Rational Synthesis Under Imperfect Information

Emmanuel Filiot
Université Libre de Bruxelles

Raffaella Gentilini
University of Perugia

Jean-François Raskin
Université Libre de Bruxelles

Abstract

In this paper, we study the rational synthesis problem for turn-based multiplayer non zero-sum games played on finite graphs for omega-regular objectives. Rationality is formalized by the concept of Nash equilibrium (NE). Contrary to previous works, we consider here the more general and more practically relevant case where players are imperfectly informed. In sharp contrast with the perfect information case, NE are not guaranteed to exist in this more general setting. This motivates the study of the NE existence problem. We show that this problem is ExpTime-C for parity objectives in the two-player case (even if both players are imperfectly informed) and undecidable for more than 2 players. We then study the rational synthesis problem and show that the problem is also ExpTime-C for two imperfectly informed players and undecidable for more than 3 players. As the rational synthesis problem considers a system (Player 0) playing against a rational environment (composed of k players), we also consider the natural case where only Player 0 is imperfectly informed about the state of the environment (and the environment is considered as perfectly informed). In this case, we show that the ExpTime-C result holds when k is arbitrary but fixed. We also analyse the complexity when k is part of the input.

1 Introduction

Context Reactive synthesis aims at producing automatically a correct reactive system from a specification and a model of the environment that interacts with the system. The system should enforce its specification no matter how the environment behaves. Two-player *zero-sum* games played on graphs is the classical mathematical model proposed to formalize the reactive synthesis problem [32] and the main solution concept for those games is the notion of winning strategy. This model is appropriate to model the situation where a monolithic controller has to be designed to interact with a monolithic and *fully antagonistic* environment.

A fully antagonistic environment is often a bold abstraction as the environment usually has its own goal which, in general, is *not* the negation of the specification of the reactive system. Nevertheless, it is a popular abstraction because it is simple and sound: a winning strategy against an antagonistic environment is clearly winning against an environment that strives for its own objective. However this abstraction may fail to find a winning strategy even if solutions exist when the objective of the environment is taken into account. As a consequence, several more recent works consider the more general model of *non* zero-sum games played on graphs in order to overcome the weakness of the simpler zero-sum model, see e.g. [5, 7, 8, 12, 14, 15, 18, 19, 23, 33].

Rational synthesis under imperfect information In this paper, we build on the notion of rational synthesis that has been introduced in [19] and [23]. In rational synthesis, the synthesis algorithm takes into account that the system is executed within an environment that is assumed to be *rational*, and not fully antagonistic as in the classical two player zero-sum setting. Rationality of the environment is modeled by assuming that the components behave according to a Nash equilibrium (NE). More precisely, in the *(non-cooperative) rational synthesis* problem [23], the components of the environment may follow any strategy, provided it is a NE: we search for a strategy σ_0 for the system winning against all the possible strategy profiles that include σ_0 for Player 0 and which are NE. In all previous works on rational synthesis, including [17], players are assumed to be perfectly informed. We lift this important restriction here as we consider the more general and more practically relevant case where players are *imperfectly informed*. This generalization has important consequences: in sharp contrast with the perfect information case, NE are not guaranteed to exist in this more general setting. So in addition to the study of the (non-cooperative) rational synthesis problem, this motivates the study of the NE existence problem for games with imperfect information.

Contributions We show that the decision problem of the existence of a Nash equilibrium is ExpTime-C for omega-regular objectives defined by parity conditions in the two-player case (Theorem 3.9), even if both players are imperfectly informed, and undecidable for 3 and more players (Theorem 3.11), even if only two players are imperfectly informed. We then study the rational synthesis problem and show that the problem is also solvable for two imperfectly informed players and ExpTime-C (Theorem 4.2), while we establish undecidability for more than 3 players (Theorem 4.7). As the rational synthesis problem considers a system (Player 0) playing against a rational environment composed of k players, we also consider the natural case where only Player 0 is imperfectly informed about the state of the environment and the environment is considered to be *omniscient*. In this case, we show that the ExpTime-C result holds when k is arbitrary but fixed (Theorem 4.6). This parametric analysis makes sense as the number of players can be expected to be low and much smaller than the number of states in the game arena. When k is not fixed, we provide a 2ExpTime solution together with an ExpSpace lower bound for omega-regular objectives defined by Rabin conditions (Theorem 5.4). To establish this last result, we show that the universality problem for nondeterministic Multi-Rabin automata on infinite words is ExpSpace-C (Theorem 5.3), this result strengthens a result by Safra and Vardi

This work was partially supported by InDAM-GNCS, the ERC Starting grant 279499 (inVEST), the ARC projects "Transform" and "Non-Zero Sum Game Graphs: Applications to Reactive Synthesis and Beyond" (Fédération Wallonie-Bruxelles), the EOS project "Verifying Learning Artificial Intelligence Systems" (FNRS-FWO) and the Belgian FNRS CDR project "Flare" J013116F. J.-F. Raskin is Professeur Francqui de Recherche funded by the Francqui foundation. E. Filiot is a research associate of F.R.S.-FNRS. .

LICS '18, July 9–12, 2018, Oxford, United Kingdom

ACM ISBN 978-1-4503-5583-4/18/07…$15.00
https://doi.org/10.1145/3209108.3209164

in [31] for Muller acceptance conditions expressed as Emerson-Lei formulas.

Two-player imperfect information zero-sum games are classically solved by reduction to the perfect information setting, via a knowledge construction [28]. In the multiplayer setting, and even in the case where only Player 0 is imperfectly informed, it is not clear whether a reduction to multiplayer perfect information rational synthesis exists. Instead, we provide a reduction to a two-player imperfect information zero-sum game, where the protagonist (Prover) wants to exhibit a solution (a strategy for Player 0) to the rational synthesis problem while the adversary tries to prove that such a strategy does not exist. To prove that there exists a solution, Prover chooses actions for Player 0 and Challenger actions for the environment players. The objective of Prover is to construct a play which is either winning for Player 0, or which is not the outcome of any NE. To check the latter, Prover may also declare, along the play, that some players have profitable deviations. The additional role of Challenger is also to force Prover to show that those players had indeed a good deviation, by taking the role of the coalition of all players against the deviating player (controlled by Prover). This results in a complex winning objective for Prover which we analyse carefully to get our complexity results.

Related works Non-zero games for synthesis have attracted a large attention recently, see e.g. [5] for a survey. *Secure equilibria* were introduced in [15] and their use for synthesis was established in [14]. Secure equilibria are refinement of NE [33]. Subgame perfect equilibria, that also refines NE, were first studied in [33, 34] and more recently in [9, 10]. To model rationality of players, the notion of *admissible strategy* is used in [2, 18] instead of the notion of NE, and computational aspects are studied in [8], potential for synthesis is studied in [7]. Rational synthesis was introduced first in [19, 23] and only LTL objectives were studied. More recently, we have study the complexity of all the classical winning conditions in [17]. In all those works, in contrast to this paper, the players have perfect information about the game.

Games played on graphs with imperfect information were studied in e.g. [3, 13, 24, 27–29]. All those works consider zero-sum games only. More recently, some works have started the study of non zero-sum games with imperfect information. In [20], Gutierrez et al. study the existence problem for NE in concurrent games with imperfect information, and LTL objectives. First, they show that, similarly to the constrained existence problem [26, 27], the unconstrained existence problem is undecidable for N-player games. Their undecidability result can be applied to our setting (see Theorem 3.11). Second, they show that the (unconstrained) existence problem is decidable for two players and 2ExpTime-C for LTL objectives. As we consider parity objectives, our ExpTime-C results for the constrained and unconstrained existence problem for NE in the two player setting is not a consequence of their result. Finally, they do not study the rational synthesis problem which is our main contribution.

The notion of admissible strategy was studied for observation-based strategies in [6], and *Doomsday equilibria*, that are an extension of secure equilibria to the k player case [12], are also considered for imperfect information. In [1], Berthon et al. study an extension of strategy logic [16, 25] in which quantification over strategies can be restricted to *observation-based* strategies. This opens the possibility to reason on strategies in games with imperfect information. The model-checking of the Strategy Logic with imperfect information is undecidable in general and the authors have identified a fragment of the logic that has a decidable model-checking problem but they do not provide complexity results. Their fragment does not cover the positive results that we have for the two-player case where both players are imperfectly informed, while their fragment can be used to show the decidability (but not the complexity bound) of the case where only the system is imperfectly informed. In [4], Bouyer studies the existence of NE in multi-player games where the information of the player is imperfect but symmetric: all the players receive the same information provided by *public signals*. This model with public signals is incomparable to our setting in which every player has his own imperfect view of the system state. In general, the NE existence problem is undecidable in the setting of public signals and the author identifies sub-cases that can be solved algorithmically.

Structure of the paper In Sec. 2, we define the necessary preliminaries. In Sec. 3, we study the existence problem for Nash equilibria. In Sec. 4, we study the rational synthesis problem. In Sec. 5, we provide an additional discussion of lower-bounds for our problems.

2 Preliminaries and Notation

2.1 Multiplayer Games with Imperfect Information

Definition 2.1 (Multiplayer Arena with Imperfect Information)**.** Given $k \in \mathbb{N}$, a *multiplayer arena* with $k+1$ players is a tuple $\mathcal{A} = \langle \Omega, V, v_0, (V_i)_{i\in\Omega}, (Act_i)_{i\in\Omega}, (\delta_i)_{i\in\Omega}, (\sim_i)_{i\in\Omega}\rangle$, where:

- $\Omega = \{0, 1, ..., k\}$ is the finite set of players,
- V is a finite set of states and $v_0 \in V$ is the initial state,
- $(V_i)_{i\in\Omega}$ is a partition of V where V_i is the set of states controlled by Player $i \in \Omega$,
- for every $i \in \Omega$, Act_i is the set of actions of Player i,
- for every $i \in \Omega$, $\delta_i : (V_i \times Act_i) \to V$ is a transition function specifying the next state in the game,
- for every $i \in \Omega$, $\sim_i$ is an equivalence relation on V such that $\{[v]_{\sim_i} \mid v \in V\}$ refines $(V_i)_{i\in\Omega}$ i.e. players observe the turns, but have imperfect information on the actual state of the game, this is assumed w.l.o.g. If $\forall j \in \Omega \cdot \forall v, v' \in V_j : v \sim_i v'$, we say that Player i is *blind*.

A *play* in a $(k+1)$-players arena $\mathcal{A}$ starts in the the initial state v_0 and proceeds in rounds. At each round, the player controlling the current state $v \in V_i$, i.e. Player i, who can observe his turn and the equivalence class $[v]_{\sim_i}$ but not the actual state $v \in V_i$ of the game, chooses the next action $a \in Act_i$ to play, and δ_i determines the next state of the play $v' = \delta_i(v)$. Formally, a play $\pi = v_0a_0v_1a_1\dots$ is an infinite sequence of states and actions such that for all positions $j \geq 0$, $v_{j+1} = \delta_i(v_j, a_j)$ where $v_j \in V_i$ and $a_j \in Act_i$, and where v_0 is the initial state of the game arena. We note $\pi_v[j]$ the state v_j in position j along π and $\pi_a[j]$ the action a_j in position j along π.

Let $\pi = v_0a_0v_1a_1\dots$ be a play, a prefix α of π is a finite sequence $\alpha = v_0a_0v_1a_1\dots v_na_n$ or $\alpha = v_0a_0v_1a_1\dots v_n$. We denote by $\alpha_v[j]$ the state v_j in position j within α, and by $\alpha_a[j]$ the action a_j in position j within α. Sometimes, we do not need explicit references to actions in plays and prefixes, then we consider that plays or prefixes are simply sequences of states. We reflect this into our notations as follows:

- $\mathsf{Plays}(\mathcal{A})$ and $\mathsf{Prefs}(\mathcal{A})$ denote the set of plays and the set of prefixes where the actions are omitted. If $\pi = v_0v_1\dots v_n\dots$

is a play where action have been omitted, we write $\pi[j]$ for $\pi_v[j]$, and $\alpha[j]$ for $\alpha_v[j]$, and $\pi[:n]$ for the sequence of states in π up to position n, i.e. $\pi[:n] = v_0 v_1 \dots v_n$.

- $\mathsf{Plays}^{\mathsf{Act}}(\mathcal{A})$ and $\mathsf{Prefs}^{\mathsf{Act}}(\mathcal{A})$ denote the set of plays and the set of prefixes where the actions are *not* omitted.

We omit the reference to $\mathcal{A}$ in those notations when it is clear from the context. For every $i \in \Omega$, we extend the equivalence relation $\sim_i$ from V to prefixes and plays as follows:

- when actions are omitted: let $\alpha, \beta \in \mathsf{Prefs}$ (resp. $\pi, \pi' \in \mathsf{Plays}$) are said *equivalent* for Player i, denoted $\alpha \sim_i \beta$ (resp. $\pi \sim_i \pi'$) if and only if they have the same length n and for every $0 \leq m \leq n$, $\alpha[m] \sim_i \beta[m]$ (resp. for every $0 \leq m$, $\pi[m] \sim_i \pi'[m]$).
- when actions are *not* omitted: let $\alpha, \beta \in \mathsf{Prefs}^{\mathsf{Act}}$ (resp. $\pi, \pi' \in \mathsf{Plays}^{\mathsf{Act}}$) are said *action equivalent* for Player i, denoted $\alpha \sim_i^a \beta$ (resp. $\pi \sim_i^a \pi'$) if and only if they have the same length n and for every $0 \leq m \leq n$, $\alpha_v[m] \sim_i \beta_v[m]$, and additionally if $v_m \in V_i$ then $\alpha_a[m] = \beta_a[m]$ (resp. for every $0 \leq m$, $\pi_v[m] \sim_i \pi'_v[m]$ and if $\pi_v[m] \in V_i$ then $\pi_a[m] = \pi'_a[m]$). So, when defining the equivalence for prefixes (resp. plays) with actions, we require that the actions played by Player i are the same in both prefixes (resp. plays).

Given a play $\pi \in \mathsf{Plays}^{\mathsf{Act}}$, we write $[\pi]^a_{\sim_i}$ for all the plays π that are action equivalent to π for Player i. Similarly, given a prefix of play α, we write $[\alpha]^a_{\sim_i}$ for the set of prefixes β that are action equivalent to α for Player i .

A *strategy* for Player $i \in \Omega$ in a $(k+1)$-players arena $\mathcal{A}$ is a total function $\sigma_i : \mathsf{Prefs} \cdot \mathsf{V_i} \mapsto \mathsf{Act_i}$. A play π is *consistent* with σ_i if $\pi[n+1] = \delta_i(\pi[n], \sigma_i(\pi[:n]))$ for all $n \geq 0$ s.t. $\pi[n] \in V_i$. The *outcome* of σ_i is the set of plays $out(\sigma_i) \subseteq \mathsf{Plays}(\mathcal{A})$ that are consistent with σ_i. A strategy σ_i for Player i is an *observation-based strategy* if for every $\alpha, \beta \in \mathsf{Prefs} \cdot V_i$, if $\alpha \sim_i \beta$, then $\sigma_i(\alpha) = \sigma_i(\beta)$. We denote by Σ_i (resp. Σ_i^{obs}) the set of strategies (resp. observation-based strategies) of Player i. Given a prefix $\alpha = v_0 v_1 \dots v_n$, the *observation* of α by Player i is the sequence $[v_0]_{\sim_i}[v_1]_{\sim_i} \dots [v_n]_{\sim_i}$, i.e. the sequence of observations that Player i receives along α.

A (observation-based) *strategy profile* in a $(k+1)$-players arena $\mathcal{A}$ is a tuple $\bar{\sigma} = \langle \sigma_0, \dots, \sigma_k \rangle$ of (observation-based) strategies, one for each of the $k+1$ involved players. The outcome of a strategy profile $\bar{\sigma}$ is the unique play that is consistent with all the strategies of the strategy profile. This play is denoted by $out(\bar{\sigma})$.

A *winning objective* (or just objective) is a set $O \subseteq V^\omega$. We say that an objective is *observable* by Player i if for each pair of plays π, π', if $\pi \sim_i \pi'$, then $\pi \in O$ if and only if $\pi' \in O$. In this paper, we consider (omega-regular) objectives defined by *parity* conditions. For a play π, we note $\inf(\pi)$ the set of states v that appear infinitely many times along π. A parity mapping is a function $p{:}V{\to}\mathbb{N}$ that assigns a color (also called priority) to each state in the arena. Given a parity mapping p, we define the associated objective as $\mathsf{Parity}(p){=}\{\pi{\in}V^\omega \mid \min\{p(v) | v \in \inf(\pi)\} \text{ is even}\}$. We define the complement of p, noted $\overline{p} : V{\to}\mathbb{N}$, as $\overline{p}(v) = p(v) + 1$ for all $v \in V$, and we have that $\mathsf{Parity}(\overline{p}) = \overline{\mathsf{Parity}(p)}$. A strategy σ_i for Player i is *winning*[1] for objective O if $out(\sigma_i){\subseteq}O$.

[1]Here we implicitly consider a two-player zero-sum game in which Player i has objective O and plays against all the other players in $\Omega \setminus \{i\}$ who have objective $\overline{O}$.

A *multiplayer parity game* with imperfect information is a pair $\mathcal{G} = \langle \mathcal{A}, (O_i)_{i \in \Omega} \rangle$ where $\mathcal{A}$ is an imperfect information multiplayer arena and $(O_i)_{i \in \Omega}$ are parity objectives for each Player $i{\in}\Omega$. The notations Plays and Prefs carries over to $\mathcal{G}$ by considering its underlying arena. We often directly write $\mathcal{G} = \langle \mathcal{A}, (p_i)_{i \in \Omega} \rangle$, when the objective $(O_i)_{i \in \Omega}$ are given by parity conditions.

2.2 Nash Equilibria and Rational Synthesis

Given a *strategy profile* $\bar{\sigma}$ in a multiplayer game $\mathcal{G} = \langle \mathcal{A}, (O_i)_{i \in \Omega} \rangle$ and a strategy τ for Player i, we write $(\bar{\sigma}_{-i}, \tau)$ for the strategy profile obtained by replacing σ_i with τ in $\bar{\sigma}$. Given winning objectives $(O_i)_{i \in \Omega}$ for each player, the *payoff* of a strategy profile $\bar{\sigma}$ is the vector $pay(\bar{\sigma}) \in \{0,1\}^{k+1}$ defined by $pay(\bar{\sigma})[i] = 1$ iff $out(\bar{\sigma}) \in O_i$. We write $pay_i(\bar{\sigma})$ for Player i's payoff in $pay(\bar{\sigma})$. Payoffs are compared by the pairwise natural order on their bits, denoted by $\leq$, i.e. $pay(\bar{\sigma}) \leq pay(\bar{\beta})$ if $pay_i(\bar{\sigma}) \leq pay_i(\bar{\beta})$ for all $i \in \Omega$. If a profile of strategy $\bar{\sigma}$ is a NE with payoff $\bar{b}$, then we say that $\bar{\sigma}$ is a $\bar{b}$−NE.

A *Nash equilibrium* (NE) for an imperfect information game $\mathcal{G} = \langle \mathcal{A}, (O_i)_{i \in \Omega} \rangle$ is an *observation-based strategy profile* $\bar{\sigma} = (\sigma_i)_{i \in \Omega}$ for $\mathcal{G}$ such that no player can improve his payoff by (unilaterally) switching to a different observation-based strategy: $\bar{\sigma}$ is a NE if for all players $i \in \Omega$ and all observation-based strategies τ of Player i, $pay(\bar{\sigma}_{-i}, \tau) \leq pay(\bar{\sigma})$. We say that $\bar{\sigma}$ is a 0-*fixed* Nash equilibrium (0NE) if $pay(\bar{\sigma}_{-i}, \tau) \leq pay(\bar{\sigma})$ for all players $i \in \Omega \setminus \{0\}$ and all observation-based strategies τ of i. In other words, it is a Nash equilibrium in which Player 0 is *not* allowed to deviate. Any NE is 0-fixed, but the converse may not hold.

Rational synthesis aims at finding an observation-based strategy for the system (Player 0) that is winning for the objective of Player 0 whenever the environment composed of several components (Players 1 to k) plays rationally. Rationality of the environment is modeled by NE. The formal definition of the *rational synthesis* problem, RS for short, is as follows. Given a multi-player game $\mathcal{G}$ with imperfect information, the observation-based strategy σ_0 for Player 0 is a solution to the *rational synthesis* problem if for all 0-fixed Nash equilibria $\bar{\sigma} = \langle \sigma_0, \dots, \sigma_k \rangle$, we have $pay_0(\bar{\sigma}) = 1$. We study in the rest of the paper the existence problems for NE and RS.

3 Nash Equilibria

In this section, we study the existence problem for NE for omega-regular objectives defined by parity conditions. First, we show that in sharp contrast to the perfect information sub-case, NE need not to exist in the imperfect information setting and so this provides motivations for the study of the existence problem for NE. Second, as a main result for this subsection, we provide an ExpTime algorithm to solve the existence decision problem together with a matching lower-bound for the two-player case. On the way to obtain our solution, we show that we can also solve the existence of *constrained* NE with the same upper bound. We close the section by showing that the problem of determining the existence of a NE is undecidable for games with more than 2 players.

3.1 Games without NE

The next proposition establishes that NE needs not to exist in games with imperfect information. The example used in the proof will be used in several other proofs in the paper.

Proposition 3.1. *There exists a 2-player imperfect information game which does not admit a Nash equilibrium.*

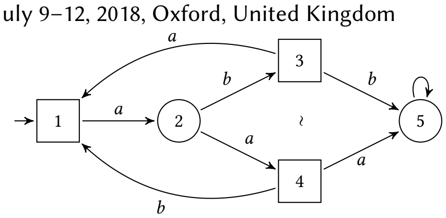

Figure 1. Imperfect information 2-player game where no NE exists. This game is called the game $\mathcal{H}$. The parity function of Player $\square$ is $p(v) = 1$ for $v \in \{1, 2, 3, 4\}$, and $p(v) = 0$ for $v = 5$. The parity function for Player $\bigcirc$ is $\overline{p}$, the complement of p. It is easy to see that Player $\square$ wants to reach 5 while Player $\bigcirc$ wants to avoid it. States 3 and 4 are indistinguishable for Player $\square$, i.e. $3 \sim_{\square} 4$.

Proof. Consider the 2-player game of Fig. 1. The parity objective of Player square is equivalent to the reachability objective "reach state 5", while the parity objective of the circle player is the complement and so it is equivalent to the safety objective "stay forever in the set of states $\{1, 2, 3, 4\}$. Suppose there exists a NE $\overline{\sigma} = \langle \sigma_{\bigcirc}, \sigma_{\square} \rangle$. We consider two cases:

- if $out(\overline{\sigma})$ reaches 5, then knowing the strategy $\sigma_{\square}$, Player $\bigcirc$ can avoid reaching 5 by changing his strategy into $\sigma'_{\bigcirc}$ as follows: take a history h which ends up in state 2, then, whatever Player $\bigcirc$ does, Player $\square$ does not see it, and gets the observation $\{3, 4\}$. Now, if $\sigma_{\square}(h\{3, 4\}) = a$, set $\sigma'_{\bigcirc}(h) = b$, otherwise set $\sigma'_{\bigcirc}(h) = a$. Clearly, $out(\langle \sigma'_{\bigcirc}, \sigma_{\square} \rangle)$ is winning for Player $\bigcirc$, hence Player $\bigcirc$ by deviating gets a better payoff. Contradiction.
- if $out(\overline{\sigma})$ does not reach 5, then Player $\square$ has a spoiling strategy $\sigma'_{\square}$ with strictly better payoff: let $h\{3, 4\}$ be a history, then set $\sigma'_{\square}(h\{3, 4\}) = \sigma_{\bigcirc}(h)$. Contradiction.

□

The main property we have exploited in the proof of Prop. 3.1 is a weak form of determinacy: the existence of spoiling observation-based strategy profiles as stated in the following lemma.

Lemma 3.2. *Let $\alpha \in \mathsf{Prefs}(\mathcal{A})$, if a strategy σ_i is not winning from α for parity objective p_i then there exists a profile of observation-based strategies $\bar{\sigma}_{-i}$ for the other players such that the only outcome compatible with both σ_i and $\bar{\sigma}_{-i}$ and that starts in α satisfies $\overline{p_i}$.*

3.2 The NE existence problem for 2 player games

For the rest of this section, we fix a two-player game $\mathcal{G} = \langle \mathcal{A}, p_0, p_1 \rangle$ with parity objectives, and let $i \in \{0, 1\}$. To introduce our solutions, we need the following additional notions and intermediary results.

Lemma 3.3. *Let σ_i be an observation-based strategy, let $\pi \in \mathsf{Plays}^{\mathsf{Act}}(\mathcal{A})$ be a play that is compatible with σ_i. For all $\pi' \in \mathsf{Plays}^{\mathsf{Act}}(\mathcal{A})$ such that $\pi' \sim^a_i \pi$, we have that:*

1. *π' is compatible with the strategy σ_i*
2. *there exists an observation-based strategy σ_{1-i} for Player $1 - i$ s.t. π' is the only outcome compatible with both σ_i and σ_{1-i}.*

Let $\alpha \in \mathsf{Prefs}^{Act}(\mathcal{A})$, we define $K^{\sim_i}(\alpha)$ as the set of states $\{last(\beta) \mid \beta \in [\alpha]^a_{\sim_i}$ and $\beta \in \mathsf{Prefs}(\mathcal{A})\}$, i.e. the set of states that can be reached in $\mathcal{A}$ while Player i observes a sequence which is $\sim^a_i$-equivalent to α. This set is called the *knowledge* of Player i after prefix α. We write $\alpha \sqsubset^a_{\sim_i} \pi$ when there exists $\pi' \sim^a_i \pi$ and $\alpha \sqsubset \pi'$, i.e. when α is a prefix which Player i cannot distinguish from π.

The following lemma tells us that all the states of $K^{\sim_i}(\alpha)$ are reachable when Player i plays an observation-based strategy σ_i that is compatible with α by choosing an adequate observation-based strategy for Player $1 - i$.

Lemma 3.4. *Let $\alpha \in \mathsf{Prefs}^{Act}(\mathcal{A})$ be a prefix compatible with strategy σ_i of Player i. For all states $v \in K^{\sim_i}(\alpha)$, there is a observation-based strategy σ_{1-i} of Player $1 - i$ and there is a prefix $\beta \in [\alpha]^a_{\sim_i}$ compatible with both σ_i and σ_{1-i} and such that $last(\beta) = v$.*

A play π is an outcome of a NE if no player has an incentive to deviate along π. To avoid deviation of Player i, Player $1 - i$ must be in position to threaten Player i of retaliation. A prefix $\alpha \in \mathsf{Prefs}^{Act}(\mathcal{A})$ is *good for retaliation* for Player i if there is an observation-based strategy σ_i of Player i such σ_i wins the objective $\overline{p_{1-i}}$ from α. Note that as parity objectives are prefix independent, this is equivalent to say that Player i has a strategy to force an outcome that falsifies the objective of the other player, for all the states in $K^{\sim_i}(\alpha)$.

We are now equipped to establish the correctness of a lemma that identifies a necessary and sufficient condition for the existence of a $(1, 0) - \mathsf{NE}$, i.e. a NE where Player 0 wins and Player 1 fails to win. We then show that this condition can be decided in ExpTime.

Lemma 3.5. *There exists a $(1, 0){-}\mathsf{NE}$ in $\mathcal{A}$ for parity objectives p_0 and p_1 iff there exists $\pi \in \mathsf{Plays}(\mathcal{A})$ s.t. the following properties hold:*

(P_1) *$\pi \models p_0$ and for all $\pi' \sim^a_0 \pi$, $\pi' \not\models p_1$,*
(P_2) *for all prefixes $\alpha \sqsubset^a_{\sim_0} \pi$, α is* good for retaliation *for Player* 0.

The existence of such a play π can be decided in ExpTime.

Proof. Let us first assume that there exists a $(1, 0) - \mathsf{NE}$ and let us establish the existence of a play with properties (P_1) and (P_2). Let $\bar{\sigma} = (\sigma_0, \sigma_1)$ be such a $(1, 0) - \mathsf{NE}$, and let $\pi = out(\bar{\sigma})$. Let us show that π satisfies the two properties.

We start with property (P_1). As π is the outcome of $\bar{\sigma}$ that is a $(1, 0) - \mathsf{NE}$, then we know that $\pi \models p_0$ and $\pi \not\models p_1$. Let us establish that for all $\pi' \sim^a_0 \pi$, we have also that $\pi' \not\models p_1$. For that let us consider an observation-based strategy σ'_1 of Player 1 such that the unique outcome of (σ_0, σ'_1) is π'. The existence of σ'_1 is guaranteed by lemma 3.3. Now assume for the sake of contradiction that $\pi' \models p_1$. If so, then σ'_1 would be a profitable deviation for Player 1. But as $\bar{\sigma}$ is a NE, this cannot be the case and so we need to conclude that $\pi' \not\models p_1$, and so (P_1) is satisfied.

We establish the property (P_2) by contradiction. Assume that there is a prefix $\alpha \sqsubset^a_{\sim_0} \pi$, that is not *good for retaliation* for Player 0. This implies that there exists a state $v \in K^{\sim_0}(\alpha)$ from which σ_0 is not winning for $\overline{p_1}$. As a consequence, by lemmas 3.2, there is an observation-based Player 1 strategy σ'_1 such the unique play that is compatible with both σ_0 and σ'_1 satisfies p_1 from state v. From σ'_1, we construct the observation-based strategy σ''_1 for Player 1 as follows. The strategy σ''_1 is chosen such that when played against σ_0 in $\mathcal{A}$, we reach state v. Such a strategy is guaranteed to exist according to lemma 3.4. Now, when v is reached (Player 1 knows it as he only needs to count the number of turns that he plays against σ_0 to determine when he reaches v), σ''_1 behaves as σ'_1 from state v. As parity objectives are prefix independent, we know that the unique outcome compatible with σ_0 and σ''_1 satisfies p_1. This shows that Player 1 has a profitable deviation from the profile $\bar{\sigma}$ and so this profile cannot be a NE, and so we have obtained our contradiction.

Let us now consider the other direction. Let π be an infinite path from v_0 that satisfies properties (P_1) and (P_2). From π, we design a $(1,0)-\mathsf{NE}$ as follows. We first define the observation-based strategy σ_1 of Player 1. This strategy σ_1 simply follows the actions as prescribed by π for all prefixes $\alpha \sqsubset^a_{\sim_1} \pi$. For all histories β such that $\beta \not\sqsubset^a_{\sim_1} \pi$, σ_1 plays arbitrarily. The strategy σ_0 of Player 0 is defined as follows. It plays according to π on all prefixes α such that $\alpha \sqsubset^a_{\sim_0} \pi$, otherwise Player 0 observes a deviation from π and plays according to an observation-based strategy that enforces $\overline{p_1}$, i.e. the complement of objective p_1. Such a strategy is guaranteed to exist thanks for property (P_2). Now, let us show that this strategy profile $\bar{\sigma} = (\sigma_0, \sigma_1)$ witnesses the existence of a $(1,0)-\mathsf{NE}$. Indeed, we first note that if both players follows their strategies in $\bar{\sigma}$ then the outcome of the game is π, and π satisfies p_0 and falsifies p_1. Second, we note that as π satisfies p_0, Player 0 has no incentive to deviate, and we only need to consider deviations of Player 1. Assume that Player 1 plays an observation-based strategy σ_1' instead of σ_1. Let the outcome of the profile (σ_0, σ_1') be π'. Then there are two cases to consider. First, assume that $\pi \sim^a_0 \pi'$, then in that case, by condition (P_1), we know that $\pi' \not\models p_1$ and so the deviation is not profitable to Player 1. Second, let us consider the case $\pi' \not\sim^a_0 \pi'$. Let α the the longest prefix such that $\alpha \sqsubset^a_{\sim_i} \pi$ and $\alpha \sqsubset^a_{\sim_i} \pi'$. By condition (P_2), we know that α is good for retaliation, and by construction of σ_0, when the next observation is given to Player 0, he observes a prefix β that is such that $\beta \not\sqsubset^a_{\sim_0} \pi$ and so he switches to its retaliation strategy. This ensure that the outcome falsifies p_1. So, again the deviation is not profitable to Player 1.

Now that we have established the correctness and completeness of properties (P_1) and (P_2), let us show that we can check them in ExpTime. To test the existence of π with the good properties, we proceed as follows:

- We construct two nondeterministic parity automata A_{p_0} and $A_{\overline{p_1}}$ that accept exactly all the sequences of observations for Player 0 such that all the underlying plays are satisfying p_0 and $\overline{p_1}$, respectively. This can be done in linear time by a construction given in Sect.3 of [11]. By taking the product of the two automata A_{p_0} and $A_{\overline{p}_1}$, we obtain a nondeterministic Streett automaton B that accepts exactly all observation sequences of paths that satisfy condition (P_1).
- We determine the set of knowledges $K \subseteq V$ of Player 0 from which Player 0 has a winning observation-based strategy for the objective $\overline{p_1}$, we denote this set $\mathcal{R}$. This can be done in ExpTime [28]. From this set it is now easy to define a deterministic safety automaton, that we note R, that keeps tracks of all sequences of observations that are good for retaliation as we simply have to make sure that we never leave the set of knowledge in $\mathcal{R}$.

Taking the product of B with R, we get an exponentially large nondeterministic Street automaton whose language is non-empty iff there exists a play π that respects properties (P_1) and (P_2). Checking emptiness of nondeterministic Street automata can be done in PTime [21]. All this gives us the claimed ExpTime procedure. ◻

As our games are symmetric and we can exchange the role of Player 0 and Player 1 in the previous lemma, we obtain that:

Corollary 3.6. *Given a two-player imperfect information game $(\mathcal{A}, p_0, p_1)$ with parity objectives p_0 and p_1, we can decide in ExpTime is there exists a $(1,0)-\mathsf{NE}$, and we can decide in ExpTime is there exists a $(0,1)-\mathsf{NE}$.*

Now, we turn our attention to the $(0,0)-\mathsf{NE}$ case.

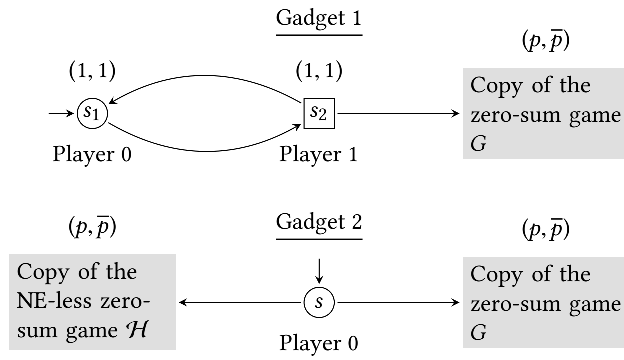

Figure 2. Two gadgets for the hardness of deciding the existence of NE and constrained NE.

Lemma 3.7. *Given a two-player imperfect information game $(\mathcal{A}, p_0, p_1)$ with parity objectives p_0 and p_1, we can decide in ExpTime if there exists a $(0,0)-\mathsf{NE}$ equilibrium.*

Proof. By inspecting the proof of lemma 3.5, it is not difficult to see that, to establish the existence of a $(0,0)-\mathsf{NE}$ equilibrium, we need to determine the existence of a path π that falsifies both p_0 and p_1, and in which both players can retaliate in case of deviation from the expected sequence of observations. This is expressed by the following four properties:

- (P_{0a}) $\pi \not\models p_0$ and for all $\pi' \sim_0 \pi$, $\pi' \not\models p_1$,
- (P_{0b}) $\pi \not\models p_1$ and for all $\pi' \sim_1 \pi$, $\pi' \not\models p_0$,
- (P_{1a}) for all prefixes $\alpha \sqsubset_{\sim_0} \pi$, α is *good for retaliation* for Player 0.
- (P_{1b}) for all prefixes $\alpha \sqsubset_{\sim_1} \pi$, α is *good for retaliation* for Player 1.

Such properties can be verified in ExpTime using a construction similar to the one used in the proof of lemma 3.5. ◻

We finish by the $(1,1)-\mathsf{NE}$ case, which is the simplest case.

Lemma 3.8. *Given a 2-player imperfect information game $(\mathcal{A}, p_0, p_1)$ for parity objectives p_0 and p_1, we can decide in PTime if there exists a $(1,1)-\mathsf{NE}$ equilibrium.*

Proof. This case is easy as we simply need to determine if there is a play π in the arena such that $\pi \models p_0 \wedge p_1$. Indeed, in such a π, no player has an incentive to deviate. As the conjunction of two parity conditions can be seen as a Streett condition, solving our problem reduces to the emptiness problem of nondeterministic Street automata, which in turn can be solved in PTime [21]. ◻

The following theorem summarizes the complexity of checking the existence of NE in two-player non-zero sum imperfect information games with parity objectives, and provides a matching lower bound for the problems that need exponential time to be solved.

Theorem 3.9. *Given a 2-player imperfect information game $(\mathcal{A}, p_0, p_1)$ with parity objectives p_0 and p_1, the problems of deciding the existence of a* NE*, of a $(1,0)-\mathsf{NE}$, of a $(0,1)-\mathsf{NE}$, and of $(0,0)-\mathsf{NE}$ are* ExpTime-C*, and the problem of deciding the existence a $(1,1)-\mathsf{NE}$ can be solved in* PTime.

Proof. First, we note that there exists a NE in $(\mathcal{A}, p_0, p_1)$ if and only if there exists $(0,0)-\mathsf{NE}$, $(1,0)-\mathsf{NE}$, $(0,1)-\mathsf{NE}$, or $(1,1)-\mathsf{NE}$.

As all four cases can be checked in ExpTime or PTime, we then conclude that our decision problem is solvable in ExpTime. Second, we prove completeness for ExpTime for NE, $(0,0)-$NE, $(1,0)-$NE, and $(0,1)-$NE. Those results are obtained as follows.

For $(1,0)$NE, and $(0,1)-$NE, this is a direct consequence of the ExpTime completeness of determining if Player 0 (or Player 1) has a winning strategy in a game graph of imperfect information with parity objectives are ExpTime-C as established in [28]. Let $\mathcal{G} = (\mathcal{A}, p, \overline{p})$ be such a zero sum game. Deciding the existence of a $(1,0)-$ NE in $\mathcal{G}$ is equivalent to decide the existence of an observation-based winning strategy for Player 0 in that game.

Let us now consider the case $(0,0)-$ NE. To show the hardness, we use the reduction from the problem of deciding the existence of a wining strategy in $\mathcal{G}$ to our problem. This is done with the gadget of Fig. 2 (gadget 1). This game contains a copy of the zero sum game $\mathcal{G}$ and two additional states s_1 and s_2. The game can loop between those two states or Player 1 can decide to move the game to the copy of $\mathcal{G}$. The parity colors of the states s_1 and s_2 are bad for both player and in the copy of $\mathcal{G}$, each play π is either winning for Player 0 (satisfying p) or for Player 1 (satisfying $\overline{p}$) as in $\mathcal{G}$, i.e. the winning condition is zero-sum. So the only possibility for the existence of a $(0,0)-$ NE is the strategy of Player 1 that loops for ever between states s_1 and s_2. But clearly Player 1 has no incentive to deviate only if Player 0 does have a winning observation-based strategy in $\mathcal{G}$. Indeed, if Player 0 does not have such a strategy, then deviating from the loop is profitable for Player 1 as the strategy of Player 0 being fixed and not winning in $\mathcal{G}$, then by lemma 3.2, we know that Player 1 can design an observation-based strategy to win $\overline{p}$ against the strategy of Player 0. On the contrary, if Player 0 has an observation-based strategy to win p in $\mathcal{G}$, then he should play this strategy to retaliate if Player 1 deviates from the loop implying that Player 1 has no incentive to deviate. So in that case, the pair of strategies that loop forever is a $(0,0)-$ NE.

Let us now consider the hardness of the only remaining case: the problem of deciding the existence of a (unconstrained) NE. Our proof relies also on a reduction from the problem of deciding the existence of an observation-based winning strategy in the zero sum case but we need another gadget which is depicted in Fig. 2 (gadget 2). This gadget is composed of a state s, a copy of $\mathcal{G}$ and a copy of a game $\mathcal{H}$ that is zero-sum and for which we know that there is no equilibrium. The example in the proof of Proposition 3.1 and given in Fig. 1 can be used for that. Let us now establish that the game defined by the gadget has a NE if and only if Player 0 has a winning strategy in the game $\mathcal{G}$.

First, we consider the case where Player 0 has a winning strategy in $\mathcal{G}$. In that case, let $\bar{\sigma} = (\sigma_0, \sigma_1)$ be any profile of strategies such that σ_0 decides in s to go to the copy of $\mathcal{G}$ and then behaves there by playing a winning observation-based strategy σ_0' for Player 0 in $\mathcal{G}$ which exists by hypothesis. The outcome of this profile of strategies is π and $\pi \models p$. As σ_0' is winning for Player 0 in $\mathcal{G}$, Player 1 has no way to deviate profitably from σ_1, and so $\bar{\sigma}$ is a NE.

Second, we consider the case where Player 0 has no winning strategy in $\mathcal{G}$. We need to consider the following possible scenarios:

1. Assume that Player 0 decides in s to enter $\mathcal{G}$. Let $\bar{\sigma} = (\sigma_0, \sigma_1)$ be such a profile. Let us show that $\bar{\sigma}$ is not a NE. For that we examine the following two sub-cases:
   a. assume that the outcome π of $\bar{\sigma}$ is such that $\pi \models \overline{p}$, i.e. it is winning for Player 1, and so losing for Player 0. In that case, there is a profitable deviation for Player 0. Indeed, as σ_1 is fixed and is not winning in $\mathcal{H}$ as none of the players have a winning strategy in $\mathcal{H}$ otherwise there would exist a NE in $\mathcal{H}$, then by lemma 3.2, there is an observation-based strategy σ_0' that in $\mathcal{H}$ wins against σ_1. So the profitable deviation for Player 0 is to go to subgame $\mathcal{H}$ and play σ_0'.
   b. assume that the outcome π of $\bar{\sigma}$ is such that $\pi \models p$, i.e. it is winning for Player 0, and so losing for Player 1. In that case, there is a profitable deviation for Player 1. Indeed, as Player 0 does not have an observation-based winning strategy in $\mathcal{G}$, this means that, by lemma 3.2, there is an observation-based strategy σ_1' that in $\mathcal{G}$ wins against σ_0, and so σ_1' is a profitable deviation for Player 1.
2. Assume that Player 0 decides in s to enter $\mathcal{H}$. Let (σ_0, σ_1) be such a profile. We consider the following two sub-cases:
   a. assume that the outcome π of $\bar{\sigma}$ is such that $\pi \models p$. In that case, let us show that there is a profitable deviation for Player 1. As in $\mathcal{H}$, Player 0 does not have a winning strategy then by lemma 3.2, Player 1 has a strategy σ_1' such that (σ_0, σ_1') is such that $\pi \models \overline{p}$ and thus a profitable deviation.
   b. assume that the outcome π of (σ_0, σ_1) is such that $\pi \models \overline{p}$. Symmetrically, Player 0 has a profitable deviation from σ_0.

So, determining the existence of a NE is ExpTime-Hard. □

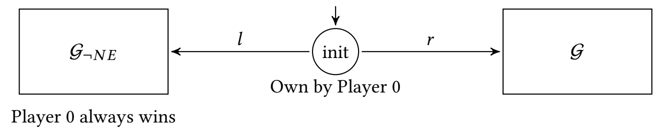

Figure 3. Reduction of $(1,1,0)$-NE existence to NE existence.

3.3 More than 2 players

The existence of constrained NE was already known to be undecidable by a reduction from the distributed synthesis problem [27]:

Theorem 3.10. *[Folklore result [26, 27]] The problem of checking the existence of a* $(1,1,0)$*-NE in a given 3-player imperfect information game with objectives* $(O, O, V^{\omega} \setminus O)$ *is undecidable.*

The unconstrained existence problem was considered in [20]:

Theorem 3.11. *[Undecidability of NE [20]] The problem of checking the existence of a NE in a given* k*-player game with imperfect information and parity objectives is undecidable for* $k \geq 3$.

4 Rational Synthesis

In this section we consider the rational synthesis problem. We start with the two-player case and provide an ExpTime solution together with a matching lower bound. Then, we show that the decidability also holds for the k-player case if only the system (Player 0) is imperfectly informed. Finally, we show the undecidability of the general case.

4.1 The case of 2 players

We provide a two steps solution. The first step, developed in Lemma 4.1 below, shows that RS on 2-player games can be reduced to the same problem over a game arena in which only Player 0 (the system) has imperfect information. Let $\mathcal{G} = \langle \mathcal{A}, O_{i \in \{0,1\}} \rangle$ be a 2-player game where both players have imperfect information. We denote by $\mathcal{G}^{\sim_1/id}$ the game $\mathcal{G}$ where the equivalence relation $\sim_1$ over V, defining the observations for player 1, is replaced by the identity relation over V (i.e. Player 1 has perfect information).

Lemma 4.1. *The two-player game* $\mathcal{G} = \langle \mathcal{A}, O_{i\in\{0,1\}}\rangle$ *admits a solution to the* RS *problem if and only if* $\mathcal{G}^{\sim_1/id}$ *admits a solution to the* RS *problem.*

Proof. ($\Leftarrow$) Let $\widehat{\sigma}_0^{obs} \in \Sigma_0^{obs}$ be a solution to the RS on $\mathcal{G}^{\sim_1/id}$. By definition of RS and 0-fixed NE, the following holds for any strategy in Σ_1, and in particular any strategy in Σ_1^{obs}:

$$\forall \sigma_1^{obs} \in \Sigma_1^{obs} : out(\widehat{\sigma}_0^{obs}, \sigma_1^{obs}) \in O_0 \vee \\ (out(\widehat{\sigma}_0^{obs}, \sigma_1^{obs}) \notin O_1 \wedge \exists \varrho_1 \in \Sigma_1(out(\widehat{\sigma}_0^{obs}, \varrho_1) \in O_1)) \quad (1)$$

Let $\sigma_1^{obs} \in \Sigma_1^{obs}$ such that $out(\widehat{\sigma}_0^{obs}, \sigma_1^{obs}) \notin (O_0 \cup O_1)$. We show how to use the strategy $\varrho_1 \in \Sigma_1$ such that $out(\widehat{\sigma}_0^{obs}, \varrho_1) \in O_1$ (existing by Equation 1) to define an observation based strategy $\varrho_1^{obs} : [V]^*_{\sim_1}[V_1]_{\sim_1} \to Act_1$ for Player 1 such that $out(\widehat{\sigma}_0^{obs}, \varrho_1^{obs}) \in O_1$. Given $\alpha = [v_0]_{\sim_1} \dots [v_n]_{\sim_1} \in [V]^*_{\sim_1}[V_1]_{\sim_1}$, if $v_i \sim_1 out(\widehat{\sigma}_0^{obs}, \varrho_1)[i]$ (the ith states of $out(\widehat{\sigma}_0^{obs}, \varrho_1)$) for all $0 \leq i \leq n$, then $\varrho_1^{obs}(\alpha) = a \in Act_1$ such that the action a satisfies $\delta_1(out(\widehat{\sigma}_0^{obs}, \varrho_1)[n], a) = out(\widehat{\sigma}_0^{obs}, \varrho_1)[n+1]$. Otherwise $\varrho_1^{obs}(\alpha) = a \in Act_1$, where a is the first action in Act_1 (assumed to be arbitrarily ordered).

Using Equation 1 and the previous transformation of strategies ϱ_1 into observation-based strategies ϱ_1^{obs}, we obtain that $\widehat{\sigma}_0^{obs} \in \Sigma_0^{obs}$ is a solution to the RS problem for $\mathcal{G} = \langle \mathcal{A}, O_{i\in\{0,1\}}\rangle$, i.e.

$$\forall \sigma_1^{obs} \in \Sigma_1^{obs} : out(\widehat{\sigma}_0^{obs}, \sigma_1^{obs}) \in O_0 \vee \\ (out(\widehat{\sigma}_0^{obs}, \sigma_1^{obs}) \notin O_1 \wedge \exists \varrho_1^{obs} \in \Sigma_1^{obs}(out(\widehat{\sigma}_0^{obs}, \varrho_1^{obs}) \in O_1)) \quad (2)$$

($\Rightarrow$) Let $\widehat{\sigma}_0^{obs} \in \Sigma_0^{obs}$ be a solution to the RS problem on $\mathcal{G} = \langle \mathcal{A}, O_{i\in\{0,1\}}\rangle$, i.e. an observation based strategy for Player 0 satisfying Equation 2. Since $\Sigma_1^{obs} \subseteq \Sigma_1$, we get

$$\forall \sigma_1^{obs} \in \Sigma_1^{obs} : out(\widehat{\sigma}_0^{obs}, \sigma_1^{obs}) \in O_0 \vee \\ (out(\widehat{\sigma}_0^{obs}, \sigma_1^{obs}) \notin O_1 \wedge \exists \varrho_1 \in \Sigma_1(out(\widehat{\sigma}_0^{obs}, \varrho_1) \in O_1)) \quad (3)$$

By contradiction, suppose there exists $\sigma_1 \in \Sigma_1 \setminus \Sigma_1^{obs}$ such that $out(\widehat{\sigma}_0^{obs}, \sigma_1) \notin O_0$ and such that $out(\widehat{\sigma}_0^{obs}, \sigma_1) \in O_1 \vee \forall \varrho_1 \in \Sigma_1(out(\widehat{\sigma}_0^{obs}, \varrho_1) \notin O_1)$. Using the same technique adopted in the other direction of this proof, we can use $out(\widehat{\sigma}_0^{obs}, \sigma_1)$ to build an observation-based strategy σ_1^{obs} for which Equation 4 holds, contradicting Equation 3.

$$out(\widehat{\sigma}_0^{obs}, \sigma_1^{obs}) \notin O_0 \wedge \\ (out(\widehat{\sigma}_0^{obs}, \sigma_1^{obs}) \in O_1 \vee \forall \varrho_1 \in \Sigma_1(out(\widehat{\sigma}_0^{obs}, \varrho_1) \notin O_1)) \quad (4)$$

□

The second step is based on the results developed in Subsection 4.2, where we provide a general decidability procedure for the RS problem on multiplayer games in which only Player 0 is imperfectly informed. Such a procedure is ExpTime when the number of players is a fixed constant k. Therefore, we get:

Theorem 4.2. *The RS problem on imperfect information* 2*-players parity games is* ExpTime-c *(even if only Player* 0 *is partially informed).*

Proof. The upper bound is a direct corollary of Theorem 4.6 (see below) and Lemma 4.1. To prove the lower-bound, it is known that deciding the winner in a two-player zero-sum parity games where Player 0 is partially informed is ExpTime-c [28]. Take such a game, and modify objective of Player 1 to an objective which is never fulfilled, by taking the parity $p_1(v) = 1$ for all states v for instance. Then, there is a solution to the RS problem in this new (2-player) game iff Player 0 has a winning strategy in the original game. □

4.2 RS on multiplayer games where only Player 0 is imperfectly informed

Our solution to the RS problem on a multiplayer game $\mathcal{G}$ where Player 0 is the only imperfectly informed player goes via a reduction to a 0-sum imperfect information 2-player game, denoted $\mathcal{G}^{\text{RS}}$, such that there is a solution to the RS problem in $\mathcal{G}$ if and only if the protagonist of $\mathcal{G}^{\text{RS}}$, called *Prover*, has an observation-based winning strategy against all strategies of the antagonist, called *Challenger*.

The role of *Prover* is to show the existence of a solution to RS, while *Challenger* tries to disprove it. *Prover* controls the actions of Player 0, and has the same partial information as Player 0 in the original game. Along the game, *Prover* builds a path in the arena of the original game $\mathcal{G}$. The path that has been constructed must be winning for Player 0 in the original game $\mathcal{G}$ or at least one other player is losing and has a profitable deviation along this path. In order to check the latter property, *Prover* has, along the play, the possibility to declare that some players have a profitable deviation, i.e. a winning strategy against all other players. These declarations are stored in a subset $D \subseteq \Omega$, i.e. $i \in D$ iff Player i has been declared along the play to have a profitable deviation. When *Prover* makes such a declaration, *Challenger* also has the possibility to check whether *Prover* has not cheated, i.e. that Player i has indeed a winning deviation. To do so, *Challenger* can also choose actions for all the players $j \neq i$ and Prover chooses Player i's actions. At any point of the game, there might also be a set of players for which Prover has to prove they have a winning strategy, against all other players. These players are stored in a subset $W \subseteq \Omega$.

More precisely, the set of states of $\mathcal{G}^{\text{RS}}$ are triples (v, W, D), where $v \in V$ is a state of $\mathcal{G}$ and $W, D \subseteq \Omega \setminus \{0\}$. A state (v, W, D) is equivalent for Prover to a state (v', W', D') in $\mathcal{G}^{\text{RS}}$ if v and v' are equivalent for Player 0 in $\mathcal{G}$. The actions of *Prover* in $\mathcal{G}^{\text{RS}}$ are the following. In each state $(v \in V_0, W, D)$ *Prover* can play an action $a \in Act_0$ to develop a strategy for Player 0 that is a solution to the RS problem on $\mathcal{G}$. The game then evolves deterministically to the next state in $\mathcal{G}$, letting W and D unchanged. On the other hand, in each state $(v \in V_i, W, D)$ for $i \neq 0$, *Prover* chooses an action which is a function ζ assigning to every triple (v', W', D') such that $v' \in V_i$, either an action $*$ or an action $\langle win, b \in Act_i\rangle$. Intuitively, each mapping $(v', W', D') \mapsto \langle win, b \in Act_i\rangle$ is used by *Prover* to claim that, if the game was actually in position (v', W', D'), then if player i plays b from $v' \in V_i$ in $\mathcal{G}$, then he ends up in a subgame where he has a winning strategy against all the strategies of the other players (since Player 0 has only partial information, unlike the other players, Prover must also provide actions for other possible states). In other words, $(v', W', D') \mapsto \langle win, b \in Act_i\rangle$ is used by *Prover* to claim that v' is a profitable deviation point for Player i. If *Prover* does not want to make any claim concerning the existence of profitable deviations from (v', W', D'), then he will map (v', W', D') to the special symbol $*$.

If the current node of the play is $(v \in V_i, W, D)$ and *Prover* chooses an action ζ such that $\zeta((v, W, D)) = \langle win, b \in Act_i \rangle$, then the play proceeds either into the successor state $(v' = \delta_i(v, b), W \cup \{i\}, D)$ (i is put in W to indicate that Player i is now checked for having a winning strategy against all other players), or in any successor state $(v'' = \delta_i(v, c), W, D \cup \{i\})$ such that $v'' \neq v'$ (i is put in D to remember that i has been declared to have a good deviation). It is up to *Challenger* (who has perfect information of the game) to choose any of these successor states. As a matter of fact, the two-player game we construct is a game where the transition relation is not deterministic over the actions, and it is up to Challenger to resolve non-determinism. This view is equivalent to the deterministic setting of the Preliminary section, modulo putting intermediate states controlled by Challenger, but it eases the notations. We now formally define the game arena $\mathcal{A}^{\mathsf{RS}}$ of $\mathcal{G}^{\mathsf{RS}}$.

Definition 4.3 (Two Player Game Arena $\mathcal{A}^{\mathsf{RS}}$)**.** Given a $k+1$-players game $\mathcal{G} = \langle \mathcal{A}, (O_{i\in\Omega}) \rangle$ over the arena $\mathcal{A} = \langle \Omega, V, v_0, (V_i)_{i\in\Omega}, (Act_i)_{i\in\Omega}, (\delta_i)_{i\in\Omega}, \sim_0 \rangle$, the 2-player game arena $\mathcal{A}^{\mathsf{RS}}$ is given by $\mathcal{A}^{\mathsf{RS}} = (\{Prover, Challenger\}, Q, q_0, Act^{\mathsf{RS}}_{Prover}, E, \sim^{\mathsf{RS}}_{Prover})$, where:

- $Q = \bigcup_{i=0}^{k} Q_i$ where $Q_i = V_i \times 2^{\Omega} \times 2^{\Omega}$, and $q_0 = (v_0, \varnothing, \varnothing)$
- $Act^{\mathsf{RS}}_{Prover} = Act_0 \cup \bigcup_{i=1}^{k} (\{\langle win, b \rangle \mid b \in Act_i\} \cup \{*\})^{Q_i}$.
- $((v, W, D), \zeta, (v', W', D')) \in E \subseteq Q \times Act^{\mathsf{RS}}_{Prover} \times Q$ if and only if one of the following conditions applies:
  1. $v \in V_0$, $v' = \delta_0(v, \zeta \in Act_0)$, $D = D'$, and $W = W'$.
  2. $v \in V_i$, $i \in \overline{W}$, $\zeta(v, W, D) = *$, $\exists a \in Act_i : v' = \delta_i(v, a)$, $W = W'$, and $D = D'$.
  3. $v \in V_i$, $i \in \overline{W} \cap \overline{D}$, $\zeta(v, W, D) = \langle win, a \in Act_i \rangle$, and
     a. either $v' = \delta_i(v, a)$,$W = W \cup \{i\}$, $D' = D$,
     b. or $v' = \delta_i(v, b) \neq \delta_i(v, a)$,$W = W$, $D' = D \cup \{i\}$.
  4. $v \in V_i$, $i \in W \cap \overline{D}$, $\zeta(v, W, D) = \langle win, a \in Act_i \rangle$, and
     a. either $v' = \delta_i(v, a)$, $W' = W$, $D' = D$,
     b. or $v' = \delta_i(v, b) \neq \delta_i(v, a)$, $W = W \setminus \{i\}$, $D' = D \cup \{i\}$.
- $(v, W, D) \sim^{\mathsf{RS}}_{Prover} (v', W', D')$ if and only if $v \sim_0 v'$.

Given a play $\varrho \in (V \times 2^{\Omega} \times 2^{\Omega})^{\omega}$ over the arena $\mathcal{A}^{\mathsf{RS}}_{\mathcal{G}}$, we denote by $\varrho_{|_V}$, $\varrho_{|_W}$ and $\varrho_{|_D}$ its V-projection, W-projection and D-projection, respectively. We can easily prove that each play ϱ over the arena $\mathcal{A}^{\mathsf{RS}}_{\mathcal{G}}$ have some monotonicity property with respect to the D- and W-components. Indeed, the D-components of the states can only increase (Players can only be added), while for W, Players can be added (and possibly removed from W) only once. This guarantees that along each play ϱ the D- and W-components eventually stabilize to sets that we denote by $lim_D(\varrho)$ and $lim_W(\varrho)$, respectively.

We are now ready to define the winning condition for the two-player zero-sum game between Prover and Challenger. Informally, a play ϱ is winning for *Prover* if either it satisfies Player 0's objective in $\mathcal{G}$, or some Player $i \in lim_D(\varrho)$ loses in $\mathcal{G}$ (i.e. Player i has profitable deviation). Moreover, each player in $lim_W(\varrho)$ (i.e. claimed to be winning along $\varrho_{|_V}$) needs to be indeed winning in $\mathcal{G}$. Formally:

$$O^{\mathsf{RS}}_{\mathcal{G}} = \{\varrho \mid \varrho_{|_V} \in O_0 \vee \bigvee_{i=1}^{k} (\varrho_{|_V} \notin O_i \wedge i \in lim_D(\varrho))\} \cap \cap \{\varrho \mid \bigwedge_{i \in lim_W(\varrho)} (\varrho_V \in O_i)\}$$

We let $\mathcal{G}^{\mathsf{RS}}$ be the two-player imperfect information zero-sum game on the arena $\mathcal{A}^{\mathsf{RS}}$ with objective $O^{\mathsf{RS}}_{\mathcal{G}}$ for Prover. The following theorem states the correctness of our reduction.

Theorem 4.4. *Let $\mathcal{G}$ be a $k+1$-player games where player 0 is the only player with imperfect information. $\mathcal{G}$ admits a solution to the* RS *problem if and only if* Prover *has an observation-based winning strategy on $\mathcal{G}^{\mathsf{RS}}$.*

To solve $\mathcal{G}^{\mathsf{RS}}$, we reduce it to a 0-sum perfect information game $\mathcal{G}^*$ where the protagonist has objective $\mathcal{P}^*$ we now define. Let $[Q]$ be the quotient of the set of states Q in $\mathcal{G}^{\mathsf{RS}}$ by $\sim^{RS}_{Prover}$ and let $\mu : Q^{\omega} \mapsto [Q]^{\omega}$ be the morphism defined by mapping each state $s \in Q$ to the equivalence class of s by $\sim^{RS}_{Prover}$. Then:

$$\mathcal{P}^* = \{\varrho \in [Q] \mid \forall \bar{s} \in Q^{\omega} : \mu(s) = \varrho \rightarrow \bar{s} \in O^{RS}_{\mathcal{G}}\}$$

Lemma 4.5. *$\mathcal{P}^*$ is recognizable by a deterministic parity word automaton with a number of states exponential in the number of states of $\mathcal{G}$ and doubly exponential in the number of players k, and a number of priorities that is polynomial in the number of states of $\mathcal{G}$ and exponential in number of players k.*

Based on the previous lemma, we can establish the complexity of RS for a fixed number of players.

Theorem 4.6. *The RS problem for a fixed number k of players where only player 0 is imperfectly informed is* ExpTime*-complete.*

Proof. Let P be the DPW built in Lemma 4.5, recognizing $\mathcal{P}^*$. We apply a knowledge construction to the arena $\mathcal{A}^{\mathsf{RS}}$, where we gather all the possible states of $\mathcal{A}^{\mathsf{RS}}$ in which Prover can be. This is a standard operation in imperfect information games, see for instance [28]. This results in a perfect information game arena $\mathcal{K}$, with which we take a product with P, to obtain a perfect information zero-sum 2-player parity game $\mathcal{G}^*$. The number of states of $\mathcal{G}^*$ is doubly exponential in k and exponential in $|V|$, and its number of priorities is polynomial in $|V|$ and exponential in k. The latter can be solved in time polynomial w.r.t. the number of states and exponential w.r.t. the number of priorities [22], leading to our upper bound. The lower bound comes from the lower bound of Theorem 4.2 (2 players). □

We conclude this section with an undecidability result.

Theorem 4.7. *The RS problem is undecidable for imperfect information games (with at least 4 players).*

Proof. This comes from the undecidability of NE existence in 3 players imperfect information games (Theorem 3.11). Given such a game G with 3 players, we construct a game G' with 4 players. The game arena is the same as G, with the only difference that Player i in G' controls the states of Player $i-1$ in G, for $i = 1, 2, 3$. The objectives of Player $1, 2, 3$ are unchanged, and the objective of Player 0 is false. Hence, Player 0 can never win. By definition of the RS, there is a solution iff there is no NE in G', iff there is no NE in G, and thus we get undecidability. □

5 Complexity Results for an Unfixed Number of Players

In this section, we analyse the complexity of the RS when the number of players is not fixed, and only Player 0 is imperfectly informed. We also show the robustness of the upper bound we obtain in the case of parity objectives, by showing that it does not change for other classical objectives, such as Rabin objectives.

5.1 Upper bound

As shown in the proof of Theorem 4.6, The RS problem for any k-players parity games with n states reduces to a two-player zero-sum perfect information parity game with a number of states which is doubly exponential in k and exponential in n, and a number of priorities exponential in k and polynomial in n. Overall, this gives a 2ExpTime upper bound, because two-player parity games can be solved in time polynomial in the number of states and exponential in the number of priorities [22]. A natural question is whether this upper bound is robust to other classical objectives such as safety, reachability, Büchi, co-Büchi, Rabin, Streett and Emerson-Lei (EL). We do not define all these objectives, and rather refer the reader to [17]. Nevertheless, since we use them later on, we define Rabin and EL objectives.

Let Γ be some finite set. An EL condition over Γ is a Boolean formula φ over variables $\{x_\gamma \mid \gamma \in \Gamma\}$. Such a formula defines a set $[\![\varphi]\!] \subseteq 2^\Gamma$, defined by $\Gamma' \in [\![\varphi]\!]$ iff putting all variables $x_\gamma, \gamma \in \Gamma'$ to true and all variables $x_\gamma, \gamma \notin \Gamma'$ to false, satisfies φ. An infinite sequence $u \in \Gamma^\omega$ satisfies an EL condition φ if $\inf(u) \in [\![\varphi]\!]$. A *Rabin constraint* over Γ is a set of n pairs $R = \{(F_1, I_1), \ldots, (F_n, I_n)\}$ such that $F_i, I_i \subseteq \Gamma$ for all i. A sequence $u \in \Gamma^\omega$ satisfies R iff it satisfies the formula $\bigvee_{(F,I)\in R}(\bigwedge_{\gamma\in F} \neg x_\gamma) \wedge (\bigvee_{\gamma\in I} x_\gamma)$, i.e., u visits finitely often the elements of F_i and infinitely often some element of I_i, for some i. The complexity bound given in the next theorem can be derived from Thm 4.6 for parity objectives. The other objectives can be reduced to the latter, using known translations.

Theorem 5.1. *The RS problem for multiplayer games where only Player 0 is imperfectly informed, and where the players have either safety, reachability, Büchi, coBüchi, parity, Streett, Rabin or EL objectives is in 2ExpTime.*

5.2 Lower bound

By Theorem 4.6, the RS problem for a fixed number of players (and only Player 0 partially informed) is ExpTime-hard for parity objectives, and so we also get this lower bound in the case of an unfixed number of players. This is shown by reducing two-player zero-sum imperfect information parity games (with only Player 0 uninformed) to RS. This reduction is generic and carries over to the classical types of objectives in reachability, safety, Büchi, coBüchi, parity, Rabin, Street and Muller. Since deciding the winner in a two-player zero-sum imperfect information games with such objectives is as well ExpTime-hard, we also get the same lower bound for RS problem and all the classical objectives. According to Theorem 5.1, this leaves an exponential gap between the lower and upper bound. In this section, we reduce this gap for Rabin objectives, by showing an ExpSpace-hardness lower bound. This lower bound is shown by reducing the universality problem for *non-deterministic multi-Rabin word automata*, a model that we now introduce. We believe this class of automata has also its own interest.

5.2.1 Multi-Rabin Word Automata

A *non-deterministic multi-Rabin automaton* (NMRA) is a tuple $A = (\Sigma, Q, Q_0, \Delta, k, R_1, \ldots, R_k)$ where Σ is a finite alphabet, Q a finite set of states with initial states $Q_0 \subseteq Q$, $\Delta \subseteq Q \times \Sigma \times Q$ a transition function, $k \in \mathbb{N}$ and $R_1, \ldots, R_k$ are Rabin constraints over Q.

A *run* of A on a word $w = w_1 w_2 \cdots \in \Sigma^\omega$ is a word $r = r_0 r_1 \cdots \in Q^\omega$ such that $r_0 \in Q_0$ and for all $i \geq 0$, $(r_i, w_{i+1}, r_{i+1}) \in \Delta$. It is *accepting* if for all $i \in \{1, \ldots, k\}$, r satisfies the Rabin constraint R_i. The language $L(A)$ accepted by A is the set of words w which admits an accepting run. The *emptiness* problem asks whether, given an NMRA, $L(A) = \varnothing$ and the *universality problem* asks, whether $L(A) = \Sigma^\omega$.

Proposition 5.2. *The emptiness problem for NMRA is NP-complete.*

In contrast, the universality problem is much harder. NMRA are a particular case of Emerson-Lei automata (modulo polynomial translation). The universality problem for EL-automata has been shown to be in ExpSpace-c [30]. The following result strengthen this ExpSpace-hardness result to NMRA universality. To prove the following result, we reduce the word problem of an ExpSpace Turing machine.

Theorem 5.3. *The universality problem is ExpSpace-c for NMRA.*

5.2.2 Lower bound for the RS problem with Rabin objectives, and only Player 0 imperfectly informed

Theorem 5.4. *The rational synthesis problem for k players all perfectly informed but Player 0 (who is blind), and Rabin objectives for each player, is ExpSpace-hard.*

Proof. We reduce the universality problem of non-deterministic multi-Rabin automata, which is ExpSpace-hard by Theorem 5.3. Let $A = (\Sigma, Q, Q_0, \alpha, R_1, \ldots, R_\alpha)$ be some NMRA. We construct a game G with $\alpha+2$ players $P_0, P_1, \ldots, P_\alpha, P_{\alpha+1}$ and Rabin objectives $R'_0, R'_1, \ldots, R'_\alpha, R'_{\alpha+1}$ such that G has a solution to the RSP iff A is not universal.

The game G has a main part G_A which simulates the automaton A, and α sink states denoted by $s_1, \ldots, s_\alpha$. Player 0 is blind in this game, i.e. $\sim_0 = V \times V$ where V are the states of G. In the part G_A, the actions of Player 0 are symbols in Σ, while the actions of player $P_{\alpha+1}$ are states. Hence, Player 0 chooses some word and Player $P_{\alpha+1}$ chooses a run. When both these players have played one round each, the players 1 to α play in turn. Player $i \in \{1, \ldots, \alpha\}$ can decide to either stay in the part G_A (action C) or to leave it for its sink state s_i (action L) where both Player i and 0 win. The game G is depicted on Fig. 4. Formally, the set of states V of the game is $V_A \cup \{s_1, \ldots, s_k\}$ where $V_A = Q \times \{0, 1, \ldots, \alpha\} \cup Q \times \Sigma$ (Player i controls any state (q, i) while Player $\alpha + 1$ controls the states (q, σ)). From a state $(q, 0)$, Player 0 can choose an action $\sigma \in \Sigma$ and move to state (q, σ). Then, Player $\alpha + 1$ can choose an action $q' \in Q$ such that $(q, \sigma, q') \in \Delta$ and move to state $(q', 1)$. From a state (q, i), Player i has two possible actions, either "stay" or "leave". If it stays, then the game move to state $(q, (i + 1)\%\alpha + 1)$. If it leaves, the game move to state s_i, on which there is only a self loop.

The winning objectives $O_0, O_1, \ldots, O_\alpha, O_{\alpha+1}$ are the following:

- if the game stay forever in G_A then Player 0 loses, otherwise he wins: $O_0 = \{(V_A, \{s_1, \ldots, s_\alpha\})\}$,
- Player P_{run} always loses: $O_{\alpha+1} = \{(\varnothing, \varnothing)\}$,
- Player $i \in \{1, \ldots, \alpha\}$ either satisfy the Rabin constraint R_i in G_A or goes to the sink s_i, i.e. $O_i = R'_i \cup \{(\varnothing, \{s_i\})\}$ where $R'_i = \{(F \times (\{0, 1, \ldots, \alpha\} \cup \Sigma), G \times (\{0, 1, \ldots, \alpha\} \cup \Sigma)) \mid (F, G) \in R_i\}$.

Intuitively, Player 0 wants to eventually leave G_A and for this, some player $i \in \{1, \ldots, \alpha\}$ has to lose if the game stays in G_A, forcing her to deviate to her sink state. Therefore, Player 0 has to choose a sequence of σ-symbols, i.e. a word, such that whatever run is chosen by Player $\alpha + 1$, it falsifies the condition of some player. If the automaton is not universal, then such a word exists. □

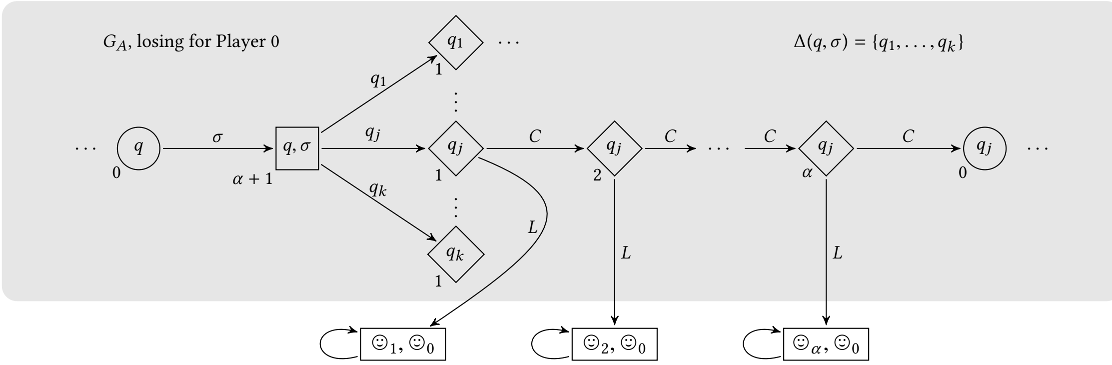

Figure 4. Reduction from NMRA to RS. The game G_A simulates the automaton A.

Remark A similar reduction to that of Thm 5.4 can be used to show that the RS problem for 2 players where only Player 0 is partially informed (and blind) and has an EL-objective is ExpSpace-hard, by reducing the universality problem of EL-automata, which is ExpSpace-c [30]. Player 0 chooses a word and Player 1 a run (which is invisible to Player 0). Player 0's objective is the formula $\neg\varphi$, where φ is the EL-formula of the automaton, and Player 1's objective is always false. Clearly in this game, any strategy profile is a 0-fixed NE since Player 1 always loses. Hence, there exists a solution to the RS problem iff there exists a strategy for Player 0 winning against all strategies of Player 1, iff there exists a word such that all runs are non-accepting, iff the automaton is not universal.

Paper Conclusion We have studied the existence of NE and of solutions to the rational synthesis problem for games played on graphs with non-zero sum ω-regular objectives and with imperfect information. In sharp contrast with the perfect information case, Nash equilibria need not to exist in games with imperfect information, and we have studied their existence problem. While this problem is ExpTime-C for two-player games and parity objectives (even if both players are imperfectly informed), it becomes undecidable for 3 and more players (even if only two players are imperfectly informed). We have obtained similar results for the rational synthesis problem: it is decidable for two players and undecidable in general. We have also identified an other interesting case: when only Player 0 is imperfectly informed, the problem remains decidable in the general case of k players. When this k is fixed, the problem remains ExpTime-C and we have shown a 2ExpTime upper bound when k is not fixed. In the later case, we leave open the exact complexity of the problem as we have only established an ExpSpace lower bound for Rabin objectives.

Acknowledgements We would like to thank Rodica Bozianu Condurache for fruitful discussions on rational synthesis, and Moshe Vardi for pointing us to the EL automata and their complexity.

References

[1] R. Berthon, B. Maubert, A. Murano, S. Rubin, and M. Y. Vardi. 2017. Strategy logic with imperfect information. In *LICS*. IEEE, 1–12.
[2] D. Berwanger. 2007. Admissibility in Infinite Games. In *STACS (LNCS)*, Vol. 4393. 188–199.
[3] D. Berwanger and L. Doyen. 2008. On the Power of Imperfect Information. In *FSTTCS (LIPIcs)*, Vol. 2. 73–82.
[4] Patricia Bouyer. 2017. Games on graphs with a public signal monitoring. *CoRR* abs/1710.07163 (2017). arXiv:1710.07163 http://arxiv.org/abs/1710.07163
[5] R. Brenguier, L. Clemente, P. Hunter, G. A. Pérez, M. Randour, J.-F. Raskin, O. Sankur, and M. Sassolas. 2015. Non-Zero Sum Games for Reactive Synthesis. *CoRR* abs/1512.05568 (2015). http://arxiv.org/abs/1512.05568
[6] R. Brenguier, A. Pauly, J.-F. Raskin, and O. Sankur. 2017. Admissibility in Games with Imperfect Information (Invited Talk). In *CONCUR (LIPIcs)*, Vol. 85. 2:1–2:23.
[7] R. Brenguier, J.-F. Raskin, and O. Sankur. 2015. Assume-Admissible Synthesis. In *CONCUR (LIPIcs)*, Vol. 42. 100–113.
[8] R. Brenguier, J.-F. Raskin, and M. Sassolas. 2014. The complexity of admissibility in Omega-regular games. In *CSL-LICS*. ACM, 23:1–23:10.
[9] T. Brihaye, V. Bruyère, N. Meunier, and J.-F. Raskin. 2015. Weak Subgame Perfect Equilibria and their Application to Quantitative Reachability. In *CSL (LIPIcs)*, Vol. 41. 504–518.
[10] V. Bruyère, S. Le Roux, A. Pauly, and J.-F. Raskin. 2017. On the Existence of Weak Subgame Perfect Equilibria. In *FOSSACS (LNCS)*, Vol. 10203. 145–161.
[11] K. Chatterjee and L. Doyen. 2010. The Complexity of Partial-Observation Parity Games. In *LPAR (LNCS)*, Vol. 6397. 1–14.
[12] K. Chatterjee, L. Doyen, E. Filiot, and J.-F. Raskin. 2014. Doomsday Equilibria for Omega-Regular Games. In *VMCAI (LNCS)*, Vol. 8318. 78–97.
[13] Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. 2013. A survey of partial-observation stochastic parity games. *Formal Methods in System Design* 43, 2 (2013), 268–284.
[14] K. Chatterjee and T. A. Henzinger. 2007. Assume-Guarantee Synthesis. In *TACAS (LNCS)*, Vol. 4424. 261–275.
[15] K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. 2006. Games with secure equilibria. *TCS* 365, 1-2 (2006), 67–82.
[16] Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. 2010. Strategy logic. *Inf. Comput.* 208, 6 (2010), 677–693. https://doi.org/10.1016/j.ic.2009.07.004
[17] R. Condurache, E. Filiot, R. Gentilini, and J.-F. Raskin. 2016. The Complexity of Rational Synthesis. In *ICALP (LIPIcs)*, Vol. 55. 121:1–121:15.
[18] M. Faella. 2009. Admissible Strategies in Infinite Games over Graphs. In *MFCS (LNCS)*, Vol. 5734. 307–318.
[19] Dana Fisman, Orna Kupferman, and Yoad Lustig. 2010. Rational Synthesis. In *TACAS (LNCS)*, Vol. 6015. Springer, 190–204.
[20] J. Gutierrez, G. Perelli, and M. Wooldridge. 2016. Imperfect Information in Reactive Modules Games. In *Principles of Knowledge Representation and Reasoning (KR)*. AAAI Press, 390–400. http://www.aaai.org/Library/KR/kr16contents.php
[21] M. Henzinger and J. A. Telle. 1996. Faster Algorithms for the Nonemptiness of Streett Automata and for Communication Protocol Pruning. In *SWAT*. 16–27.
[22] M. Jurdzinski. 2011. Algorithms for solving parity games. In *Lectures in Game Theory for Computer Scientists*. 74–98.
[23] O. Kupferman, G. Perelli, and M. Y. Vardi. 2014. Synthesis with Rational Environments. In *EUMAS*. 219–235.
[24] O. Kupferman and M.Y. Vardi. 1997. Synthesis with incomplete information. In *2nd International Conference on Temporal Logic*. Manchester, 91–106.
[25] F. Mogavero, A. Murano, and M. Y. Vardi. 2010. Reasoning About Strategies. In *FSTTCS (LIPIcs)*, Vol. 8. 133–144.
[26] G.L. Peterson and J.H. Reif. 1979. Multiple person alternation. In *FOCS*. 348–363.
[27] A. Pnueli and R. Rosner. 1990. Distributed Reactive Systems are Hard to Synthesize. In *FOCS*. 746–757.
[28] J.-F. Raskin, K. Chatterjee, L. Doyen, and T. A. Henzinger. 2007. Algorithms for Omega-Regular Games with Imperfect Information. *LMCS* 3, 3 (2007).
[29] John H. Reif. 1984. The Complexity of Two-Player Games of Incomplete Information. *J. Comput. Syst. Sci.* 29, 2 (1984), 274–301.
[30] S. Safra and M.Y. Vardi. 1989. On ω automata and temporal logic. In *21th Annual Symp. on Theory of Computing*. 127–137.
[31] S. Safra and M. Y. Vardi. 1989. On omega-Automata and Temporal Logic. In *STOC*. ACM, 127–137.
[32] Wolfgang Thomas. 1995. On the Synthesis of Strategies in Infinite Games. In *STACS*. 1–13. https://doi.org/10.1007/3-540-59042-0_57
[33] M. Ummels. 2006. Rational Behaviour and Strategy Construction in Infinite Multiplayer Games. In *FSTTCS (LNCS)*, Vol. 4337. 212–223.
[34] M. Ummels. 2008. The Complexity of Nash Equilibria in Infinite Multiplayer Games. In *FOSSACS (LNCS)*, Vol. 4962. 20–34.