The Complexity of Partial-Observation Parity Games

Krishnendu Chatterjee[1] and Laurent Doyen[2]

[1] IST Austria (Institute of Science and Technology Austria)
[2] LSV, ENS Cachan & CNRS, France

Abstract. We consider two-player zero-sum games on graphs. On the basis of the information available to the players these games can be classified as follows: (a) partial-observation (both players have partial view of the game); (b) one-sided partial-observation (one player has partial-observation and the other player has complete-observation); and (c) complete-observation (both players have complete view of the game). We survey the complexity results for the problem of deciding the winner in various classes of partial-observation games with ω-regular winning conditions specified as parity objectives. We present a reduction from the class of parity objectives that depend on sequence of states of the game to the sub-class of parity objectives that only depend on the sequence of observations. We also establish that partial-observation acyclic games are PSPACE-complete.

1 Introduction

Games on graphs. Games played on graphs provide the mathematical framework to analyze several important problems in computer science as well as mathematics. In particular, when the vertices and edges of a graph represent the states and transitions of a reactive system, then the synthesis problem (Church's problem) asks for the construction of a winning strategy in a game played on a graph [5, 21, 20, 18]. Game-theoretic formulations have also proved useful for the verification [1], refinement [13], and compatibility checking [9] of reactive systems. Games played on graphs are dynamic games that proceed for an infinite number of rounds. In each round, the players choose moves which, together with the current state, determine the successor state. An outcome of the game, called a *play*, consists of the infinite sequence of states that are visited.

Strategies and objectives. A strategy for a player is a recipe that describes how the player chooses a move to extend a play. Strategies can be classified as follows: *pure* strategies, which always deterministically choose a move to extend the play, and *randomized* strategies, which may choose at a state a probability distribution over the available moves. Objectives are generally Borel measurable functions [17]: the objective for a player is a Borel set B in the Cantor topology on S^ω (where S is the set of states), and the player satisfies the objective iff the outcome of the game is a member of B. In verification, objectives are usually *ω-regular languages*. The ω-regular languages generalize the classical regular languages to infinite strings; they occur in the low levels of the Borel hierarchy (they lie in $\Sigma_3 \cap \Pi_3$) and they form a robust and expressive language for determining payoffs for commonly used specifications. We consider parity objectives and its sub-classes that are canonical forms to express objectives in verification.

Classification of games. Games played on graphs can be classified according to the knowledge of the players about the state of the game. Accordingly, there are (a) *partial-observation* games, where each player only has a partial or incomplete view about the state and the moves of the other player; (b) *one-sided partial-observation* games, where one player has partial knowledge and the other player has complete knowledge about the state and moves of the other player; and (c) *complete-observation* games, where each player has complete knowledge of the game.

Analysis. The analysis of games can be classified as *qualitative* and *quantitative* analysis. The qualitative analysis consists of the following questions: given an objective and a state of the game, (a) can Player 1 ensure the objective with certainty against all strategies of Player 2 (*sure winning* problem); (b) can Player 1 ensure the objective with probability 1 against all strategies of Player 2 (*almost-sure winning* problem); and (c) can Player 1 ensure the objective with probability arbitrarily close to 1 against all strategies of Player 2 (*limit-sure winning* problem). Given an objective, a state of the game, and a rational threshold ν, the quantitative analysis problem asks whether the maximal probability with which Player 1 can ensure the objective against all Player 2 strategies is at least ν.

Organization. The paper is organized as follows: In Section 3 we show a new result that presents a reduction of general parity objectives that depend on state sequences to *visible* objectives that only depend on the sequence of observations (rather than the sequence of states). In Section 4 we survey the complexity of solving the three classes of partial-observation games with parity objectives and its sub-classes both for qualitative and quantitative analysis. In Section 5 we show that for the special case of *acyclic* games the qualitative analysis problem is PSPACE-complete both for one-sided partial-observation and partial-observation games. The PSPACE-completeness result for acyclic games is in contrast to general games where the complexities are EXPTIME-complete, 2EXPTIME-complete, and undecidable (depending on the objective and the specific qualitative analysis question).

2 Definitions

In this section we present the definition of partial-observation games and their sub-classes, and the notions of strategies and objectives. A *probability distribution* on a finite set A is a function $\kappa : A \to [0,1]$ such that $\sum_{a \in A} \kappa(a) = 1$. We denote by $\mathcal{D}(A)$ the set of probability distributions on A. We focus on partial-observation turn-based games, where at each round one of the players is in charge of choosing the next action.

Partial-observation games. A *partial-observation game* (or simply a *game*) is a tuple $G = \langle S_1 \cup S_2, A_1, A_2, \delta_1 \cup \delta_2, \mathcal{O}_1, \mathcal{O}_2 \rangle$ with the following components:

1. *(State space).* $S = S_1 \cup S_2$ is a finite set of states, where $S_1 \cap S_2 = \emptyset$ (i.e., S_1 and S_2 are disjoint), states in S_1 are Player 1 states, and states in S_2 are Player 2 states.
2. *(Actions).* A_i ($i = 1, 2$) is a finite set of actions for Player i.
3. *(Transition function).* For $i \in \{1, 2\}$, the transition function for Player i is the function $\delta_i : S_i \times A_i \to S_{3-i}$ that maps a state $s_i \in S_i$, and action $a_i \in A_i$ to the successor state $\delta_i(s_i, a_i) \in S_{3-i}$ (i.e., games are alternating).

4. *(Observations).* $\mathcal{O}_i \subseteq 2^{S_i}$ $(i = 1, 2)$ is a finite set of observations for Player i that partition the state space S_i. These partitions uniquely define functions $\mathsf{obs}_i : S_i \to \mathcal{O}_i$ $(i = 1, 2)$ that map each Player i state to its observation such that $s \in \mathsf{obs}_i(s)$ for all $s \in S_i$.

Special cases. We consider the following special cases of partial-observation games, obtained by restrictions in the observations:

– *(Observation restriction).* The games with *one-sided partial-observation* are the special case of games where $\mathcal{O}_2 = \{\{s\} \mid s \in S_2\}$ (Player 2 has complete observation), i.e., only Player 1 has partial-observation. The *games of complete-observation* are the special case of games where $\mathcal{O}_1 = \{\{s\} \mid s \in S_1\}$ and $\mathcal{O}_2 = \{\{s\} \mid s \in S_2\}$, i.e., every state is visible to each player and hence both players have complete observation. If a player has complete observation we omit the corresponding observation sets from the description of the game.

Classes of game graphs. We use the following abbreviations: Pa for partial-observation, Os for one-sided complete-observation, Co for complete-observation. For $\mathcal{C} \in \{\mathsf{Pa}, \mathsf{Os}, \mathsf{Co}\}$, we denote by $\mathcal{G}_{\mathcal{C}}$ the set of all $\mathcal{C}$ games. Note that the following strict inclusions hold: partial-observation (Pa) is more general than one-sided partial-observation (Os) and Os is more general than complete-observation (Co).

Plays. In a game, in each turn, for $i \in \{1, 2\}$, if the current state s is in S_i, then Player i chooses an action $a \in A_i$, and the successor state is $\delta_i(s, a)$. A *play* in G is an infinite sequence of states and actions $\rho = s_0 a_0 s_1 a_1 \ldots$ such that for all $j \geq 0$, if $s_j \in S_i$, for $i \in \{1, 2\}$, then there exists $a_j \in A_i$ such that $\delta_i(s_j, a_j) = s_{j+1}$. The *prefix up to* s_n of the play ρ is denoted by $\rho(n)$, its *length* is $|\rho(n)| = n + 1$ and its *last element* is $\mathsf{Last}(\rho(n)) = s_n$. The set of plays in G is denoted by $\mathsf{Plays}(G)$, and the set of corresponding finite prefixes is denoted $\mathsf{Prefs}(G)$. For $i \in \{1, 2\}$, we denote by $\mathsf{Prefs}_i(G)$ the set of finite prefixes in G that end in a state in S_i. The *observation sequence* of $\rho = s_0 a_0 s_1 a_1 \ldots$ for Player i $(i = 1, 2)$ is the unique infinite sequence of observations and actions of Player i, i.e., $\mathsf{obs}_i(\rho) \in (\mathcal{O}_i A_i)^\omega$ defined as follows: (i) if $s_0 \in S_i$, then $\mathsf{obs}_i(\rho) = o_0 a_0 o_2 a_2 o_4 \ldots$ such that $s_j \in o_j$ for all even $j \geq 0$; (ii) if $s_0 \in S_{3-i}$, then $\mathsf{obs}_i(\rho) = o_1 a_1 o_3 a_3 o_5 \ldots$ such that $s_j \in o_j$ for all odd $j \geq 1$. The observation sequence for finite sequences (prefix of plays) is defined analogously.

Strategies. A *pure strategy* in G for Player 1 is a function $\sigma : \mathsf{Prefs}_1(G) \to A_1$. A *randomized strategy* in G for Player 1 is a function $\sigma : \mathsf{Prefs}_1(G) \to \mathcal{D}(A_1)$. A (pure or randomized) strategy σ for Player 1 is *observation-based* if for all prefixes $\rho, \rho' \in \mathsf{Prefs}(G)$, if $\mathsf{obs}_1(\rho) = \mathsf{obs}_1(\rho')$, then $\sigma(\rho) = \sigma(\rho')$. We omit analogous definitions of strategies for Player 2. We denote by Σ_G, Σ_G^O, Σ_G^P, Π_G, Π_G^O and Π_G^P the set of all Player-1 strategies in G, the set of all observation-based Player-1 strategies, the set of all pure Player-1 strategies, the set of all Player-2 strategies in G, the set of all observation-based Player-2 strategies, and the set of all pure Player-2 strategies, respectively. Note that if Player 1 has complete observation, then $\Sigma_G^O = \Sigma_G$.

Objectives. An *objective* for Player 1 in G is a set $\phi \subseteq S^\omega$ of infinite sequences of states. A play $\rho \in \mathsf{Plays}(G)$ *satisfies* the objective ϕ, denoted $\rho \models \phi$, if $\rho \in \phi$. Objectives are generally Borel measurable: a Borel objective is a Borel set in the Cantor topology on S^ω [15]. We specifically consider ω-regular objectives specified as parity objectives (a canonical form to express all ω-regular objectives [25]). For a play $\rho = s_0 a_0 s_1 a_1 \dots$ we denote by ρ_k the k-th state s_k of the play and denote by $\mathrm{Inf}(\rho)$ the set of states that occur infinitely often in ρ, that is, $\mathrm{Inf}(\rho) = \{s \mid s_j = s \text{ for infinitely many } j\text{'s}\}$. We consider the following classes of objectives.

1. *Reachability and safety objectives.* Given a set $\mathcal{T} \subseteq S$ of target states, the *reachability* objective $\mathsf{Reach}(\mathcal{T})$ requires that a state in $\mathcal{T}$ be visited at least once, that is, $\mathsf{Reach}(\mathcal{T}) = \{\rho \mid \exists k \geq 0 \cdot \rho_k \in \mathcal{T}\}$. Dually, the *safety* objective $\mathsf{Safe}(\mathcal{T})$ requires that only states in $\mathcal{T}$ be visited. Formally, $\mathsf{Safe}(\mathcal{T}) = \{\rho \mid \forall k \geq 0 \cdot \rho_k \in \mathcal{T}\}$.
2. *Büchi and coBüchi objectives.* The *Büchi* objective $\mathsf{Buchi}(\mathcal{T})$ requires that a state in $\mathcal{T}$ be visited infinitely often, that is, $\mathsf{Buchi}(\mathcal{T}) = \{\rho \mid \mathrm{Inf}(\rho) \cap \mathcal{T} \neq \emptyset\}$. Dually, the *coBüchi* objective $\mathsf{coBuchi}(\mathcal{T})$ requires that only states in $\mathcal{T}$ be visited infinitely often. Formally, $\mathsf{coBuchi}(\mathcal{T}) = \{\rho \mid \mathrm{Inf}(\rho) \subseteq \mathcal{T}\}$
3. *Parity objectives.* For $d \in \mathbb{N}$, let $p : S \to \{0, 1, \dots, d\}$ be a *priority function*, which maps each state to a nonnegative integer priority. The *parity* objective $\mathsf{Parity}(p)$ requires that the minimum priority that occurs infinitely often be even. Formally, $\mathsf{Parity}(p) = \{\rho \mid \min\{p(s) \mid s \in \mathrm{Inf}(\rho)\} \text{ is even}\}$. The Büchi and coBüchi objectives are the special cases of parity objectives with two priorities, $p : S \to \{0, 1\}$ and $p : S \to \{1, 2\}$, respectively.
4. *Visible objectives.* We say that an objective ϕ is *visible* for Player i if for all $\rho, \rho' \in S^\omega$, if $\rho \models \phi$ and $\mathsf{obs}_i(\rho) = \mathsf{obs}_i(\rho')$, then $\rho' \models \phi$. For example if the priority function maps observations to priorities (i.e., $p : \mathcal{O}_i \to \{0, 1, \dots, d\}$), then the parity objective is visible for Player i.

Outcomes. The *outcome* of two randomized strategies σ (for Player 1) and π (for Player 2) from a state s in G is the set of plays $\rho = s_0 s_1 \dots \in \mathsf{Plays}(G)$, with $s_0 = s$, where for all $j \geq 0$, if $s_j \in S_1$ (resp. $s_j \in S_2$), then there exists an action $a_j \in A_1$ (resp. $a_j \in A_2$), such that $\sigma(\rho(j))(a_j) > 0$ (resp. $\pi(\rho(j))(a_j) > 0$) and $\delta_1(s_j, a_j) = s_{j+1}$ (resp. $\delta_2(s_j, a_j) = s_{j+1}$). This set is denoted $\mathrm{Outcome}(G, s, \sigma, \pi)$. The outcome of two pure strategies is defined analogously by viewing pure strategies as randomized strategies that play their chosen action with probability one. The *outcome set* of the pure (resp. randomized) strategy σ for Player 1 in G is the set $\mathsf{Outcome}_1(G, s, \sigma)$ of plays ρ such that there exists a pure (resp. randomized) strategy π for Player 2 with $\rho \in \mathrm{Outcome}(G, s, \sigma, \pi)$. The outcome set $\mathsf{Outcome}_2(G, s, \pi)$ for Player 2 is defined symmetrically.

Sure winning, almost-sure winning, limit-sure winning and value function. An *event* is a measurable set of plays, and given strategies σ and π for the two players, the probabilities of events are uniquely defined [26]. For a Borel objective ϕ, we denote by $\mathrm{Pr}_s^{\sigma,\pi}(\phi)$ the probability that ϕ is satisfied by the play obtained from the starting state s when the strategies σ and π are used. Given a game G, an objective ϕ, and a state s, we consider the following winning modes: (1) an observation-based strategy σ for Player 1 is *sure winning* for the objective ϕ from s if $\mathrm{Outcome}(G, s, \sigma, \pi) \subseteq \phi$ for all observation-based strategies π for Player 2; (2) an observation-based strategy σ for Player 1 is *almost-sure winning* for the objective ϕ from s if $\mathrm{Pr}_s^{\sigma,\pi}(\phi) = 1$ for all observation-based strategies π for Player 2; and (3) a family $(\sigma_\varepsilon)_{\varepsilon>0}$ of observation-based strategies for Player 1 is *limit-sure winning* for the objective ϕ from s if $\mathrm{Pr}_s^{\sigma_\varepsilon,\pi}(\phi) \geq 1 - \varepsilon$, for all $\varepsilon > 0$ and all observation-based strategies π for Player 2. The *value function* $\langle\langle 1 \rangle\rangle_{val}^G : S \to \mathbb{R}$ for objective ϕ for Player 1 assigns to every state the maximal probability with which Player 1 can guarantee the satisfaction of ϕ with an observation-based strategy, against all observation-based strategies for Player 2. Formally we have

$$\langle\langle 1 \rangle\rangle_{val}^G(\phi)(s) = \sup_{\sigma \in \Sigma_G^O} \inf_{\pi \in \Pi_G^O} \mathrm{Pr}_s^{\sigma,\pi}(\phi).$$

For $\varepsilon \geq 0$, an observation-based strategy is *ε-optimal* for ϕ from s if we have $\inf_{\pi \in \Pi_G^O} \mathrm{Pr}_s^{\sigma,\pi}(\phi) \geq \langle\langle 1 \rangle\rangle_{val}^G(\phi)(s) - \varepsilon$. An *optimal* strategy is a 0-optimal strategy. Given a rational value $0 \leq \nu \leq 1$ and a state s, the *value decision problem* asks whether the value of the game at s is at least ν. The qualitative analysis consists of the sure, almost-sure and limit-sure winning problems, and the quantitative analysis is the value decision problem.

3 Reduction of Objectives to Visible Objectives

The complexity lower bounds in this paper are given for visible objectives, while upper bounds are given for general objectives. In [23, 7], algorithms based on a subset construction are given for visible objective, establishing upper bounds (namely EXPTIME) for visible objectives only.

We show that games with general parity objectives can be reduced to games with visible parity objective with an exponential blow-up. However, this blow-up has no impact on the complexity upper bounds because from a game G, the reduction constructs a game G' as the product of G with an exponentially large automaton M, such that the further subset construction of [7] applied to G' induces an exponential blow-up only with respect to G (the subset construction for G' has size $O(2^{|G|} \cdot |M|) = O(2^{|G|} \cdot 2^{|G|})$ which is simply exponential). This is because M is a deterministic automaton.

We give the idea of the construction. Assume that we have a game G with parity objective given by the priority function $p : S \to \{0, 1, \ldots, d\}$. We construct a game G' with visible objective as a product $G \times M$ where M is a finite automaton with parity condition that "synchronizes" with G on observations and actions of Player 1. We construct M as the complement of the automaton M' that we define as follows.

The automaton M' has state space S_1 and alphabet $\Sigma = \mathcal{O}_1 \times A_1$ that accepts the observations of the plays that are losing for Player 1. An observation sequence is losing if it is the observation of a losing play. The initial state of M' is the initial state of the game (we assume w.l.o.g that the game starts in a Player 1 state). The transitions of M' are $(s, (\mathsf{obs}_1(s), a), s'')$ for all $s, s'' \in S_1$ and $a \in A_1$ such that $\delta_1(s, a) = s'$ and $\delta_2(s', b) = s''$ for some $s' \in S_2$ and $b \in A_2$. The priority assigned to this transition is $1 + \min\{p(s), p(s')\}$. Note that M' has at most one run over each infinite word, and that a run in M' corresponds to a play in G. The language of M' is the set of infinite words over $\Sigma = \mathcal{O}_1 \times A_1$ that have a run in M' in which the least priority visited infinitely often is even, i.e. such that the least priority (according to p) visited infinitely often is odd (and thus the corresponding run violates the winning condition of the game G). By complementing M', we get an exponentially larger automaton M that accepts the winning observation sequences [24]. We can assume that M is deterministic and that the states rather than the transitions are labeled by priorities and letters. The game G' is obtained by a synchronized product of G and M in which Player 1 can see the state of M (i.e., the observation of a state (s, u) where s is a state of G and u is a state of M is $(\mathsf{obs}_1(s), u)$). The priority of a state (s, u) depends only on u and therefore defines a visible parity objective. Transitions in G and M are synchronized on the obervations and actions of Player 1.

Note that for reachability and safety objectives, there exists a reduction to a visible objective in polynomial time. First, we can assume that the target states $\mathcal{T}$ defining the objective are sink states (because once $\mathcal{T}$ is reached, the winner of the game is fixed). Second, we make the sink states visible, which makes the objective vissible, and does not change the winner of the game (observing the sink states is of no help since the game is over when this happens).

Theorem 1. *Given a game $G \in \mathcal{G}_{\mathsf{Os}}$ with parity objective, one can construct a game G' as a product of G with an exponentially large automaton M with a* visible *parity objective such that the following assertions hold:*

1. *G and G' have the same answer to the sure winning problem; and*
2. *the sure winning problem for G' can be solved in time exponential in the size of G.*

4 Complexity of Partial-Observation Parity Games

In this section we present a complete picture of the complexity results for the three different classes of partial-observation games, with different classes of parity objectives, both for qualitative and quantitative analysis.

4.1 Complexity of sure winning

We first show that for sure winning, pure strategies are sufficient for all partial-observation games.

Lemma 1 (Pure strategies suffice for sure winning). *For all games $G \in \mathcal{G}_{\mathsf{Pa}}$ and all objectives ϕ, if there is a sure winning strategy, then there is a pure sure winning strategy.*

Proof. Consider a randomized strategy σ for Player 1, let σ^P be the pure strategy such that for all $\rho \in \mathsf{Prefs}_1(G)$, the strategy $\sigma^P(\rho)$ chooses an action from $\mathrm{Supp}(\sigma(\rho))$. Then for all s we have $\mathsf{Outcome}_1(G, s, \sigma^P) \subseteq \mathsf{Outcome}_1(G, s, \sigma)$, and thus, if σ is sure winning, then so is σ^P. The result also holds for observation-based strategies. ∎

Spoiling strategies. To spoil a strategy of Player 1 (for sure-winning), Player 2 does not need the full memory of the history of the play, but only needs counting strategies [7]. We say that a pure strategy $\pi : \mathsf{Prefs}_2(G) \to A_2$ for Player 2 is *counting* if for all prefixes $\rho, \rho' \in \mathsf{Prefs}_2(G)$ such that $|\rho| = |\rho'|$ and $\mathsf{Last}(\rho) = \mathsf{Last}(\rho')$, we have $\pi(\rho) = \pi(\rho')$. Let Π_G^C be the set of counting strategies for Player 2. The memory needed by a counting strategy is only the number of turns that have been played. This type of strategy is sufficient to spoil the non-winning strategies of Player 1.

Lemma 2 (Counting spoiling strategies suffice). *Let G be a partial-observation game and ϕ be an objective. There exists a pure observation-based strategy $\sigma^o \in \Sigma_G^O$ such that for all $\pi^o \in \Pi_G^O$ we have $\mathrm{Outcome}(G, s, \sigma^o, \pi^o) \in \phi$ if and only if there exists a pure observation-based strategy $\sigma^o \in \Sigma_G^O$ such that for all counting strategies $\pi^c \in \Pi_G^C$ we have $\mathrm{Outcome}(G, s, \sigma^o, \pi^c) \in \phi$.*

Proof. We prove the equivalent statement that: for all pure observation-based strategies $\sigma^o \in \Sigma_G^O$ there exists $\pi^o \in \Pi_G^O$ such that $\mathrm{Outcome}(G, s, \sigma^o, \pi^o) \subsetneq \phi$ iff for all pure observation-based strategies $\sigma^o \in \Sigma_G^O$ there exists $\pi^c \in \Pi_G^C$ such that $\mathrm{Outcome}(G, s, \sigma^o, \pi^c) \subsetneq \phi$. The right implication ($\leftarrow$) is trivial. For the left implication ($\rightarrow$), let $\sigma^o \in \Sigma_G^O$ be an arbitrary pure observation-based strategy for Player 1 in G. Let $\pi^o \in \Pi_G^O$ be a strategy for Player 2 such that there exists $\rho^* \in \mathrm{Outcome}(G, s, \sigma^o, \pi^o)$ and $\rho^* \notin \phi$. Let $\rho^* = s_0 a_0 s_1 a_1 \ldots a_{n-1} s_n a_n \ldots$ and define a counting strategy π^c for Player 2 such that for all $\rho \in \mathsf{Prefs}_2(G)$ if $\mathsf{Last}(\rho) = s_{n-1}$ for $n = |\rho|$, then $\pi^c(\rho) = s_n$, and otherwise $\pi^c(\rho)$ is fixed arbitrarily. Clearly, π^c is a counting strategy and we have $\rho^* \in \mathrm{Outcome}(G, s, \sigma^o, \pi^o)$. Hence it follows that $\mathrm{Outcome}(G, s, \sigma^o, \pi^c) \subsetneq \phi$, and we obtain the desired result. ∎

Sure winning coincide for Pa **and** Os **games.** For all $\mathcal{O}_2$ partitions of a partial-observation game, a counting strategy is an observation-based strategy. From Lemma 1 it follows that pure strategies suffice for sure winning, and Lemma 2 shows that counting strategies suffice for spoiling pure strategies. Hence it follows that for spoiling strategies in sure winning games, the observation for Player 2 does not matter, and hence for sure winning, Pa and Os games coincide.

Lemma 3. *For a partial-observation game $G = \langle S_1 \cup S_2, A_1, A_2, \delta_1 \cup \delta_2, \mathcal{O}_1, \mathcal{O}_2 \rangle$ with an objective ϕ, consider the one-sided partial-observation game $G' = \langle S_1 \cup S_2, A_1, A_2, \delta_1 \cup \delta_2, \mathcal{O}_1, \mathcal{O}_2' \rangle$ such that $\mathcal{O}_2' = \{\{s\} \mid s \in S_2\}$. The answer to the sure winning questions in G and G' coincide for objective ϕ.*

Complexity of sure winning. The results for complete-observation games are as follows: (1) safety and reachability objectives can be solved in linear-time (this is alternating reachability in AND-OR graphs) [14]; (2) Büchi and coBüchi objectives can be solved in quadratic time [25]; and (3) parity objectives lie in NP $\cap$ coNP [10] and no polynomial time algorithm is known. The results for one-sided partial-observation games are as follows: (1) the EXPTIME-completeness for reachability objectives follows from the results of [22]; (2) the EXPTIME-completeness for safety objectives follows from the results of [4]; and (3) the EXPTIME-upper bound for all parity objective follows from the results of [7] and hence it follows that for all Büchi, coBüchi and parity objectives we have EXPTIME-complete bound. From Lemma 3 the results follow for partial-observation games. The results are summarized in the following theorem and shown in Table 1.

Theorem 2 (Complexity of sure winning). *The following assertions hold:*

1. *The sure winning problem for complete-observation games (i) with reachability and safety objectives can be solved in linear time; (ii) with Büchi and coBüchi objectives can be solved in quadratic time; and (iii) with parity objectives is in NP $\cap$ coNP.*

	Complete-observation	One-sided	Partial-observation
Safety	Linear-time	EXPTIME-complete	EXPTIME-complete
Reachability	Linear-time	EXPTIME-complete	EXPTIME-complete
Büchi	Quadratic-time	EXPTIME-complete	EXPTIME-complete
coBüchi	Quadratic-time	EXPTIME-complete	EXPTIME-complete
Parity	NP $\cap$ coNP	EXPTIME-complete	EXPTIME-complete

Table 1. Complexity of sure winning.

	Complete-observation	One-sided	Partial-observation
Safety	Linear-time	EXPTIME-complete	EXPTIME-complete
Reachability	Linear-time	EXPTIME-complete	2EXPTIME-complete
Büchi	Quadratic-time	EXPTIME-complete	2EXPTIME-complete
coBüchi	Quadratic-time	Undecidable	Undecidable
Parity	NP $\cap$ coNP	Undecidable	Undecidable

Table 2. Complexity of almost-sure winning.

2. *The sure winning problem for partial-observation and one-sided partial-observation games with reachability, safety, Büchi, coBüchi and parity objectives are EXPTIME-complete.*

4.2 Complexity of almost-sure winning

In contrast to sure winning (Lemma 1), for almost-sure winning, randomized strategies are more powerful than pure strategies (for example see [7]) for one-sided partial-observation games. The celebrated determinacy result of Martin [16] shows that for complete-observation games either there is a sure winning strategy for Player 1, or there is a pure strategy for Player 2 that ensures against all Player 1 strategies the objective is not satisfied. It follows that for complete-observation games, the almost-sure, limit-sure winning, and value decision problems coincide with the sure winning problem. For safety objectives, the counter-examples are always finite prefixes, and it can be shown that for a given observation-based strategy for Player 1 if there is a strategy for Player 2 to produce a finite counter-example, then the finite counter-example is produced with some constant positive probability. It follows that for partial-observation games and one-sided partial-observation games with safety objectives, the almost-sure and the limit-sure winning problems coincide with the sure winning problem.

Lemma 4. *The following assertions hold:*

1. *For complete-observation games, the almost-sure, limit-sure winning, and value decision problems coincide with the sure winning problem.*
2. *For safety objectives, the almost-sure and the limit-sure winning problems coincide with the sure winning problem for partial-observation and one-sided partial-observation games.*

	Complete-observation	One-sided	Partial-observation
Safety	Linear-time	EXPTIME-complete	EXPTIME-complete
Reachability	Linear-time	Undecidable	Undecidable
Büchi	Quadratic-time	Undecidable	Undecidable
coBüchi	Quadratic-time	Undecidable	Undecidable
Parity	NP $\cap$ coNP	Undecidable	Undecidable

Table 3. Complexity of limit-sure winning.

Complexity of almost-sure winning. In view of Lemma 4 the almost-sure winning analysis for complete-observation games with all classes of objectives follow from Theorem 2. Similarly due to Lemma 4 the results for partial-observation games and one-sided partial-observation games with safety objectives follow from Theorem 2. The EXPTIME-completeness for almost-sure winning with reachability and Büchi objectives for one-sided partial-observation games follows from [7]; and the 2EXPTIME-completeness for almost-sure winning with reachability and Büchi objectives for partial-observation games follows from [3, 12]. The undecidability result for almost-sure winning for coBüchi objectives for one-sided partial-observation games is obtained as follows: (i) in [2] it was shown that for probabilistic automata with coBüchi conditions, the problem of deciding if there exists a word that is accepted with probability 1 is undecidable and from this it follows that for one-sided partial-observation games with probabilistic transitions, the problem of deciding the existence of a pure observation-based almost-sure winning strategy is undecidable; (ii) it was shown in [6] that probabilistic transitions can be removed from the game graph, and the problem remains undecidable under randomized observation-based strategies. The undecidability for the more general parity objectives, and partial-observation games follows. This gives us the results for almost-sure winning, and they are summarized in the theorem below (see also Table 2).

Theorem 3 (Complexity of almost-sure winning). *The following assertions hold:*

1. *The almost-sure winning problem for one-sided partial-observation games (i) with safety, reachability and Büchi objectives are EXPTIME-complete, and (ii) is undecidable for coBüchi and parity objectives.*
2. *The almost-sure winning problem for partial-observation games (i) with safety, reachability and Büchi objectives are 2EXPTIME-complete, and (ii) is undecidable for coBüchi and parity objectives.*

4.3 Complexity of limit-sure winning and value decision problems

The complexity results for limit-sure winning and value decision problems are as follows.

Complexity of limit-sure winning. In view of Lemma 4 the results for (i) limit-sure winning and value decision problem for complete-observation games with all classes

	Complete-observation	One-sided	Partial-observation
Safety	Linear-time	Undecidable	Undecidable
Reachability	Linear-time	Undecidable	Undecidable
Büchi	Quadratic-time	Undecidable	Undecidable
coBüchi	Quadratic-time	Undecidable	Undecidable
Parity	NP $\cap$ coNP	Undecidable	Undecidable

Table 4. Complexity of value decision.

of objectives, and (ii) for partial-observation games and one-sided partial-observation games with safety objectives with limit-sure winning, follow from Theorem 2. It follows from the results of [11] that the following question is undecidable for probabilistic automata with reachability condition: for all $\varepsilon > 0$ is there a word w_ε that is accepted with probability greater than $1 - \varepsilon$? It follows that for one-sided partial-information games with probabilistic transitions, the problem of deciding the existence of a family of pure observation-based limit-sure winning strategies is undecidable; and again it follows from [6] that the problem is undecidable by removing probabilistic transitions from the game graph, and also for randomized observation-based strategies. Since (i) reachability objectives are special cases of Büchi, coBüchi and parity objectives, and (ii) one-sided partial-observation games are special cases of partial-observation games, the undecidability results for the more general cases follow. This gives us the results for limit-sure winning, and they are summarized in the theorem below (see also Table 3).

Theorem 4 (Complexity of limit-sure winning). *The following assertions hold:*

1. *The limit-sure winning problem for one-sided partial-observation games (i) with safety objectives are EXPTIME-complete, and (ii) with reachability, Büchi, coBüchi, and parity objectives are undecidable.*
2. *The limit-sure winning problem for partial-observation games (i) with safety objectives are EXPTIME-complete, and (ii) with reachability, Büchi, coBüchi, and parity objectives are undecidable.*

Complexity of the value decision problems. Since the limit-sure winning problem is a special case of the value decision problem (with $\nu = 1$), the undecidability results for all objectives other than safety objectives follow from Theorem 4. The undecidability of the value decision problem for probabilistic safety automata was shown in [8], and from [6] the undecidability follows for the value decision problem of one-sided partial-observation games with safety objectives. We summarize the results in Theorem 5 and Table 4.

Theorem 5 (Complexity of value decision problems). *The value decision problems for partial-observation and one-sided partial-observation games with safety, reachability, Büchi, coBüchi, and parity objectives are undecidable.*

5 The Complexity of Acyclic Games

We show that partial-observation games with reachability and safety objective played on acyclic graphs are PSPACE-complete. Note that for such games, the notion of sure-winning, almost-sure winning, and limit-sure winning coincide, and that randomized strategies are no more powerful than pure strategies.

A partial-observation game is *acyclic* if there are two distinguished sink states s_{acc} and s_{rej} (accepting and rejecting states) such that the transition relation is acyclic over $S \setminus \{s_{acc}, s_{rej}\}$. The objective is $\mathsf{Reach}(\{s_{acc}\})$ or equivalently $\mathsf{Safe}(S \setminus \{s_{rej}\})$. Clearly the winner of an acyclic game is known after at most $|S|$ rounds of playing. We claim that the qualitative analysis of acyclic partial-observation games (with reachability or safety objective) is PSPACE-complete. Since for acyclic games parity objectives reduce to safety or reachability objectives, the PSPACE-completeness follows for all parity objectives.

PSPACE upper bound. A PSPACE algorithm to solve acyclic games is as follows. Starting from $t_0 = \{s_0\}$, we choose an action $a \in A_1$ and we check that Player 1 is winning from each set $t_1 = \mathsf{Post}_a(t_0) \cap o_1$ for each $o_1 \in \mathcal{O}_1$ where $\mathsf{Post}_a(t) = \{s'' \mid \exists s \in t, b \in A_2 : \delta_2(\delta_1(s, a), b) = s''\}$. For each observation $o_1 \in \mathcal{O}_1$, we can reuse the space used by previous checks. Since the number of rounds is at most $|S_1|$, we can check if Player 1 is winning using a recursive procedure that tries out all choices of actions (the stack remains bounded by a polynomial in $|S_1|$).

PSPACE lower bound. We prove PSPACE-hardness using a reduction from QBF, which is the problem of evaluating a quantified boolean formula and is known to be PSPACE-complete [19]. A formula is defined over a finite set X of boolean variables, and is of the form $\varphi \equiv Q_1 x_1 \dots Q_n x_n \bigwedge_i c_i$, where $Q_k \in \{\exists, \forall\}$, $x_k \in X$ ($k = 1 \dots n$) and each clause c_i is of the form $u_1 \vee u_2 \vee u_3$ and u_j are literals (i.e., either a variable or the negation of a variable). We assume without loss of generality that all variables occurring in φ are quantified. Given a formula φ, we construct an acyclic game G_φ and state s_I such that Player 1 has a sure winning strategy in G_φ from s_I if and only if the formula φ is true.

The idea of the construction is as follows. Let us call Player 1 the $\exists$player and Player 2 the $\forall$player. In the game G_φ, the $\forall$player chooses a valuation of the universally quantified variables, and the $\exists$player chooses a valuation of the existentially quantified variables. The choices are made in alternation, according to the structure of the formula. Moreover, the $\forall$player (secretly) chooses one clause that he will monitor. Since the $\exists$player does not know which clause is chosen by the $\forall$player, she has to ensure that all clauses are satisfied by her choice of valuation.

To be fair, when the $\exists$player is asked to choose a value for an existentially quantified variable x, the $\forall$player should have announced the value he has chosen for the variables that are quantified before x in the formula. We use observations to simulate this.

Note that, having chosen a clause, the $\forall$player has a unique clever choice of the value of the universally quantified variables (namely such that the corresponding literals in the clause are all false). Indeed, for any other choice, the clause would be satisfied no matter the $\exists$player's choice, and there would be nothing to check.

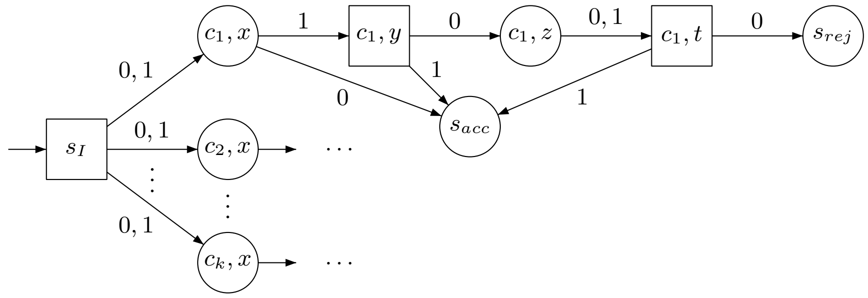

Fig. 1. Reduction of QBF to acyclic games for $\varphi = \exists x \forall y \exists z \forall t \ (x \vee \bar{y} \vee \bar{t}) \wedge \ldots$ Circles are states of $\exists$player, boxes are states of $\forall$player.

The reduction is illustrated in Fig.1. We formally describe below the game G_φ. W.l.o.g. we assume that the quantifiers in φ are alternating, i.e. φ is of the form $\exists x_1 \forall x_2 \ldots \exists x_{2n-1} \forall x_{2n} \bigwedge_i c_i$. The set of actions in G_φ is $A_1 = A_2 = \{0,1\}$ and the state space is $S_1 \cup S_2 \cup \{s_{acc}, s_{rej}\}$ where $S_1 = \{(c,x) \mid c \text{ is a clause in } \varphi \text{ and } x \text{ is an existentially quantified variable}\}$ and $S_2 = \{s_I\} \cup \{(c,x) \mid c \text{ is a clause in } \varphi \text{ and } x \text{ is a universally quantified variable}\}$. The transitions are as follows, for each clause c of φ:

- $(s_I, a, (c, x_1))$ for each $a \in \{0,1\}$. Intuitively, Player 2 initially chooses which clause he will track;
- $((c, x_i), a, s_{acc})$ for all $1 \leq i \leq 2n$ if $a = 0$ and $\bar{x}_i \in c$, or if $a = 1$ and $x_i \in c$. Intuitively, the state s_{acc} is reached if the assignment to variable x_i makes the clause c true;
- $((c, x_i), a, (c, x_{i+1}))$ for all $1 \leq i \leq 2n$ if $a = 0$ and $\bar{x}_i \notin c$, or if $a = 1$ and $x_i \notin c$ (and we assume that (c, x_{2n+1}) denotes s_{rej}). Intuitively, the state s_{rej} is reached if no literal in c is set to true by the players.

The set of observations for Player 1 is $\mathcal{O}_1 = \{\mathsf{init}\} \cup \{x = 0 \mid x \in X\} \cup \{x = 1 \mid x \in X\}$, and the observation function is defined by $\mathsf{obs}_1(c, x_1) = \mathsf{init}$ for all clauses c in φ, and $\mathsf{obs}_1(c, x_i) = \begin{cases} x_i = 1 & \text{if } x_{i-1} \notin c \\ x_i = 0 & \text{otherwise} \end{cases}$ for all clauses c in φ, and all $1 < i \leq n$. Intuitively, the $\exists$player does not know which clause is monitored by the $\forall$player, but knows the value assigned by the $\forall$player to the universally quantified variables.

The correctness of this construction is established as follows. First, assume that $\exists$player has a sure winning strategy in G_φ. Since strategies are observation-based, the action choice after a prefix of a play $s_I a_0 (c, x_1) a_1 \ldots (c, x_k)$ is independent of c and depends only on the sequence of previous actions and observations which provide the value of variables $x_1, \ldots, x_{k-1}$ only. Therefore we can view the winning strategy as a function that assigns to each existentially quantified variable a value that depends on the value of the variables quantified earlier in the formula. This function is a witness for showing that φ holds, since the state s_{rej} is avoided.

Conversely, if φ holds, then there is a strategy to assign a value to the existentially quantified variables given the value of the variables quantified earlier in the formula, from which it is easy to construct a winning strategy in G_φ to reach s_{acc}.

Thus PSPACE-completeness follows for one-sided partial-observation games for sure winning. Since sure, almost-sure, and limit-sure winning coincide for acyclic games, and for sure winning partial-observation games coincide with one-sided partial-observation games (Lemma 3), the PSPACE-completeness for all the qualitative analysis problems follow.

Theorem 6 (Complexity of acyclic games). *The sure, almost-sure, and limit-sure winning problems for acyclic games of partial observation and one-sided partial observation with all parity objectives are PSPACE-complete.*

References

1. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.
2. C. Baier, N. Bertrand, and M. Größer. On decision problems for probabilistic Büchi automata. In *FoSSaCS*, LNCS 4962, pages 287–301. Springer, 2008.
3. N. Bertrand, B. Genest, and H. Gimbert. Qualitative determinacy and decidability of stochastic games with signals. In *LICS*, pages 319–328. IEEE Computer Society, 2009.
4. D. Berwanger and L. Doyen. On the power of imperfect information. In *FSTTCS*, Dagstuhl Seminar Proceedings 08004. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), 2008.
5. J.R. Büchi and L.H. Landweber. Solving sequential conditions by finite-state strategies. *Transactions of the AMS*, 138:295–311, 1969.
6. K. Chatterjee, L. Doyen, H. Gimbert, and T.A. Henzinger. Randomness for free. In *MFCS*. Springer, 2010.
7. K. Chatterjee, L. Doyen, T. A. Henzinger, and J.-F. Raskin. Algorithms for omega-regular games of incomplete information. *Logical Methods in Computer Science*, 3(3:4), 2007.
8. K. Chatterjee and T.A. Henzinger. Probabilistic automata on infinite words: Decidability and undecidability results. In *ATVA*. Springer, 2010.
9. L. de Alfaro and T.A. Henzinger. Interface theories for component-based design. In *EMSOFT*, LNCS 2211, pages 148–165. Springer, 2001.
10. E.A. Emerson and C. Jutla. Tree automata, mu-calculus and determinacy. In *FOCS*, pages 368–377. IEEE, 1991.
11. Hugo Gimbert and Youssouf Oualhadj. Probabilistic automata on finite words: Decidable and undecidable problems. In *ICALP*, LNCS 6199, pages 527–538. Springer, 2010.
12. V. Gripon and O. Serre. Qualitative concurrent stochastic games with imperfect information. In *ICALP(2)*, LNCS 5556, pages 200–211. Springer, 2009.
13. T. A. Henzinger, O. Kupferman, and S. Rajamani. Fair simulation. *Information and Computation*, 173:64–81, 2002.
14. N. Immerman. Number of quantifiers is better than number of tape cells. *Journal of Computer and System Sciences*, 22:384–406, 1981.
15. A. Kechris. *Classical Descriptive Set Theory*. Springer, 1995.
16. D. A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.
17. D. A. Martin. The determinacy of Blackwell games. *The Journal of Symbolic Logic*, 63(4):1565–1581, 1998.

18. R. McNaughton. Infinite games played on finite graphs. *Annals of Pure and Applied Logic*, 65:149–184, 1993.
19. C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1993.
20. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190. ACM Press, 1989.
21. P. J. Ramadge and W. M. Wonham. Supervisory control of a class of discrete-event processes. *SIAM Journal of Control and Optimization*, 25(1):206–230, 1987.
22. J. H. Reif. Universal games of incomplete information. In *STOC*, pages 288–308. ACM Press, 1979.
23. J. H. Reif. The complexity of two-player games of incomplete information. *Journal of Computer and System Sciences*, 29(2):274–301, 1984.
24. S. Safra. On the complexity of ω-automata. In *FOCS*, pages 319–327. IEEE Computer Society Press, 1988.
25. W. Thomas. Languages, automata, and logic. In *Handbook of Formal Languages*, volume 3, Beyond Words, chapter 7, pages 389–455. Springer, 1997.
26. M. Y. Vardi. Automatic verification of probabilistic concurrent finite-state systems. In *FOCS*, pages 327–338. IEEE Computer Society Press, 1985.