

Minimizing Expected Cost Under Hard Boolean Constraints, with Applications to Quantitative Synthesis*

Shaull Almagor¹, Orna Kupferman², and Yaron Velner³

1 School of Computer Science and Engineering, The Hebrew University, Israel.

2 School of Computer Science and Engineering, The Hebrew University, Israel.

3 School of Computer Science and Engineering, The Hebrew University, Israel.

Abstract

In Boolean synthesis, we are given an LTL specification, and the goal is to construct a transducer that realizes it against an adversarial environment. Often, a specification contains both Boolean requirements that should be satisfied against an adversarial environment, and multi-valued components that refer to the quality of the satisfaction and whose expected cost we would like to minimize with respect to a probabilistic environment.

In this work we study, for the first time, mean-payoff games in which the system aims at minimizing the expected cost against a probabilistic environment, while surely satisfying an ω -regular condition against an adversarial environment. We consider the case the ω -regular condition is given as a parity objective or by an LTL formula. We show that in general, optimal strategies need not exist, and moreover, the limit value cannot be approximated by finite-memory strategies. We thus focus on computing the limit-value, and give tight complexity bounds for synthesizing ϵ -optimal strategies for both finite-memory and infinite-memory strategies.

We show that our game naturally arises in various contexts of synthesis with Boolean and multi-valued objectives. Beyond direct applications, in synthesis with costs and rewards to certain behaviors, it allows us to compute the minimal sensing cost of ω -regular specifications – a measure of quality in which we look for a transducer that minimizes the expected number of signals that are read from the input.

1998 ACM Subject Classification F.4.3 : Formal Languages, B.8.2 Performance Analysis and Design Aids, F.1.1 Models of Computation

Keywords and phrases Stochastic and Quantitative Synthesis, Mean Payoff Games, Sensing.

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2016.9

1 Introduction

Synthesis is the automated construction of a system from its specification: given a linear temporal logic (LTL) formula ψ over sets I and O of input and output signals, we synthesize a system that *realizes* ψ [11, 17]. At each moment in time, the system reads a truth assignment, generated by the environment, to the signals in I , and it generates a truth assignment to the signals in O . The system realizes ψ if all the computations that are generated by the above interaction satisfy ψ .

* The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no 278410, from the Israel Science Foundation (grant no 1229/10), and from the US-Israel Binational Science Foundation (grant no 2010431).



One weakness of automated synthesis in practice is that it pays no attention to the quality of the synthesized system. Indeed, the classical setting is Boolean: a computation satisfies a specification or does not satisfy it. Accordingly, while the synthesized system is correct, there is no guarantee about its quality. In recent years, researchers have considered extensions of the classical Boolean setting to a quantitative one, which takes quality into account. Quality measures refer to the system itself, examining parameters like size or consumption of memory, sensors, voltage, bandwidth, etc., or refer to the way the system satisfies the specification. In the latter, we allow the designer to specify the quality of a behavior using quantitative specification formalisms [1, 6, 13]. For example, rather than the Boolean specification requiring all requests to be followed by a grant, a quantitative specification formalism gives a different satisfaction value to a computation in which requests are responded immediately and one in which requests are responded after long delays.¹

Solving the synthesis problem in the Boolean setting amounts to solving a two-player zero-sum game between the system and the environment. The goal of the system is to satisfy the (Boolean) specification, and the environment is adversarial. Then, a winning strategy for the system corresponds to a transducer that realizes the specification. In the quantitative setting, the goal of the system is no longer Boolean, as every play is assigned a cost by the specification. In the classical quantitative approach, we measure the satisfaction value in the worst-case semantics. Thus, the value of a strategy for the system is the maximal cost of a play induced by this strategy, and the goal of the system is to minimize this value. Recently, there is a growing interest also in the expected cost of a play, under a probabilistic environment. The motivation behind this approach is that the quality of satisfaction is a “soft constraint”, and should not be measured in a worst-case semantics. Then, the game above is replaced by a mean-payoff Markov Decision Process (MDP): a game in which each state has a cost, inducing also costs to infinite plays (essentially, the cost of an infinite play is the limit of the average cost of prefixes of increased lengths). The goal is to find a strategy that minimizes the expected cost [10, 12].

While quantitative satisfaction refines the Boolean one, often a specification contains both Boolean conditions that should be satisfied against all environments, and multi-valued components that refer to the quality of the satisfaction and whose expectation we would like to minimize with respect to a probabilistic environment. Accordingly, researchers have suggested the *beyond worst-case* approach, where a specification has both hard and soft constraints, and the goal is to realize the hard constraints, while maximizing the expected satisfaction value of the soft constraints.

In this work, we consider, for the first time, mean-payoff MDPs equipped with a parity winning condition (parity-MDPs, for short). The goal is to find a strategy that surely wins the parity game (that is, against an adversarial environment), while minimizing the expected cost of a play against a probabilistic environment. While the starting point in earlier related work is the MDP itself, possibly augmented by different objectives, our starting point depends on the application, and we view the construction of the MDP as an integral part of our contribution. We focus on two applications: synthesis with penalties to undesired scenarios and synthesis with minimal sensing.

Let us describe the two applications. We start with penalties to scenarios. Consider

¹ Note that the polarity of some quality measures is negative, as we want to minimize size, consumption, costs, etc., whereas the polarity of other measures is positive, as we want to maximize performance and satisfaction value. For simplicity, we assume that all measures are associated with costs, which we want to minimize.

an LTL specification ψ over I and O . Activating an output signal may have a cost; for example, when the activation involves a use of a resource. Taking these costs into account, the input to the synthesis problem includes, in addition to ψ , a cost function γ assigning cost to some assignments to output signals. The cost of a computation is then the mean cost of assignments in it. While the specification ψ is a hard constraint, as we only allow correct computations, minimizing the expected cost of computations with respect to γ is a soft constraint. More elaborated cost functions refer to on-going regular scenarios. Power consumption, for example, is an important consideration in modern chip design. As the chips become more complex, the cost of powering a server farm can easily outweigh the cost of the servers themselves, thus design teams go to great lengths in order to reduce power consumption in their designs. The most widely researched logical power saving techniques are *clock gating*, in which a clock is prevented from making a “tick” if it is redundant [5], and *power gating*, in which whole sections of the chip are powered off when not needed and then powered on again [14]. The goal of these techniques is to reduce power consumption and the number of changes in the values of signals, the main source of power consumption in chips. The input to the problem of synthesis with penalties to scenarios includes, in addition to ψ , a set of deterministic automata on finite words, each describing a undesired scenario and its cost. For example, it is easy to specify the scenario of “value flip” with a two-state deterministic automaton. We show how the setting can be translated into solving our parity-MDPs, thus generating systems that realize ψ with minimal expected cost.

Our primary application considers activation of sensors. The quality measure of sensing was introduced in [2, 3], as a measure for the detail with which a random input word needs to be read in order to realize the specification. In the context of synthesis, our goal is to construct a transducer that realizes the specification and minimizes the expected average number of sensors (of input signals) that are used along the interaction. Thus, the hard constraint is the LTL specification, and the soft one is the expected number of active sensors. Giving up sensing has a flavor of synthesis with incomplete information [15]: the transducer has to realize the specification no matter what the incomplete information is. Thus, as opposed to the examples above, the modeling of cost involves a careful construction of the MDP to be analyzed, and also involves an exponential blow-up, which we show to be unavoidable. In [3], the problem was solved for safety specifications. Our solution to the parity-MDP problem enables a solution for full LTL. We also study the complexity of the problem when the input is an LTL formula, rather than a deterministic automaton.

Back to parity-MDPs, we show that in general, optimal strategies need not exist. That is, there are parity-MDPs in which an infinite-state strategy can get arbitrarily close to some limit optimal value, but cannot attain it. Moreover, the limit value cannot be approximated by finite-memory strategies. Accordingly, our solution to parity-MDPs suggests two algorithms. The first, described in Section 3.1, finds the limit value of all possible strategies, which corresponds to infinite-state transducers. The second, described in Section 3.2, computes the limit value over all finite-memory strategies. The complexity of both algorithms is $\text{NP} \cap \text{coNP}$. Moreover, they are computable in polynomial time when an oracle to a two-player parity game is given. Hence, our complexity upper bounds match the trivial lower bounds that arise from the fact that every solution to a parity-MDP is also a solution to a parity game. For our applications, we show that the complexity of the synthesis problem for LTL specifications stays doubly-exponential, as in the Boolean setting, even when we minimize penalties to undesired scenarios or minimize sensing.

Related Work. The combination of worst-case synthesis with expected-cost synthesis, dubbed *beyond worst-case synthesis*, was studied in [7, 12] for models that are closely related to ours. In [7] the authors study mean-payoff MDPs, where both the hard constraints and the soft constraints are quantitative. Thus, a system needs to ensure a strict upper bound on the mean-payoff cost, while minimizing the expected cost. In [12], multidimensional mean-payoff MDPs are considered. Thus, the MDP is equipped with several mean-payoff costs, and the goal is to find a system that ensures the mean-payoff in some of the mean-payoffs is below an upper bound, while minimizing the expected mean-payoffs (or rather, approximating their Pareto-curve).

In comparison, our work is the first to consider a hard Boolean constraint (namely the parity condition). This poses both a conceptual and a technical difference. Conceptually, when quantitative synthesis is taken as a refinement of Boolean synthesis, it is typically meant as a ranking of different systems that satisfy a Boolean specification. Thus, it makes sense for the hard constraint to be Boolean as well. Technically, combining Boolean and quantitative constraints gives rise to some subtleties that do not exist in the pure-quantitative setting. Specifically, when both the hard and the soft constraints are quantitative, a strategy can intuitively “alternate” between satisfying them. Thus, if while trying to meet the soft constraint the hard constraint is violated, we can switch to a worst-case strategy until the hard constraint is satisfied, and go back to trying to minimize the soft constraint. This alternation can be done infinitely often. In the Boolean setting, however, this alternation can violate the Boolean constraint. We note that unlike classical parity games, where the parity winning condition can be translated to a richer mean-payoff objective, the parity winning condition in our parity-MDPs does not admit a similar translation.

Other works on MDPs and mean-payoff objectives tackle different aspects of quantitative analysis. In [18], a solution to the expected mean-payoff value over MDPs is presented. In [8] and [9], the authors study a combination of mean-payoff and parity objectives over MDPs and over stochastic two-player games. There, the goal of the system is to ensure with probability 1 that the parity condition holds and that the mean-payoff is below a threshold. This differs from our work in that the parity condition is not a hard constraint, as it is met only almost-surely, and in that the expected mean-payoff is not guaranteed to be minimized. As detailed in the paper, these differences make the technical challenges very different.

Due to lack of space, some of the proofs are omitted and can be found in the full version [4].

2 Parity-MDPs

A *parity Markov decision process* (Parity-MDP, for short) combines a parity game with a mean-payoff MDP. The game is played between Player 1, who models a system, and Player 2, who models the environment. The environment is adversarial with respect to the parity winning condition and is stochastic with respect to the mean-payoff objective. Formally, a parity-MDP is a tuple $\mathcal{M} = \langle S_1, S_2, s_0, A_1, A_2, \delta_1, \delta_2, P, cost, \alpha \rangle$, with the following components. The sets S_1 and S_2 are finite set of states, for Players 1 and 2, respectively. Let $S = S_1 \cup S_2$. Then, $s_0 \in S$ is an initial state, and A_1 and A_2 are sets of actions for the players. Not all actions are available in all states: for every state $s \in S_i$, for $i \in \{1, 2\}$, we use $A_i(s)$ to denote the finite set of actions available to Player i in the state s . For $i \in \{1, 2\}$, the (partial) transition function $\delta_i : S_i \times A_i \rightarrow S$ is such that $\delta_i(s, a)$ is defined iff $a \in A_i(s)$. Let $\delta = \delta_1 \cup \delta_2$. Note that δ_2 gets an action of Player 2 as a parameter. We distinguish between two approaches to the way an action is chosen in Player 2 states. In the

adversarial approach, it is Player 2 who chooses the action. In the stochastic approach, the choice depends on the (partial) function $P : S_2 \times A_2 \rightarrow [0, 1]$, where for every state $s \in S_2$ and $a \in A_2$, we have that $P(s, a) > 0$ only if $a \in A_2(s)$. Also, $\sum_{a \in A_2(s)} P(s, a) = 1$. Finally, $cost : S \rightarrow \mathbb{N}$ is a cost function, and $\alpha : S \rightarrow \{0, \dots, d\}$, for some $d \in \mathbb{N}$, is a parity winning condition.

The parity-MDP \mathcal{M} induces a *parity game* $\mathcal{M}^P = \langle S_1, S_2, s_0, A_1, A_2, \delta_1, \delta_2, \alpha \rangle$, obtained by omitting P and $cost$. In this game, we follow the adversarial approach to the environment. Thus, both players choose their actions. Formally, a *strategy* for Player i in \mathcal{M} , for $i \in \{1, 2\}$ is a function $f_i : S^* \times S_i \rightarrow A_i$ such that for s_0, \dots, s_n , we have $f(s_0, \dots, s_n) \in A_i(s_n)$. Thus, a strategy suggests to Player i an available action given the history of the states traversed so far. Note that we do not consider randomized strategies, but rather deterministic ones. Our results in Section 3 show that this is sufficient, in the sense that the players cannot gain by using randomization.

Given strategies f_1 and f_2 for Players 1 and 2, the *play* induced f_1 and f_2 is the infinite sequence of states s_0, s_1, \dots such that for every $j \geq 0$, if $s_j \in S_i$, for $i \in \{1, 2\}$, then $s_{j+1} = \delta_i(s_j, f(s_0, \dots, s_j))$. For an infinite play r , we denote by $\text{inf}(r)$ the set of states that r visits infinitely often. The play $r = s_0, s_1, \dots$ of \mathcal{M} is *parity winning* if $\max \{\alpha(s) : s \in \text{inf}(r)\}$ is even.

The parity-MDP \mathcal{M} also induces an *MDP* $\mathcal{M}^{\text{MDP}} = \langle S_1, S_2, s_0, A_1, A_2, \delta_1, \delta_2, P, cost \rangle$, obtained by omitting α . In this game, we follow the stochastic approach to the environment and consider the distribution of plays when only a strategy for Player 1 is given. Formally, we first extend P to transitions as follows: For states $s \in S_2$ and $s' \in S$, we define $P(s, s') = \sum_{a \in A_2(s) : \delta_2(s, a) = s'} P(s, a)$. Then, a play of \mathcal{M} with strategy f_1 for Player 1 is an infinite sequence of states s_0, s_1, \dots such that for every $j \geq 0$, if $s_j \in S_1$, then $s_{j+1} = \delta_1(s_j, f_1(s_0, \dots, s_j))$, and if $s_j \in S_2$, then $P(s_j, s_{j+1}) > 0$. The *cost* of a strategy f_1 is the expected average cost of a random walk in \mathcal{M} in which Player 1 proceeds according to f_1 . Formally, for $m \in \mathbb{N}$ and for a prefix $\tau = s_0, s_1, \dots, s_m$ of a play, let $I_2 = \{j : j < m \text{ and } s_j \in S_2\}$. Then, we define $P_{f_1}(\tau) = \prod_{j \in I_2} P(s_j, s_{j+1})$ and $cost_m(f_1, \tau) = \frac{1}{m+1} \sum_{j=0}^m cost(s_j)$. The cost of the strategy f_1 is then $cost(f_1) = \liminf_{m \rightarrow \infty} \sum_{\tau : |\tau|=m} cost_m(f_1, \tau) \cdot P_{f_1}(\tau)$. We denote by $\text{inf}(f_1)$ the random variable that associates $\text{inf}(\rho)$ with a sequence of states $\rho = s_0, s_1, \dots$, under the probability space induced by \mathcal{M} with f_1 .

A *finite memory* strategy for \mathcal{M} is described by a finite set M called *memory*, an initial memory $init \in M$, a *memory update function* $next : S_1 \times M \rightarrow M$, and an *action function* $act : S_1 \times M \rightarrow A_1$ such that $act(s, m) \in A_1(s)$ for every $s \in S_1$ and $m \in M$.

A strategy is *memoryless* if it has finite memory M with $|M| = 1$. Note that a memoryless strategy depends only on the current state. Thus, we can describe a memoryless strategy by $f_1 : S_1 \rightarrow A_1$. Let $cost(\mathcal{M}) = \inf \{cost(f_1) : f_1 \text{ is a strategy for } \mathcal{M}\}$. That is, $cost(\mathcal{M})$ is the expected cost of a game played on \mathcal{M} in which Player 1 uses an optimal strategy.

The following is a basic property of MDPs.

► **Theorem 1.** *Consider an MDP \mathcal{M} . Then, $cost(\mathcal{M})$ can be attained by a memoryless strategy, which can be computed in polynomial time.*

Recall that a strategy f_1 for player 1 is winning in M^P if every play of \mathcal{M} with f_1 satisfies the parity condition α . Note that we require *sure* winning, in the sense that all plays must be winning, rather than winning with probability 1 (*almost-sure* winning). On the other hand, the definition of cost in M^{MDP} considered strategies for Player 1 and ignore the parity winning condition. We now define the *sure cost* of the parity-MDP, which does take them into account. For a strategy f_1 for Player 1, the sure cost of f_1 , denoted $cost_{\text{sure}}(f_1)$, is

$cost(f_1)$, if f_1 is winning, and is ∞ otherwise. The sure cost of \mathcal{M} is then $cost_{\text{sure}}(\mathcal{M}) = \inf \{cost_{\text{sure}}(f_1) : f_1 \text{ is a strategy for } \mathcal{M}\}$.

Consider a parity-MDP $\mathcal{M} = \langle S_1, S_2, s_0, A_1, A_2, \delta_1, \delta_2, P, cost, \alpha \rangle$. An *end component* (EC, for short) is a set $U \subseteq S$ such that for every state $s \in U$, the following hold.

1. If $s \in S_1$, then there exists an action $a \in A_s$ such that $\delta_1(s, a) \in U$.
2. If $s \in S_2$, then for every $a \in A_2(s)$ such that $P(s, a) > 0$, it holds that $\delta_2(s, a) \in U$.
3. For every $t, t' \in U$, there exist a path $t = t_0, t_1, \dots, t_k = t'$ and actions a_1, \dots, a_t such that for every $0 \leq i < t$, it holds that $t_i \in U$, and there exists an action a such that $\delta(t_i, a) = t_{i+1}$.

Intuitively, the probabilistic player cannot force the play to leave U , and Player 1 has positive probability of reaching every state in U from every other state.


For an EC U and a state $s \in U$, we can consider the parity-MDP $\mathcal{M}|_U^s$, in which the states are U , the initial state is s , and all the components are naturally restricted to U . Since U is an EC, then this is indeed a parity-MDP. An EC U is *maximal* if for every nonempty $U' \subseteq S \setminus U$, we have that $U \cup U'$ is not an EC.

3 Solving Parity MDPs

In this section we study the problem of finding the sure cost for an MDP. Recall that for MDPs, there always exists an optimal memoryless strategy. We start by demonstrating that for the sure cost of parity-MDPs, the situation is much more complicated.

► **Theorem 2.** *There is a parity-MDP \mathcal{M} in which Player 1 does not have an optimal strategy (in particular, not a memoryless one) for attaining the sure cost of \mathcal{M} . Moreover, for every $\epsilon > 0$, Player 1 may need infinite memory in order to ϵ -approximate $cost_{\text{sure}}(\mathcal{M})$.*

Proof. Consider the parity-MDP \mathcal{M} appearing in Figure 1. Player 1 can decrease the cost of a play towards 1 by staying in the initial state s_0 . However, in order to ensure an even parity rank, Player 1 must either play b and reach s_3 w.p. 0.5, or play c but incur cost 10. A finite memory strategy for Player 1 must eventually play c from s_0 in almost every play,² thus the cost of every winning finite-memory strategy is 10. On the other hand, for every $\epsilon > 0$, there exists an infinite memory strategy f that gets cost at most $1 + \epsilon$. Essentially (see Lemma 4 for a formal proof of the general case), the strategy f plays b for a long time. If s_3 is reached, it plays b for even longer, and otherwise plays c .



■ **Figure 1** The Parity MDP \mathcal{M} . States of Player 1 are circles, these of Player 2 are squares, with outgoing edges marked by their probability. Each state is labeled by its parity rank (left) and cost (right). Player 1 has no optimal strategy and needs infinite memory for an ϵ approximation.

Finally, there is no optimal strategy for Player 1, as every strategy that plays c from s_0 eventually (i.e., as a response to some strategy of Player 2) gets cost 10 with some positive probability. However, a strategy that never plays c is not parity-winning. ◀

² Note that this also implies that randomized strategies could not be of help here.

Following Theorem 2, our solution to parity MDPs suggests two algorithms. The first, described in Section 3.1, finds the limit value of all possible strategies, which corresponds to infinite-state transducers. The second, described in Section 3.2, computes the limit value over all finite-memory strategies. The complexity of both algorithms is $\text{NP} \cap \text{coNP}$. Moreover, they are computable in polynomial time when an oracle to a two-player parity game is given. Hence, our complexity upper bounds match the trivial lower bounds that arise from the fact that every solution to a parity-MDP is also a solution to a parity game.

3.1 Infinite-Memory Strategies

In this section we study the problem of finding the sure cost of a parity-MDP when infinite-memory strategies are allowed. We prove the upper bound in the following theorem. As stated above, the lower bound is trivial.

► **Theorem 3.** *Consider a parity-MDP \mathcal{M} . Then, $\text{cost}_{\text{sure}}(\mathcal{M})$ can be computed in $\text{NP} \cap \text{coNP}$, and is parity-games hard.*

Consider a parity-MDP $\mathcal{M} = \langle S_1, S_2, s_0, A_1, A_2, \delta_1, \delta_2, P, \text{cost}, \alpha \rangle$. We first remove from \mathcal{M} all states that are not sure-winning for Player 1 in \mathcal{M}^P . Clearly, every strategy that attains $\text{cost}_{\text{sure}}(\mathcal{M})$ cannot visit a state that is losing in \mathcal{M}^P . Thus, we henceforth assume that all states in \mathcal{M} are winning for Player 1 in \mathcal{M}^P . We say that an EC C of \mathcal{M} is *good* (GEC, for short) if its maximal rank is even. That is, $\max_{s \in C} \{\alpha(s)\}$ is even.

The idea behind our algorithm is as follows. W.p. 1, each play in \mathcal{M} eventually reaches and visits infinitely often all states of some EC. Hence, when restricting attention to plays that are winning for Player 1 in \mathcal{M}^P , it must be the case that this EC is good. It follows that the sure cost of \mathcal{M} is affected only by the properties of its GECs. Moreover, since the minimal expected mean-payoff value is the same in all the states of an EC, we can consider only maximal GECs and refer to the value of an EC, namely the minimal expected value that Player 1 can ensure while staying in the EC. Our algorithm constructs a new MDP (without ranks) \mathcal{M}' in which the cost of a state is the value of the maximal GEC it belongs to. If a state does not belong to a GEC, then we assign it a very high cost in \mathcal{M}' , where the intuition is that Player 1 cannot benefit from visiting this state infinitely often. We claim that the sure cost in the parity-MDP \mathcal{M} coincides with the cost of the MDP \mathcal{M}' .

Formally, for an EC C , let C^{\max} be the set of the states of C with the maximal parity rank in C . By definition, this rank is even when C is a GEC. Note that if C and C' are GECs and $C \cap C' \neq \emptyset$, then $C \cup C'$ is also a GEC. Thus, we can restrict attention to maximal GEC. For a GEC C , there exists a memoryless strategy f^C that maximizes the probability of reaching C^{\max} from every state $s \in C$ while staying in C . Moreover, since C is an EC, the probability of reaching C^{\max} by playing f^C is strictly positive from every state $s \in C$. Let t be a state in C . Consider the MDP $\mathcal{M}^{\text{MDP}}|_C^t$. Since C is EC, we have that $\text{cost}(\mathcal{M}^{\text{MDP}}|_C^t)$ is independent of the initial state t . Thus, we can define $\text{cost}(\mathcal{M}^{\text{MDP}}|_C)$ as $\text{cost}(\mathcal{M}^{\text{MDP}}|_C^t)$ for some $t \in C$.

Recall that our algorithm starts by a preprocessing step that removes all states that are not sure-winning for Player 1 in \mathcal{M}^P . It then finds the maximal GECs of \mathcal{M} (using a polynomial-time procedure that we describe in the full version [4]), and obtain an MDP \mathcal{M}' by assigning every state within a GEC C the cost $\text{cost}(\mathcal{M}^{\text{MDP}}|_C)$, and assigning every state that is not inside a GEC cost $W + 1$, where W is the maximal cost that appears in \mathcal{M} . We claim that $\text{cost}_{\text{sure}}(\mathcal{M}) = \text{cost}(\mathcal{M}')$.

Before proving the claim, note that all the steps of the algorithm except for the preprocessing step that involves a solution of parity game require polynomial time. In particular,

the strategies f^C above are computable in polynomial time by solving a reachability MDP, and, by Theorem 1, so does the final step of finding $\text{cost}(\mathcal{M}')$.

Proving that $\text{cost}_{\text{sure}}(\mathcal{M}) = \text{cost}(\mathcal{M}')$ involves the following steps (see the full version [4] for the full proof). First, proving $\text{cost}_{\text{sure}}(\mathcal{M}) \geq \text{cost}(\mathcal{M}')$ is not hard, as a play with a winning strategy f for Player 1 in \mathcal{M} reaches and stays in some GEC C w.p. 1, and within C , the best expected cost one can hope for is $\text{cost}(\mathcal{M}^{\text{MDP}}|_C)$, which is exactly what the strategy f attains when played in \mathcal{M}' .

Next, proving $\text{cost}_{\text{sure}}(\mathcal{M}) \leq \text{cost}(\mathcal{M}')$, we show how an optimal strategy f' in \mathcal{M}' induces an ϵ -optimal strategy f in \mathcal{M} . We start with Lemma 4, which justifies the costs within a GEC.

► **Lemma 4.** *Consider a GEC C in \mathcal{M} , and $s \in C$. Let $v(s) = \text{cost}(\mathcal{M}^{\text{MDP}}|_C^s)$, then for every $\epsilon > 0$ there exists a strategy f of \mathcal{M}^s with $\text{cost}_{\text{sure}}(f) \leq v(s) + \epsilon$.*

Intuitively, in a good EC, f minimizes the expected mean-payoff and *once in a while* it plays reachability, aiming to visit to a state with the maximal rank in the EC. Since the EC is good, this rank is even. If reachability is not obtained after N rounds, for a parameter N , then f gives up and aims at only surely satisfying the parity objective (our preprocessing step ensures that this is possible). Otherwise, after reaching the maximal rank, f switches again to minimizing mean-payoff. This process is repeated forever, increasing N in each iteration. Hence, the probability that Player 1 eventually gives up can be bounded from above by an arbitrarily small $\epsilon > 0$. Accordingly, Player 1 can achieve a value that is arbitrarily close to $\text{cost}(\mathcal{M}^{\text{MDP}}|_C)$.

Finally, we construct the ϵ -optimal strategy f in \mathcal{M} as follows. The strategy f first mimics f' for a large number of steps k , or until an EC (in which f' stays forever) is reached. If a good EC is not reached, then f aims at only surely satisfying the parity objective. If a good EC is reached, then f behaves as prescribed above, per Lemma 4. Since the probability of f' reaching a good EC within k steps tends to 1, then Player 1 can achieve a value within ϵ of $\text{cost}(\mathcal{M}')$.

► **Example 5.** Recall the parity-MDP \mathcal{M} in Figure 1. The GECs are $C_1 = \{s_0, s_2, s_3\}$ and $C_2 = \{s_1\}$, and their costs as MDPs are 1 and 10, respectively. Since Player 1 can force the play to stay in C_1 , then $\text{cost}_{\text{sure}}(\mathcal{M}) = 1$, and an ϵ approximation is obtained as per Lemma 4, and as described in the proof of Theorem 2. ◀

3.2 Finite-Memory Strategies

In this section we study the problem of finding the sure cost of a parity-MDP, when restricted to finite memory strategies. For a parity-MDP \mathcal{M} , we define $\text{cost}_{\text{sure}, < \infty}(\mathcal{M}) = \inf\{\text{cost}_{\text{sure}}(f) : f \text{ is a finite memory strategy for } \mathcal{M}\}$. We prove the upper bound in the following theorem. As stated above, the lower bound is trivial.

► **Theorem 6.** *Consider a parity-MDP \mathcal{M} . Then, $\text{cost}_{\text{sure}, < \infty}(\mathcal{M})$ can be computed in $\text{NP} \cap \text{co-NP}$, and is parity-games hard.*

The general approach is similar to the one we took in Section 3.1. That is, we remove from \mathcal{M} all states that are not sure-winning for Player 1 in \mathcal{M}^P , and proceed by reasoning about a certain type of ECs. However, for finite-memory strategies, we need a more restricted class of ECs than the GECs that were used in Section 3.1. Indeed, a finite-memory strategy might not suffice to win the sure-parity condition in a GEC.

For a GEC C , let k be the maximal odd priority in C , with $k = -1$ if there are no odd priorities. We define $C_{\text{even}}^{\max} = \{q \in C : \alpha(q) > k \text{ and } \alpha(q) \text{ is even}\}$. We say that a GEC C in

\mathcal{M} is *super good* (SGEC, for short) if from every state $s \in C$, there exists a finite-memory strategy f for $\mathcal{M}|_C^s$ such that the play of \mathcal{M} under f reaches C_{even}^{\max} w.p. 1, and if the play does not reach C_{even}^{\max} , then it is parity winning. We refer to f as a *witness* to C being a SGEC. If C is not a SGEC, we call its states that do satisfy the above *super-good states*.

We argue that SGECs are the proper notion for reasoning about finite-memory strategies. Specifically, we show that in a SGEC, Player 1 can achieve ϵ -optimal expected cost with a finite-memory strategy, and that every finite-memory winning strategy reaches a SGEC w.p. 1.

Our algorithm finds the maximal SGECs of \mathcal{M} and obtain an MDP \mathcal{M}' in the same manner we did in Section 3.1, namely by assigning high weights to states not in SGECs, and the optimal mean-payoff MDP value to states in SGECs. As there, we claim that $\text{cost}(\mathcal{M}') = \text{cost}_{\text{sure}, < \infty}(\mathcal{M})$. The analysis of the algorithm as well as its concrete details, are, however, more intricate.

► **Example 7.** Recall again the parity-MDP \mathcal{M} in Figure 1. The only SGEC is $\{s_1\}$, and thus we have that $\text{cost}_{\text{sure}, < \infty}(\mathcal{M}) = 10$ (see the proof of Theorem 2). ◀

We start by proving that the notion of maximal SGECs is well defined. To this end, we present the following lemma, whose proof appears in the full version [4]. Note that in the case of GECs, the lemma was trivial.

► **Lemma 8.** *Consider SGEC C and D , such that $C \cap D \neq \emptyset$, then $C \cup D$ is also a SGEC.*

Intuitively, we prove this by considering witnesses f, g for C and D being SGECs. We then modify f such that from every state in C , it tries to reach D for N steps, for some parameter N . Once D is reached, g takes over. If D is not reached, f attempts to reach C_{even}^{\max} . Thus, w.p. 1, the strategy reaches D_{even}^{\max} , and if it does not, it either reaches C_{even}^{\max} infinitely often, or wins the parity condition.

Next, we note that unlike the syntactic definition of GECs, the definition of SGECs is semantic, as it involves a strategy. Thus, finding the maximal SGECs adds another complication to the algorithm. In fact, it is not hard to see that even checking whether an EC is a SGEC is parity-games hard. Using techniques from [8], we show in the full version [4] that we can reduce the latter to the problem of solving a parity-Büchi game. We thus have the following lemma.

► **Lemma 9.** *Consider an EC C in a parity-MDP \mathcal{M} . We can decide whether C is a SGEC in $\text{NP} \cap \text{co-NP}$, as well as compute a witness strategy and, if C is not a SGEC, find the set of super-good states.*

Next, we show how to find the maximal SGECs of \mathcal{M} . Essentially, for every odd rank k , we can find the SGECs whose maximal odd rank is k by removing all states with higher odd ranks, and recursively refining ECs by keeping only super-good states, using Lemma 9. Thus, we have the following (see the full version [4] for complete details).

► **Theorem 10.** *Consider a parity-MDP \mathcal{M} . We can find the maximal SGECs of \mathcal{M} in $\text{NP} \cap \text{co-NP}$.*

Theorem 10 shows that our algorithm for computing $\text{cost}_{\text{sure}, < \infty}(\mathcal{M})$ solves the problem in $\text{NP} \cap \text{co-NP}$. It remains to prove its correctness. First, Lemma 11 justifies the assignment of costs within a SGEC.

► **Lemma 11.** *Consider a SGEC C in \mathcal{M} and a state s in C . Let $v(s) = \text{cost}(\mathcal{M}^{\text{MDP}}|_C^s)$. Then, for every $\epsilon > 0$, there exists a finite-memory strategy f of $\mathcal{M}|_C^s$ with $\text{cost}_{\text{sure}}(f) \leq v(s) + \epsilon$.*

Proof. Let g be a memoryless strategy such that $\text{cost}(g) = \text{cost}(\mathcal{M}^{\text{MDP}}|_C^s)$. By Theorem 1 such a strategy exists. Let h be a finite-memory strategy that witnesses C being a SGEC. For every $k \in \mathbb{N}$, consider the strategy f_k that repeatedly plays g for k steps and then plays h until C_{even}^{\max} is reached. Since g and h are finite-memory, then f_k is finite memory. In addition, observe that h reaches C_{even}^{\max} w.p. 1, and if C_{even}^{\max} is not reached, then h is parity-winning. Thus f_k is parity-winning, and it reaches Step 1 infinitely often w.p. 1. Moreover, since h has finite memory, then for every $n \in \mathbb{N}$, there is a bounded probability $0 < p(n) \leq 1$ that f reaches C_{even}^{\max} within n steps, with $\lim_{n \rightarrow \infty} p(n) = 1$. Thus, we get that $\lim_{k \rightarrow \infty} \text{cost}_{\text{sure}}(f_k) = \text{cost}_{\text{sure}}(g) = \text{cost}(\mathcal{M}^{\text{MDP}}|_C^s)$, which concludes the proof. ◀

Lemma 11 implies that we can approximate the optimal value of SGECs with finite-memory strategies. It remains to show that it is indeed enough to consider SGECs. Consider a finite-memory strategy f . Then, w.p. 1, f reaches an EC. Let C be an EC with $\Pr_{\mathcal{M}}(\text{inf}(f) = C) > 0$. The following lemma characterizes an assumption we can make on the behavior of f in such an EC.

► **Lemma 12.** *Consider a parity-MDP \mathcal{M} and an EC C . For every finite-memory strategy f , if $\Pr_{\mathcal{M}}(\text{inf}(f) = C) > 0$, then there exists a finite-memory strategy g such that for every $s \in C$, we have that $\Pr_{\mathcal{M}^s}(\text{inf}(g) = C) = 1$ and every play of g from s stays in C . Moreover, if f is parity winning, then so is g .*

Intuitively, we show that there exists some finite history h such that the strategy f_h , which is f played after seeing the history h , has the following property: f_h reaches and stays in C , and w.p. 1 visits infinitely often all the states in C , and in particular C_{even}^{\max} . For the proof, we consider the set $F = \{f_h : h \text{ is a finite history}\}$. Since f has finite memory, it follows that this set is finite. Using this, we show that if $\Pr_{\mathcal{M}}(\text{inf}(g) = C) < 1$ for every $g \in F$, then $\Pr_{\mathcal{M}}(\text{inf}(f) = C) = 0$, which is a contradiction. Finally, since f is also parity winning, it follows that f_h above is also parity-winning, and is thus a witness for C being a SGEC. The full proof appears in the full version [4].

Finally, by Lemma 13, we can assume that once f reaches an EC C , it stays in C and visits all its states infinitely often w.p. 1. Since f is parity-winning, it follows that C has a maximal even rank, and that f reaches C_{even}^{\max} w.p. 1. Moreover, in every play that does not reach C_{even}^{\max} , f wins the parity condition. We can thus conclude with the following Lemma, which completes the correctness proof of our algorithm for computing $\text{cost}_{\text{sure}, < \infty}(\mathcal{M})$. See the full version [4] for the proof.

► **Lemma 13.** *Consider a parity-MDP \mathcal{M} and an EC C . For every finite-memory strategy f , if f is parity winning and $\Pr_{\mathcal{M}}(\text{inf}(f) = C) > 0$, then C is a SGEC.*

3.3 Comparison with Related Work

Both our work and [7, 12] analyze ECs and reduce the problem to reasoning about an MDP that ignores the hard constraints. The main difference with [7] is that there, the hard and soft constraints have the same objective (i.e., worst-case mean-payoff value and expected-case mean-payoff value). In [7], the strategy played for N rounds to satisfy the soft objective and then at most M rounds to satisfy the hard objective, for some constants N and M . In our setting, we cannot bound M , and in fact it might be the case that Player 1 would play to satisfy the parity objective for the rest of the game (i.e., forever) even after reaching a super-good end component.

The difference with [12] is twofold. First, technically, the type of hard constraints in [12] is worst-case mean-payoff, whereas our setting uses the Boolean parity condition. In classical parity games, the parity condition can be reduced to a mean-payoff objective. Similar reductions, however, do not work in order to reduce our setting to the setting of [12]. Thus, our contribution is orthogonal to [12]. Secondly, Boolean constraints are conceptually different than quantitative constraints, and as we demonstrate in Section 4, they arise naturally in quantitative extensions of Boolean paradigms.

We note that [12] also study a relaxation in which almost-sure winning is allowed for the hard constraints. An analogue in our setting is to consider an almost-sure parity condition. We note that in such a setting, GECs are sufficient for reasoning both about finite-memory and infinite-memory strategies. Moreover, the preprocessing involves solving an almost-sure parity MDP (without mean-payoff constraints), which can be done in polynomial time. Thus, as is the case in [12], we can compute the cost of an MDP with almost-sure hard constraints in polynomial time.

4 Applications

In this section we study two applications of parity-MDPs. Both extend the Boolean synthesis problem. Due to lack of space, our description is only an overview. Full definitions and details can be found in the full version [4]. We start with some basic definitions.

For finite sets I and O of input and output signals, respectively, an *I/O transducer* is $\mathcal{T} = \langle I, O, Q, q_0, \delta, \rho \rangle$, where Q is a set of states, $q_0 \in Q$ is an initial state, $\delta : Q \times 2^I \rightarrow Q$ is a total (deterministic) transition function, and $\rho : Q \rightarrow 2^O$ is a labeling function on the states. The run of \mathcal{T} on a word $w = i_0 \cdot i_1 \cdots \in (2^I)^\omega$ is the sequence of states q_0, q_1, \dots such that $q_{k+1} = \delta(q_k, i_k)$ for all $k \geq 0$. The *output* of \mathcal{T} on w is then $o_1, o_2, \dots \in (2^O)^\omega$ where $o_k = \rho(q_k)$ for all $k \geq 1$. Note that the first output assignment is that of q_1 , and we do not consider $\rho(q_0)$. This reflects the fact that the environment initiates the interaction. The *computation of \mathcal{T} on w* is then $\mathcal{T}(w) = i_0 \cup o_1, i_1 \cup o_2, \dots \in (2^{I \cup O})^\omega$. When Q is a finite set, we say that the transducer is finite.

The synthesis problem gets as input a specification $L \subseteq (2^{I \cup O})^\omega$ and generates a transducer \mathcal{T} that realizes L ; namely, all the computations of \mathcal{T} are in L . The language L is typically given by an LTL formula [16] or by means of an automaton of infinite words.

4.1 Penalties on Undesired Scenarios

Recall that in Boolean synthesis, the goal is to generate a transducer that associates with each infinite sequence of inputs an infinite sequence of outputs so that the result computation satisfies a given specification. Typically, some behaviors generated by the transducers may be less desired than others. For example, as discussed in Section 1, designs that use fewer resources or minimize expensive activities are preferable. The input to the *synthesis with penalties* problem includes, in addition to the Boolean specification, languages of finite words that describe undesired behaviors, and their costs. The goal is to generate a transducer that realizes the specification and minimizes cost due to undesired behaviors.

Formally, the input to the problem includes languages L_1, \dots, L_m of finite words over the alphabet $2^{I \cup O}$ and a penalty function $\gamma : \{1, \dots, m\} \rightarrow \mathbb{N}$ specifying for each $1 \leq i \leq m$ the penalty that should be applied for generating a behavior in L_i . As described in Section 1, the language L_i may be local (that is, include only words of length 1) and thus refer only to activation of output signals, may specify short scenarios like flips of output signals, and may also specify rich regular scenarios. Note that we allow penalties also for behaviors that

depend on the input signals. Intuitively, whenever a computation π includes a behavior in L_i , a penalty of $\gamma(i)$ is applied. Formally, if $\pi = \sigma_1, \sigma_2, \dots$, then for every position $j \geq 1$, we define $\text{penalty}(j) = \{i : \text{there is } k \leq i \text{ such that } \sigma_k \cdot \sigma_{k+1} \cdots \sigma_j \in L_i\}$. That is, $\text{penalty}(j)$ points to the subset of languages L_i such that a word in L_i ends in position j . Then, the cost of position j , denoted $\text{cost}(j)$, is $\sum_{i \in \text{penalty}(j)} \gamma(i)$. Finally, for a finite computation $\pi = \sigma_1, \sigma_2, \dots$, we define its cost, denoted $\text{cost}(\pi)$, as $\limsup_{m \rightarrow \infty} \frac{1}{m} \sum_{j=1}^m \text{cost}(j)$.

Let \mathcal{A} be a deterministic parity automaton (DPW, for short) over the alphabet $2^{I \cup O}$ that specifies the specification ψ . We describe a parity-MDP whose solution is a transducer that realizes \mathcal{A} with the minimal cost for penalties. The idea is simple: on top of the parity game \mathcal{G} described above, we compose monitors that detect undesired scenarios. We assume that the languages L_1, \dots, L_m and are given by means of deterministic automata on finite words (DFWs) $\mathcal{U}_1, \dots, \mathcal{U}_m$ where for every $1 \leq i \leq m$, we have that $L(\mathcal{U}_i) = (2^{I \cup O})^* \cdot L_i$. That is, \mathcal{U}_i accepts $\sigma_1 \cdots \sigma_n$ iff there exists $k \leq n$ such that $\sigma_k \cdots \sigma_n \in L_i$. Essentially, we turn \mathcal{A} into a parity-MDP by composing it with the DFWs $\mathcal{U}_1, \dots, \mathcal{U}_m$. Then, \mathcal{U}_i reaching an accepting state indicates that the penalty for L_i should be applied, which induces the costs in the parity-MDP. The probabilities in the parity-MDP are induced from the distribution of the assignments to input signals. The full construction can be found in the full version [4]. We note that an alternative definition can replace the DFWs $\mathcal{U}_1, \dots, \mathcal{U}_m$ and the cost function γ by a single weighted automaton that can be composed with \mathcal{A} .

4.2 Sensing

Consider a transducer $\mathcal{T} = \langle I, O, Q, q_0, \delta, \rho \rangle$. For a state $q \in Q$ and a signal $p \in I$, we say that p is *sensed in* q if there exists a set $S \subseteq I$ such that $\delta(q, S \setminus \{p\}) \neq \delta(q, S \cup \{p\})$. Intuitively, a signal is sensed in q if knowing its value may affect the destination of at least one transition from q . We use $\text{sensed}(q)$ to denote the set of signals sensed in q . The *sensing cost* of a state $q \in Q$ is $\text{scost}(q) = |\text{sensed}(q)|$. For a finite run $r = q_1, \dots, q_m$ of \mathcal{T} , we define the sensing cost of r , denoted $\text{scost}(r)$, as $\frac{1}{m} \sum_{i=0}^{m-1} \text{scost}(q_i)$. That is, $\text{scost}(r)$ is the average number of sensors that \mathcal{T} uses during r . For a finite input sequence $w \in (2^I)^*$, we define the sensing cost of w in \mathcal{T} , denoted $\text{scost}_{\mathcal{T}}(w)$, as the sensing cost of the run of \mathcal{T} on w . Finally, the sensing cost of \mathcal{T} is the expected sensing cost of input sequences of length that tends to infinity, which is parameterized by a distribution on $(2^I)^\omega$ given by a sequence of distributions D_1, D_2, \dots such that $D_t : 2^I \rightarrow [0, 1]$ describes the distribution over 2^I at time $t \in \mathbb{N}$. For simplicity, we assume that the distribution is uniform. Thus, $D_t(i) = 2^{-|I|}$ for every $t \in \mathbb{N}$. For the uniform distribution we have $\text{scost}(\mathcal{T}) = \lim_{m \rightarrow \infty} \frac{1}{|2^I|^m} \sum_{w \in (2^I)^m} \text{scost}_{\mathcal{T}}(w)$.

Note that this definition also applies when the transducer is infinite. However, for infinite transducers, the limit in the definition of $\text{scost}(\mathcal{T})$ might not exist, and we therefore define $\text{scost}(\mathcal{T}) = \limsup_{m \rightarrow \infty} \frac{1}{|2^I|^m} \sum_{w \in (2^I)^m} \text{scost}_{\mathcal{T}}(w)$. Finally, for a realizable specification $L \in 2^{I \cup O}$, we define $\text{scost}_{I/O}(L) = \inf\{\text{scost}(\mathcal{T}) : \mathcal{T} \text{ is an } I/O \text{ transducer that realizes } L\}$.

In [3], we study the sensing cost of *safety* properties. We show that there, a finite, minimally-sensing transducer, always exists (albeit of exponential size), and the problem of computing the sensing cost is EXPTIME-complete. In our current setting, however, a minimally-sensing transducer need not exist, and any approximation may require infinite memory. We demonstrate this with an example.

► **Example 14.** Let $I = \{a\}$ and $O = \{b\}$, and consider the specification $\psi = (\text{GF}a \wedge \text{Gb}) \vee \text{G}(\neg b \rightarrow \text{XG}(a \leftrightarrow b))$. Thus, ψ states that either a holds infinitely often and b always holds, or, if b does not holds at a certain time, then henceforth, a holds iff b holds. Observe that once the system outputs $\neg b$, it has to always sense a in order to determine the output. The

system thus has an incentive to always output b . This, however, may render ψ false, as a need not hold infinitely often.

We start by claiming that every finite-memory transducer \mathcal{T} that realizes ψ has sensing cost 1. Indeed, let n be the number of states in \mathcal{T} . A random input sequence contains the infix $(\neg a)^{n+1}$ w.p. 1. Upon reading such an infix, \mathcal{T} has to output $\neg b$, as otherwise it would not realize ψ on a computation with suffix $(\neg a)^\omega$. Thus, from then on, \mathcal{T} senses a in every state. So, $\text{scost}(\mathcal{T}) = 1$.

However, by using infinite-memory transducers, we can follow the construction in Section 3.1 and reduce the sensing cost arbitrarily close to 0. Let $M \in \mathbb{N}$. We construct a transducer \mathcal{T}' as follows. After initializing i to 1, the transducer \mathcal{T}' senses a and outputs b for iM steps. If a does not hold during this time, then \mathcal{T}' outputs $\neg b$ and starts sensing a and outputting b accordingly. Otherwise, if a holds during this time, then \mathcal{T}' stops sensing a for 2^i steps, while outputting b . It then increases i by 1 and repeats the process. Note that \mathcal{T}' outputs $\neg b$ iff a does not hold for iM consecutive positions at the i -th round (which happens w.p. 2^{-iM}). Thus, the probability of \mathcal{T}' outputting $\neg b$ in a random computation is bounded from above by $\sum_{i=1}^{\infty} 2^{-iM} = 2^{-M}$, which tends to 0 as M tends to ∞ . Note that in the i -th round, \mathcal{T}' senses a for only iM steps, and then does not sense anything for 2^{iM} steps, so if \mathcal{T}' does not output $\neg b$, the sensing cost is 0. Thus, we have $\lim_{M \rightarrow \infty} \text{scost}(\mathcal{T}') = 0$. ◀

We proceed by describing the general solution to computing the sensing cost of a specification. Recall that synthesis of a DPW \mathcal{A} is reduced to solving a parity game. When sensing is introduced, it is not enough for the system to win this game, as it now has to win while minimizing the sensing cost. Intuitively, not sensing some inputs introduces incomplete information to the game: once the system gives up sensing, it may not know the state in which the game is and knows instead only a set of states in which the game may be. Technically (see the full version for the detailed proof), we force the system to satisfy the specification with respect to all assignments to the un-sensed inputs by converting the DPW \mathcal{A} into a *universal parity automaton* (UPW) – an automaton in which a the transition function maps each state and letter to a set of successor states, and a word is accepted if all the runs on it are accepting.

► **Theorem 15.** *Consider a DPW specification \mathcal{A} over $2^{I \cup O}$. There exists a parity-MDP \mathcal{M} such that $\text{cost}_{\text{sure}}(\mathcal{M}) = \text{scost}_{I/O}(L(\mathcal{A}))$. Moreover, the number of states of \mathcal{M} is singly exponential in that of \mathcal{A} , and the number of parity ranks on \mathcal{M} is polynomial in that of \mathcal{A} .*

► **Theorem 16.** *Consider a DPW specification \mathcal{A} over $2^{I \cup O}$. We can compute $\text{scost}_{I/O}(L(\mathcal{A}))$ in singly-exponential time. Moreover, the problem of deciding whether $\text{scost}_{I/O}(L(\mathcal{A})) > 0$ is EXPTIME-complete.*

Proof. We obtain from \mathcal{A} a parity-MDP \mathcal{M} as per Theorem 15. Observe that the algorithm in the proof of Theorem 3 essentially runs in polynomial time, apart from solving a parity game, which is done in $\text{NP} \cap \text{co-NP}$. However, deterministic algorithms for solving parity games run in time polynomial in the number of states, and singly-exponential in the number of parity ranks. Since the number of parity ranks in \mathcal{M} is polynomial in that of \mathcal{A} , we can find $\text{cost}_{\text{sure}}(\mathcal{M})$ in time singly-exponential in the size of \mathcal{A} . Since $\text{cost}_{\text{sure}}(\mathcal{M}) = \text{scost}_{I/O}(L(\mathcal{A}))$, we are done.

For the lower bound, we note that the problem of deciding whether $\text{scost}_{I/O}(L(\mathcal{A})) > 0$ is EXPTIME-hard even for a restricted class of automata, namely looping automata [3]. ◀

The input to the synthesis problem is typically given as an LTL formula, rather than a DPW. Then, the translation from LTL to a DPW involves a doubly-exponential blowup.

Thus, a naive solution for computing the sensing cost of a specification given by an LTL formula is in 3EXPTIME. However, by translating the formula to a UPW, rather than a DPW, we show how we can avoid one exponent, thus matching the 2EXPTIME complexity of standard Boolean synthesis.

► **Theorem 17.** *Consider an LTL specification ψ over $I\cup O$. We can compute $\text{scost}_{I/O}(L(\psi))$ in doubly-exponential time.*

Proof. We start by translating ψ to a UPW \mathcal{A} of size single-exponential in the size of ψ . This can be done, for example, by translating $\neg\psi$ to a nondeterministic Büchi automaton [19] and dualizing it. We then follow the proof of Theorem 15, by adding the universal transitions described there directly to the UPW \mathcal{A} . Thus, when we finally determinize the UPW to a DPW, the size of the DPW is doubly-exponential, so computing the sensing cost can also be done in doubly-exponential time. ◀

References

- 1 S. Almagor, U. Boker, and O. Kupferman. Formalizing and reasoning about quality. *Journal of the ACM*, 63(3), 2016.
- 2 S. Almagor, D. Kuperberg, and O. Kupferman. Regular sensing. In *34th FSTTCS*, volume 29 of *LIPICs*, pages 161–173, 2014.
- 3 S. Almagor, D. Kuperberg, and O. Kupferman. The sensing cost of monitoring and synthesis. In *35th FSTTCS*, volume 35 of *LIPICs*, pages 380–393, 2015.
- 4 Shaul 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>.
- 5 E. Arbel, O. Rokhlenko, and K. Yorav. Sat-based synthesis of clock gating functions using 3-valued abstraction. In *FMCAD*, pages 198–204, 2009.
- 6 Patricia Bouyer, Nicolas Markey, and Raj Mohan Matteplackel. Averaging in LTL. In *25th CONCUR*, pages 266–280, 2014.
- 7 V. Bruyère, E. Filiot, M. Randour, and J-F. Raskin. Meet your expectations with guarantees: Beyond worst-case synthesis in quantitative games. In *31st STACS*, volume 25 of *LIPICs*, pages 199–213, 2014.
- 8 K. Chatterjee and L. Doyen. Energy and mean-payoff parity markov decision processes. In *36th MFCS*, pages 206–218, 2011.
- 9 K. Chatterjee, L. Doyen, H. Gimbert, and Y. Oualhadj. Perfect-information stochastic mean-payoff parity games. In *17th FOSSACS*, pages 210–225, 2014.
- 10 K. Chatterjee, Z. Komárková, and J. Kretínský. Unifying two views on multiple mean-payoff objectives in markov decision processes. In *30th LICS*, pages 244–256, 2015.
- 11 A. Church. Logic, arithmetics, and automata. In *Proc. Int. Congress of Mathematicians, 1962*, pages 23–35. Institut Mittag-Leffler, 1963.
- 12 L. Clemente and J-F. Raskin. Multidimensional beyond worst-case and almost-sure problems for mean-payoff objectives. In *30th LICS*, pages 257–268, 2015.
- 13 L. de Alfaro, M. Faella, T.A. Henzinger, R. Majumdar, and M. Stoelinga. Model checking discounted temporal properties. *Theoretical Computer Science*, 345(1):139–170, 2005.
- 14 M. Keating, D. Flynn, R. Aitken, A. Gibbons, and K. Shi. *Low Power Methodology Manual*. Springer, 2007.
- 15 O. Kupferman and M.Y. Vardi. Synthesis with incomplete information. In *Advances in Temporal Logic*, pages 109–127. Kluwer Academic Publishers, 2000.
- 16 A. Pnueli. The temporal semantics of concurrent programs. *Theoretical Computer Science*, 13:45–60, 1981.

- 17 A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *16th POPL*, pages 179–190, 1989.
- 18 M.L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- 19 M.Y. Vardi and P. Wolper. Reasoning about infinite computations. *IEC*, 115(1):1–37, 1994.