

A Survey of Partial-Observation Stochastic Parity Games

Krishnendu Chatterjee · Laurent Doyen · Thomas
A. Henzinger

the date of receipt and acceptance should be inserted later

Abstract We consider two-player zero-sum stochastic games on graphs with ω -regular winning conditions specified as parity objectives. These games have applications in the design and control of reactive systems. We survey the complexity results for the problem of deciding the winner in such games, and in classes of interest obtained as special cases, based on the information and the power of randomization available to the players, on the class of objectives and on the winning mode.

On the basis of information, 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). The one-sided partial-observation games have two important subclasses: the one-player games, known as partial-observation Markov decision processes (POMDPs), and the blind one-player games, known as probabilistic automata.

On the basis of randomization, (a) the players may not be allowed to use randomization (pure strategies), or (b) they may choose a probability distribution over actions but the actual random choice is external and not visible to the player (actions invisible), or (c) they may use full randomization.

Finally, various classes of games are obtained by restricting the parity objective to a reachability, safety, Büchi, or coBüchi condition. We also consider several winning modes, such as sure-winning (i.e., all outcomes of a strategy have to satisfy the winning condition), almost-sure winning (i.e., winning with probability 1), limit-sure winning (i.e., winning with probability arbitrarily close to 1), and value-threshold winning (i.e., winning with probability at least ν , where ν is a given rational).

K. Chatterjee
krishnendu.chatterjee@ist.ac.at
IST Austria (Institute of Science and Technology Austria)

L. Doyen
doyen@lsv.ens-cachan.fr
LSV, ENS Cachan & CNRS, France

T. A. Henzinger
tah@ist.ac.at
IST Austria (Institute of Science and Technology Austria)

1 Introduction

Games on graphs. Games played on graphs provide the mathematical framework to analyze several important problems in computer science as well as in mathematics. In particular, the synthesis problem (Church’s problem) can be reduced to the construction of a winning strategy in a game played on a graph where the vertices and edges of the graph represent the states and transitions of a reactive system [7,36,34,32]. Game-theoretic formulations have also proved useful for the verification [1], refinement [26], and compatibility checking [22] of reactive systems. Games played on graphs are dynamic games that proceed for an infinite number of rounds. In each round, one of the players chooses a move which, together with the current state, determine a probability distribution over 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 prefix. 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. If the randomized strategy is allowed to observe the result of the move chosen from the probability distribution, then it is a randomized *action-visible* strategy; otherwise it is a randomized *action-invisible* strategy. Objectives are generally Borel measurable functions [31]: the objective for a player is a Borel set B in the Cantor topology on Q^ω (where Q is the set of states), and the objective is satisfied if 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 $\Sigma_3 \cap \Pi_3$, called the first $2^{1/2}$ -levels of the Borel hierarchy) and they form a robust and expressive language for determining payoffs for commonly used specifications. We consider the parity objective and its subclasses, which is a canonical form to express ω -regular 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 states visited and the actions played in the game; (b) *one-sided partial-observation* games, where one player has partial knowledge and the other player has complete knowledge; and (c) *complete-observation* games, where each player has complete knowledge of the play. Partial-observation games are useful to model many important problems related to reactive systems, for example, the interaction of a plant and controller when the plant has private variables not accessible to the controller. The class of one-sided partial-observation games has two important subclasses when there is only one player: (a) one-player partial-observation games, called partial-observation Markov decision processes (POMDPs); and (b) the special case of POMDPs with a single observation, which are called blind POMDPs or probabilistic automata [35].

Analysis. The analysis of games can be classified into *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 that the objective is satisfied with certainty against all strategies of player 2 (*sure winning* problem); (b) can player 1 ensure that the objective is satisfied with probability 1 against all strategies of player 2 (*almost-sure winning* problem); and (c) can player 1 ensure that the objective is satisfied with probability arbitrarily close to 1 against all strategies of player 2 (*limit-sure winning* problem). Given an objective, an initial 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 .

Structure of the paper. We survey the main complexity results about various classes of partial-observation games, for different classes of parity objectives, for all three types of strategies (pure, randomized action-visible, and randomized action-invisible). We organize the results according to sure winning (see Table 1), almost-sure winning (see Table 2 and Table 3), limit-sure winning (see Table 4), and quantitative analysis (see Table 5).

2 Definitions

We define the class of two-player partial-observation stochastic games, and we consider some relevant subclasses.

A *probability distribution* on a finite set S is a function $\kappa : S \rightarrow [0, 1]$ such that $\sum_{s \in S} \kappa(s) = 1$. The *support* of κ is the set $\text{Supp}(\kappa) = \{s \in S \mid \kappa(s) > 0\}$. We denote by $\mathcal{D}(S)$ the set of probability distributions on S .

Stochastic games. Given finite alphabets A_i of actions for player i ($i = 1, 2$), a *two-player stochastic game* (or simply a *game*) is a tuple $G = \langle Q_1, Q_2, \delta \rangle$ where

- Q_1 is a finite state of player-1 states, Q_2 is a finite state of player-2 states, with $Q_1 \cap Q_2 = \emptyset$.
- $\delta : Q_i \times A_i \rightarrow \mathcal{D}(Q_{3-i})$ is a probabilistic transition function.

Let $Q = Q_1 \cup Q_2$. Games are played for infinitely many rounds from an initial state $q_0 \in Q$. In each round, if the current state is $q \in Q_i$ ($i \in \{1, 2\}$), then player i chooses an action $a \in A_i$ and the game proceeds to the next round in state q' with probability $\delta(q, a)(q')$. In this way, the players construct an infinite path called a *play*.

The game is *deterministic* if for $i = 1, 2$, for all states $q \in Q_i$ and all actions $a \in A_i$, there exists a state $q' \in Q_{3-i}$ such that $\delta(q, a)(q') = 1$.

Partial-observation stochastic games. In partial-observation games, the players have a partial or incomplete view of the states visited in the game, and the views of the two players are in general different. A *partial-observation two-player stochastic game* consists of a game $G = \langle Q_1, Q_2, \delta \rangle$ and two sets $\mathcal{O}_i \subseteq 2^Q$ ($i = 1, 2$) of *observations* for player i that define two partitions of Q . The states in an observation of \mathcal{O}_i are indistinguishable for player i . These partitions uniquely define two functions $\text{obs}_i : Q \rightarrow \mathcal{O}_i$ ($i = 1, 2$) that map each state $q \in Q$ to its observation for player i such that $q \in \text{obs}_i(q)$. A model of games where states may have multiple observations (i.e., the observation sets can overlap) can be encoded in the model where observations form a partition of the state space [17].

Fig. 1 shows an example of a two-player stochastic game where the player-1 states $Q_1 = \{q_1, q_2, q_5\}$ are circles, and the player-2 states $Q_2 = \{q_0, q_3, q_4, \odot\}$ are boxes. The transition from q_1 on a is such that $\delta(q_1, a)(q_3) = \delta(q_1, a)(\odot) = \frac{1}{2}$. Observations (not depicted) are $\mathcal{O}_1 = \{\{q_0\}, \{q_i \mid 1 \leq i \leq 5\}, \{\odot\}\}$ for player 1, and $\mathcal{O}_2 = \{\{q_0\}, \{q_1\}, \{q_2\} \dots\}$ for player 2. Hence, player 1 can observe the initial state and the state \odot , while player 2 can distinguish all states.

Special cases. We consider the following subclasses of partial-observation stochastic games:

- (*Observation restriction*). The games with *one-sided partial-observation* are the special case of games where $\mathcal{O}_2 = \{\{q\} \mid q \in Q\}$ (equivalently, $\text{obs}_2(q) = \{q\}$ for all $q \in Q$), i.e. player 2 has complete observation, and only player 1 has partial-observation.




Fig. 1 A two-player stochastic game.

The *games of complete observation* are the special case of games where $\mathcal{O}_1 = \mathcal{O}_2 = \{\{q\} \mid q \in Q\}$, i.e. every state is visible to each player and hence both players have complete observation. Two-player games of complete observation are also called perfect-information stochastic games, or $2^{1/2}$ -player games.

- (*Player restriction*). The *partial-observation Markov decision processes* (POMDP) are the special case of games where A_2 is a singleton, i.e. only player 1 has multiple choice. An important subclass of POMDP is obtained when player 1 has only one observation (i.e., $\mathcal{O}_1 = \{Q\}$) and is called blind POMDPs or *probabilistic automata* [35].

Strategies. Strategies are recipes used by the players to choose their action. In partial-observation games, it is more powerful to use randomization in strategies, and when randomization is used it is more powerful to observe which action was (randomly) chosen, as compared to the case where the action is not visible. Therefore, we consider the following classes of strategies.

- *Randomized action-visible*. Let $A_i^\varepsilon = A_i \cup \{\varepsilon\}$. A *randomized action-visible strategy* in G for player i is a function $\sigma_i : (\mathcal{O}_i \cdot A_i^\varepsilon)^* \cdot \mathcal{O}_i \rightarrow \mathcal{D}(A_i)$.
- *Randomized action-invisible*. A *randomized action-invisible strategy* in G for player i is a function $\sigma_i : \mathcal{O}_i^+ \rightarrow \mathcal{D}(A_i)$ where \mathcal{O}_i^+ is the set of nonempty sequences of observations.
- *Pure*. A *pure strategy* in G for player i is a randomized action-invisible strategy such that for all observation sequences $\rho \in \mathcal{O}_i^+$, there exists an action $a \in A_i$ such that $\sigma_i(\rho)(a) = 1$.

For pure strategies, it does not matter whether the actions are visible or not because given a pure strategy σ_i , and a finite prefix $\rho \in \mathcal{O}_i^+$, the actions played by player i along ρ can be completely reconstructed.

Note that games where the players use randomized action-visible strategies can be encoded into games where the actions are invisible, in which the last action played is stored in the state space, and the observations on states are enriched to reveal the last action. This encoding has a polynomial blow-up. The same trick can be used to encode games where the actions can be partially observed (e.g., through observations on the action sets) into games with actions invisible [25]. Hence the model with action invisible is more general.

Memory of strategies. A randomized action-invisible strategy for player i uses *finite-memory* if it can be encoded by a deterministic transducer $\langle \text{Mem}, m_0, \alpha_u, \alpha_n \rangle$ where Mem is a finite set (the memory of the strategy), $m_0 \in \text{Mem}$ is the initial memory value, $\alpha_u : \text{Mem} \times \mathcal{O}_i \rightarrow \text{Mem}$ is an update function, and $\alpha_n : \text{Mem} \times \mathcal{O}_i \rightarrow \mathcal{D}(A_i)$ is a next-move function. The *memory size* of the strategy is the number $|\text{Mem}|$ of memory values. The strategy is *memoryless* if $|\text{Mem}| = 1$. If the current observation is $o \in \mathcal{O}_i$, and the current memory value is m , then the strategy chooses the next action according to the probability distribution $\alpha_n(m, o)$, and the memory is updated to $\alpha_u(m, o)$. Formally, $\langle \text{Mem}, m_0, \alpha_u, \alpha_n \rangle$ defines the strategy σ such that $\sigma(\rho \cdot o) = \alpha_n(\hat{\alpha}_u(m_0, \text{obs}_1(\rho)), o)$ for all $\rho \in \mathcal{O}_i^*$ and $o \in \mathcal{O}_i$, where $\hat{\alpha}_u$ extends α_u to sequences of observations as expected. The memory used by randomized action-visible strategies can be defined analogously.

Outcome. Strategies induce a set of plays and a (probability) measure μ over play prefixes, defined as follows:

- For randomized action-visible strategies, we lift the functions obs_i to sequences of interleaved states and actions: for $i = 1, 2$, denote by $\text{obs}_i(q_0 a_0 q_1 a_1 \dots q_k)$ the sequence $\text{obs}_i(q_0) a'_0 \text{obs}_i(q_1) a'_1 \dots \text{obs}_i(q_k)$ where $a'_j = a_j$ if $q_j \in Q_i$, and $a'_j = \varepsilon$ if $q_j \in Q_{3-i}$. Let $\text{state}(q_0 a_0 q_1 a_1 \dots q_k)$ be the sequence $q_0 q_1 \dots q_k$ (i.e., the projection that removes all actions).

A play $\rho = q_0 q_1 \dots$ is an *outcome* of the randomized action-visible strategies σ_1 for player 1 and σ_2 for player 2 from a state $q \in Q$ if $q_0 = q$ and there exists a sequence $\pi = q_0 a_0 q_1 a_1 \dots q_k$ such that $\rho = \text{state}(\pi)$ and for all $j \geq 0$, if $q_j \in Q_i$ then $\sigma_i(\text{obs}_i(q_0 a_0 q_1 \dots q_j))(a_j) > 0$, and $\delta(q_j, a_j)(q_{j+1}) > 0$. We denote by $\text{Outcome}(G, q, \sigma_1, \sigma_2)$ the set of all outcomes of the strategies σ_1 and σ_2 from q . Define the function $\xi_q^{\sigma_1, \sigma_2} : (Q \cdot (A_1 \cup A_2))^* \cdot Q \rightarrow [0, 1]$ inductively as follows. Let $\xi_q^{\sigma_1, \sigma_2}(q_0) = 1$ if $q_0 = q$, and $\xi_q^{\sigma_1, \sigma_2}(q_0) = 0$ if $q_0 \neq q$. For all $\pi \in (Q \cdot (A_1 \cup A_2))^* \cdot Q$, and $q' \in Q$, let $i \in \{1, 2\}$ such that $\text{Last}(\pi) \in Q_i$ (where $\text{Last}(\pi)$ is the last state in the sequence π), and define

$$\xi_q^{\sigma_1, \sigma_2}(\pi \cdot a \cdot q') = \begin{cases} \xi_q^{\sigma_1, \sigma_2}(\pi) \cdot \sigma_i(\text{obs}_i(\pi))(a) \cdot \delta(\text{Last}(\pi), a)(q') & \text{if } a \in A_i \\ 0 & \text{if } a \in A_{3-i} \end{cases}$$

From $\xi_q^{\sigma_1, \sigma_2}$, define the function $\mu_q^{\sigma_1, \sigma_2} : Q^+ \rightarrow [0, 1]$ by

$$\mu_q^{\sigma_1, \sigma_2}(\rho) = \sum_{\pi \in (Q \cdot (A_1 \cup A_2))^* \cdot Q \mid \text{state}(\pi) = \rho} \xi_q^{\sigma_1, \sigma_2}(\pi).$$

- For randomized action-invisible strategies, we lift the functions obs_i to sequences of states and denote by $\text{obs}_i(q_0 q_1 \dots q_k)$ the sequence $\text{obs}_i(q_0) \text{obs}_i(q_1) \dots \text{obs}_i(q_k)$. A play $\rho = q_0 q_1 \dots$ is an *outcome* of the strategies σ_1 for player 1 and σ_2 for player 2 from a state $q \in Q$ if $q_0 = q$ and for all $j \geq 0$, if $q_j \in Q_i$ then there exists $a \in A_i$ such that $\sigma_i(\text{obs}_i(q_0 \dots q_j))(a) > 0$, and $\delta(q_j, a)(q_{j+1}) > 0$. We denote by $\text{Outcome}(G, q, \sigma_1, \sigma_2)$ the set of all outcomes of the strategies σ_1 and σ_2 from q . Define the function $\mu_q^{\sigma_1, \sigma_2} : Q^+ \rightarrow [0, 1]$ inductively as follows. Let $\mu_q^{\sigma_1, \sigma_2}(q_0) = 1$ if $q_0 = q$, and $\mu_q^{\sigma_1, \sigma_2}(q_0) = 0$ if $q_0 \neq q$. For all $\rho \in Q^+$ and $q \in Q$, let $i \in \{1, 2\}$ such that $\text{Last}(\rho) \in Q_i$ (where $\text{Last}(\rho)$ is the last state in the sequence ρ), and define

$$\mu_q^{\sigma_1, \sigma_2}(\rho \cdot q) = \mu_q^{\sigma_1, \sigma_2}(\rho) \cdot \sum_{a_i \in A_i} \sigma_i(\text{obs}_i(\rho))(a_i) \cdot \delta(\text{Last}(\rho), a_i)(q).$$

- The definition of outcome for pure strategies is derived as a special case of the definition of outcome for randomized action-invisible strategies.

By Caratheodary’s extension theorem, the function $\mu_q^{\sigma_1, \sigma_2}$ can be uniquely extended to a probability measure $\Pr_q^{\sigma_1, \sigma_2}(\cdot)$ over Borel sets of infinite plays [39, 6].

Objectives. An *objective* in G is a Borel set $\varphi \subseteq Q^\omega$ of plays in the Cantor topology on Q^ω [28]. All objectives we consider in this paper lie in the first $2^{1/2}$ -levels of the Borel hierarchy. We specifically consider the parity objective, which is a canonical form to express all ω -regular objectives [38]. For a play $\rho = q_0q_1q_2 \dots$, denote by $\text{Inf}(\rho) = \{q \in Q \mid \forall i \geq 0 \cdot \exists j \geq i : q_j = q\}$ the set of states that occur infinitely often in ρ . We consider the following objectives.

- *Reachability and safety objectives.* Given a set $T \subseteq Q$ of target states, the *reachability* objective $\text{Reach}(T) = \{q_0q_1q_2 \dots \in Q^\omega \mid \exists k \geq 0 : q_k \in T\}$ requires that a target state in T be visited at least once. Dually, the *safety* objective $\text{Safe}(T) = \{q_0q_1q_2 \dots \in Q^\omega \mid \forall k \geq 0 : q_k \in T\}$ requires that only states in T be visited.
- *Büchi and coBüchi objectives.* Given a set $T \subseteq Q$ of target states The *Büchi* objective $\text{Büchi}(T) = \{\rho \in Q^\omega \mid \text{Inf}(\rho) \cap T \neq \emptyset\}$ requires that a state in T be visited infinitely often. Dually, the *coBüchi* objective $\text{coBüchi}(T) = \{\rho \in Q^\omega \mid \text{Inf}(\rho) \subseteq T\}$ requires that only states in T be visited infinitely often.
- *Parity objectives.* For $d \in \mathbb{N}$, let $p : Q \rightarrow \{0, 1, \dots, d\}$ be a *priority function* that maps each state to a nonnegative integer priority. The *parity* objective $\text{Parity}(p) = \{\rho \in Q^\omega \mid \min\{p(q) \mid q \in \text{Inf}(\rho)\} \text{ is even}\}$ requires that the smallest priority that appears infinitely often be even.

Note that a reachability objective $\text{Reach}(T)$ can be viewed as a special case of Büchi objectives, and safety objectives are special cases of coBüchi objectives.

And the objectives $\text{Büchi}(T)$ and $\text{coBüchi}(T)$ are special cases of parity objectives defined by respective priority functions p_1, p_2 such that $p_1(q) = 0$ and $p_2(q) = 2$ if $q \in T$, and $p_1(q) = p_2(q) = 1$ otherwise.

Winning modes. Given a game structure G , a state q , and an objective φ , and classes $\mathcal{C}_1, \mathcal{C}_2$ of strategies, we say that:

- a strategy $\sigma_1 \in \mathcal{C}_1$ for player 1 is *sure winning* if $\text{Outcome}(G, q, \sigma_1, \sigma_2) \subseteq \varphi$ for all strategies $\sigma_2 \in \mathcal{C}_2$ for player 2,
- a strategy $\sigma_1 \in \mathcal{C}_1$ for player 1 is *almost-sure winning* if $\Pr_q^{\sigma_1, \sigma_2}(\varphi) = 1$ for all strategies $\sigma_2 \in \mathcal{C}_2$ for player 2,
- player 1 is *limit-sure winning* if for all $\epsilon > 0$ there exists a strategy $\sigma_1 \in \mathcal{C}_1$ for player 1 such that for all strategies $\sigma_2 \in \mathcal{C}_2$ for player 2, we have $\Pr_q^{\sigma_1, \sigma_2}(\varphi) \geq 1 - \epsilon$,
- the *value function* $\langle\langle 1 \rangle\rangle_{\text{val}}^G(\varphi) : Q \rightarrow \mathbb{R}$ for player 1 assigns to every state the maximal probability with which player 1 can guarantee that φ using a strategy in \mathcal{C}_1 , against all strategies of \mathcal{C}_2 for player 2. Formally, let

$$\langle\langle 1 \rangle\rangle_{\text{val}}^G(\varphi)(q) = \sup_{\sigma_1 \in \mathcal{C}_1} \inf_{\sigma_2 \in \mathcal{C}_2} \Pr_q^{\sigma_1, \sigma_2}(\varphi).$$

In the above definitions, the class \mathcal{C}_2 of player-2 strategies can be taken as any of the three classes of strategies we have defined (i.e., pure, randomized action-invisible, or randomized action-visible) without changing the definition. This is because once a strategy of player 1 is fixed, we obtain a (possibly infinite-state) POMDP in which pure strategies are sufficient for parity objectives [15].

For limit-sure and the value function, the complexity and decidability results presented in this survey are independent of the choice of the class \mathcal{C}_1 among the three classes we have defined.

Qualitative and quantitative analysis. We are interested in the problems of deciding, given a game G , a state q , and an objective φ , whether there exists a {pure, randomized action-invisible, randomized action-visible} strategy for player 1 that is {sure, almost-sure} winning from q for the objective φ , or whether player 1 is limit-sure winning from q for the objective φ .

Given a rational value $0 \leq \nu \leq 1$, the *value decision problem* asks whether the value $\langle\langle 1 \rangle\rangle_{val}^G(\varphi)(q)$ of the game at q 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.

Consider the objective $\text{Buchi}(\{\odot\})$ in the example of Fig. 1. There is no sure winning strategy for player 1 for this objective, but player 1 has an almost-sure winning strategy that uses randomization (play everywhere a and b with probability $\frac{1}{2}$), as well as a pure almost-sure winning strategy (alternate playing a and b). Hence for the three classes of strategies, we have $\langle\langle 1 \rangle\rangle_{val}^G(\text{Buchi}(\{\odot\}))(q_0) = 1$.

3 The Complexity of Partial-Observation Stochastic Parity Games

In this section we present a survey of results about the computational complexity and strategy complexity of different classes of partial-observation games, with different classes of parity objectives and strategies, both for qualitative and quantitative analysis. We organize the results as follows: we first present the results for sure winning, then for almost-sure winning, followed by limit-sure winning, and finally the quantitative analysis (value problem).

3.1 Complexity of sure winning

In this section we consider partial-observation games and its subclasses with parity objectives and the winning mode is sure winning. We first present a simple result from the literature that shows that for sure winning, pure strategies are sufficient for all partial-observation games.

Pure strategies suffice for sure winning. The key argument to show that pure strategies suffice for sure winning is as follows. Consider a randomized action-visible strategy σ_1 for player 1, let σ_1^P be a pure strategy such that for all finite prefixes ρ , the strategy $\sigma_1^P(\rho)$ chooses an action from $\text{Supp}(\sigma_1(\rho))$. Then for all states q and all strategies σ_2 for player 2, we have $\text{Outcome}(G, q, \sigma_1^P, \sigma_2) \subseteq \text{Outcome}(G, q, \sigma_1, \sigma_2)$, and thus if σ_1 is sure winning, then so is σ_1^P .

Counting 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 [17]. We say that a pure strategy σ_2 for player 2 is *counting* if for all finite prefixes ρ, ρ' such that $|\rho| = |\rho'|$, we have $\sigma_2(\rho) = \sigma_2(\rho')$. 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 [17].

Sure winning coincide for partial-observation and one-sided games. For all partitions \mathcal{O}_2 of a partial-observation game, counting strategies are valid strategies. Since pure strategies

suffice for sure winning and counting strategies suffice for spoiling pure strategies, it follows that for spoiling strategies in sure winning games, the observation for player 2 does not matter, and hence for sure winning, partial-observation and one-sided games coincide.

Computational complexity of sure winning. We now summarize the results related to the computational complexity of sure winning in various classes of partial-observation games. The following basic facts follow from the above argument: (1) For sure winning, pure strategies are as powerful as randomized action-visible strategies, and hence we only focus on pure strategies. (2) For sure winning, partial-observation games are equivalent as the special class of one-sided partial-observation games.

1. *Complete-observation games.* 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) [27]; (2) Büchi and coBüchi objectives can be solved in quadratic time [38]; and (3) parity objectives lie in $NP \cap coNP$ and no polynomial time algorithm is known [23]. The results for sure winning are derived as follows: for sure winning since all paths need to satisfy the objective, the probabilistic choices can be interpreted as the choice of the adversary (player 2) and then we obtain complete-observation deterministic games, and all results follow from the results for complete-observation deterministic games.
2. *Probabilistic automata and POMDPs.* The sure winning problem for probabilistic automata is the same as the blind games problem considered by Reif [37], and the POMDP problem is same as the sure winning problem for one-sided deterministic games (again interpreting the probabilistic choice as adversarial choice). The results for blind games are as follows: using the subset construction technique, Reif showed that the blind games problem can be solved as a graph problem on an exponentially large graph for reachability objectives [37]. It follows from [17] that the same (subset construction) technique also works for parity objectives, and [5] presents a reduction of parity objectives to safety objectives. As a consequence it follows that the sure winning problem for probabilistic automata with parity objectives can be solved in PSPACE, and the PSPACE lower bound follows from the universality problem for non-deterministic finite automata [37]. Hence we have PSPACE-completeness for the sure winning problem for probabilistic automata. The result for POMDPs is the same as for one-sided games discussed below.
3. *One-sided and partial-observation games.* The results for one-sided partial-observation games are as follows: (1) the EXPTIME-completeness for reachability objectives follows from the results of [37]; (2) the EXPTIME-completeness for safety objectives follows from the results of [5]; and (3) the EXPTIME-upper bound for all parity objectives follows from the results of [17] and hence it follows that for all Büchi, coBüchi and parity objectives we have EXPTIME-complete bound. Again the EXPTIME upper bound follows by constructing and solving an exponential-size game of complete observation. The lower bound follows from a reduction of the membership problem of alternating polynomial space Turing machines. The results follow for partial-observation games as sure winning for partial-observation games coincides with sure winning for one-sided partial-observation games.

The results are summarized in Theorem 1 and shown in Table 1.

Strategy complexity of sure winning. The sure winning problem for complete-observation games is the same as complete-observation deterministic games, and hence the existence of pure memoryless sure winning strategies for complete-observation games with parity objectives follows from the results of [23]. For probabilistic automata, POMDP, and partial-observation games, sure-winning strategies of exponential size suffices: the witness sure

	Complete-observation	Prob. Automata	POMDP	One-sided	Partial-observation
Safety	Linear-time	PSPACE-comp.	EXPTIME-comp.	EXPTIME-comp.	EXPTIME-comp.
Reachability	Linear-time	PSPACE-comp.	EXPTIME-comp.	EXPTIME-comp.	EXPTIME-comp.
Büchi	Quadratic-time	PSPACE-comp.	EXPTIME-comp.	EXPTIME-comp.	EXPTIME-comp.
coBüchi	Quadratic-time	PSPACE-comp.	EXPTIME-comp.	EXPTIME-comp.	EXPTIME-comp.
Parity	$NP \cap coNP$	PSPACE-comp.	EXPTIME-comp.	EXPTIME-comp.	EXPTIME-comp.

Table 1 Complexity of sure winning.

	Complete-observation	Prob. Automata	POMDP	One-sided	Partial-observation
Safety	Linear-time	PSPACE-comp.	EXPTIME-comp.	EXPTIME-comp.	EXPTIME-comp.
Reachability	Linear-time	PSPACE-comp.	EXPTIME-comp.	EXPTIME-comp.	2EXPTIME-comp.
Büchi	Quadratic-time	PSPACE-comp.	EXPTIME-comp.	EXPTIME-comp.	2EXPTIME-comp.
coBüchi	Quadratic-time	Undecidable	Undecidable	Undecidable	Undecidable
Parity	$NP \cap coNP$	Undecidable	Undecidable	Undecidable	Undecidable

Table 2 Complexity of almost-sure winning: randomized action-visible strategies.

winning strategy is based on the subset construction where the subset represents the *belief* of the player (where belief represents the possible current states of the game). It follows that pure belief-based sure winning strategies exist. The exponential lower bound for memory follows from the fact that the shortest word to witness the non-universality of a non-deterministic finite automata is of exponential size.

Theorem 1 (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 $NP \cap coNP$.*
2. *The sure winning problem for probabilistic automata with reachability, safety, Büchi, coBüchi and parity objectives are PSPACE-complete.*
3. *The sure winning problem for POMDPs, one-sided partial-observation games, and partial-observation games with reachability, safety, Büchi, coBüchi and parity objectives are EXPTIME-complete.*
4. *Pure memoryless sure-winning strategies exist in complete-observation parity games. Sure-winning strategies with exponential memory is sufficient for probabilistic automata, POMDP, one-sided partial-observation games, and partial-observation games with parity objectives, and exponential memory is necessary for sure-winning strategies in probabilistic automata with safety and reachability objectives.*

3.2 Complexity of almost-sure winning

In this section, we discuss the results for almost-sure winning. In contrast to sure winning in partial-observation games where all problems are decidable, the almost-sure winning problem becomes undecidable in many cases. We start with the main undecidability result related to almost-sure winning.

Undecidability result for almost-sure winning. The main undecidability result for the almost-sure winning problem was established in [2] where it was shown that the problem of deciding the existence of a pure almost-sure winning strategy for probabilistic automata

with coBüchi objectives is undecidable. The key steps of this deep result are: (1) First the authors consider the undecidability of the approximation problem for probabilistic automata over finite words which shows that given a probabilistic automata on finite words and a constant $0 < \varepsilon < 1$ such that either there is a word accepted with probability $1 - \varepsilon$ or all words are accepted with probability at most ε , decide which is the case, is undecidable [30]. (2) Then the approximation problem is reduced to the problem of checking, given two probabilistic Büchi automata, whether there is a word that is accepted with positive probability by both the automata. (3) Using the closure properties of probabilistic Büchi automata [3], it follows that the problem of deciding the existence of a word (which is same as a pure strategy for probabilistic automata) that is accepted with positive probability is undecidable for probabilistic automata with Büchi conditions. As the coBüchi condition is dual to the Büchi condition, and almost-sure acceptance criteria is dual to acceptance with positive probability, it follows that the problem of deciding the existence of pure almost-sure winning strategy for probabilistic automata with coBüchi objectives is undecidable.

Undecidability result for randomized strategies. As discussed above, the existence of pure almost-sure winning strategy for probabilistic automata with coBüchi objectives is undecidable. It was established in [15] that for POMDPs pure strategies are as powerful as randomized action-visible strategies for all objectives (i.e., for one-player games randomization is not useful for strategies). As a consequence it follows that the existence of randomized action-visible almost-sure winning strategy for probabilistic automata coincides with the existence of pure almost-sure winning strategies. Hence it follows that the existence of almost-sure winning strategy for probabilistic automata, POMDPs, one-sided partial-observation games, and partial observation games is undecidable for coBüchi (and hence also for parity) objectives for all classes (pure, randomized action-visible, and randomized action-invisible) of strategies.

Some basic facts. In contrast to sure winning in all classes of games and almost-sure winning for probabilistic automata and POMDPs, where pure strategies are as powerful as randomized action-visible strategies, the scenario is quite different for one-sided partial-observation games. It was shown in [17, Example 2.3] that for almost-sure winning, randomized action-visible strategies are more powerful than pure strategies for one-sided partial-observation games with reachability objectives (and the example of [17] also shows that randomized action-invisible strategies are more powerful than pure strategies for almost-sure winning in one-sided partial-observation games). It was shown in [21, Section 3] that for one-sided partial-observation games, randomized action-visible strategies are more powerful than randomized action-invisible strategies for almost-sure winning with reachability objectives. For safety objectives, the counter-examples are always finite prefixes, and it can be shown that for a given 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, and all the subclasses with safety objectives, the almost-sure problem coincide with the sure winning problem.

Complexity of almost-sure winning: randomized action-visible strategies. We now summarize the results related to the computational complexity of deciding the existence of randomized action-visible almost-sure winning strategies in various classes of partial-observation games. As for safety objectives the results are exactly the same as in the case of sure winning, we omit the discussion of safety objectives below.

1. *Complete-observation games.* A quadratic-time reduction of the almost-sure winning problem to the sure winning problem for parity objectives was established in [18, 12].

The reduction is linear-time for reachability, Büchi and coBüchi objectives. It follows that reachability objectives can be solved in linear time, Büchi and coBüchi objectives in quadratic time, and parity objectives in $\text{NP} \cap \text{coNP}$.

2. *Probabilistic automata*. The PSPACE-completeness of the almost-sure winning problem for reachability and Büchi objectives was established in [9, 10] for pure strategies. Since for probabilistic automata, the problem for pure strategies coincides with randomized action-visible strategies, the result of PSPACE-completeness follows for all classes of strategies. The almost-sure winning problem is undecidable for coBüchi and parity objectives [2].
3. *POMDPs*. The EXPTIME upper bound for the almost-sure winning problem for reachability and Büchi objectives was established in [2] for pure strategies (and hence also for randomized action-visible strategies), and the EXPTIME lower bound was shown in [16]. The almost-sure winning problem is undecidable for coBüchi and parity objectives [2].
4. *One-sided games*. The EXPTIME-completeness for almost-sure winning with reachability and Büchi objectives for one-sided partial-observation games was established in [17]. The standard subset construction does not work for almost-sure winning, as the standard subset construction yields a complete-observation game where randomization is not useful, whereas in one-sided partial-observation games randomized strategies are more powerful than pure strategies for almost-sure winning with reachability objectives. A slightly more involved exponential-size game can be constructed to solve the almost-sure winning problem, and an EXPTIME lower bound was established through a reduction of the membership problem of alternating polynomial-space Turing machine [17]. The almost-sure winning problem is undecidable for coBüchi and parity objectives [2].
5. *Partial-observation games*. The 2EXPTIME-completeness for almost-sure winning with reachability and Büchi objectives for partial-observation games was shown in [4]: the upper bound was established generalizing the result of [17] and the lower bound was established through a reduction of the membership problem of alternating exponential-space Turing machine. The almost-sure winning problem is undecidable for coBüchi and parity objectives [2].

This gives the results for almost-sure winning, summarized in Theorem 2 below (see also Table 2).

Strategy complexity of almost-sure winning: randomized action-visible strategies. The existence of pure memoryless almost-sure winning strategy for complete-observation games with parity objectives was established in [18, 12], and a consequence of the result also shows that for complete-observation games with parity objectives pure memoryless strategies are as powerful as randomized action-visible strategies. It was shown in [17] that in one-sided partial-observation games belief-based randomized action-visible almost-sure winning strategies exist for reachability and Büchi objectives. As a consequence, an exponential upper bound for memory follows for all decidable problems. The exponential lower bound for memory for POMDPs with reachability objectives was established in [16], and the construction can be adapted to show the same lower bound for probabilistic automata. It was shown in [4] that even in partial-observation games belief-based randomized action-visible almost-sure winning strategies exist for reachability and Büchi objectives. As a consequence optimal exponential memory bound follows for randomized action-visible almost-sure winning strategies for reachability and Büchi objectives for probabilistic automata, POMDPs, one-sided partial observation games, and partial-observation games. For coBüchi and parity objectives, no bound on memory of almost-sure winning strategies can be established

as the problem is undecidable. In the sequel, we do not discuss the memory bounds for problems where undecidability has been established as the undecidability result implies that infinite memory is necessary (as sufficiency of finite-memory strategies would imply semi-decidability, contradicting the undecidability results).

Theorem 2 (Complexity of almost-sure winning: randomized action-visible strategies)

The following assertions hold for randomized action-visible strategies:

1. *The almost-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 $NP \cap coNP$.*
2. *The almost-sure winning problem for probabilistic automata with reachability, safety, Büchi objectives are PSPACE-complete, and undecidable for coBüchi and parity objectives.*
3. *The almost-sure winning problem for POMDPs and one-sided partial-observation games with reachability, safety, and Büchi objectives are EXPTIME-complete, and undecidable for coBüchi and parity objectives.*
4. *The almost-sure winning problem for partial-observation games is EXPTIME-complete for safety objectives, 2EXPTIME-complete for reachability, and Büchi objectives, and undecidable for coBüchi and parity objectives.*
5. *Pure memoryless almost-sure winning strategies exist in complete-observation parity games. Exponential memory is sufficient for almost-sure winning strategies in probabilistic automata, POMDP, one-sided partial-observation games, and partial-observation games with safety, reachability and Büchi objectives; and exponential memory is necessary for almost-sure winning strategies in probabilistic automata with safety and reachability objectives. For probabilistic automata, infinite memory may be necessary for almost-sure winning with coBüchi and parity objectives.*

Computational and strategy complexity of almost-sure winning: pure and randomized action-invisible strategies.

We now summarize the results related to the computational complexity of pure and randomized action-invisible almost-sure winning strategies. We start with some basic facts. (1) For complete-observation games, probabilistic automata, and POMDPs, the results for randomized action-visible, randomized action-invisible and pure strategies are the same for almost-sure winning, as for all these classes pure strategies are as powerful as randomized action-visible strategies [13]. (2) For almost-sure winning for reachability and Büchi objectives, the equivalence of the pure and randomized action-invisible strategies has been established in [14] (polynomial-time reduction in both directions for equivalence). Note that randomized action-invisible strategies are more powerful than pure strategies for almost-sure winning; however, the reduction of [14] shows that given a game G , a game H (polynomial in the size of G) can be constructed such that there is a randomized action-invisible almost-sure winning strategy in G iff there is a pure one in H ; and similarly a reduction in the other direction. (3) For safety objectives, all results are similar to the case of sure winning. (4) For almost-sure winning for Büchi objectives, there is a linear-time reduction to reachability objectives for probabilistic automata [2], and the same reduction also works for partial-observation games. (5) The undecidability results for coBüchi and parity objectives follows from the results of probabilistic automata.

In view of the above facts we discuss the results for almost-sure winning strategies for reachability objectives in one-sided partial-observation games and in general partial-observation games. It was previously claimed that belief-based randomized action-invisible strategies suffices for almost-sure winning for reachability objectives in partial-observation

	Complete-observation	Prob. Automata	POMDP	One-sided	Partial-observation
Safety	Linear-time	PSPACE-comp.	EXPTIME-comp.	EXPTIME-comp.	EXPTIME-comp.
Reachability	Linear-time	PSPACE-comp.	EXPTIME-comp.	EXPTIME-comp.	??(open)
Büchi	Quadratic-time	PSPACE-comp.	EXPTIME-comp.	EXPTIME-comp.	??(open)
coBüchi	Quadratic-time	Undecidable	Undecidable	Undecidable	Undecidable
Parity	$NP \cap coNP$	Undecidable	Undecidable	Undecidable	Undecidable

Table 3 Complexity of almost-sure winning: pure and randomized action-invisible strategies.

games, and enumerating over the space of exponentially many belief-based strategies gives a $2EXPTIME$ upper bound for the problem [25]. However, it was shown that this result is not correct, and even in the special case of one-sided partial-observation games both pure and randomized action-invisible almost-sure winning strategies require more than belief memory, still exponential memory is sufficient for reachability objectives and the problem is $EXPTIME$ -complete for one-sided partial-observation games [14]. However, the scenario is much more complicated for partial-observation games, and although finite-memory pure and randomized action-invisible strategies suffice for almost-sure winning for reachability objectives, memory of non-elementary size is required in general [14], in contrast to the exponential upper bound claimed in [25]. The exact computational complexity of the problem remains open. The results are summarized in Table 3.

Theorem 3 (Complexity of almost-sure winning: pure and randomized action-invisible strategies) *The following assertions hold for pure and randomized action-invisible strategies:*

1. *The almost-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 $NP \cap coNP$.*
2. *The almost-sure winning problem for probabilistic automata with reachability, safety, and Büchi objectives are PSPACE-complete, and undecidable for coBüchi and parity objectives.*
3. *The almost-sure winning problem for POMDPs and one-sided partial-observation games with reachability, safety, and Büchi objectives are EXPTIME-complete, and undecidable for coBüchi and parity objectives.*
4. *The almost-sure winning problem for partial-observation games is EXPTIME-complete for safety objectives, and undecidable for coBüchi and parity objectives.*
5. *Pure memoryless almost-sure winning strategies exist in complete-observation parity games. Almost-sure winning strategies with exponential memory is sufficient for probabilistic automata, POMDP, and one-sided partial-observation games with safety, reachability and Büchi objectives; and exponential memory is necessary for almost-sure winning strategies in probabilistic automata with safety and reachability objectives. For probabilistic automata infinite memory may be necessary for almost-sure winning with coBüchi and parity objectives. For partial-observation games with reachability and Büchi objectives finite-memory almost-sure winning strategies exist, and in general almost-sure winning strategies require at least memory of non-elementary size.*

3.3 Complexity of limit-sure winning

In this section we discuss the results for limit-sure winning.

	Complete-observation	Prob. Automata	POMDP	One-sided	Partial-observation
Safety	Linear-time	PSPACE-comp.	EXPTIME-comp.	EXPTIME-comp.	EXPTIME-comp.
Reachability	Linear-time	Undecidable	Undecidable	Undecidable	Undecidable
Büchi	Quadratic-time	Undecidable	Undecidable	Undecidable	Undecidable
coBüchi	Quadratic-time	Undecidable	Undecidable	Undecidable	Undecidable
Parity	$NP \cap coNP$	Undecidable	Undecidable	Undecidable	Undecidable

Table 4 Complexity of limit-sure winning.

Complexity of limit-sure winning. We summarize the results for limit-sure winning. Like in the case of almost-sure winning, for safety objectives limit-sure winning coincides with sure winning (the same argument of finite counter-examples applies). Hence, all results for limit-sure winning for safety objectives follow from the results for sure winning.

1. *Complete-observation games.* It follows from the results of [18, 12] that for complete-observation games limit-sure winning coincide with almost-sure winning for all parity objectives, and hence all results follow from the results for almost-sure winning
2. *Probabilistic automata.* The main undecidability result for limit-sure winning was established for probabilistic finite automata. In [24] the authors start with the automata construction of [2] and show that the following question is undecidable for probabilistic finite automata: for all $\varepsilon > 0$ is there a word w_ε that is accepted with probability greater than $1 - \varepsilon$? The results of [24] can be easily adapted to show that the existence of pure strategies for limit-sure winning is undecidable for probabilistic automata with reachability objectives. By the result of [15] for probabilistic automata pure strategies are as powerful as randomized action-visible strategies, and hence the limit-sure winning problem is undecidable for all classes of strategies in probabilistic automata. Since (i) reachability objectives are a special case of Büchi, coBüchi and parity objectives, and (ii) probabilistic automata are a special class of POMDPs, one-sided partial-observation games, and partial-observation games, the undecidability follows for all more general problems.

This gives the results for limit-sure winning, and they are summarized in Theorem 4 (see also Table 4).

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

1. *The results for limit-sure winning for complete-observation games with all classes of parity objectives, and for all classes of games with safety objectives coincide with the corresponding results for the almost-sure winning problem.*
2. *The limit-sure winning problem for probabilistic automata, POMDPs, one-sided partial-observation games, and partial-observation games with reachability, Büchi, coBüchi, and parity objectives is undecidable.*

3.4 Complexity of value decision problem

In this section we consider the quantitative analysis problem, i.e., the value decision problem.

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 for all classes other than complete-observation games. The undecidability of the value decision problem for probabilistic safety

	Complete-observation	Prob. Automata	POMDP	One-sided	Partial-observation
Safety	$NP \cap coNP$	Undecidable	Undecidable	Undecidable	Undecidable
Reachability	$NP \cap coNP$	Undecidable	Undecidable	Undecidable	Undecidable
Büchi	$NP \cap coNP$	Undecidable	Undecidable	Undecidable	Undecidable
coBüchi	$NP \cap coNP$	Undecidable	Undecidable	Undecidable	Undecidable
Parity	$NP \cap coNP$	Undecidable	Undecidable	Undecidable	Undecidable

Table 5 Complexity of value decision (quantitative analysis).

automata with pure strategies can be derived from the results of [35, 33], and from [15] the undecidability follows for all other classes of strategies. For complete-observation games, the value decision problem was shown to be in $NP \cap coNP$ for reachability and safety objectives in [20], and parity objectives in [19, 12], and for complete-observation games with parity objectives pure memoryless optimal strategies exist [20, 19, 12]. We summarize the results in Theorem 5 and Table 5.

Theorem 5 (Complexity of value decision problems) *The value decision problem for complete-observation games with with safety, reachability, Büchi, coBüchi, and parity objectives lie in $NP \cap coNP$; and for probabilistic automata, POMDPs, one-sided partial-observation games and partial-observation games are undecidable.*

4 Other Related Results

We presented a survey of results for partial-observation stochastic games and their subclasses, with parity objectives for different classes of strategies. To keep the presentation focused we omitted several special cases (such as deterministic games), generalizations (such as arithmetic hierarchy characterizations of the undecidable problems), other variants (such as positive winning), and applications. We briefly discuss them below.

Deterministic games. In this article we considered stochastic games where the transition function is probabilistic. The special case of deterministic games consist of games graphs with deterministic transition function that given a state and an action gives a unique next state rather than a probability distribution over the next states. All the problems we studied in this work can also be studied for deterministic games, and the collection of results for deterministic games is presented in [13].

Arithmetic hierarchy characterization. As we have discussed in the survey, many problems related to qualitative and quantitative analysis are undecidable. An interesting theoretical question is to establish in which level of the arithmetic hierarchy the undecidable problems lie. The arithmetic hierarchy characterization for many of the problems we consider for the special case of probabilistic automata has been considered in [10, 11] where it has been established that the undecidable problems lie in the low level of arithmetic hierarchy.

Positive winning. Along with sure, almost-sure and limit-sure winning, there is another variant of qualitative analysis, namely positive winning, considered in literature. The positive winning problem asks for the existence of a strategy to ensure that an objective is satisfied with positive probability. Intuitively, the positive winning problem is dual to the almost-sure winning problem, and the results of [2] show that the positive winning problem for Büchi objectives is undecidable for probabilistic automata. The positive winning problem for partial-observation games with safety and reachability objectives has been considered in [4] where optimal complexity and memory bounds have been established.

Applications and future directions. Partial-observation games are an important generalization of complete-observation games, and arise naturally in many applications such as controller synthesis where the controller does not have access to private variables of the plant; in program analysis such as synchronizers for lock placement in concurrent programs [8]; and in artificial intelligence such as controllers for robot planning [29]. Thus analysis of partial-observation games is an important practical problem. However many problems are undecidable and decidable problems often have high complexity. An interesting direction of future research would be to identify natural and practically relevant subclasses of partial-observation games where the qualitative and quantitative analysis problems are more tractable.

Acknowledgements. The research was supported by Austrian Science Fund (FWF) Grant No P 23499-N23 on Modern Graph Algorithmic Techniques in Formal Verification, FWF NFN Grant No S11407-N23 (RiSE), ERC Start grant (279307: Graph Games), Microsoft faculty fellows award, ERC Advanced grant QUAREM, and FWF Grant No S11403-N23 (RiSE).

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. C. Baier and M. Größer. Recognizing omega-regular languages with probabilistic automata. In *LICS*, pages 137–146, 2005.
4. 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.
5. 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.
6. P. Billingsley, editor. *Probability and Measure*. Wiley-Interscience, 1995.
7. J.R. Büchi and L.H. Landweber. Solving sequential conditions by finite-state strategies. *Transactions of the AMS*, 138:295–311, 1969.
8. P. Cerný, K. Chatterjee, T. A. Henzinger, A. Radhakrishna, and R. Singh. Quantitative synthesis for concurrent programs. In *Proc. of CAV*, LNCS 6806, pages 243–259. Springer, 2011.
9. R. Chadha, A. P. Sistla, and M. Viswanathan. On the expressiveness and complexity of randomization in finite state monitors. *J. ACM*, 56(5), 2009.
10. R. Chadha, A. P. Sistla, and M. Viswanathan. Power of randomization in automata on infinite strings. In *Proc. of CONCUR*, LNCS 5710, pages 229–243. Springer, 2009.
11. R. Chadha, A. P. Sistla, and M. Viswanathan. Probabilistic Büchi automata with non-extremal acceptance thresholds. In *Proc. of VMCAI*, pages 103–117, 2011.
12. K. Chatterjee. *Stochastic ω -Regular Games*. PhD thesis, UC Berkeley, 2007.
13. K. Chatterjee and L. Doyen. The complexity of partial-observation parity games. In *Proc. of LPAR (Yogyakarta)*, pages 1–14, 2010.
14. K. Chatterjee and L. Doyen. Partial-observation stochastic games: How to win when belief fails. In *LICS'12*. IEEE, 2012 (CoRR abs/1107.2141, July 2011).
15. K. Chatterjee, L. Doyen, H. Gimbert, and T.A. Henzinger. Randomness for free. In *MFCS*. Springer, 2010.
16. K. Chatterjee, L. Doyen, and T. A. Henzinger. Qualitative analysis of partially-observable Markov decision processes. In *MFCS*, pages 258–269, 2010.
17. 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.
18. K. Chatterjee, M. Jurdziński, and T.A. Henzinger. Simple stochastic parity games. In *CSL'03*, volume 2803 of *LNCS*, pages 100–113. Springer, 2003.
19. K. Chatterjee, M. Jurdziński, and T.A. Henzinger. Quantitative stochastic parity games. In *SODA'04*, pages 121–130. SIAM, 2004.
20. A. Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992.

21. J. Cristau, C. David, and F. Horn. How do we remember the past in randomised strategies? In *GANDALF*, pages 30–39, 2010.
22. L. de Alfaro and T.A. Henzinger. Interface theories for component-based design. In *EMSOFT*, LNCS 2211, pages 148–165. Springer, 2001.
23. E.A. Emerson and C. Jutla. Tree automata, mu-calculus and determinacy. In *FOCS*, pages 368–377. IEEE, 1991.
24. H. Gimbert and Y. Oualhadj. Probabilistic automata on finite words: Decidable and undecidable problems. In *Proc. of ICALP*, LNCS 6199, pages 527–538. Springer, 2010.
25. V. Gripon and O. Serre. Qualitative concurrent stochastic games with imperfect information. In *ICALP(2)*, LNCS 5556, pages 200–211. Springer, 2009.
26. T. A. Henzinger, O. Kupferman, and S. Rajamani. Fair simulation. *Information and Computation*, 173:64–81, 2002.
27. N. Immerman. Number of quantifiers is better than number of tape cells. *Journal of Computer and System Sciences*, 22:384–406, 1981.
28. A. Kechris. *Classical Descriptive Set Theory*. Springer, 1995.
29. H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas. Temporal-logic-based reactive mission and motion planning. *IEEE Transactions on Robotics*, 25(6):1370–1381, 2009.
30. O. Madani, S. Hanks, and A. Condon. On the undecidability of probabilistic planning and related stochastic optimization problems. *Artificial Intelligence*, 147(1-2):5–34, 2003.
31. D. A. Martin. The determinacy of Blackwell games. *The Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
32. R. McNaughton. Infinite games played on finite graphs. *Annals of Pure and Applied Logic*, 65:149–184, 1993.
33. A. Paz. *Introduction to probabilistic automata (Computer science and applied mathematics)*. Academic Press, 1971.
34. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190. ACM Press, 1989.
35. M. O. Rabin. Probabilistic automata. *Information and Control*, 6:230–245, 1963.
36. 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.
37. J. H. Reif. Universal games of incomplete information. In *STOC*, pages 288–308. ACM Press, 1979.
38. W. Thomas. Languages, automata, and logic. In *Handbook of Formal Languages*, volume 3, Beyond Words, chapter 7, pages 389–455. Springer, 1997.
39. M. Y. Vardi. Automatic verification of probabilistic concurrent finite-state systems. In *FOCS*, pages 327–338. IEEE Computer Society Press, 1985.