

The case $> c$ is the easier. First, we compute the maximal probability $v_{s_0}^*$ to reach T and an optimal strategy λ^* , in polynomial time (Lemma 1). If $v_{s_0}^* \leq c$, then the answer is clearly NO. Otherwise, we claim it is YES. We construct a witness strategy λ for $s_0 \models \mathbf{S}(p) \wedge \mathbf{P}_{> c}(\diamond T)$ from λ^* and λ^p as follows. Starting in s_0 , the strategy λ plays as λ^* for k steps where k is taken as in Lemma 1: the probability to reach T after k steps is strictly greater than c , which implies that $s_0, \lambda \models \mathbf{P}_{> c}(\diamond T)$. Then, λ switches to λ^p . Since parity is prefix-independent, we have that $s_0, \lambda \models \mathbf{S}(p)$, and we are done. Our procedure lies in $\mathbf{P}^{\text{NP} \cap \text{coNP}} = \text{NP} \cap \text{coNP}$ [6], and λ is finite-memory since λ^* and λ^p are memoryless and k is finite.

We now turn to case $\geq c$. We compute $v_{s_0}^*$ in polynomial time. If $v_{s_0}^* > c$, then we answer YES as we apply the same reasoning as in the previous case. If $v_{s_0}^* < c$, then we trivially answer NO. The more involved case is $v_{s_0}^* = c$. We must verify that probability c is still achievable if, in addition, it is required to enforce $\mathbf{S}(p)$. To answer this, we modify \mathcal{M}^w and we reduce our problem to the almost-sure case of Theorem 2. Intuitively, we construct the MDP \mathcal{M}' as follows. (i) We enrich states with one bit that records if T has been visited. (ii) While T has not been visited, we suppress all edges controlled by \mathcal{P}_1 that are not optimal for reachability, i.e., all edges in $E^{-\text{opt}}$. (iii) While T has not been visited, we delete all states that cannot reach T and normalize the probability of the edges that survive this deletion.

We prove that $s'_0 \models_{\mathcal{M}'} \mathbf{S}(p) \wedge \mathbf{AS}(\diamond T') \iff s_0 \models_{\mathcal{M}^w} \mathbf{S}(p) \wedge \mathbf{P}_{\geq c}(\diamond T)$ holds, where s'_0 is the initial state in \mathcal{M}' and T' the translation of T . The crux is the restriction to E^{opt}



■ **Figure 2** This MDP is a UGEC: going to d satisfies (1_U) , whereas going to b satisfies (2_U) .

before visiting T : \mathcal{P}_1 must be able to ensure $S(p_1)$ while using only edges that are optimal for reachability if T cannot be forced surely. As the almost-sure problem is in $\text{NP} \cap \text{coNP}$ by Theorem 2, and \mathcal{M}' is polynomially larger than \mathcal{M}^w (and thus \mathcal{M}), we obtain the claimed complexity using $\text{P}^{\text{NP} \cap \text{coNP}} = \text{NP} \cap \text{coNP}$ [6]. Our reduction implies that the witness strategy can be finite-memory. Again, this problem generalizes parity games by taking $T = S$. ◀

4 Almost-sure parity under parity constraints

Overview and key lemma. We consider an MDP $\mathcal{M} = (G = (S, E), S_1, S_2, \delta)$ with two priority functions p_1 and p_2 . We look at the problem $s \models S(p_1) \wedge \text{AS}(p_2)$. The cornerstone of our approach is the notion of *ultra-good end-component*.

► **Definition 5.** An end-component C of \mathcal{M} is *ultra-good* (UGEC) if in the sub-MDP $\mathcal{M}_{|C}$, the following two properties hold:

- $(1_U) \forall s \in C, s \models_{\mathcal{M}_{|C}} S(p_1) \wedge \text{AS}(\diamond C_{\text{even}}^{\max}(p_1))$, where

$$C_{\text{even}}^{\max}(p_i) = \{s \in C \mid (p_i(s) \text{ is even}) \wedge (\forall s' \in C, p_i(s') \text{ is odd} \implies p_i(s') < p_i(s))\}$$

contains the states with even priorities that are larger than any odd priority in C (this set can be empty for arbitrary ECs but needs to be non-empty for UGECs);

- $(2_U) \forall s \in C, s \models_{\mathcal{M}_{|C}} \text{AS}(p_1) \wedge \text{AS}(p_2)$, or equivalently, $s \models_{\mathcal{M}_{|C}} \text{AS}(p_1 \cap p_2)$.

We introduce the following notations: $\text{UGEC}(\mathcal{M})$ is the set of all UGECs of \mathcal{M} , and $\mathcal{U} = \cup_{U \in \text{UGEC}(\mathcal{M})} U$ is the set of states that belong to a UGEC in \mathcal{M} .

Intuitively, within a UGEC, \mathcal{P}_1 has a strategy to almost-surely visit $C_{\text{even}}^{\max}(p_1)$ while guaranteeing $S(p_1)$, and he also has a (generally different) strategy that almost-surely ensures both parity objectives. Figure 2 gives an example of UGEC. This notion strengthens the concept of *super-good* EC from [1]: essentially, the super-good ECs are exactly the ECs satisfying (1_U) . Thus, every UGEC is a super-good EC, but the converse is false.

The central lemma underpinning our approach is the following.

► **Lemma 6.** *The following equivalence holds:*

$$s_0 \models S(p_1) \wedge \text{AS}(\diamond \mathcal{U}) \iff s_0 \models S(p_1) \wedge \text{AS}(p_2).$$

Essentially, this lemma permits to reduce the problem under study to the one treated in Theorem 2, provided that we are able to compute \mathcal{U} , the set of states appearing in a UGEC. The rest of this section is dedicated to the proof of this lemma and its consequences.

Left-to-right implication (sufficient condition). We first study witness strategies for conditions (1_U) and (2_U) of Definition 5. For (1_U) , it was shown in the proof of [2, Lemma 3] (extended version of [1]) that deciding if the condition holds is in $\text{NP} \cap \text{coNP}$ and that uniform finite-memory witness strategies exist. For (2_U) , we establish the following lemma.

► **Lemma 7.** *Let C be an EC of \mathcal{M} . The following assertions hold.*

1. *It can be decided in polynomial time if condition (2_U) holds.*
2. *If it holds, then there exists a (uniform randomized) memoryless witness strategy $\lambda_{2,C}$ and a sub-EC $D \subseteq C$ such that $D_{\text{even}}^{\max}(p_1) \neq \emptyset$, $D_{\text{even}}^{\max}(p_2) \neq \emptyset$, and for all $s \in C$, we have that $\mathbb{P}_{\mathcal{M}_{1,C},s}^{\lambda_{2,C}} [\{\pi \in \text{Out}^{\mathcal{M}_{1,C}}(\lambda_{2,C}) \mid \inf(\pi) = D\}] = 1$.*
3. *Furthermore, $\lambda_{2,C}$ satisfies the following property: $\forall s \in C, \forall \varepsilon > 0, \exists n \in \mathbb{N}$ such that $\mathbb{P}_{\mathcal{M}_{1,C},s}^{\lambda_{2,C}} [\{\pi \in \text{Out}^{\mathcal{M}_{1,C}}(\lambda_{2,C}) \mid \exists i, 0 \leq i \leq n, \pi(i) \in D_{\text{even}}^{\max}(p_1)\}] \geq 1 - \varepsilon$.*

Proof Sketch. For Point **2**, we resort on a classical result on almost-sure reachability of ECs and almost-sure satisfaction of both parity objectives. For Point **3**, we use Lemma 1 and Point **2**. For Point **1**, we show that (i) the existence of a sub-EC D such that $D_{\text{even}}^{\max}(p_1) \neq \emptyset$ and $D_{\text{even}}^{\max}(p_2) \neq \emptyset$ is not only necessary but also sufficient to satisfy condition (2_U) , and (ii) the existence of such a set can be decided in polynomial time. For (i), it suffices to build a uniform randomized memoryless strategy λ that reaches the sub-EC D almost-surely and then plays uniformly at random in it forever: λ will be a witness for $s \models_{\mathcal{M}_{1,C}} \text{AS}(p_1) \wedge \text{AS}(p_2)$, so condition (2_U) holds in C . For (ii), we first check if $C_{\text{even}}^{\max}(p_1) \neq \emptyset$ and $C_{\text{even}}^{\max}(p_2) \neq \emptyset$. If this holds, then $D = C$ and the answer is YES (it takes linear time obviously). If it does not hold, then we compute the sets $C_{\text{odd}}^{\max}(p_i) = \{s \in C \mid (p_i(s) \text{ is odd}) \wedge (\forall s' \in C, p_i(s') \text{ is even} \implies p_i(s') < p_i(s))\}$ and we iterate this procedure in the sub-EC $C' \subset C$ defined as $C' = C \setminus \text{Attr}_2(C_{\text{odd}}^{\max}(p_1) \cup C_{\text{odd}}^{\max}(p_2))$, where Attr_2 is the classical attractor for \mathcal{P}_2 . A suitable D exists if and only if this procedure stops before $C' = \emptyset$. In addition, this procedure takes at most $|C|$ iterations (as we remove at least one state at each step) and each iteration takes linear time. ◀

We will now prove that inside any UGEC, there is a strategy for $\mathbf{S}(p_1) \wedge \text{AS}(p_2)$. From now on, let C be a UGEC of \mathcal{M} , $\lambda_{1,C}$ be a uniform finite-memory witness strategy for (1_U) in Definition 5, and $\lambda_{2,C}$ be a uniform randomized memoryless one for (2_U) , additionally satisfying the properties of Lemma 7. We build a strategy λ_C based on $\lambda_{1,C}$ and $\lambda_{2,C}$.

► **Definition 8.** Let $C \in \text{UGEC}(\mathcal{M})$. Let $(n_i)_{i \in \mathbb{N}}$ be a sequence of naturals n_i such that $\mathbb{P}_{\mathcal{M}_{1,C},s}^{\lambda_{2,C}} [\{\pi \in \text{Out}^{\mathcal{M}_{1,C}}(\lambda_{2,C}) \mid \exists i, 0 \leq i \leq n_i, \pi(i) \in D_{\text{even}}^{\max}(p_1)\}] \geq 1 - 2^{-i}$, whose existence is guaranteed by Lemma 7. We build strategy λ_C as follows, starting with $i = 0$.

(a) Play $\lambda_{2,C}$ for n_i steps. Then $i = i + 1$ and go to (b).

(b) If $D_{\text{even}}^{\max}(p_1)$ was visited in phase a), then go to (a).

Else, play $\lambda_{1,C}$ until $C_{\text{even}}^{\max}(p_1)$ is reached and then go to (a).

Observe that λ_C requires infinite memory. In the next lemma, we prove that λ_C is a proper witness for $\mathbf{S}(p_1) \wedge \text{AS}(p_2)$ in the UGEC C .

► **Lemma 9.** *Let $C \in \text{UGEC}(\mathcal{M})$. For all $s \in C$, it holds that $s, \lambda_C \models \mathbf{S}(p_1) \wedge \text{AS}(p_2)$.*

Proof Sketch. First consider $s, \lambda_C \models \mathbf{S}(p_1)$. Fix any $\pi \in \text{Out}_s^{\mathcal{M}_{1,C}}(\lambda_C)$: we will show that $\max_{s' \in \inf(\pi)} p_1(s')$ is even. Three cases are possible: (i) λ_C switches infinitely often between $\lambda_{1,C}$ and $\lambda_{2,C}$, (ii) it eventually plays $\lambda_{1,C}$ forever, and (iii) it eventually plays $\lambda_{2,C}$ forever. In case (i), $C_{\text{even}}^{\max}(p_1)$ is visited infinitely often. Since any state in this set has an even priority higher than any odd priority in C , we are good. In case (ii), we know that $s, \lambda_{1,C} \models \mathbf{S}(p_1)$. By prefix-independence, we are also good. In case (iii), $D_{\text{even}}^{\max}(p_1)$ is visited infinitely often, and, eventually, play π never leaves the sub-EC D . Hence, we are good here too and we conclude that $s, \lambda_C \models \mathbf{S}(p_1)$.

To show that $s, \lambda_C \models \text{AS}(p_2)$, we prove that λ_C almost-surely ends up in playing only $\lambda_{2,C}$ (which ensures $\text{AS}(p_2)$). The crux here is the choice of durations n_i for the strategy: we

can show that the probability to never play $\lambda_{1,C}$ again after round i tends to one when i tends to infinity. This entails the needed result. \blacktriangleleft

We can now prove the left-to-right implication of Lemma 6. For this, assume that for $s_0 \in S$, we have that $\lambda_{\mathcal{U}}$ is a witness for $s_0 \models \mathbf{S}(p_1) \wedge \mathbf{AS}(\diamond\mathcal{U})$, where we recall that \mathcal{U} represents the union of all UGECs of the MDP \mathcal{M} . Note that such a strategy can be finite-memory w.l.o.g. as proved in Theorem 2. We build a global strategy λ as follows.

► **Definition 10.** Based on strategies $\lambda_{\mathcal{U}}$ and λ_C for all $C \in \text{UGEC}(\mathcal{M})$, we build the global strategy λ as follows.

- (a) Play $\lambda_{\mathcal{U}}$ until a UGEC C is reached, then go to (b).
- (b) Play λ_C forever.

This strategy requires infinite memory because of the strategies λ_C . We prove that λ is a witness for $s_0 \models \mathbf{S}(p_1) \wedge \mathbf{AS}(p_2)$.

► **Lemma 11.** *It holds that $s_0, \lambda \models \mathbf{S}(p_1) \wedge \mathbf{AS}(p_2)$.*

Right-to-left implication (necessary condition). We now turn to the converse implication of Lemma 6, i.e., that $s_0 \models \mathbf{S}(p_1) \wedge \mathbf{AS}(p_2)$ implies $s_0 \models \mathbf{S}(p_1) \wedge \mathbf{AS}(\diamond\mathcal{U})$. We start by an intermediate lemma regarding witness strategies: it establishes that all states reachable via such a strategy also satisfy the property.

► **Lemma 12.** *For every state $s \in S$, every strategy λ such that $s, \lambda \models \mathbf{S}(p_1) \wedge \mathbf{AS}(p_2)$, and every prefix $\rho \in \text{Pref}(\text{Out}_s^{\mathcal{M}}(\lambda))$, we have that $\text{Last}(\rho) \models \mathbf{S}(p_1) \wedge \mathbf{AS}(p_2)$.*

The next lemma establishes that at least one UGEC must exist in \mathcal{M} .

► **Lemma 13.** *The following holds: $s_0 \models \mathbf{S}(p_1) \wedge \mathbf{AS}(p_2) \implies \text{UGEC}(\mathcal{M}) \neq \emptyset$.*

Proof Sketch. Given $\Pi \subseteq \text{Plays}(G)$, we define $\text{States}(\Pi) = \{s \in S \mid \exists \pi \in \Pi, \exists n \in \mathbb{N}, \pi(n) = s\}$. We then study the set

$$\mathcal{S} = \{R \subseteq S \mid \exists s \in S, \exists \lambda \in \Lambda, (s, \lambda \models \mathbf{S}(p_1) \wedge \mathbf{AS}(p_2)) \wedge (R = \text{States}(\text{Out}_s^{\mathcal{M}}(\lambda)))\}.$$

Intuitively, it contains any subset of S that captures all states reachable by some witness strategy λ , from some state $s \in S$. First note that $s_0 \models \mathbf{S}(p_1) \wedge \mathbf{AS}(p_2)$ implies that \mathcal{S} is non-empty, as for a witness strategy λ , $R = \text{States}(\text{Out}_{s_0}^{\mathcal{M}}(\lambda)) \in \mathcal{S}$, by definition.

We then show that all minimal elements of \mathcal{S} for set inclusion \subseteq are UGECs, which suffices to establish our lemma. The most important ingredients to prove that any $R \in \min_{\subseteq}(\mathcal{S})$ is a UGEC are the following. First the existence, for any $s \in R$, of a strategy λ_R such that $s, \lambda_R \models_{\mathcal{M}_R} \mathbf{S}(p_1) \wedge \mathbf{AS}(p_2)$ (i.e., λ_R satisfies the property without leaving R), which follows from Lemma 12 and the minimality of R in \mathcal{S} . Second, proving that $R_{\text{even}}^{\max}(p_1)$ is dense in the subtree induced by $\text{Out}_s^{\mathcal{M}}(\lambda_R)$, that is, that for every prefix ρ , the subtree defined by λ_R from ρ reaches a state of $R_{\text{even}}^{\max}(p_1)$, and that this holds in all subsequent subtrees. From this density argument, we can derive a witness strategy for condition (1_U) in Definition 5. \blacktriangleleft

Collecting in $\mathcal{U}_{\min} = \bigcup_{R \in \min_{\subseteq}(\mathcal{S})} R$ all states that belong to minimal sets R of \mathcal{S} , we finally prove the implication.

► **Lemma 14.** *The following holds: $s_0 \models \mathbf{S}(p_1) \wedge \mathbf{AS}(p_2) \implies s_0 \models \mathbf{S}(p_1) \wedge \mathbf{AS}(\diamond\mathcal{U}_{\min})$.*

Algorithm. Lemma 11 and Lemma 14 prove the correctness of the reduction presented in Lemma 6. It is the cornerstone of our algorithm.

► **Theorem 15.** *Given an MDP $\mathcal{M} = (G = (S, E), S_1, S_2, \delta)$, a state $s_0 \in S$, and two priority functions $p_i: S \rightarrow \{1, \dots, d\}$, $i \in \{1, 2\}$, it can be decided in $\text{NP} \cap \text{coNP}$ if $s_0 \models \mathbf{S}(p_1) \wedge \mathbf{AS}(p_2)$. If the answer is YES, then there exists an infinite-memory witness strategy, and infinite memory is in general necessary. This decision problem is at least as hard as solving parity games.*

Proof. The algorithm can be sketched as follows:

1. Compute the set $\max_{\subseteq}(\text{SGEC}(\mathcal{M}))$ of maximal super-good ECs, using [1]. Those are the maximal ECs satisfying condition $(\mathbf{1}_U)$ in Definition 5. There are only polynomially many of them, and their computation is in $\text{NP} \cap \text{coNP}$.
2. For each of them, check if $(\mathbf{2}_U)$ holds using Lemma 7, in polynomial time. If an EC does not satisfy $(\mathbf{2}_U)$, then it is also the case of all its sub-ECs (as seen in the proof of Lemma 7). Hence, we have that $\mathcal{U} = \{C \in \max_{\subseteq}(\text{SGEC}(\mathcal{M})) \mid C \text{ satisfies } (\mathbf{2}_U)\}$.
3. Decide if $s_0 \models \mathbf{S}(p_1) \wedge \mathbf{AS}(\diamond \mathcal{U})$ using Theorem 2. This is in $\text{NP} \cap \text{coNP}$. If it holds, then answer YES, otherwise answer NO.

Its correctness was established in Lemma 6. It belongs to $\text{P}^{\text{NP} \cap \text{coNP}} = \text{NP} \cap \text{coNP}$ [6], and it trivially generalizes classical parity games (e.g., by taking $p_2: s \mapsto 0$ for all $s \in S$).

Finally, let us discuss strategies. A witness strategy λ plays as follows: (i) it plays as the finite-memory strategy witness for $s_0 \models \mathbf{S}(p_1) \wedge \mathbf{AS}(\diamond \mathcal{U})$ given by Theorem 2 until a UGEC C is reached, (ii) then it switches to the infinite-memory strategy λ_C described in Definition 8. It is clear that such a strategy is a witness for $s_0 \models \mathbf{S}(p_1) \wedge \mathbf{AS}(p_2)$, as expected.

Infinite memory is required in general, as shown in the UGEC C in Figure 2: there exists no finite-memory witness strategy in C . Indeed, assume \mathcal{P}_1 is restricted to a finite-memory strategy λ . To be able to ensure p_1 on the play in which \mathcal{P}_2 always goes to c from b , \mathcal{P}_1 must visit d infinitely often, and because of the finite memory of λ , he must do it after a bounded number of steps along which a is not visited: say n steps. Hence, the probability to do it will be bounded from below by a strictly positive constant, here $2^{-\frac{n}{2}}$ (the probability that \mathcal{P}_2 chooses c for $\frac{n}{2}$ times in a row), all along a consistent play. Therefore, \mathcal{P}_1 will almost-surely visit d infinitely often, and p_2 will actually be satisfied with probability zero. ◀

5 Parity with threshold probability under parity constraints

We now turn to the problem $s_0 \models \mathbf{S}(p_1) \wedge \text{P}_{\sim c}(p_2)$ for $\sim \in \{>, \geq\}$ and $c \in \mathbb{Q} \cap [0, 1)$.


Very-good end-components. In addition to UGECs, we need the new notion of *very-good end-component*.

► **Definition 16.** An end-component C of \mathcal{M} is *very-good* (VGEC) if the following two properties hold:

- $(\mathbf{1}_V) \forall s \in C, s \models_{\mathcal{M}} \mathbf{S}(p_1)$;
- $(\mathbf{2}_V) \forall s \in C, s \models_{\mathcal{M}_1 C} \mathbf{AS}(p_1) \wedge \mathbf{AS}(p_2)$, or equivalently, $s \models_{\mathcal{M}_1 C} \mathbf{AS}(p_1 \cap p_2)$.

We introduce the following notations: $\text{VGEC}(\mathcal{M})$ is the set of all VGECs of \mathcal{M} , and $\mathcal{V} = \cup_{V \in \text{VGEC}(\mathcal{M})} V$ is the set of states that belong to a VGEC in \mathcal{M} .

Note that in condition $(\mathbf{1}_V)$, \mathcal{P}_1 is allowed to leave C to ensure $\mathbf{S}(p_1)$: this is in contrast to condition $(\mathbf{1}_U)$ for UGECs, in Definition 5. On the contrary, condition $(\mathbf{2}_V)$ is exactly the same as $(\mathbf{2}_U)$. From these definitions, it is trivial to see that any UGEC is also a VGEC,



■ **Figure 3** The EC $\{a, b, c\}$ is very-good but not ultra-good, as \mathcal{P}_1 has to leave it to ensure $\mathbf{S}(p_1)$.

but the converse is false. Consider Figure 3: $\{a, b, c\}$ is a VGEC. The strategy ensuring $(\mathbf{2}_V)$ from a is to go to b , and the strategy ensuring $(\mathbf{1}_V)$ from a is to go to d . As we will prove in Lemma 18 and as in all VGECs, \mathcal{P}_1 can ensure $a \models \mathbf{S}(p_1) \wedge P_{>1-\varepsilon}(p_2)$ for any $\varepsilon > 0$. Still, $\{a, b, c\}$ is not a UGEC: no strategy ensures $\mathbf{S}(p_1)$ on $\mathcal{M}_{\{a,b,c\}}$, as \mathcal{P}_2 can enforce the play $(ab)^\omega$ that has odd maximal priority. This illustrates why the notion of UGEC is too strong when reasoning about threshold probability, hence why we need to introduce VGECs.

Available strategies in VGECs. As for UGECs, we will use witness strategies for $(\mathbf{1}_V)$ and $(\mathbf{2}_V)$. Deciding if $(\mathbf{1}_V)$ holds is solving a classical parity game, in $\text{NP} \cap \text{coNP}$ [19]. Uniform pure memoryless witness strategies exist: let λ_1 be such a witness. For simplicity of presentation, we assume in the following that all states of \mathcal{M} satisfy $(\mathbf{1}_V)$, as otherwise they will trivially not satisfy the properties we consider (as $\mathbf{S}(p_1)$ will not be ensured). For $(\mathbf{2}_V)$, we established in Lemma 7 that deciding if it holds is in polynomial time and that uniform randomized memoryless witness strategies exist: let $\lambda_{2,C}$ be one of them.

Reaching VGECs. We prove a strong relationship between the measure of paths that satisfy p_1 and p_2 , and the measure of paths that reach VGECs, under any strategy.

► **Lemma 17.** For all $s \in S$, and all $\lambda \in \Lambda$, the following holds: $\mathbb{P}_{\mathcal{M},s}^\lambda[\diamond \mathcal{V}] \geq \mathbb{P}_{\mathcal{M},s}^\lambda[p_1 \cap p_2]$.

Limit-sure satisfaction in VGECs. For each state in a VGEC, we claim that the parity objective p_2 can be satisfied with probability arbitrarily close to one, while ensuring p_1 surely.

► **Lemma 18.** Let $C \in \text{VGEC}(\mathcal{M})$. For all $s \in C$ and $\varepsilon \in (0, 1]$, the following property holds: $s \models \mathbf{S}(p_1) \wedge P_{>1-\varepsilon}(p_2)$.

Proof Sketch. As for UGECs, we build a witness strategy based on two simpler ones: λ_1 and $\lambda_{2,C}$. However, our strategy here depends on ε . The rough idea is as follows: the witness strategy λ_ε will play $\lambda_{2,C}$ for longer and longer rounds, and switch to λ_1 forever if $D_{\text{even}}^{\max}(p_1)$ is not visited along one of those rounds. Using Lemma 7 cleverly, we can define the sequence of round lengths in such a way that the probability to play as $\lambda_{2,C}$ forever exceeds $1 - \varepsilon$, yielding the result. ◀

The strict threshold case. We reduce the problem to a reachability problem toward \mathcal{V} . The first lemma gives a sufficient condition under which the property is satisfied. Its proof tells us how to construct witness strategies.

► **Lemma 19.** The following holds: $s_0 \models \mathbf{S}(p_1) \wedge P_{>c}(\diamond \mathcal{V}) \implies s_0 \models \mathbf{S}(p_1) \wedge P_{>c}(p_2)$.

Proof Sketch. We build a witness strategy λ based on (i) $\lambda_{\diamond \mathcal{V}}$, a strategy that ensures to reach \mathcal{V} with probability $q > c$ from s_0 , (ii) λ_1 , and (iii) $\lambda_{\varepsilon,C}$ from Lemma 18 for a well-chosen $\varepsilon > 0$ and every VGEC C . The idea is to first play $\lambda_{\diamond \mathcal{V}}$ long enough so that a VGEC C is reached with probability close to q and, if such a C is reached, to switch to $\lambda_{\varepsilon,C}$ for ε

121:12 Threshold Constraints with Guarantees for Parity Objectives in MDPs

sufficiently small so that the total probability to satisfy p_2 is higher than c . If no VGEC is reached, λ switches to λ_1 , hence ensuring $\mathbf{S}(p_1)$. ◀

This second lemma gives a necessary condition. Its proof uses Lemma 17.

► **Lemma 20.** *The following holds: $s_0 \models \mathbf{S}(p_1) \wedge P_{>c}(p_2) \implies s_0 \models \mathbf{S}(p_1) \wedge P_{>c}(\diamond \mathcal{V})$.*

The non-strict threshold case. As we solved the strict case, the only interesting remaining case is when \mathcal{P}_1 , while surely forcing p_1 , can force p_2 with probability c , but no more. The main tool here is UGECs. The first lemma gives a sufficient condition. Its proof is constructive. Recall that $\mathcal{U} = \cup_{U \in \text{UGEC}(\mathcal{M})} U$.

► **Lemma 21.** *The following holds: $s_0 \models \mathbf{S}(p_1) \wedge P_{\geq c}(\diamond \mathcal{U}) \implies s_0 \models \mathbf{S}(p_1) \wedge P_{\geq c}(p_2)$.*

Proof Sketch. We define a witness strategy based on (i) $\lambda_{\diamond \mathcal{U}}$, a witness for $s_0 \models \mathbf{S}(p_1) \wedge P_{\geq c}(\diamond \mathcal{U})$, and (ii) strategies λ_C for every UGEC C : it suffices to play $\lambda_{\diamond \mathcal{U}}$ as long as no UGEC C is reached and to switch to λ_C when reached, if ever. ◀

The next lemma gives a necessary condition, keeping in mind that we consider the case where \mathcal{P}_1 cannot ensure probability strictly larger than c .

► **Lemma 22.** *The following holds: $(s_0 \models \mathbf{S}(p_1) \wedge P_{\geq c}(p_2)) \wedge (s_0 \not\models \mathbf{S}(p_1) \wedge P_{>c}(p_2)) \implies s_0 \models \mathbf{S}(p_1) \wedge P_{\geq c}(\diamond \mathcal{U})$.*

Algorithm. Based on the reductions shown above, we can now establish an algorithm and complexity results for the threshold problem.

► **Theorem 23.** *Given an MDP $\mathcal{M} = (G = (S, E), S_1, S_2, \delta)$, a state $s_0 \in S$, and two priority functions $p_i: S \rightarrow \{1, \dots, d\}$, $i \in \{1, 2\}$, it can be decided in $\text{NP} \cap \text{coNP}$ if $s_0 \models \mathbf{S}(p_1) \wedge P_{\sim c}(p_2)$ for $\sim \in \{>, \geq\}$ and $c \in \mathbb{Q} \cap [0, 1)$. If the answer is YES, then there exists an infinite-memory witness strategy, and infinite memory is in general necessary. This decision problem is at least as hard as solving parity games.*

Proof. The algorithm can be sketched as follows:

1. Remove from \mathcal{M} all states where $\mathbf{S}(p_1)$ does not hold, as well as their attractor for \mathcal{P}_2 : if s_0 is removed, then answer NO. Let \mathcal{M}' be the remaining MDP. This operation is in $\text{NP} \cap \text{coNP}$ as it consists in solving a classical parity game [19].
2. Compute the set \mathcal{V} representing the union of VGECs in \mathcal{M}' . This can be done in polynomial time by computing the maximal ECs of \mathcal{M}' and applying Lemma 7 to check condition $(2_{\mathbf{V}})$ for each of them (condition $(1_{\mathbf{V}})$ holds thanks to the previous step).
3. Decide if $s_0 \models \mathbf{S}(p_1) \wedge P_{>c}(\diamond \mathcal{V})$ using Theorem 4. This is in $\text{NP} \cap \text{coNP}$. If it holds, then answer YES. If it does not hold and \sim is $>$, then answer NO, otherwise, i.e., if \sim is \geq , continue with the next step.
4. Use the sub-algorithm described in Theorem 15 to compute the set \mathcal{U} representing the union of UGECs in \mathcal{M}' . This is in $\text{NP} \cap \text{coNP}$.
5. Decide if $s_0 \models \mathbf{S}(p_1) \wedge P_{\geq c}(\diamond \mathcal{U})$ using Theorem 4. This is in $\text{NP} \cap \text{coNP}$. If it holds, answer YES, otherwise answer NO.

The correctness of this algorithm follows from Lemma 19, Lemma 20, Lemma 21, and Lemma 22. It belongs to $\text{P}^{\text{NP} \cap \text{coNP}} = \text{NP} \cap \text{coNP}$ [6], and it trivially generalizes classical parity games (e.g., by taking $p_2: s \mapsto 0$ for all $s \in S$).

Finally, let us discuss strategies. Witness strategies for the case $>$ (resp. \geq) were described in Lemma 19 (resp. Lemma 21). In both cases, infinite memory is in general required, because it is in general necessary to play optimally in both VGECs and UGECs. For UGECs, see Theorem 15 for an example. For VGECs, consider the VGEC $\{a, b, c\}$ in the MDP of Figure 3. We claim that for every finite-memory strategy λ ensuring $S(p_1)$, the probability to ensure p_2 is zero, hence there is no finite-memory witness for $a \models S(p_1) \wedge P_{>1-\varepsilon}(p_2)$. As argued for the UGEC case, in order to ensure p_1 on the play in which \mathcal{P}_2 always goes to a from b , \mathcal{P}_1 must go to d at some point, and because of the finite memory of λ , he must do it after a bounded number of steps along which c is not visited: say n steps. Again, the probability to do it will be bounded from below by a strictly positive constant, here $2^{-\frac{n}{2}}$ (the probability that \mathcal{P}_2 chooses a for $\frac{n}{2}$ times in a row), all along a consistent play. Therefore, \mathcal{P}_1 will almost-surely go to d , and p_2 will actually be satisfied with probability zero. ◀

6 Conclusion

We further extended the beyond worst-case synthesis framework by studying the case of two parity objectives and proved $\text{NP} \cap \text{coNP}$ membership for all considered variants.

Our algorithms can easily be generalized to more than two parity objectives as long as we consider only the S and AS operators. Indeed, we have that for any MDP \mathcal{M} , any state s in \mathcal{M} , and any number of priority functions p_1, \dots, p_n , it holds that $s \models \bigwedge_i S(p_i) \wedge \bigwedge_j AS(p_j) \iff s \models S(\bigwedge_i p_i) \wedge AS(\bigwedge_j p_j)$, and it is easy to reduce the latter problem to $s' \models S(p') \wedge AS(p'')$ on a (larger) MDP \mathcal{M}' , using classical techniques (e.g., any conjunction of parity objectives can be expressed as a Muller condition [13], that in turn can be transformed into a single parity condition on a larger graph [20]). Extending this generalization to the operator $P_{\sim c}$ is more challenging and would require to mix our techniques to methods for percentile queries [24]: an interesting direction for future work.

Another question is the limits of finite-memory strategies. We saw that in general, infinite memory is needed. We would like to investigate under which additional conditions finite-memory strategies suffice, and to develop corresponding algorithms.

References

- 1 Shaull Almagor, Orna Kupferman, and Yaron Velner. Minimizing expected cost under hard boolean constraints, with applications to quantitative synthesis. In Josée Desharnais and Radha Jagadeesan, editors, *27th International Conference on Concurrency Theory, CONCUR 2016, August 23-26, 2016, Québec City, Canada*, volume 59 of *LIPICs*, pages 9:1–9:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016. doi:10.4230/LIPICs.CONCUR.2016.9.
- 2 Shaull Almagor, Orna Kupferman, and Yaron Velner. Minimizing expected cost under hard boolean constraints, with applications to quantitative synthesis. *CoRR*, abs/1604.07064, 2016. URL: <http://arxiv.org/abs/1604.07064>.
- 3 Christel Baier, Marcus Größer, and Frank Ciesinski. Quantitative analysis under fairness constraints. In Zhiming Liu and Anders P. Ravn, editors, *Automated Technology for Verification and Analysis, 7th International Symposium, ATVA 2009, Macao, China, October 14-16, 2009. Proceedings*, volume 5799 of *Lecture Notes in Computer Science*, pages 135–150. Springer, 2009. doi:10.1007/978-3-642-04761-9_12.
- 4 Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT press, 2008.
- 5 Raphaël Berthon, Mickael Randour, and Jean-François Raskin. Threshold constraints with guarantees for parity objectives in Markov decision processes. *CoRR*, abs/1702.05472, 2017. URL: <http://arxiv.org/abs/1702.05472>.

- 6 Gilles Brassard. A note on the complexity of cryptography (corresp.). *IEEE Transactions on Information Theory*, 25(2):232–233, 1979.
- 7 Tomáš Brázdil, Antonín Kucera, and Petr Novotný. Optimizing the expected mean payoff in energy Markov decision processes. In Cyrille Artho, Axel Legay, and Doron Peled, editors, *Automated Technology for Verification and Analysis - 14th International Symposium, ATVA 2016, Chiba, Japan, October 17-20, 2016, Proceedings*, volume 9938 of *Lecture Notes in Computer Science*, pages 32–49, 2016. doi:10.1007/978-3-319-46520-3_3.
- 8 Romain Brenguier, Lorenzo Clemente, Paul Hunter, Guillermo A. Pérez, Mickael Randour, Jean-François Raskin, Ocan Sankur, and Mathieu Sassolas. Non-zero sum games for reactive synthesis. In Adrian-Horia Dediu, Jan Janousek, Carlos Martín-Vide, and Bianca Truthe, editors, *Language and Automata Theory and Applications - 10th International Conference, LATA 2016, Prague, Czech Republic, March 14-18, 2016, Proceedings*, volume 9618 of *Lecture Notes in Computer Science*, pages 3–23. Springer, 2016. doi:10.1007/978-3-319-30000-9_1.
- 9 Véronique Bruyère, Emmanuel Filiot, Mickael Randour, and Jean-François Raskin. Expectations or guarantees? I want it all! A crossroad between games and MDPs. In Fabio Mogavero, Aniello Murano, and Moshe Y. Vardi, editors, *Proceedings 2nd International Workshop on Strategic Reasoning, SR 2014, Grenoble, France, April 5-6, 2014*, volume 146 of *EPTCS*, pages 1–8, 2014. doi:10.4204/EPTCS.146.1.
- 10 Véronique Bruyère, Emmanuel Filiot, Mickael Randour, and Jean-François Raskin. Meet your expectations with guarantees: Beyond worst-case synthesis in quantitative games. In Ernst W. Mayr and Natacha Portier, editors, *31st International Symposium on Theoretical Aspects of Computer Science, STACS 2014, March 5-8, 2014, Lyon, France*, volume 25 of *LIPICs*, pages 199–213. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2014. doi:10.4230/LIPICs.STACS.2014.199.
- 11 Véronique Bruyère, Quentin Hautem, and Mickael Randour. Window parity games: an alternative approach toward parity games with time bounds. In Domenico Cantone and Giorgio Delzanno, editors, *Proceedings of the Seventh International Symposium on Games, Automata, Logics and Formal Verification, GandALF 2016, Catania, Italy, 14-16 September 2016*, volume 226 of *EPTCS*, pages 135–148, 2016. doi:10.4204/EPTCS.226.10.
- 12 Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdzinski. Mean-payoff parity games. In *20th IEEE Symposium on Logic in Computer Science (LICS 2005), 26-29 June 2005, Chicago, IL, USA, Proceedings*, pages 178–187. IEEE Computer Society, 2005. doi:10.1109/LICS.2005.26.
- 13 Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Generalized parity games. In Helmut Seidl, editor, *Foundations of Software Science and Computational Structures, 10th International Conference, FOSSACS 2007, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2007, Braga, Portugal, March 24-April 1, 2007, Proceedings*, volume 4423 of *Lecture Notes in Computer Science*, pages 153–167. Springer, 2007. doi:10.1007/978-3-540-71389-0_12.
- 14 Lorenzo Clemente and Jean-François Raskin. Multidimensional beyond worst-case and almost-sure problems for mean-payoff objectives. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015*, pages 257–268. IEEE Computer Society, 2015. doi:10.1109/LICS.2015.33.
- 15 Alexandre David, Peter Gjøøl Jensen, Kim Guldstrand Larsen, Axel Legay, Didier Lime, Mathias Grund Sørensen, and Jakob Haahr Taankvist. On time with minimal expected cost! In Franck Cassez and Jean-François Raskin, editors, *Automated Technology for Verification and Analysis - 12th International Symposium, ATVA 2014, Sydney, NSW, Australia, November 3-7, 2014, Proceedings*, volume 8837 of *Lecture Notes in Computer Science*, pages 129–145. Springer, 2014. doi:10.1007/978-3-319-11936-6_10.

- 16 E. Allen Emerson, Charanjit S. Jutla, and A. Prasad Sistla. On model-checking for fragments of μ -calculus. In Costas Courcoubetis, editor, *Computer Aided Verification, 5th International Conference, CAV'93, Elounda, Greece, June 28 - July 1, 1993, Proceedings*, volume 697 of *Lecture Notes in Computer Science*, pages 385–396. Springer, 1993. doi:10.1007/3-540-56922-7_32.
- 17 Jerzy Filar and Koos Vrieze. *Competitive Markov decision processes*. Springer, 1997.
- 18 Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002.
- 19 Marcin Jurdzinski. Deciding the winner in parity games is in $UP \cap co-UP$. *Inf. Process. Lett.*, 68(3):119–124, 1998. doi:10.1016/S0020-0190(98)00150-1.
- 20 Christof Löding. Optimal bounds for transformations of omega-automata. In C. Pandu Rangan, Venkatesh Raman, and Ramaswamy Ramanujam, editors, *Foundations of Software Technology and Theoretical Computer Science, 19th Conference, Chennai, India, December 13-15, 1999, Proceedings*, volume 1738 of *Lecture Notes in Computer Science*, pages 97–109. Springer, 1999. doi:10.1007/3-540-46691-6_8.
- 21 Mickael Randour. Automated synthesis of reliable and efficient systems through game theory: A case study. In *Proc. of ECCS 2012*, Springer Proceedings in Complexity XVII, pages 731–738. Springer, 2013. doi:10.1007/978-3-319-00395-5_90.
- 22 Mickael Randour. Reconciling rationality and stochasticity: Rich behavioral models in two-player games. *CoRR*, abs/1603.05072, 2016. *GAMES 2016, the 5th World Congress of the Game Theory Society, Maastricht, Netherlands*. URL: <http://arxiv.org/abs/1603.05072>.
- 23 Mickael Randour, Jean-François Raskin, and Ocan Sankur. Variations on the stochastic shortest path problem. In Deepak D’Souza, Akash Lal, and Kim Guldstrand Larsen, editors, *Verification, Model Checking, and Abstract Interpretation - 16th International Conference, VMCAI 2015, Mumbai, India, January 12-14, 2015. Proceedings*, volume 8931 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2015. doi:10.1007/978-3-662-46081-8_1.
- 24 Mickael Randour, Jean-François Raskin, and Ocan Sankur. Percentile queries in multi-dimensional Markov decision processes. *Formal Methods in System Design*, 50(2):207–248, 2017. doi:10.1007/s10703-016-0262-7.