Games with a Weak Adversary*,**

Krishnendu Chatterjee[1] and Laurent Doyen[2]

[1] IST Austria
[2] LSV, ENS Cachan & CNRS, France

Abstract. We consider multi-player graph games with partial-observation and parity objective. While the decision problem for three-player games with a coalition of the first and second players against the third player is undecidable in general, we present a decidability result for partial-observation games where the first and third player are in a coalition against the second player, thus where the second player is adversarial but weaker due to partial-observation. We establish tight complexity bounds in the case where player 1 is less informed than player 2, namely 2-EXPTIME-completeness for parity objectives. The symmetric case of player 1 more informed than player 2 is much more complicated, and we show that already in the case where player 1 has perfect observation, memory of size non-elementary is necessary in general for reachability objectives, and the problem is decidable for safety and reachability objectives. From our results we derive new complexity results for partial-observation stochastic games.

1 Introduction

Games on graphs. Games played on graphs are central in several important problems in computer science, such as reactive synthesis [21, 22], verification of open systems [2], and many others. The game is played by several players on a finite-state graph, with a set of angelic (existential) players and a set of demonic (universal) players as follows: the game starts at an initial state, and given the current state, the successor state is determined by the choice of moves of the players. The outcome of the game is a *play*, which is an infinite sequence of states in the graph. A *strategy* is a transducer to resolve choices in a game for a player that given a finite prefix of the play specifies the next move. Given an objective (the desired set of behaviors or plays), the goal of the existential players is to ensure the play belongs to the objective irrespective of the strategies of the universal players. In verification and control of reactive systems an objective is typically an ω-regular set of paths. The class of ω-regular languages, that extends classical regular languages to infinite strings, provides a robust specification language to express all commonly used specifications, and parity objectives are a canonical way to define such ω-regular specifications [27]. Thus games on graphs with parity objectives provide a general framework for analysis of reactive systems.

* This research was partly supported by Austrian Science Fund (FWF) Grant No P23499- N23, FWF NFN Grant No S11407-N23 (RiSE), ERC Start grant (279307: Graph Games), Microsoft Faculty Fellowship Award, and European project Cassting (FP7-601148).

** Fuller version: [1].

Perfect vs partial observation. Many results about games on graphs make the hypothesis of *perfect observation* (i.e., players have perfect or complete observation about the state of the game). In this setting, due to determinacy (or switching of the strategy quantifiers for existential and universal players) [17], the questions expressed by an arbitrary alternation of quantifiers reduce to a single alternation, and thus are equivalent to solving two-player games (all the existential players against all the universal players). However, the assumption of perfect observation is often not realistic in practice. For example in the control of physical systems, digital sensors with finite precision provide partial information to the controller about the system state [12, 14]. Similarly, in a concurrent system the modules expose partial interfaces and have access to the public variables of the other processes, but not to their private variables [25, 2]. Such situations are better modeled in the more general framework of *partial-observation* games [24–26].

Partial-observation games. Since partial-observation games are not determined, unlike the perfect-observation setting, the multi-player game problems do not reduce to the case of two-player games. Typically, multi-player partial-observation games are studied in the following setting: a set of partial-observation existential players, against a perfect-observation universal player, such as for distributed synthesis [21, 13, 23]. The problem of deciding if the existential players can ensure a reachability (or a safety) objective is undecidable in general, even for two existential players [20, 21]. However, if the information of the existential players form a chain (i.e., existential player 1 more informed than existential player 2, existential player 2 more informed than existential player 3, and so on), then the problem is decidable [21, 16, 18].

Games with a weak adversary. One aspect of multi-player games that has been largely ignored is the presence of weaker universal players that do not have perfect observation. However, it is natural in the analysis of composite reactive systems that some universal players represent components that do not have access to all variables of the system. In this work we consider games where adversarial players can have partial observation. If there are two existential (resp., two universal) players with incomparable partial observation, then the undecidability results follows from [20, 21]; and if the information of the existential (resp., universal) players form a chain, then they can be reduced to one partial-observation existential (resp., universal) player. We consider the following case of partial-observation games: one partial-observation existential player (player 1), one partial-observation universal player (player 2), one perfect-observation existential player (player 3), and one perfect-observation universal player (player 4). Roughly, having more partial-observation players in general leads to undecidability, and having more perfect-observation players reduces to two perfect-observation players. We first present our results and then discuss two applications of the model.

Results. Our main results are as follows:

1. *Player 1 less informed.* We first consider the case when player 1 is less informed than player 2. We establish the following results: (i) a 2-EXPTIME upper bound for parity objectives and a 2-EXPTIME lower bound for reachability objectives (i.e., we establish 2-EXPTIME-completeness); (ii) an EXPSPACE upper bound for parity objectives when player 1 is blind (has only one observation), and EXPSPACE lower bound for reachability objectives even when both player 1 and player 2 are

	Reachability		Parity		Parity	
Player 2	Finite- or infinite-memory strategies		Infinite-memory strategies		Finite-memory strategies	
Player 1	Perfect	More informed	Perfect	More informed	Perfect	More informed
Randomized	EXP-c [9]	EXP-c [4]	Undec. [3, 8]	Undec. [3, 8]	EXP-c [10]	**2EXP**
Pure	EXP-c [7]	**2EXP-c**	Undec. [3]	Undec. [3]	EXP-c [10]	**2EXP-c**

Table 1. Complexity of qualitative analysis (almost-sure winning) for partial-observation stochastic games with partial observation for player 1 with reachability and parity objectives. Player 2 has either perfect observation or more information than player 1(new results boldfaced). For positive winning, all entries other than the first (randomized strategies for player 1 and perfect observation for player 2) remain the same, and the complexity for the first entry for positive winning is PTIME-complete.

blind. In all these cases, if the objective can be ensured then the upper bound on memory requirement of winning strategies is at most doubly exponential.

2. *Player 1 more informed.* We consider the case when player 1 can be more informed as compared to player 2, and show that even when player 1 has perfect observation there is a non-elementary lower bound on the memory required by winning strategies. This result is also in sharp contrast with distributed games, where if only one player has partial observation then the upper bound on memory of winning strategies is exponential.

Applications. We discuss two applications of our results: the sequential synthesis problem, and new complexity results for partial-observation *stochastic* games.

1. The sequential synthesis problem consists of a set of partially implemented modules, where first a set of modules needs to be refined, followed by a refinement of some modules by an external source, and then the remaining modules are refined so that the composite open reactive system satisfies a specification. Given the first two refinements cannot access all private variables, we have a four-player game where the first refinement corresponds to player 1, the second refinement to player 2, the third refinement to player 3, and player 4 is the environment.
2. In partial-observation stochastic games, there are two partial-observation players (one existential and one universal) playing in the presence of uncertainty in the transition function (i.e., stochastic transition function). The qualitative analysis question is to decide the existence of a strategy for the existential player to ensure the parity objective with probability 1 (or with positive probability) against all strategies of the universal player. The witness strategy can be randomized or deterministic (pure). While the qualitative problem is undecidable, the practically relevant restriction to finite-memory pure strategies reduces to the four-player game problem. Moreover, for finite-memory strategies, the decision problem for randomized strategies reduces to the pure-strategy question [7]. By the results we establish in this paper, new decidability and complexity results are obtained for the qualitative analysis of partial-observation stochastic games with player 2 partially informed but more informed than player 1. The complexity results for almost-sure winning are summarized in Table 1. Surprisingly for reachability objectives, whether player 2 is perfectly informed or more informed than player 1 does not change the complexity

for randomized strategies, but it results in an exponential increase in the complexity for pure strategies.

2 Definitions

We first consider three-player (non-stochastic) games with parity objectives and we establish new complexity results in Section 3 that we later extend to four-player games in Section 5. We also present the related model of two-player stochastic games for which our contribution implies new complexity results.

Three-player games Given alphabets A_i of actions for player i ($i = 1, 2, 3$), a *three-player game* is a tuple $G = \langle Q, q_0, \delta \rangle$ where:
- Q is a finite set of states with $q_0 \in Q$ the initial state; and
- $\delta : Q \times A_1 \times A_2 \times A_3 \to Q$ is a deterministic transition function that, given a current state q, and actions $a_1 \in A_1$, $a_2 \in A_2$, $a_3 \in A_3$ of the players, gives the successor state $q' = \delta(q, a_1, a_2, a_3)$.

The games we consider are sometimes called *concurrent* because all three players need to choose simultaneously an action to determine a successor state. The special class of *turn-based* games corresponds to the case where in every state, one player has the turn and his sole action determines the successor state. In our framework, a turn-based state for player 1 is a state $q \in Q$ such that $\delta(q, a_1, a_2, a_3) = \delta(q, a_1, a_2', a_3')$ for all $a_1 \in A_1$, $a_2, a_2' \in A_2$, and $a_3, a_3' \in A_3$. We define analogously turn-based states for player 2 and player 3. A game is turn-based if every state of G is turn-based (for some player). The class of two-player games is obtained when A_3 is a singleton. In a game G, given $s \subseteq Q$, $a_1 \in A_1$, $a_2 \in A_2$, let $\mathsf{post}^G(s, a_1, a_2, -) = \{q' \in Q \mid \exists q \in s \cdot \exists a_3 \in A_3 : q' = \delta(q, a_1, a_2, a_3)\}$.

Observations For $i = 1, 2, 3$, a set $\mathcal{O}_i \subseteq 2^Q$ of *observations* (for player i) is a partition of Q (i.e., $\mathcal{O}_i$ is a set of non-empty and non-overlapping subsets of Q, and their union covers Q). Let $\mathsf{obs}_i : Q \to \mathcal{O}_i$ be the function that assigns to each state $q \in Q$ the (unique) observation for player i that contains q, i.e. such that $q \in \mathsf{obs}_i(q)$. The functions obs_i are extended to sequences $\rho = q_0 \dots q_n$ of states in the natural way, namely $\mathsf{obs}_i(\rho) = \mathsf{obs}_i(q_0) \dots \mathsf{obs}_i(q_n)$. We say that player i is *blind* if $\mathcal{O}_i = \{Q\}$, that is player i has only one observation; player i has *perfect information* if $\mathcal{O}_i = \{\{q\} \mid q \in Q\}$, that is player i can distinguish each state; and player 1 is *less informed* than player 2 (we also say player 2 is more informed) if for all $o_2 \in \mathcal{O}_2$, there exists $o_1 \in \mathcal{O}_1$ such that $o_2 \subseteq o_1$.

Strategies For $i = 1, 2, 3$, let Σ_i be the set of *strategies* $\sigma_i : \mathcal{O}_i^+ \to A_i$ of player i that, given a sequence of past observations, give an action for player i. Equivalently, we sometimes view a strategy of player i as a function $\sigma_i : Q^+ \to A_i$ satisfying $\sigma_i(\rho) = \sigma_i(\rho')$ for all $\rho, \rho' \in Q^+$ such that $\mathsf{obs}_i(\rho) = \mathsf{obs}_i(\rho')$, and say that σ_i is *observation-based*.

Outcome Given strategies $\sigma_i \in \Sigma_i$ ($i = 1, 2, 3$) in G, the *outcome play* from a state q_0 is the infinite sequence $\rho_{q_0}^{\sigma_1,\sigma_2,\sigma_3} = q_0 q_1 \dots$ such that for all $j \geq 0$, we have $q_{j+1} = \delta(q_j, a_1^j, a_2^j, a_3^j)$ where $a_i^j = \sigma_i(q_0 \dots q_j)$ (for $i = 1, 2, 3$).

Objectives An *objective* is a set $\alpha \subseteq Q^\omega$ of infinite sequences of states. A play ρ *satisfies* the objective α if $\rho \in \alpha$. An objective α is *visible* for player i if for all $\rho, \rho' \in Q^\omega$, if $\rho \in \alpha$ and $\mathsf{obs}_i(\rho) = \mathsf{obs}_i(\rho')$, then $\rho' \in \alpha$. We consider the following objectives:

- *Reachability*. Given a set $\mathcal{T} \subseteq Q$ of target states, the *reachability* objective $\mathsf{Reach}(\mathcal{T})$ requires that a state in $\mathcal{T}$ be visited at least once, that is, $\mathsf{Reach}(\mathcal{T}) = \{\rho = q_0 q_1 \cdots \mid \exists k \geq 0 : q_k \in \mathcal{T}\}$.
- *Safety*. Given a set $\mathcal{T} \subseteq Q$ of target states, the *safety* objective $\mathsf{Safe}(\mathcal{T})$ requires that only states in $\mathcal{T}$ be visited, that is, $\mathsf{Safe}(\mathcal{T}) = \{\rho = q_0 q_1 \cdots \mid \forall k \geq 0 : q_k \in \mathcal{T}\}$.
- *Parity*. For a play $\rho = q_0 q_1 \dots$ we denote by $\mathsf{Inf}(\rho)$ the set of states that occur infinitely often in ρ, that is, $\mathsf{Inf}(\rho) = \{q \in Q \mid \forall k \geq 0 \cdot \exists n \geq k : q_n = q\}$. For $d \in \mathbb{N}$, let $p : Q \to \{0, 1, \dots, d\}$ be a priority function, which maps each state to a nonnegative integer priority. The parity objective $\mathsf{Parity}(p)$ requires that the minimum priority occurring infinitely often be even. Formally, $\mathsf{Parity}(p) = \{\rho \mid \min\{p(q) \mid q \in \mathsf{Inf}(\rho)\}$ is even$\}$. Parity objectives are a canonical way to express ω-regular objectives [27]. If the priority function is constant over observations of player i, that is for all observations $\gamma \in \mathcal{O}_i$ we have $p(q) = p(q')$ for all $q, q' \in \gamma$, then the parity objective $\mathsf{Parity}(p)$ is visible for player i.

Decision problem Given a game $G = \langle Q, q_0, \delta \rangle$ and an objective $\alpha \subseteq Q^\omega$, the *three-player decision problem* is to decide if $\exists \sigma_1 \in \Sigma_1 \cdot \forall \sigma_2 \in \Sigma_2 \cdot \exists \sigma_3 \in \Sigma_3 : \rho_{q_0}^{\sigma_1, \sigma_2, \sigma_3} \in \alpha$.

The results for the three-player decision problem have implications for decision problems on partial-observation stochastic games that we formally define below.

Two-player partial-observation stochastic games Given alphabet A_i of actions, and set $\mathcal{O}_i$ of observations (for player $i \in \{1, 2\}$), a *two-player partial-observation stochastic game* (for brevity, two-player stochastic game) is a tuple $G = \langle Q, q_0, \delta \rangle$ where Q is a finite set of states, $q_0 \in Q$ is the initial state, and $\delta : Q \times A_1 \times A_2 \to \mathcal{D}(Q)$ is a probabilistic transition where $\mathcal{D}(Q)$ is the set of probability distributions $\kappa : Q \to [0, 1]$ on Q, such that $\sum_{q \in Q} \kappa(q) = 1$. Given a current state q and actions a, b for the players, the transition probability to a successor state q' is $\delta(q, a, b)(q')$. Observation-based strategies are defined as for three-player games. An *outcome play* from a state q_0 under strategies σ_1, σ_2 is an infinite sequence $\rho = q_0\, a_0 b_0\, q_1 \dots$ such that $a_i = \sigma_1(q_0 \dots q_i)$, $b_i = \sigma_2(q_0 \dots q_i)$, and $\delta(q_i, a_i, b_i)(q_{i+1}) > 0$ for all $i \geq 0$.

Qualitative analysis Given an objective α that is Borel measurable (all Borel sets in the Cantor topology and all objectives considered in this paper are measurable [15]), a strategy σ_1 for player 1 is *almost-sure winning* (resp., *positive winning*) for the objective α from q_0 if for all observation-based strategies σ_2 for player 2, we have $\Pr_{q_0}^{\sigma_1, \sigma_2}(\alpha) = 1$ (resp., $\Pr_{q_0}^{\sigma_1, \sigma_2}(\alpha) > 0$) where $\Pr_{q_0}^{\sigma_1, \sigma_2}(\cdot)$ is the unique probability measure induced by the natural probability measure on finite prefixes of plays (i.e., the product of the transition probabilities in the prefix).

3 Three-Player Games with Player 1 Less Informed

We consider the three-player (non-stochastic) games defined in Section 2. We show that for reachability and parity objectives the three-player decision problem is decidable when player 1 is less informed than player 2. The problem is EXPSPACE-complete when player 1 is blind, and 2-EXPTIME-complete in general.

Remark 1. Observe that once the strategies of the first two players are fixed we obtain a graph, and in graphs perfect-information coincides with blind for construction of a path (see [6, Lemma 2] that counting strategies that count the number of steps are sufficient which can be ensured by a player with no information). Hence without loss of generality we consider that player 3 has perfect observation, and drop the observation for player 3.

Theorem 1 (Upper bounds). *Given a three-player game $G = \langle Q, q_0, \delta \rangle$ with player 1 less informed than player 2 and a parity objective α, the problem of deciding whether $\exists \sigma_1 \in \Sigma_1 \cdot \forall \sigma_2 \in \Sigma_2 \cdot \exists \sigma_3 \in \Sigma_3 : \rho_{q_0}^{\sigma_1,\sigma_2,\sigma_3} \in \alpha$ can be solved in 2-EXPTIME. If player 1 is blind, then the problem can be solved in EXPSPACE.*

Proof. The proof is by a reduction of the decision problem for three-player games to a decision problem for partial-observation two-player games with the same objective. We present the reduction for parity objectives that are visible for player 2 (defined by priority functions that are constant over observations of player 2). The general case of not necessarily visible parity objectives can be solved using a reduction to visible objectives, as in [6, Section 3].

Given a three-player game $G = \langle Q, q_0, \delta \rangle$ over alphabet of actions A_i ($i = 1, 2, 3$), and observations $\mathcal{O}_1, \mathcal{O}_2 \subseteq 2^Q$ for player 1 and player 2, with player 1 less informed than player 2, we construct a two-player game $H = \langle Q_H, \{q_0\}, \delta_H \rangle$ over alphabet of actions A'_i ($i = 1, 2$), and observations $\mathcal{O}'_1 \subseteq 2^{Q_H}$ and perfect observation for player 2, where (intuitive explanations follow):

- $Q_H = \{s \in 2^Q \mid s \neq \varnothing \land \exists o_2 \in \mathcal{O}_2 : s \subseteq o_2\}$;
- $A'_1 = A_1 \times (2^Q \times A_2 \to \mathcal{O}_2)$, and $A'_2 = A_2$;
- $\mathcal{O}'_1 = \big\{\{s \in Q_H \mid s \subseteq o_1\} \mid o_1 \in \mathcal{O}_1\big\}$, and let $\mathsf{obs}'_1 : Q_H \to \mathcal{O}'_1$ be the corresponding observation function;
- $\delta_H(s, (a_1, f), a_2) = \mathsf{post}^G(s, a_1, a_2, -) \cap f(s, a_2)$.

Intuitively, the state space Q_H is the set of knowledges of player 2 about the current state in G, i.e., the sets of states compatible with an observation of player 2. Along a play in H, the knowledge of player 2 is updated to represent the set of possible current states in which the game G can be. In H player 2 has perfect observation and the role of player 1 in the game H is to simulate the actions of both player 1 and player 3 in G. Since player 2 fixes his strategy before player 3 in G, the simulation should not let player 2 know player-3's action, but only the observation that player 2 will actually see while playing the game. The actions of player 1 in H are pairs $(a_1, f) \in A'_1$ where a_1 is a simple action of player 1 in G, and f gives the observation $f(s, a_2)$ received by player 2 after the response of player 3 to the action a_2 of player 2 when the knowledge of player 2 is s. In H, player 1 has partial observation, as he cannot distinguish knowledges of player 2 that belong to the same observation of player 1 in G. The transition relation updates the knowledges of player 2 as expected. Note that $|\mathcal{O}_1| = |\mathcal{O}'_1|$, and therefore if player 1 is blind in G then he is blind in H as well.

Given a visible parity objective $\alpha = \mathsf{Parity}(p)$ where $p : Q \to \{0, 1, \ldots, d\}$ is constant over observations of player 2, let $\alpha' = \mathsf{Parity}(p')$ where $p'(s) = p(q)$ for all $q \in s$ and $s \in Q_H$. Note that the function p' is well defined since s is a subset of an observation of player 2 and thus $p(q) = p(q')$ for all $q, q' \in s$. However, the parity objective $\alpha' = \mathsf{Parity}(p')$ may not be visible to player 1 in G. We establish that given witness strategies in G we can construct witness strategies in H and vice-versa, and the details of the strategy constructions are presented in [1]. $\square$

Theorem 2 (Lower bounds). *Given a three-player game $G = \langle Q, q_0, \delta \rangle$ with player 1 less informed than player 2 and a reachability objective α, the problem of deciding whether $\exists \sigma_1 \in \Sigma_1 \cdot \forall \sigma_2 \in \Sigma_2 \cdot \exists \sigma_3 \in \Sigma_3 : \rho_{q_0}^{\sigma_1,\sigma_2,\sigma_3} \in \alpha$ is 2-EXPTIME-hard. If player 1 is blind (and even when player 2 is also blind), then the problem is EXPSPACE-hard.*

Proof. The proof of 2-EXPTIME-hardness is obtained by a polynomial-time reduction of the membership problem for exponential-space *alternating* Turing machines to the three-player problem. The same reduction for the special case of exponential-space *nondeterministic* Turing machines shows EXPSPACE-hardness when player 1 is blind (because our reduction yields a game in which player 1 is blind when we start from a nondeterministic Turing machine). The membership problem for Turing machines is to decide, given a Turing machine M and a finite word w, whether M accepts w. The membership problem is 2-EXPTIME-complete for exponential-space alternating Turing machines, and EXPSPACE-complete for exponential-space nondeterministic Turing machines [19].

An alternating Turing machine is a tuple $M = \langle Q_\vee, Q_\wedge, \Sigma, \Gamma, \Delta, q_0, q_{acc}, q_{rej} \rangle$ where the state space $Q = Q_\vee \cup Q_\wedge$ consists of the set $Q_\vee$ of or-states, and the set $Q_\wedge$ of and-states. The input alphabet is Σ, the tape alphabet is $\Gamma = \Sigma \cup \{\#\}$ where $\#$ is the blank symbol. The initial state is q_0, the accepting state is q_{acc}, and the rejecting state is q_{rej}. The transition relation is $\Delta \subseteq Q \times \Gamma \times Q \times \Gamma \times \{-1, 1\}$, where a transition $(q, \gamma, q', \gamma', d) \in \Delta$ intuitively means that, given the machine is in state q, and the symbol under the tape head is γ, the machine can move to state q', replace the symbol under the tape head by γ', and move the tape head to the neighbor cell in direction d. A configuration c of M is a sequence $c \in (\Gamma \cup (Q \times \Gamma))^\omega$ with exactly one symbol in $Q \times \Gamma$, which indicates the current state of the machine and the position of the tape head. The initial configuration of M on $w = a_0 a_1 \dots a_n$ is $c_0 = (q_0, a_0) \cdot a_1 \cdot a_2 \cdot \dots \cdot a_n \cdot \#^\omega$. Given the initial configuration of M on w, it is routine to define the execution trees of M where at least one successor of each configuration in an or-state, and all successors of the configurations in an and-state are present (and we assume that all branches reach either q_{acc} or q_{rej}), and to say that M accepts w if all branches of some execution tree reach q_{acc}. Note that $Q_\wedge = \varnothing$ for nondeterministic Turing machines, and in that case the execution tree reduces to a single path. A Turing machine M uses exponential space if for all words w, all configurations in the execution of M on w contain at most $2^{O(|w|)}$ non-blank symbols.

We present the key steps of our reduction from alternating Turing machines. Given a Turing machine M and a word w, we construct a three-player game with reachability objective in which player 1 and player 2 have to simulate the execution of M on w, and player 1 has to announce the successive configurations and transitions of the machine along the execution. Player 1 announces configurations one symbol at a time, thus the alphabet of player 1 is $A_1 = \Gamma \cup (Q \times \Gamma) \cup \Delta$. In an initialization phase, the transition relation of the game forces player 1 to announce the initial configuration c_0 (this can be done with $O(n)$ states in the game, where $n = |w|$). Then, the game proceeds to a loop where player 1 keeps announcing symbols of configurations. At all times along the execution, some finite information is stored in the finite state space of the game: a window of the last three symbols $\mathsf{z}_1, \mathsf{z}_2, \mathsf{z}_3$ announced by player 1, as well as the last symbol $\mathsf{head} \in Q \times \Gamma$ announced by player 1 (that indicates the current machine state and the position of the tape head). After the initialization phase, we should have $\mathsf{z}_1 = \mathsf{z}_2 = \mathsf{z}_3 = \#$ and $\mathsf{head} = (q_0, a_0)$. When player 1 has announced a full configuration, he moves to a state of the game where either player 1 or player 2 has to announce a transition of the machine: for $\mathsf{head} = (p, a)$, if $p \in Q_\vee$, then player 1 chooses the next transition, and if $p \in Q_\wedge$, then player 2 chooses. Note that the transitions chosen by player 2 are visible to player 1 and this is the only information that player 1 observes. Hence player 1 is less informed than player 2, and both player 1 and player 2 are blind when the machine is nondeterministic. If a transition $(q, \gamma, q', \gamma', d)$ is chosen by player i, and either $p \neq q$ or $a \neq \gamma$, then player i loses (i.e., a sink state is reached to let player 1 lose, and the target state of the reachability objective is reached to let player 2 lose). If at some point player 1 announces a symbol (p, a) with $p = q_{acc}$, then player 1 wins the game.

The role of player 2 is to check that player 1 faithfully simulates the execution of the Turing machine, and correctly announces the configurations. After every announcement of a symbol by player 1, the game offers the possibility to player 2 to compare this symbol with the symbol at the same position in the next configuration. We say that player 2 *checks* (and whether player 2 checks or not is not visible to player 1), and the checked symbol is stored as z_2. Note that player 2 can be blind to check because player 2 fixes his strategy after player 1. The window $\mathsf{z}_1, \mathsf{z}_2, \mathsf{z}_3$ stored in the state space of the game provides enough information to update the middle cell z_2 in the next configuration, and it allows the game to verify the check of player 2. However, the distance (in number of steps) between the same position in two consecutive configurations is exponential (say 2^n for simplicity), and the state space of the game is not large enough to check that such a distance exists between the two symbols compared by player 2. We use player 3 to check that player 2 makes a comparison at the correct position. When player 2 decides to check, he has to count from 0 to 2^n by announcing after every symbol of player 1 a sequence of n bits, initially all zeros (again, this can be enforced by the structure of the game with $O(n)$ states). It is then the responsibility of player 3 to check that player 2 counts correctly. To check this, player 3 can at any time choose a bit position $p \in \{0, \dots, n-1\}$ and store the bit value b_p announced by player 2 at position p. The value of b_p and p is not visible to player 2. While player 2 announces the bits $b_{p+1}, \dots, b_{n-1}$ at position $p+1, \dots, n-1$, the finite state of the game is used to flip the value of b_p if all bits $b_{p+1}, \dots, b_{n-1}$ are equal to 1, hence updating b_p to the value of the p-th bit in what should be the next announcement of player 2. In the next bit sequence announced by player 2, the p-th bit is compared with b_p. If they match, then the game goes to a sink state (as player 2 has faithfully counted), and if they differ then the game goes to the target state (as player 2 is caught cheating). It can be shown that this can be enforced by the structure of the game with $O(n^2)$ states, that is $O(n)$

states for each value of p. As before, whether player 3 checks or not is not visible to player 2.

Note that the checks of player 2 and player 3 are one-shot: the game will be over (either in a sink or target state) when the check is finished. This is enough to ensure a faithful simulation by player 1, and a faithful counting by player 2, because (1) partial observation allows to hide to a player the time when a check occurs, and (2) player 2 fixes his strategy after player 1 (and player 3 after player 2), thus they can decide to run a check exactly when player 1 (or player 2) is not faithful. This ensures that player 1 does not win if he does not simulate the execution of M on w, and that player 2 does not win if he does not count correctly.

Hence this reduction ensures that M accepts w if and only if the answer to the three-player game problem is YES, where the reachability objective is satisfied if player 1 eventually announces that the machine has reached q_{acc} (that is if M accepts w), or if player 2 cheats in counting, which can be detected by player 3. □

4 Three-Player Games with Player 1 Perfect

When player 2 is less informed than player 1, we show that three-player games get much more complicated (even in the special case where player 1 has perfect information). We note that for reachability objectives, the three-player decision problem is equivalent to the qualitative analysis of positive winning in two-player stochastic games, and we show that the techniques developed in the analysis of two-player stochastic games can be extended to solve the three-player decision problem with safety objectives as well.

For reachability objectives, the three-player decision problem is equivalent to the problem of positive winning in two-player stochastic games where the third player is replaced by a probabilistic choice over the action set with uniform probability. Intuitively, after player 1 and player 2 fixed their strategy, the fact that player 3 can construct a (finite) path to the target set is equivalent to the fact that such a path has positive probability when the choices of player 3 are replaced by uniform probabilistic transitions. Given a three-player game $G = \langle Q, q_0, \delta \rangle$, let $\mathsf{Uniform}(G) = \langle Q, q_0, \delta' \rangle$ be the two-player partial-observation *stochastic* game (with same state space, action sets, and observations for player 1 and player 2) where $\delta'(q, a_1, a_2)(q') = \frac{\left|\{a_3 \mid \delta(q,a_1,a_2,a_3)=q'\}\right|}{|A_3|}$ for all $a_1 \in A_1$, $a_2 \in A_2$, and $q, q' \in Q$. Formally, the equivalence result is presented in Lemma 1, and the equivalence holds for all three-player games (not restricted to three-player games where player 1 has perfect information). However, we will use Lemma 1 to establish results for three-player games where player 1 has perfect information.

Lemma 1. *Given a three-player game G and a reachability objective α, the answer to the three-player decision problem for $\langle G, \alpha \rangle$ is* YES *if and only if player* 1 *is positive winning for α in the two-player partial-observation stochastic game* $\mathsf{Uniform}(G)$.

Reachability objectives. Even in the special case where player 1 has perfect information, and for reachability objectives, non-elementary memory is necessary in general for player 1 to win in three-player games. This result follows from Lemma 1 and from the result of [7, Example 4.2 Journal version] showing that non-elementary memory is necessary to win with positive probability in two-player stochastic games. It also follows from Lemma 1 and the result of [7, Corollary 4.9 Journal version] that the three-player decision problem for reachability games is decidable. The decidability result can be extended to safety objectives [1].

Theorem 3. *When player 1 has perfect information, the three-player decision problem is decidable for both reachability and safety games, and for reachability games memory of size non-elementary is necessary in general for player* 1.

5 Four-Player Games

We show that the results presented for three-player games extend to games with four players (the fourth player is universal and perfectly informed). The definition of four-player games and related notions is a straightforward extension of Section 2.

In a four-player game with player 1 less informed than player 2, and perfect information for both player 3 and player 4, consider the *four-player decision problem* which is to decide if $\exists \sigma_1 \in \Sigma_1 \cdot \forall \sigma_2 \in \Sigma_2 \cdot \exists \sigma_3 \in \Sigma_3 \cdot \forall \sigma_4 \in \Sigma_4 : \rho_{q_0}^{\sigma_1,\sigma_2,\sigma_3,\sigma_4} \in \alpha$ for a parity objective α (also see [1, Remark 2] for further discussion). Since player 3 and player 4 have perfect information, we assume without loss of generality that the game is turn-based for them, that is there is a partition of the state space Q into two sets Q_3 and Q_4 (where $Q = Q_3 \cup Q_4$) such that the transition function is the union of $\delta_3 : Q_3 \times A_1 \times A_2 \times A_3 \to Q$ and $\delta_4 : Q_4 \times A_1 \times A_2 \times A_4 \to Q$. Strategies and outcomes are defined analogously to three-player games. A strategy of player $i \in \{3, 4\}$ is of the form $\sigma_i : Q^* \cdot Q_i \to A_i$.

We present a polynomial reduction of the problem for four-player games to solving a three-player game with the first player less informed than the second player [1]. Hardness follows from the special case of three-player games.

Theorem 4. *The four-player decision problem with player* 1 *less informed than player* 2*, and perfect information for both player* 3 *and player* 4 *is 2-EXPTIME-complete for parity objectives.*

6 Applications

We now discuss applications of our results in the context of synthesis and qualitative analysis of two-player partial-observation stochastic games.

Sequential synthesis. The *sequential synthesis* problem consists of an open system of partially implemented modules (with possible non-determinism or choices) $M_1, M_2, \ldots, M_n$ that need to be refined (i.e., the choices determined by strategies) such that the composite system after refinement satisfy a specification. The system is open in the sense that after the refinement the composite system is reactive and interact with an environment. Consider the problem where first a set $M_1, \ldots, M_k$ of modules are refined, then a set $M_{k+1}, \ldots, M_\ell$ are refined by an external implementor, and finally the remaining set of modules are refined. In other words, the modules are refined sequentially: first a set of modules whose refinement can be controlled, then a set of modules whose refinement cannot be controlled as they are implemented externally, and finally the remaining set of modules. If the refinements of modules $M_1, \ldots, M_\ell$ do not have access to private variables of the remaining modules we obtain a partial-observation game with four players: the first (existential) player corresponds to the refinement of modules $M_1, \ldots, M_k$, the second (universal) player corresponds to the refinement of modules $M_{k+1}, \ldots, M_\ell$, the third (existential) player corresponds to the refinement of the remaining modules, and the fourth (adversarial) player is the environment. If the second player has access to all the variables visible to the first player, then player 1 is less informed.

Two-player partial-observation stochastic games. Our results for four-player games imply new complexity results for two-player stochastic games. For qualitative analysis (positive and almost-sure winning) under finite-memory strategies for the players the following reduction has been established in [10, Lemma 1] (see Lemma 2.1 of the arxiv version): the probabilistic transition function can be replaced by a turn-based gadget consisting of two perfect-observation players, one angelic (existential) and one demonic (universal). The turn-based gadget is the same as used for perfect-observation stochastic games [5, 11]. In [10], only the special case of perfect observation for player 2 was considered, and hence the problem reduced to three-player games where only player 1 has partial observation and the other two players have perfect observation. In case where player 2 has partial observation, the reduction of [10] requires two perfect-observation players, and gives the problem of four-player games (with perfect observation for player 3 and player 4). Hence when player 1 is less informed, we obtain a 2-EXPTIME upper bound from Theorem 4, and obtain a 2-EXPTIME lower bound from Theorem 2 and Lemma 1 (see [1] for lower bound for almost-sure winning). Thus we obtain the following result.

Theorem 5. *The qualitative analysis problems (almost-sure and positive winning) for two-player partial-observation stochastic parity games where player 1 is less informed than player 2, under finite-memory strategies for both players, are 2-EXPTIME-complete.*

Remark 2. Note that the lower bounds for Theorem 5 are established for reachability objectives. Moreover, it was shown in [7, Section 5] that for qualitative analysis of two-player partial-observation stochastic games with reachability objectives, finite-memory strategies suffice, i.e., if there is a strategy to ensure almost-sure (resp., positive) winning, then there is a finite-memory strategy. Thus the results of Theorem 5 hold for reachability objectives even without the restriction of finite-memory strategies, and it extends the result of [7, Theorem 1] which showed EXPTIME-completeness for reachability objectives when player 2 has perfect observation.

References

1. Full version. In *ArXiv CoRR: http://arxiv.org/abs/1404.5453*, 2014.

2. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.
3. C. Baier, N. Bertrand, and M. Größer. On decision problems for probabilistic Büchi automata. In *Proc. of FoSSaCS*, LNCS 4962, pages 287–301. Springer, 2008.
4. N. Bertrand, B. Genest, and H. Gimbert. Qualitative determinacy and decidability of stochastic games with signals. In *Proc. of LICS*, pages 319–328, 2009.
5. K. Chatterjee. *Stochastic ω-Regular Games*. PhD thesis, UC Berkeley, 2007.
6. K. Chatterjee and L. Doyen. The complexity of partial-observation parity games. In *Proc. of LPAR*, LNCS 6397, pages 1–14. Springer, 2010.
7. K. Chatterjee and L. Doyen. Partial-observation stochastic games: How to win when belief fails. In *Proc. of LICS 2012; Journal version ACM ToCL*, pages 175–184. IEEE, 2012.
8. K. Chatterjee, L. Doyen, H. Gimbert, and T. A. Henzinger. Randomness for free. In *Proc. of MFCS*, LNCS 6281, pages 246–257. Springer, 2010.
9. 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.
10. K. Chatterjee, L. Doyen, S. Nain, and M. Y. Vardi. The complexity of partial-observation stochastic parity games with finite-memory strategies. In *Proc. of FoSSaCS*, LNCS. Springer, 2014.
11. K. Chatterjee, M. Jurdziński, and T. A. Henzinger. Simple stochastic parity games. In *CSL'03*, volume 2803 of *LNCS*, pages 100–113. Springer, 2003.
12. M. De Wulf, L. Doyen, and J.-F. Raskin. A lattice theory for solving games of imperfect information. In *Proc. of HSCC*, LNCS 3927, pages 153–168. Springer, 2006.
13. B. Finkbeiner and S. Schewe. Coordination logic. In *Proc. of CSL*, LNCS 6247, pages 305–319. Springer, 2010.
14. T. A. Henzinger and P.W. Kopke. Discrete-time control for rectangular hybrid automata. *Theor. Comp. Science*, 221:369–392, 1999.
15. A. Kechris. *Classical Descriptive Set Theory*. Springer, 1995.
16. P. Madhusudan and P. S. Thiagarajan. Distributed controller synthesis for local specifications. In *ICALP*, pages 396–407, 2001.
17. D.A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.
18. S. Mohalik and I. Walukiewicz. Distributed games. In *FSTTCS*, pages 338–351, 2003.
19. C. H. Papadimitriou. *Computational complexity*. Addison-Wesley, 1994.
20. G. L. Peterson and J. H. Reif. Multiple-person alternation. In *FOCS*, pages 348–363, 1979.
21. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. of POPL*, pages 179–190. ACM Press, 1989.
22. 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.
23. R. Ramanujam and S. E. Simon. A communication based model for games of imperfect information. In *Proc. of CONCUR*, LNCS 6269, pages 509–523. Springer, 2010.
24. J. H. Reif. Universal games of incomplete information. In *Proc. of STOC*, pages 288–308, 1979.
25. J. H. Reif. The complexity of two-player games of incomplete information. *JCSS*, 29:274–301, 1984.
26. J. H. Reif and G. L. Peterson. A dynamic logic of multiprocessing with incomplete information. In *Proc. of POPL*, pages 193–202. ACM, 1980.
27. W. Thomas. Languages, automata, and logic. In *Handbook of Formal Languages*, volume 3, Beyond Words, chapter 7, pages 389–455. Springer, 1997.