

Admissibility in Games with Imperfect Information*

Romain Brenguier¹, Arno Pauly², Jean-François Raskin³, and Ocan Sankur⁴

- 1 Oxford University, UK
- 2 Université Libre de Bruxelles, U.L.B., Belgium
- 3 Université Libre de Bruxelles, U.L.B., Belgium
- 4 CNRS – IRISA Rennes, France

Abstract

In this invited paper, we study the concept of *admissible strategies* for two player win/lose infinite sequential games with *imperfect information*. We show that in stark contrast with the perfect information variant, admissible strategies are *only* guaranteed to exist when players have objectives that are *closed sets*. As a consequence, we also study decision problems related to the existence of admissible strategies for regular games as well as finite duration games.

1998 ACM Subject Classification F.1.1 Models of Computation, B.1.2 [Automatic Synthesis] Formal Models

Keywords and phrases Admissibility, non-zero sum games, reactive synthesis, imperfect information

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2017.2

Category Invited Talk

1 Introduction

Two-player zero-sum perfect information games played on finite (directed) graphs are the canonical model to formalize the reactive synthesis problem [24, 1]. Unfortunately, this mathematical model is often a too coarse abstraction of reality. First, realistic systems are usually made up of several components, each of them with its own objective. These objectives are not necessarily antagonistic. Hence, the setting of non-zero sum graph games needs to be investigated, see [9] and additional references therein. Second, in systems made of several components, each component has usually a partial view on the entire system. Hence it is natural to study games with imperfect information [25, 13]. In this paper, we investigate the notion of *admissible strategies* for infinite duration non-zero sum games played on graphs in which players have *imperfect information*.

The *objective* W of a player in such a game is a set of infinite paths, those that model behaviors of the system that are regarded as satisfactory by the player. During the course of the game, players move a token from vertices to adjacent vertices. Each player applies one *strategy* which is a function that maps histories of plays that ends in a vertex owned by the player to adjacent vertices. The strategy instructs the player where to move the token

* Work partially supported by the ERC Starting grant 279499 (inVEST) and the ARC project “Non-Zero Sum Game Graphs: Applications to Reactive Synthesis and Beyond” (Fédération Wallonie-Bruxelles). J.-F. Raskin is Professeur Francqui de Recherche.



according to the prefix of play constructed so far. For a player with objective W , a strategy σ is said to be *dominated* by a strategy σ' if σ' does as well as σ with respect to W against all the strategies of the other players and strictly better for some of them. A strategy σ is *admissible* for a player if it is not dominated by any other of his strategies. Clearly, playing a strategy which is not admissible is sub-optimal, and as a consequence a rational player should favor admissible strategies.

While admissibility is a classical concept for games in normal form, i.e. matrix games, it was first studied in this context of infinite duration games by Berwanger in [6]. In that paper, it was shown that admissibility, i.e. the avoidance of dominated strategies, is well-behaved in infinite duration n -player non-zero sum turn-based games with *perfect information* and *Boolean outcomes* (two possible payoffs: win or lose). This framework encompasses games with omega-regular objectives. The main contributions of Berwanger were to show that:

- (i) In all n -player game structures, for all objectives, players have admissible strategies. Berwanger even shows the existence of strategies that survive the iterated elimination of dominated strategies.
- (ii) Every strategy that is dominated by a strategy is dominated by an admissible strategy.
- (iii) For finite game structures, the set of admissible strategies forms a regular set.

Contributions

In this paper, we study the notion of admissible strategies in the more general setting of infinite duration games with *imperfect* information. We obtain results that are in stark contrast with the results obtained by Berwanger:

- (i) In a 2-player game with imperfect information, players may have no admissible strategies at all, even for reachability objectives (Example 9). In particular, there are strategies that are dominated but not dominated by any admissible strategy.
- (ii) Admissible strategies always exist for the class of closed objectives, i.e. safety winning conditions, and their existence is only guaranteed for this class of objectives (Theorem 10).
- (iii) The set of admissible strategies of a player depends on the informedness of the other players: this set is the largest when the other players are perfectly informed (Theorem 17).
- (iv) For the special case of finite duration games, we know by point (ii) that admissible strategies always exist, but we show that simple queries about this set are NP-complete (Theorem 28).
- (v) For infinite duration 2-player games, we characterize precisely the set of admissible strategies when the second player is fully informed and we show how to decide the existence of admissible strategies (Theorem 20). When the second player is not fully informed but more informed than the first player, then we show how to decide the existence of admissible strategies for the first player by a reduction to the model-checking problem of Strategy Logic with (hierarchical) imperfect information (Theorem 32). In the general case, we show how to decide, given a regular observation-based strategy if this strategy is admissible or not (Theorem 35). For this last case, we left open the decidability status of the non-emptiness problem for the set of admissible strategies.

Related works

The iterated elimination of dominated strategies formalizes a strong notion of rationality [2]. In [18], Faella considers games played on finite graphs and focuses on the vertices from which one designated player cannot force a win. He compares several criteria for establishing what is the preferable behavior of this player from those vertices, eventually settling on the notion of admissible strategy. In [11], starting from the notion of admissible strategy, we have defined a novel rule for the compositional synthesis of reactive systems, applicable to systems made of n components which have each their own objective. We have shown that this synthesis rule leads to solutions which are robust and resilient. In [12], we have shown how to solve the central algorithmic questions related to admissibility for omega-regular objectives. In all those works, players are assumed to have perfect information.

In [10], we have studied the notion of admissible strategies for non-zero quantitative games with perfect information. Similarly to games with imperfect information, in the quantitative case, admissible strategies may not exist. Nevertheless, in all games where the *adversarial* and the *cooperative* values can be realized, i.e. games in which there exist worst-case optimal strategies and cooperative optimal profiles, players are guaranteed to have admissible strategies.

In [4], we have studied the notion of admissibility for concurrent games. Concurrent games [17] are n -player games played on graphs in which players take moves concurrently: at each round, all players choose an action at the same time and without informing the other players. This set of actions determines the next vertex. Because players choose actions without being informed of the concurrent choices of the other players, concurrent games are a special case of imperfect information games. In concurrent games, in contrast with general games with imperfect information studied here, admissible strategies are guaranteed to exist. While in this paper, we limit our study to deterministic strategies, in [4], we have also considered the more general case of randomized strategies and the notion of almost-sure winning.

In [16], Damm and Finkbeiner use the notion of *dominant strategy* to provide a compositional semi-algorithm for the (undecidable) distributed synthesis problem. So while we use the notion of admissible strategy, they use a notion of dominant strategy. The notion of dominant strategy is *strictly stronger*: every dominant strategy is admissible but an admissible strategy is not necessary dominant.

In normal form games it is trivial to decide whether a strategy dominates another one, or is dominated by an arbitrary strategy. Iterated elimination of dominated strategies does yield a relevant computational problem here, though. Depending on the type of dominance used (weak dominance, dominance, strict dominance) and the number of distinct payoffs, these problems were shown to be complete for L, NL, P or NP in [23, 15, 20, 8]. See [23] for an overview.

Structure of the paper

In Section 3, we discuss the existence of admissible strategies in two player win/lose infinite sequential games with imperfect information. In Section 4, we discuss the impact of informedness of the opponent on the set of admissible strategies. In Section 5, we characterize the set of admissible strategies when the opponent is perfectly informed. In Section 6, we consider the special case of games on finite trees. In Section 7, we investigate decision problems related to dominance in regular games of infinite duration.

2 Definitions

Games of imperfect information

Given some finite word w , let $\ell(w)$ denote its last element and $|w|$ its length. For an finite or infinite word w , if $|w| \geq n$, let $w_{\leq n}$ denote the prefix of w of length n , and for i , $0 \leq i < |w|$, let $w(i)$ denote the letter in position i in w . Given an infinite word w , let $\text{inf}(w)$ denote the letters in w that appear infinitely often along w .

► **Definition 1.** A two player win/lose infinite sequential game with imperfect information (short: infinite game with imperfect information) between Player 1 (sometimes called the protagonist) and Player 2 (sometimes called the opponent) $G = (d, W, \cong_1, \cong_2)$ is given by the following:

1. a function $d : \{0, 1\}^* \rightarrow \{1, 2\}$ assigning a player to each vertex in the full infinite binary tree.
2. a winning condition $W \subseteq \{0, 1\}^\omega$ for the first player.
3. equivalence relations \cong_i on $\{0, 1\}^*$ for $i \in \{1, 2\}$ satisfying the following properties:
 - a. $v \cong_i u \Rightarrow d(v) = d(u)$ (knowledge of who is playing is guaranteed).
 - b. If $v \not\cong_i u$, then $vb \not\cong_i ub'$ for any $b, b' \in \{0, 1\}$ (perfect recall).
 - c. $v \cong_i u \Rightarrow |v| = |u|$ (ability to count moves).

Player i is *perfectly informed* if \cong_i is the identity noted id_i . We only specify an objective for Player 1 as we will characterize the admissible strategies of this player, and to do this, the objective of Player 2 is irrelevant.

► **Definition 2.** Elements of $\{0, 1\}^*$ (resp. $\{0, 1\}^\omega$) are called *finite plays* or histories (resp. *plays*). A strategy of Player i is a function $\sigma : d^{-1}(i) \rightarrow \{0, 1\}$. It is *observation-based*, if $v \cong_i u$ implies $\sigma(v) = \sigma(u)$. A play $p \in \{0, 1\}^\omega$ is *compatible* with a strategy σ of Player i , if for all $n \in \mathbb{N}$, if $d(p_{\leq n}) = i$, then $\sigma(p_{\leq n}) = p(n+1)$. We write Σ_1 ($\Sigma_1^{\cong_1}$) and Σ_2 ($\Sigma_2^{\cong_2}$) for the set of all (observation-based) strategies for Player 1 and Player 2 respectively.

► **Definition 3.** Given strategies $\sigma \in \Sigma_1^{\cong_1}$, $\tau \in \Sigma_2^{\cong_2}$, we let $\text{Out}(\sigma, \tau)$ be the unique $p \in \{0, 1\}^\omega$ compatible with both σ and τ . A strategy $\sigma \in \Sigma_1^{\cong_1}$ of Player 1 is an *adversarially winning strategy* in G if $\forall \tau \in \Sigma_2^{\cong_2} : \text{Out}(\sigma, \tau) \in W$, we denote this by $G, \sigma \models W$. A strategy $\sigma \in \Sigma_1^{\cong_1}$ is *cooperatively winning* if $\exists \tau \in \Sigma_2^{\cong_2} : \text{Out}(\sigma, \tau) \in W$.

► **Definition 4.** A *regular game* $R(v_0) = (V, E, V_1, V_2, \mathcal{O}^1, \mathcal{O}^2, \Omega)$ is given by

1. a finite directed graph (V, E) with a designated starting vertex $v_0 \in V$ where all vertices have out-degree 1 or 2, moreover, the successors of each vertex $v \in V$ are ordered, and denoted by $s(v, 0)$ and $s(v, 1)$ (if v has out-degree 1, they are equal).
2. a control partition $V = V_1 \cup V_2$, V_i being the vertices controlled by Player i , v_0 is the starting vertex.
3. two observation partitions $V = \bigcup_{j \in I^i} \mathcal{O}_j^i$ for $i \in \{1, 2\}$ which are both refinements of the control partition. We denote by $\mathcal{O}_{\text{turn}}^i$ the set of observations of Player i that contain vertices that belong to Player i .
4. an ω -regular winning condition $\Omega \subseteq V^\omega$ defined by a parity condition $pr : V \rightarrow \{1, 2, \dots, d\}$ and such that $\rho \in \Omega$ if and only if $\min\{k \mid v \in \text{inf}(\rho) \wedge pr(v) = k\}$ is even.

► **Remark.** For convenience, we will give examples of regular games violating the condition that the out-degree is 1 or 2. In this case, we understand implicitly that the following

modifications are employed to satisfy this condition: If 2^d is an upper bound for the out-degree, then any non-sink vertex is replaced by a full binary tree of height d , with the partitions and winning condition being extended accordingly.

► **Definition 5.** From a regular game $R(v_0)$, we obtain an infinite sequential game $R^U(v_0) = (d, W, \cong_0, \cong_1)$ by unfolding as follows:

1. Let $\iota : \{0, 1\}^* \rightarrow V^*$ be defined via $\iota(\langle \rangle) = v_0$ and $\iota(wb) = \iota(w) \cdot s(\ell(\iota(w)), b)$. Let $d(w) = i$ if $\ell(\iota(w)) \in V_i$.
2. Let $\widehat{\iota} : \{0, 1\}^\omega \rightarrow V^\omega$ be the limit induced by ι . Let $W = \widehat{\iota}^{-1}(\Omega)$.
3. Let \cong_i be the smallest equivalence relation that satisfies the following constraints:
 - a. If $w \cong_i w'$ and $d(w) = 3 - i$ and $\exists j \ell(\iota(wb)) \in \mathcal{O}_j^i \wedge \ell(\iota(w'b')) \in \mathcal{O}_j^i$, then $wb \cong_i w'b'$.
 - b. If $w \cong_i w'$ and $d(w) = i$ and $\exists j \ell(\iota(wb)) \in \mathcal{O}_j^i \wedge \ell(\iota(w'b)) \in \mathcal{O}_j^i$, then $wb \cong_i w'b$.

We also consider the special case in which Player 2 is *perfectly informed*. In that case, we do not take into account the partition \mathcal{O}^2 and $R^U(v_0) = (d, W, \cong_1, \cong_2)$ is defined as above but $\cong_2 = id_2$ is the identity relation.

To ease presentation, we make a limited use of *computation tree logic (CTL)* to state simple facts on strategies in regular games. We refer the interested reader for instance to [3] for formal definitions. Recall that for subsets $A, B \subseteq V$, **FA** (resp. **GA**) is the set of plays that visit A (resp. that always stay in A). We use the following CTL formulas. $G, \sigma \models \mathbf{AFA}$ means that all plays compatible with σ eventually reach A . Furthermore $G, \sigma \models \mathbf{EFA}$ means that there exists $\tau \in \Sigma_{3-i}^{\cong_{3-i}}$ such that $\iota(\text{Out}(\sigma, \tau))$ eventually reaches A .

For a subset $A \subseteq V$ of vertices, $G(A)$ denotes the game that starts with nondeterministically moving to any vertex of A , while both players receive respective observations of the chosen vertex. Strategy σ is *winning from* $v \in V$ (resp. $A \subseteq V$) if it is winning in $G(\{v\})$ (resp. $G(A)$).


► **Definition 6.** We consider a regular game $R(v_0)$. Given some finite sequence or infinite sequence $w \in \{0, 1\}^{\leq \omega}$, the *derived observation sequence* for Player i is $\mathcal{O}^i(w) = \mathcal{O}_{j_0}^i \mathcal{O}_{j_1}^i \dots$ where $\iota(w)(k) \in \mathcal{O}_{j_k}^i$ for all $k \leq |w|$. An observation-based strategy for Player i in $R(v_0)$ is equivalent to a function $\sigma : (\mathcal{O}^i)^* \cdot \mathcal{O}_{\text{turn}}^i \rightarrow \{0, 1\}$. Now a strategy σ of Player i is *regular*, if there exists a finite deterministic automaton $\mathcal{A} = (Q^A, q_0^A, \mathcal{O}^i, \delta^A, F^A)$, where Q^A is a finite set of states, q_0^A is the initial state, \mathcal{O}^i is the alphabet, $\delta : Q^A \times \mathcal{O}^i \rightarrow Q^A$, and $F^A \subseteq Q^A$ is the set of accepting states, and the observation-based strategy σ is encoded by the language of the automaton as follows: when reading words $\mathcal{O}_{j_0}^i \mathcal{O}_{j_1}^i \dots$, and whenever $d(w) = i$, then the automaton accepts the observation sequence $\mathcal{O}^i(w)$ induced by w iff $\sigma(\mathcal{O}^i(w)) = 0$.

► **Definition 7.** Given a regular game $R(v_0)$, the *knowledge* of Player i playing the observation-based strategy σ after a sequence of observations $\rho = o_1 o_2 \dots o_n \in (\mathcal{O}^i)^*$, denoted $K(\sigma, \rho)$, is the set of vertices v such that there exists $w \in \{0, 1\}^*$:

1. $\mathcal{O}^i(\iota(w)) = \rho$
2. for all j , $1 \leq j < n$, if $o_j \in \mathcal{O}_{\text{turn}}^i$, then $\sigma(\rho_{\leq j}) = w_j$
3. v is equal to $\ell(\iota(w))$.

Dominance and admissibility

► **Definition 8.** For $\sigma, \sigma' \in \Sigma_1^{\cong_1}$, we say that σ' *weakly dominates* σ and write $\sigma' \succeq \sigma$ if $\forall \tau \in \Sigma_2^{\cong_2}$, $\text{Out}(\sigma, \tau) \in W \Rightarrow \text{Out}(\sigma', \tau) \in W$. If in addition $\exists \tau \in \Sigma_2^{\cong_2}$, $\text{Out}(\sigma, \tau) \notin W$, and $\text{Out}(\sigma', \tau) \in W$, then σ' *dominates* σ and write $\sigma' \succ \sigma$. We say that a strategy $\sigma \in \Sigma_1^{\cong_1}$ is *admissible* if there is no $\sigma' \in \Sigma_1^{\cong_1}$ such that $\sigma' \succ \sigma$. Note that if $\sigma \preceq \sigma'$ but $\sigma \not\succeq \sigma'$, then $\sigma' \preceq \sigma$, justifying the notation.



■ **Figure 1** Player 1 can not differentiate between s_1 and s_2 , and not between t_1 and t_2 . In this example, no strategy is admissible for Player 1.

3 Existence of Admissible Strategies

A starting observation for our investigation is that the existence of admissible strategies in infinite sequential games with imperfect information is not guaranteed, even for regular games and reachability objectives. This is in stark contrast to the case where both players are perfectly informed [6].

► **Example 9.** We exhibit a regular game with a reachability objective (to reach the state marked w) lacking observation-based admissible strategies for Player 1. The graph is depicted in Figure 1. Player 1 controls circle vertices, Player 2 controls box vertices. Player 1 wins all plays that eventually enter the vertex w . Player 1 has imperfect information: he can not differentiate between s_1 and s_2 , and not between t_1 and t_2 . As a consequence, the observation-based strategies available to Player 1 are essentially the strategies σ_n , one for each $n \in \mathbb{N}$, which play the action "0" for n consecutive steps followed by the action "1", and σ_∞ , which always plays the action "0". Then it is easy to show that $\sigma_\infty \prec \sigma_0 \prec \sigma_1 \dots$, hence there is no observation-based admissible strategy for Player 1 in this game.

We can characterize exactly the properties of the winning condition W that ensures the existence of admissible strategies, provided we use colours to define the winning condition. Let C be a finite set of colours, and $c : \{0, 1\}^* \rightarrow C$ be a colouring function. Let $\widehat{c} : \{0, 1\}^\omega \rightarrow C^\omega$ be defined via $\widehat{c}(p)(n) = c(p_{\leq n})$. We say that $W \subseteq \{0, 1\}^\omega$ is induced by $W_C \subseteq C^\omega$ if there exists a colouring function c such that $W = \widehat{c}^{-1}(W_C)$. Note that functions of the form \widehat{c} are exactly the 1-Lipschitz functions from $\{0, 1\}^\omega$ to C^ω .

- **Theorem 10.** *The following are equivalent for $W_C \subseteq C^\omega$:*
1. W_C is closed (i.e. corresponds to a safety condition).
 2. All imperfect information games with winning set induced by W_C have admissible strategies.

Our proof of this theorem heavily relies on topological arguments. Note that we can conceive of Σ_i^{\cong} as a metric space by setting $d(\sigma, \sigma') = 2^{-k}$ for the least $k \in \mathbb{N}$ such that $\exists w \in \{0, 1\}^k \cap d^{-1}(i)$ with $\sigma(w) \neq \sigma'(w)$. With this metric, Σ_i^{\cong} is a compact metric space, hence every sequence of strategies has an accumulation point. Moreover, the map $(\sigma, \tau) \mapsto \text{Out}(\sigma, \tau) : \Sigma_1^{\cong} \times \Sigma_2^{\cong} \rightarrow \{0, 1\}^\omega$ is continuous. We will write $B(\sigma, k) := \{\sigma' \in \Sigma_i^{\cong} \mid d(\sigma, \sigma') < 2^{-k}\}$ and utilize the fact that there are only countably many distinct sets $B(\sigma, k)$.

► **Lemma 11.** *Let W be closed. Let $(\sigma_n)_{n \in \mathbb{N}}$ be a sequence in Σ_1^{\cong} with $\sigma_n \preceq \sigma_{n+1}$ for all n , and let σ be an accumulation point of $(\sigma_n)_{n \in \mathbb{N}}$. Then $\forall n \in \mathbb{N} : \sigma_n \preceq \sigma$.*

Proof. Assume that for some $i \in \mathbb{N}$ we have $\sigma_i \not\preceq \sigma$. Then there is some $\tau \in \Sigma_2^{\cong 2}$ such that $\text{Out}(\sigma, \tau) \notin W$ and $\text{Out}(\sigma_n, \tau) \in W$. As W is closed, $\text{Out}(\sigma, \tau) \notin W$ depends only on some finite prefix of σ , i.e. whenever $d(\sigma, \sigma') < 2^{-k}$ for a suitably large k , then also $\text{Out}(\sigma', \tau) \notin W$. As σ is an accumulation point of $(\sigma_n)_{n \in \mathbb{N}}$, there is some $j \geq i$ with $d(\sigma, \sigma_j) < 2^{-k}$.

By assumption and transitivity of \preceq , we find that $\sigma_i \preceq \sigma_j$. But $\text{Out}(\sigma_i, \tau) \in W$, whereas $\text{Out}(\sigma_j, \tau) \notin W$ – contradiction. \blacktriangleleft

► **Corollary 12.** *Let $(\sigma_n)_{n \in \mathbb{N}}$ be a sequence of protagonist strategies with $\sigma_n \prec \sigma_{n+1}$ for all n , and let σ be an accumulation point of $(\sigma_n)_{n \in \mathbb{N}}$. Then $\forall n \in \mathbb{N} \sigma_n \prec \sigma$.*

► **Lemma 13.** *Let $(\sigma_\beta)_{\beta < \alpha}$ be an ordinal-indexed sequence of protagonist strategies such that $\sigma_\beta \prec \sigma_\gamma$ for any $\beta < \gamma < \alpha$. Then α is countable.*

Proof. By assumption, $p_\gamma \not\preceq p_\beta$ for any $\gamma > \beta$. This is witnessed by some $\tau \in \Sigma_2^{\cong 2}$ such that $\text{Out}(\sigma_\gamma, \tau) \in W$ and $\text{Out}(\sigma_\beta, \tau) \notin W$. As W is closed, there is some $k \in \mathbb{N}$ such that $\text{Out}(\sigma_\beta, \tau') \notin W$ for all $\tau' \in B_{\beta, \gamma} := B(\tau, k)$.

As there are only countably many distinct sets of the form $B(\tau, k)$, if α were uncountable, there would have to be $\beta < \gamma < \beta' < \gamma'$ with $B_{\beta, \gamma} = B_{\beta', \gamma'}$. By construction we have that $\exists \tau \in B_{\beta, \gamma} : \text{Out}(\sigma_\gamma, \tau) \in W$ and $\forall \tau \in B_{\beta', \gamma'} : \text{Out}(\sigma_{\beta'}, \tau) \notin W$. But this contradicts that $\sigma_\gamma \preceq \sigma_{\beta'}$, hence α has to be countable. \blacktriangleleft

► **Lemma 14.** *If $W_C \subseteq C^\omega$ is not closed, then there is an imperfect information game with winning set induced by W_C without an admissible strategy.*

Proof. This is a generalization of the Example 9. As W_C is not closed, we can pick a sequence $(p_n)_{n \in \mathbb{N}}$ and path p_∞ in C^ω such that $d(p_n, p_\infty) < 2^{-2n}$, $p_n \in W_C$ and $p_\infty \notin W_C$.

The two players take turns, i.e. $d(w) = 1$ iff $|w|$ is even. It is unknown to Player 1 which moves Player 2 has taken, i.e. \cong_1 is the coarsest equivalence relation satisfying the criteria. The colouring ensures that any sequence of the form $0^{2j+1}1w1q$ is mapped to p_{j+l+1} where $|w| = 2l$ and w contains no 1 in an odd position, while any other sequence gets mapped to p_∞ . The sequence $(p_n)_{n \in \mathbb{N}}$ and p_∞ were chosen in a way to make this possible.

Intuitively, this means that Player 1 wins iff he plays a 1 for the first time at some turn after Player 2 already has played a 1. As Player 1 cannot observe the moves of Player 2, we find that as in Example 9 the observation based strategies available to Player 1 are essentially σ_n , which plays 0 the first n times it is Player 1's turn and then 1, and σ_∞ , which always plays 1. Then again $\sigma_\infty \prec \sigma_0 \prec \sigma_1 \dots$, hence there is no admissible strategy. \blacktriangleleft

Proof of Theorem 10. If W_C is closed, then any W induced by it is closed, too. We can start with some strategy p_0 , and if it is dominated, move to some strategy p_1 it is dominated by, etc., and create a strictly increasing sequence $(p_n)_{n \in \mathbb{N}}$, unless we hit a non-dominated strategy. Then pick some accumulation point p_ω (as the space of strategies is compact). If p_ω is dominated, pick a witness dominating strategy $p_{\omega+1}$, etc. If we would never reach a non-dominated strategy, then this would by Lemma 11 create an increasing Ω_1 -sequence (where Ω_1 is the first uncountable ordinal), but this would contradict Lemma 13.

The second implication is the statement of Lemma 14. \blacktriangleleft

Some further remarks

► **Observation 15.** *In infinite trees, there are strictly increasing sequences of any countable length.*

Proof. By induction. If T is a tree with a strictly increasing sequence $(p_\beta)_{\beta < \alpha}$ of length α , let opponent choose between playing in T or letting protagonist choose between losing and winning. Starting with protagonist choosing to lose outside of T , and playing p_0 inside T , he can improve inside T for α steps and then decide to win if outside of T , yielding an improvement sequence of length $\alpha + 1$.

If for each $i \in \mathbb{N}$, T_i is a tree with an improvement sequence $(p_\beta^i)_{\beta < \alpha_i}$, then by letting opponent choose in which tree to play (or non at all), we can obtain an improvement sequence of length $\sup_{i \in \mathbb{N}} \alpha_i$. ◀

► **Observation 16.** *The construction in Lemma 14 is sufficiently uniform that the preceding argument also shows that for some ω -regular winning condition W all imperfect information games played on finite graphs have admissible strategies iff W is a safety condition.*

4 The Impact of the Informedness of the Opponent

The following theorem characterizes how the information available to the opponent (Player 2) impacts dominance between strategies of Player 1. When the second player is less informed then there are more dominated strategies for Player 1 and so less admissible strategies for Player 1.

We consider two infinite sequential games with imperfect information that only differ in the equivalence relation of Player 2, with one having \cong_2 and the other having \cong'_2 . The set of (observation-based) strategies for Player 1 in both games is the same. We assume that \cong_2 is coarser than \cong'_2 , i.e. $w \cong'_2 w' \Rightarrow w \cong_2 w'$. Then $\Sigma_2^{\cong'_2} \supseteq \Sigma_2^{\cong_2}$. We write $\sigma \prec \sigma'$ ($\sigma \prec' \sigma'$) if σ is dominated by σ' in the game built with \cong_2 (with \cong'_2).

► **Theorem 17.** *$G = (d, W, \cong_1, \cong_2)$ and $G' = (d, W, \cong_1, \cong'_2)$, where \cong_2 is coarser than \cong'_2 , let \prec be the dominance ordering in G and \prec' in G' , for all observation-based strategies σ and σ' of Player 1 if $\sigma \prec' \sigma'$, then $\sigma \prec \sigma'$.*

Proof. By the definition of \prec and \prec' , we see that $\Sigma_2^{\cong'_2} \supseteq \Sigma_2^{\cong_2}$ implies that if $\sigma \prec' \sigma'$, then $\sigma \preceq \sigma'$. To extend this to \prec and \prec' , we need show that if there is some strategy $\tau' \in \Sigma_2^{\cong'_2}$ such that $\text{Out}(\sigma, \tau') \notin W$ but $\text{Out}(\sigma', \tau') \in W$, then there already is some $\tau \in \Sigma_2^{\cong_2}$ with this property.

Consider only word-strategies played by Player 2, i.e. strategies not depending on the actions of the opponent at all. If there is a word-strategy obtaining different outcomes against σ and σ' , which by assumption of $\sigma \preceq' \sigma'$ can only mean that the outcome of the word strategy against σ is not in W while the outcome of the word strategy against σ' is in W , we would be done – as all word-strategies are in $\Sigma_2^{\cong_2}$ by the requirement that players can count the number of moves played so far.

If τ is played against σ , it will act like some particular word strategy (denoted by τ^σ), likewise τ acts like the word strategy $\tau^{\sigma'}$ when played against σ' . If $\tau^\sigma = \tau^{\sigma'}$, then this is the witness we are looking for. Otherwise, $\tau^\sigma = h\rho$ and $\tau^{\sigma'} = h\rho'$ for some maximal common prefix h , and after playing h against either σ or σ' , Player 2 can distinguish w.r.t. \cong'_2 which of the two strategies he is playing against. Thus, the following strategy τ' is also in $\Sigma_2^{\cong'_2}$: Play h . If faced against σ , play ρ' , if faced against σ' , play ρ . Now τ' behaves like $\tau^{\sigma'}$ if played against σ , and like τ^σ if played against σ' . If all word strategies would yield the same outcomes against σ and σ' , then τ' wins against σ' and loses against σ , a contradiction to $\sigma \prec' \sigma'$. Thus, there has to be a word strategy yielding different outcomes, which as explained above, provides the witness for $\sigma \prec \sigma'$. ◀

► **Corollary 18.** *If a strategy is not admissible w.r.t. perfect information strategies of the opponent, then it is not admissible w.r.t. observation-based strategies of the opponent.*

To conclude this section, we shall demonstrate that Player 2 having imperfect information can indeed lead to strictly less admissible strategies for Player 1 compared to a perfectly informed Player 2:

► **Example 19.** Consider a game where Player 1 moves first and plays **a** or **b**. Then Player 2 responds with **x**, **y** or **z**. Player 1 wins with the combinations **ax**, **bx** and **by**. If Player 2 can observe the move of Player 1, then both strategies of Player 1 are admissible. If Player 2 cannot observe the first move, then only playing **b** is admissible for Player 1.

5 Characterizing Admissibility with Perfectly Informed Opponents

We will provide a characterization of the admissible strategies of Player 1 under the assumption that Player 2 is perfectly informed. Corollary 18 shows that this is in a sense a conservative information, as any strategy found to be not admissible here will not be admissible against arbitrary opponents.

Fix some $\sigma \in \Sigma_1^{\cong 1}$. We say that $M \subseteq \{0, 1\}^*$ is a σ -monochromatic set, if M is maximal under the two following constraints:

1. For any $v, u \in M$ we find that $v \cong_1 u$.
2. Any $v \in M$ is compatible with σ .

Given some σ -monochromatic set M , we partition M into $M_w^\sigma \uplus M_c^\sigma \uplus M_\ell^\sigma$ such that:

- $M_w^\sigma = \{v \in M \mid \forall \tau \in \Sigma_2 : v \text{ is compatible with } \tau \Rightarrow \text{Out}(\sigma, \tau) \in W\}$, i.e. the set of histories in M from which the strategy σ is adversarially winning;
- $M_c^\sigma = \{v \in M \mid v \notin M_w^\sigma \mid \exists \tau \in \Sigma_2 : v \text{ is compatible with } \tau \wedge \text{Out}_\subseteq(\sigma, \tau)W\}$, i.e. the set of histories in M from which the strategy σ is cooperatively winning;
- $M_\ell^\sigma = \{v \in M \mid \forall \tau \in \Sigma_2 : v \text{ is compatible with } \tau \Rightarrow \text{Out}(\sigma, \tau) \notin W\}$, i.e. the set of histories in M from which the strategy σ is losing against all the strategies of Player 2.

Note that in the definition of monochromatic set, we quantify over the *entire* set of strategies of Player 2 as Player 2 is assumed to be perfectly informed.

► **Theorem 20.** *An observation-based strategy $\sigma \in \Sigma_1^{\cong 1}$ of Player 1 is dominated, when Player 2 is perfectly informed, if and only if there exists a σ -monochromatic set K such that*

1. *there exists another strategy $\sigma' \in \Sigma_1^{\cong 1}$ such that K is also σ' -monochromatic,*
2. *the strategies σ and σ' induce partitions $(K_w^\sigma, K_c^\sigma, K_\ell^\sigma)$ and $(K_w^{\sigma'}, K_c^{\sigma'}, K_\ell^{\sigma'})$ of K such that:*
 - (a) $K_w^\sigma \cup K_c^\sigma \subseteq K_w^{\sigma'}$;
 - (b) *and either*
 - (i) $K_\ell^\sigma \cap (K_w^{\sigma'} \cup K_c^{\sigma'}) \neq \emptyset$,
 - (ii) *or* $K_c^\sigma \neq \emptyset$.

Proof. We first establish the right to left direction. We construct an observation-based strategy σ'' that dominates σ when Player 2 is perfectly informed. The strategy σ'' is defined as follows. In all histories that are not extensions of those in K , σ'' plays exactly as σ . For all histories extending some history in K , then σ'' plays as σ' . Let us prove that σ'' has all the required properties to dominate σ :

- First, we note that σ'' is observation-based as σ and σ' are both observation-based, and extensions of K cannot be \cong_1 -equivalent to non-extensions of K .

- Second, let π be a strategy of Player 2, and assume that $\text{Out}(\sigma, \pi) \in W$. Then we need to show that $\text{Out}(\sigma'', \pi) \in W$. We consider two different cases:
 1. If no prefix of $\text{Out}(\sigma, \pi)$ is in K then we have that $\text{Out}(\sigma, \pi) = \text{Out}(\sigma'', \pi)$ so $\text{Out}(\sigma'', \pi) \in W$.
 2. If a prefix w of $\text{Out}(\sigma, \pi)$ is in K , then once we reach this prefix σ'' is now acting as σ' . As $\text{Out}(\sigma, \pi) \in W$, w is such that $w \in K_w^\sigma \cup K_c^\sigma$, then we can conclude that $w \in K_w^{\sigma'}$, and so $\text{Out}(\sigma'', \pi) \in W$.
- Third, we need to establish the existence of one strategy π of Player 2 such that $\text{Out}(\sigma'', \pi) \in W$ and $\text{Out}(\sigma, \pi) \notin W$. For that we consider a history $w \in K$ such that either: (i) $w \in K_c^\sigma$ (and so $w \notin K_w^\sigma$), or (ii) $w \in K_\ell^\sigma \wedge w \in K_c^{\sigma''} \cup K_w^{\sigma''}$. Such a w always exists by hypothesis. We design π as follows. First, π plays compatible with w . As σ and σ'' disagree at the history w , w.l.o.g. let us consider that σ plays 0 while σ'' plays 1. As Player 2 is perfectly informed, he observes this and behaves in a way that:
 - if (i) $w \in K_c^\sigma$ and thus $w \in K_w^{\sigma'}$, then π does not help σ (this is possible as $w \notin K_w^\sigma$), and ensures that the outcome is outside W , while σ'' behaves like σ' and is thus winning from w no matter what π is.
 - if (ii) $w \in K_\ell^\sigma \wedge w \in K_c^{\sigma''} \cup K_w^{\sigma''}$, then π helps Player 1 when he plays σ'' to ensure that the outcome is in W and as $w \in K_\ell^\sigma$, we know that the outcome when σ is played is not in W .

We now establish the left to right direction. Assume that σ is dominated by some other strategy σ'' . By definition of dominance, there exists a strategy π of Player 2 such that $\text{Out}(\sigma, \pi) \notin W$ and $\text{Out}(\sigma'', \pi) \in W$. Let w be the longest common prefix of $\text{Out}(\sigma, \pi)$ and $\text{Out}(\sigma'', \pi)$, and let K be the σ -monochromatic set containing w (K is then also σ'' -monochromatic). Let us consider the partition of K for σ and σ'' . First, we know that $w \notin K_w^\sigma$ and $w \in K_c^{\sigma''} \cup K_w^{\sigma''}$. Also, because σ is dominated by σ'' then we have that:

- for all $v' \in K_w^\sigma : v' \in K_w^{\sigma''}$,
- for all $v' \in K_c^\sigma : v' \in K_w^{\sigma''}$,
- for all $v' \in K_\ell^{\sigma''} : v' \in K_\ell^\sigma$,

Otherwise, it is easy to obtain a contradiction with the fact that σ is dominated by σ'' . All this implies the right properties on the respective partitions. ◀

► **Remark.** This characterization of dominance is particularly useful in the case of regular games. There the question whether for some monochromatic K there exists a strategy σ inducing a particular partition $(K_w^\sigma, K_c^\sigma, K_\ell^\sigma)$ depends only on the set of last vertices of the histories of K . This is because winning sets in regular games are defined with parity conditions that lead to prefix independence. In particular, there are only finitely many cases to check. We will exploit this algorithmically in Subsection 7.2.

6 Games on Finite Trees

The simplest non-trivial subclass of infinite sequential games with imperfect information are those with clopen winning sets W . Essentially, this means that whether or not $p \in W$ for some run p only depends on some fixed-length prefix of p . Hence, we can consider these as *finite tree games*. Alternatively, we could conceive of these as regular games where the underlying graph is a tree.

The initial vertex is the root of the tree, and we denote by $\text{Leaves}(G)$ the leaves of the tree, i.e. the finite histories determining membership in W of any run passing through them. W.l.o.g., we assume that all players know when a leaf is reached, and that both players can distinguish winning from losing leaves; in other terms, no \cong_i -equivalence class contains both

a leaf and a non-leaf, and no class contains both a winning and a losing leaf. We denote the set of winning leaves by Ω_T .

In this section, for convenience, we assume that each vertex of Player i has k_i successors for a given $k_i \geq 1$, and that these successors are labeled with numbers between 0 and $k_i - 1$. We sometimes call these numbers *actions*, and denote $\text{Act}_i = \{0, 1, \dots, k_i\}$, and $\text{Act} = (\text{Act}_1, \text{Act}_2)$. For any vertex $v \in V$, and $a \in \text{Act}$, let $\delta(v, a)$ denote the a -th successor of v . Games with $|\text{Act}_i| > 2$ can be equivalently modeled using two actions by adding intermediate states (see the remark after Definition 4).

We denote the finitely many \cong_i -equivalence classes that matter (i.e. are not proper extensions of the leaves) by $\mathcal{O}^i = \{\mathcal{O}_j^i \mid j \in I^i\}$. An observation-based strategy then is, up to irrelevant moves, simply a function $\sigma : \mathcal{O}^i \rightarrow \text{Act}_i$, which can be stored in linear space.

6.1 Characterization of Domination

We are going to give an algorithm to decide whether a given strategy is dominated in a finite tree game. To do so, we are going to simultaneously simulate σ and another strategy that is to be chosen by Player i , in a product construction. We need the following additional definition:

► **Definition 21** (Switching strategies). For any player i , strategies $\tau, \tau' \in \Sigma_i^{\cong_i}$ and an observation $o \in \mathcal{O}_i$, we denote $\tau[o/\tau']$ the strategy that plays τ but upon visiting o switches to τ' . Formally,

$$\tau[o/\tau'](o_1 \cdots o_n) = \begin{cases} \tau'(o_1 \cdots o_n) & \text{if } \exists i, o_i = o \\ \tau(o_1 \cdots o_n) & \text{otherwise} \end{cases}$$

Let us consider a finite tree game $G = (V, E, V_1, V_2, \mathcal{O}^1, \mathcal{O}^2, \Omega_T)$ and let us note $\Sigma_i^{\cong_i}(G)$ the observation-based strategies of Player i in G . For strategy $\sigma \in \Sigma_i^{\cong_i}(G)$, and vertex $v \in V_i$, let $\delta^*(v, \sigma)$ be the vertex obtained by repeatedly applying σ as long as the vertex stays in V_i . Thus, $\delta^*(v, \sigma)$ is either a leaf, or a vertex of V_{3-i} . For $v \notin V_i$, let $\delta^*(v, \sigma) = v$. Given action a and strategy σ , let us define $a \cdot \sigma$ that plays a in the first step, and then switches to σ .

► **Definition 22.** Given a finite tree game $G(v_{\text{init}}) = (V, E, V_1, V_2, \mathcal{O}^1, \mathcal{O}^2, \Omega)$ with actions Act , and strategy $\sigma \in \Sigma_i^{\cong_i}(G)$, define $G_\sigma(v_{\text{init}}^\sigma) = (V^\sigma, E^\sigma, V_1^\sigma, V_2^\sigma, \mathcal{O}_\sigma^1, \mathcal{O}_\sigma^2, \Omega^\sigma)$ with actions Act^σ , where

- $V^\sigma = V \times V$, $V_i^\sigma = V \times V_i$, and $V_{3-i}^\sigma = V^\sigma \setminus V_i^\sigma$,
- $v_{\text{init}}^\sigma = (\delta^*(v_{\text{init}}, \sigma), v_{\text{init}})$,
- $\text{Act}_i^\sigma = \text{Act}_i$, and $\text{Act}_{3-i}^\sigma = \text{Act}_{3-i} \times \text{Act}_{3-i}$,
- $\mathcal{O}_i^\sigma = \{V \times \mathcal{O}_j^i \mid \mathcal{O}_j \in \mathcal{O}^i\}$, $\mathcal{O}_{3-i}^\sigma = \{\mathcal{O}_j^{3-i} \times \mathcal{O}_{j'}^{3-i} \mid \mathcal{O}_j^{3-i}, \mathcal{O}_{j'}^{3-i} \in \mathcal{O}^{3-i}\}$.
- For all $(v, v') \in V_i^\sigma$, $a \in \text{Act}_i^\sigma$,

$$\delta^\sigma((v, v'), a) = (v, \delta(v', a)).$$

For all $(v, v') \in V_{3-i}^\sigma$ with $\mathcal{O}^{3-i}(v) = \mathcal{O}^{3-i}(v')$, for all $(a, b) \in \text{Act}_{3-i}^\sigma$,

$$\delta^\sigma((v, v'), (a, b)) = (\delta^*(v, a \cdot \sigma), \delta(v', a)).$$

For all $(v, v') \in V_{3-i}^\sigma$, if $\mathcal{O}^{3-i}(v) \neq \mathcal{O}^{3-i}(v')$, for all $(a, b) \in \text{Act}_{3-i}^\sigma$,

$$\delta^\sigma((v, v'), (a, b)) = \begin{cases} (\delta^*(v, a \cdot \sigma), \delta(v', b)) & \text{if } v, v' \notin \text{Leaves}(G), \\ (\delta^*(v, a \cdot \sigma), v') & \text{if } v \notin \text{Leaves}(G), v' \in \text{Leaves}(G), \\ (v, \delta(v', b)) & \text{if } v \in \text{Leaves}(G), v' \notin \text{Leaves}(G), \end{cases}$$

Informally, Player i only sees the second component and plays as in G ; in fact, the second component of G_σ reproduces precisely G . The first component always moves according to σ from Player- i vertices, and according to Player $3-i$'s actions otherwise. Player $3-i$ plays the same actions in both components as long as her observations match, but can choose different actions otherwise. Observe that we “accelerate” the transitions in the first component in Player- i vertices, so we ensure that the first component is always either a Player $3-i$ vertex, or a leaf.

Note that $\Sigma_i^{\cong}(G) = \Sigma_i^{\cong}(G_\sigma)$ since the observations and actions for Player i are identical in both games.

Consider game G , player i , and $\sigma, \sigma' \in \Sigma_i^{\cong}(G)$. For any strategy $\underline{\tau} \in \Sigma_{3-i}^{\cong}(G)$, let us call $n_{\underline{\tau}}^{\sigma, \sigma'}$ the smallest integer such that the observation of Player $3-i$ at the $n_{\underline{\tau}}^{\sigma, \sigma'}$ -th step is different between the plays $\rho = \iota(\text{Out}_G(\sigma, \tau))$ and $\rho' = \iota(\text{Out}_G(\sigma', \tau))$, that is, $n_{\underline{\tau}}^{\sigma, \sigma'} = \min\{n \mid \mathcal{O}^{3-i}(\rho_n) \neq \mathcal{O}^{3-i}(\rho'_n)\}$ when this minimum is finite, and $n_{\underline{\tau}}^{\sigma, \sigma'} = \infty$ otherwise. Note that the lengths of the plays in the latter case must be equal since observations distinguish leaves from non-leaves. Let $\text{dist}_1^{G, \underline{\tau}}(\sigma, \sigma') = \rho_{n_{\underline{\tau}}^{\sigma, \sigma'}}$ and $\text{dist}_2^{G, \underline{\tau}}(\sigma, \sigma') = \rho'_{n_{\underline{\tau}}^{\sigma, \sigma'}}$ be these vertices where Player $3-i$ distinguishes both plays for the first time.

► **Lemma 23.** *Consider any game G , player i , $\sigma_1, \sigma_2 \in \Sigma_i^{\cong}(G)$, and $\underline{\tau} \in \Sigma_{3-i}^{\cong}(G_{\sigma_1})$. If we write $(t_1, t_2) = \ell(\iota(\text{Out}_{G_{\sigma_1}}(\sigma_2, \tau)))$, there exists $\tau \in \Sigma_{3-i}^{\cong}(G)$ such that $t_j = \ell(\iota(\text{Out}_G(\sigma_j, \tau)))$ for all $j \in \{1, 2\}$.*

Proof. If $n_{\underline{\tau}}^{\sigma_1, \sigma_2} = \infty$, then Player $3-i$ receives the same observation in both components. Let us define τ by $\tau(o) = a$ where $(a, b) = \underline{\tau}((o, o))$. Then choosing $\tau_1 = \tau_2 = \tau$ yields the result.

Assume $n_{\underline{\tau}}^{\sigma_1, \sigma_2} < \infty$, and write $(t_1, t_2) = \ell(\iota(\text{Out}_{G_{\sigma_1}}(\sigma_2, \underline{\tau})))$. We define $\tau(o)$ as the first component of $\underline{\tau}((o, o))$ for all observations $o \in \mathcal{O}_{3-i}$ such that $o \notin \{\text{dist}_j^{G, \underline{\tau}}(\sigma_1, \sigma_2)\}_{j \in \{1, 2\}}$ or o is not a descendant of these observations. Let (s_1, s_2) be the $n_{\underline{\tau}}^{\sigma_1, \sigma_2}$ -th vertex of $\iota(\text{Out}_{G_{\sigma_1}}(\sigma_2, \underline{\tau}))$. Notice that the run $\iota(\text{Out}_G(\sigma_j, \tau))$ visits s_j , for each $j \in \{1, 2\}$. For each j , in $\text{dist}_j^{G, \underline{\tau}}(\sigma_1, \sigma_2)$, there exists strategy τ_j such that $\ell(\iota(\text{Out}_{G_{s_j}}(\sigma_j, \tau_j))) = t_j$. We complete the definition of τ as follows: at each $\text{dist}_j^{G, \underline{\tau}}(\sigma_1, \sigma_2)$, it switches to τ_j . ◀

Let us define $P_\sigma = \{(t, t') \in \text{Leaves}(G) \times \text{Leaves}(G) \mid t \in \Omega_T \Rightarrow t' \in \Omega_T\}$.

► **Lemma 24.** *For any $\sigma' \in \Sigma_i^{\cong}(G)$, σ is weakly dominated by σ' if, and only if $G_\sigma, \sigma' \models \mathbf{AFP}_\sigma$.*

Proof. Assume $G_\sigma, \sigma' \models \mathbf{AFP}_\sigma$. Let $\tau \in \Sigma_{3-i}^{\cong}$ such that $G, \sigma, \tau \models \Omega$; we will show that $G, \sigma', \tau \models \Omega$.

Let $\underline{\tau}$ defined by $\underline{\tau}((o, o')) = (\tau(o), \tau(o'))$ for all $o, o' \in \mathcal{O}_{3-i}$. We have, by assumption, $G_\sigma, \sigma', \underline{\tau} \models \mathbf{FP}_\sigma$. Moreover, the projection of this run to the first component is exactly $\iota(\text{Out}_G(\sigma, \tau))$. The projection to the second component is $\iota(\text{Out}_G(\sigma', \tau))$. By the definition of P_σ , and since $\text{Out}_G(\sigma, \tau)$ is in Ω , $\text{Out}_G(\sigma', \tau)$ is in Ω too.

Conversely, assume that $G_\sigma, \sigma', \underline{\tau} \models \mathbf{G-P}_\sigma$. Let $(t, t') = \ell(\text{Out}(\sigma', \underline{\tau})_{G_\sigma})$, with $t \in \Omega_T$, $t' \notin \Omega_T$. Let $\tau \in \Sigma_{3-i}^{\cong}(G)$ be as given by Lemma 23. We get that $\iota(\text{Out}_G(\sigma, \tau)) = t$ and $\iota(\text{Out}_G(\sigma', \tau)) = t'$, which gives the desired result. ◀

We let $Q_\sigma = \{(t, t') \in \text{Leaves}(G) \times \text{Leaves}(G) \mid t \notin \Omega_T \wedge t' \in \Omega_T\}$.

► **Lemma 25.** *For any $\sigma' \in \Sigma_i^{\cong}(G)$, σ is dominated by σ' if, and only if $G_\sigma, \sigma' \models \mathbf{AFP}_\sigma \wedge \mathbf{EFQ}_\sigma$.*

1 \ 2	x	y	z
a a	1	0	0
a b	0	1	1
b a	1	1	0
b b	0	0	1

■ **Figure 2** An admissible strategy of a sub-game may not be a part of an admissible strategy of the whole game.

Proof. Assume $G_\sigma, \sigma' \models \mathbf{AFP}_\sigma \wedge \mathbf{EF}Q_\sigma$. The previous lemma shows that σ' weakly dominates σ . Let τ such that G_σ, σ', τ ends in $(t, t') \in Q_\sigma$. Let τ given by Lemma 23. We get that $\iota(\text{Out}_G(\sigma, \tau)) = t$ and $\iota(\text{Out}_G(\sigma', \tau)) = t'$, which means that σ is dominated by σ' . ◀

► **Lemma 26.** *Given a game G , player i , and subsets $P, Q \subseteq \text{Leaves}(G)$, one can decide in polynomial time if there is a strategy $\sigma \in \Sigma_i^{\cong i}(G)$ such that $G, \sigma \models \mathbf{AFP} \wedge \mathbf{EF}Q$.*

Proof. We first show that there exist maximally permissive strategies for objectives of the type \mathbf{AFP} . In fact, define

$$\text{Safe} = \{(o, a) \in \mathcal{O}_i \times \text{Act}_i \mid \exists \sigma \in \Sigma_i^{\cong i}, G, \sigma \models \mathbf{AFP} \wedge \mathbf{EF}o, \sigma(o) = a\}.$$

This is the set of pairs of observation-actions (o, a) such that some Player- i strategy that is winning for \mathbf{AFP} is compatible with o , and chooses a from o . We claim that any $\sigma \in \Sigma_i^{\cong i}(G)$ such that $\forall o \in \mathcal{O}_i, (o, \sigma(o)) \in \text{Safe}$ is winning for objective \mathbf{AFP} . This can be proved by induction on the length of plays, and using perfect recall.

Now, the set Safe can be computed bottom-up. Let G' be the game G where we only keep Player- i actions that conform to Safe . There exists a strategy in G winning for $\mathbf{AFP} \wedge \mathbf{EF}Q$ if, and only if some vertex of Q is reachable in G' . ◀

► **Theorem 27.** *Given game G on a finite tree, player i and strategy $\sigma \in \Sigma_i^{\cong i}(G)$, one can decide in polynomial time whether σ is admissible.*

6.2 Hardness

Let us start with a remark showing that one cannot hope to prune some actions of the game locally, to obtain a description of the admissible plays as it is the case in the perfect-information setting, see section on safety games in [12]. In fact, a strategy that only uses actions that appear in (the reachable part of) some admissible strategy may not be admissible itself.

Consider the game in normal form in Figure 2. It is easy to encode it as a game on a tree by ensuring that Player 2 is blind (i.e. all states in V_2 leads to the same observation) and players have perfect recall (in particular, they remember which actions they have played). Player 1 starts choosing a or b , then Player 2 plays (x , y or z), and Player 1 plays again. We have that ab is admissible, so there is an admissible strategy playing a in the first vertex. Furthermore, in the subgame where Player 1 has already played a , playing a is admissible (to see this, compare the first two lines in the figure). However strategy aa is dominated by ba .

Similarly there is an admissible strategy of Player 1 playing b in the first step. Playing b is admissible in the subgame that starts with history b , but the strategy bb is dominated by ab .

We now show that deciding the existence of an admissible strategy choosing a particular action at a given state is NP-complete.

► **Theorem 28.** *Given two-player game G on a finite tree, and an action a , deciding whether there is an admissible strategy σ with $\sigma(s_{\text{init}}) = a$ is NP-complete.*

Proof. NP-membership holds because strategies of Player 1 can be represented in polynomial size, and having guessed some strategy σ_1 which uses a , whether σ_1 is admissible can be verified in polynomial time by Corollary 27.

For NP-hardness, we encode an instance of the 3-SAT problem of the following form

$$\psi = \exists x_1, \dots, x_n. C_1 \wedge \dots \wedge C_m,$$

where $C_i = \ell_{i,1} \vee \ell_{i,2} \vee \ell_{i,3}$ and $\ell_{i,j} \in \{x_k, \neg x_k \mid k \in [1, n]\}$, for each $1 \leq i \leq m$, and $1 \leq j \leq 3$.

We define a game G_Ω as follows. From the initial vertex Player 1 either chooses a clause among C_1, \dots, C_m , or goes to a special vertex C_0 . Then Player 2 blindly selects a literal $\ell \in \{x_k, \neg x_k \mid k \in [1, n]\}$. At this point, Player 1 only sees the index k such that $\ell \in \{x_k, \neg x_k\}$. Let us write $o^{i,k}$ for this observation, where i is the index of the clause, or 0, and k is the index of the variable chosen by Player 2.

From $o^{0,k}$, Player 1 chooses T or F meaning, intuitively, that x_k is evaluated to true and false respectively. All other observations $o^{i,k}$ with $i > 0$ are leaves. The reachability condition Ω is satisfied when either Player 1 selects C_i with $i > 0$ and ℓ is a literal which does not appear in C_i or Player 1 selected C_0 and the valuation chosen in the last vertex is such that ℓ is evaluated to true.

Assume there is a strategy σ_1 choosing C_0 which is not dominated by any strategy not choosing C_0 . Now, define the valuation v as follows. For each variable x_k , we set $v(x_k)$ to true if, and only $\sigma_1(o^{0,k}) = T$. Let σ'_1 be any strategy of Player 1 choosing some clause C_i with $i > 0$. Because σ_1 is not dominated by σ'_1 , either: (A) there is a strategy τ^i of Player 2 which makes σ_1 win but not σ'_1 ; or (B) for all strategies τ of Player 2, σ'_1, τ wins implies σ_1, τ wins.

Let us first show that case (B) cannot happen. Recall that σ'_1, τ is winning when τ selects a literal outside C_i . If there are more than 2 variables ($n > 2$), which we can assume without loss of generality, then there is one variable x_k and its negation that do not appear in C_i . Consider the strategies τ and τ' that select x_k and $\neg x_k$ respectively. Since they make σ'_1 win, they must also make σ_1 win, which means that the valuation defined by σ_1 makes both x_k and $\neg x_k$ true, which is a contradiction.

Thus, only the case (A) is possible. For each $1 \leq i \leq m$, let σ'_i be the strategy that chooses C_i . Let τ^i be a Player-2 strategy such that σ_1, τ^i wins and σ'_i, τ^i loses. Let us fix $1 \leq i \leq m$, and let ℓ denote the literal chosen by τ^i . The run of σ'_i, τ^i is losing, which means that $\ell \in C_i$. Since σ_1, τ^i wins, ℓ is satisfied by v . Since $\ell \in C_i$, C_i is satisfied by v . Since i was arbitrary, all clauses C_i are satisfied. Hence ψ is satisfiable.

In the other direction, assume the formula is satisfiable. Let v be a valuation satisfying the formula, and σ_v be the strategy that first plays C_0 , and then plays according to the valuation v . Let σ'_1 be any strategy that first plays C_i for any $i \in [1, n]$. As v satisfies the formula, there is a literal ℓ in C_i such that $v(\ell) = T$. Let τ_ℓ be the strategy of Player 2 that chooses this literal. As $\ell \in C_i$, τ_ℓ makes σ'_1 lose, and as $v(\ell) = T$, τ_ℓ makes σ_1 win. So σ_1 is not dominated by σ'_1 . ◀

7 Decision problems for dominance in regular games

In this section, we study decision problems related to admissible strategies for regular infinite sequential games with imperfect information. We have seen in Theorem 17 that a strategy of Player 1 that is dominated when Player 2 is perfectly informed is also dominated when Player 2 is imperfectly informed. The special case of a perfectly informed Player 2 is thus important. In this section, we start by considering this case, exploiting the characterization obtained in Theorem 20. Then we turn to the case when Player 2 is not perfectly informed. Before presenting our results, we need to recall some notions related to infinite trees and automata on infinite trees.

7.1 Automata on infinite trees and tree encodings of strategies

Infinite trees

Following e.g. [21], given a finite set \mathcal{D} , a \mathcal{D} -tree is a prefix closed language $T \subseteq \mathcal{D}^*$. When \mathcal{D} is clear from the context, we call T a tree. Elements of T are called *nodes*, and the empty word ϵ is the *root* of the tree. For every $x \in T$, the nodes $x \cdot d \in T$ for $d \in \mathcal{D}$ are the *children* of x , the children $x \cdot d$ is called the child of direction d of x . A \mathcal{D} -tree is a full infinite tree if $T = \mathcal{D}^*$. A *branch* b of T is a set $b \subseteq T$ such that $\epsilon \in b$ and for every $x \in b$ there exists a unique $d \in \mathcal{D}$ such that $x \cdot d \in b$. Given two sets \mathcal{D} and \mathcal{L} , a \mathcal{L} -labelled \mathcal{D} -tree is a pair $\langle T, L \rangle$ where T is a \mathcal{D} -tree and $L : T \rightarrow \mathcal{L}$ labels each node of T with an element of \mathcal{L} .

Tree encoding of observation-based strategies

We encode observation-based strategies $\sigma \in \Sigma_1^{\cong 1}$, i.e. $\sigma : (\mathcal{O}^1)^* \cdot \mathcal{O}_{\text{turn}}^1 \rightarrow \{0, 1\}$, as full infinite $\{0, 1, *\}$ -labelled \mathcal{O}^1 -trees with the convention that:

- the root node is always labelled with $*$;
- every node that follows an observation $o \in \mathcal{O}_{\text{turn}}^1$ is labeled with 0 or 1 according to σ : if the node is reached with a sequence of directions $\rho = o_1 o_2 \dots o_n$ that corresponds to observations in \mathcal{O}^i and $o_n \in \mathcal{O}_{\text{turn}}^1$, then the node is labelled by $\sigma(\rho)$;
- every node that follows an observation $o \in \mathcal{O}^1 \setminus \mathcal{O}_{\text{turn}}^1$ is labelled with $*$ as in that case, it is the turn of Player 2 to play and the choice of Player 2 is not visible for Player 1.

Alternating tree automata

We use *alternating tree automata* (AT) to recognize regular sets of infinite trees that encode observation-based strategies. A *AT* $P = (Q, q_0, \Delta, \Psi)$ that operates on \mathcal{L} -labelled \mathcal{D} -trees is defined by:

- its finite non-empty set of states Q ;
- its initial state q_0 ;
- its transition function $\Delta : Q \times \mathcal{L} \rightarrow \mathcal{B}^+(Q \times \mathcal{D})$, where $\mathcal{B}^+(Q \times \mathcal{D})$ is the set of (positive) Boolean formula built from elements in $Q \times \mathcal{D}$ using \vee and \wedge . For a set $E \subseteq Q \times \mathcal{D}$ and a formula $\psi \in \mathcal{B}^+(Q \times \mathcal{D})$, we say that E satisfies ψ iff assigning true to elements in E and assigning false to elements in $(Q \times \mathcal{D}) \setminus E$ makes ψ true;
- Ψ is the acceptance condition which is a set of infinite sequences of states in Q that are accepting. Typically, we consider either parity functions, giving alternating parity tree automata (APT), or Muller conditions (that subsume Boolean combinations of parity conditions), giving alternating Muller tree automata (AMT), to define Ψ .

A run of $P = (Q, q_0, \Delta, \Psi)$ on a \mathcal{L} -labelled \mathcal{D} -tree $\langle T, L \rangle$ is $Q \times \mathcal{D}^*$ -labelled tree $\langle T_r, r \rangle$ such that:

1. $\epsilon \in T_r$ and $r(\epsilon) = (q_0, \epsilon)$;
2. Let $y \in T_r$ with $r(y) = (q, x)$ and $\Delta(q, L(x)) = \psi$. Then there exists a (possibly empty) set $E = \{(q_0, d_0), (q_1, d_1), \dots, (q_n, d_n)\} \subseteq Q \times D$ and:
 - a. E satisfies ψ ;
 - b. for all i , $0 \leq i \leq n$, we have $y \cdot i \in T_r$ and $r(y \cdot i) = (q_i, x \cdot d_i)$.

A run is *accepting* iff all its infinite branches b^∞ are such that $b^\infty \in \Psi$, i.e. they satisfy the acceptance condition.

7.2 Player 2 perfectly informed

As we mentioned following Theorem 20, in regular games the validity of the condition in the theorem depends only on the set of last vertices in the histories forming a monochromatic set. In slight abuse of notation, we call a set of vertices arising in this way a monochromatic set, too.

► **Lemma 29.** *Let $M = M_w \uplus M_c \uplus M_\ell$ be a monochromatic set. We can construct in PTIME an APT that recognizes the set of trees that encode observation-based strategies $\sigma : (\mathcal{O}^1)^* \cdot \mathcal{O}_{\text{turn}}^1 \rightarrow \{0, 1\}$ of Player 1 and such that $M_w = M_w^\sigma$, $M_c = M_c^\sigma$, and $M_\ell = M_\ell^\sigma$.*

Proof. The main idea of the construction is as follows. The APT has four parts:

1. one that checks that σ ensures a win from all vertices in M_w no matter what are the choices made by Player 2;
2. one that checks that σ can cooperate with Player 2 to win from all vertices in M_c ;
3. one that checks that σ can cooperate with Player 2 to lose from all the vertices in M_c ;
4. one that checks that all outcomes that are compatible with σ are losing from all the vertices in M_ℓ no matter what are the choices made by Player 2.

Accordingly, its state space is defined as $Q = \{q_0\} \cup V \times \{w, c_w, c_\ell, \ell\}$. The transitions are defined as follows:

- for states in $V \times \{w\}$. If a node n of the tree is reached by a run of the APT that ends up in state (v, w) then we must verify that the strategy encoded in the tree ensures to win no matter what are the choices of Player 2. So, the automaton follows the choices prescribed by the strategy for nodes annotated by $i \in \{0, 1\}$ and branches universally on the choices of Player 2 for nodes annotated by $*$. Accordingly:
 - for $i \in \{0, 1\}$: $\Delta((v, w), i) = ((s(v, i), w), \mathcal{O}^1(s(v, i)))$
 - for $*$: $\Delta((v, w), *) = ((s(v, 0), w), \mathcal{O}^1(s(v, 0))) \wedge ((s(v, 1), w), \mathcal{O}^1(s(v, 1)))$
- for states in $V \times \{c_w\}$. If a node n of the tree is reached by a run of the APT that ends up in state (v, c_w) then the automaton must verify that the strategy σ encoded in the tree ensures that at least one outcome is winning. So, the automaton follows the choices of Player 1 as encoded in the tree and for nodes labelled with $*$ that denote choices of Player 2, the automaton *nondeterministically* chooses one choice for Player 2. Accordingly:
 - for $i \in \{0, 1\}$: $\Delta((v, c_w), i) = ((s(v, i), c_w), \mathcal{O}^1(s(v, i)))$
 - for $*$: $\Delta((v, c_w), *) = ((s(v, 0), c_w), \mathcal{O}^1(s(v, 0))) \vee ((s(v, 1), c_w), \mathcal{O}^1(s(v, 1)))$
- for states in $V \times \{c_\ell\}$. The approach is similar. If a node n of the tree is reached by a run of the APT that ends up in state (v, c_ℓ) then the automaton must verify that the strategy σ encoded in the tree ensures that at least one outcome compatible with σ is losing. So, the automaton follows the choices of Player 1 as encoded in the tree and for nodes labelled with $*$ that denote choices of Player 2, the automaton *nondeterministically*

chooses one choice for Player 2. The difference with the previous case (c_w) is handled by the acceptance condition. Accordingly:

- for $i \in \{0, 1\}$: $\Delta((v, c_\ell), i) = ((s(v, i), c_\ell), \mathcal{O}^1(s(v, i)))$
- for $*$: $\Delta((v, c_\ell), *) = ((s(v, 0), c_\ell), \mathcal{O}^1(s(v, 0))) \vee ((s(v, 1), c_\ell), \mathcal{O}^1(s(v, 1)))$
- for states in $V \times \{\ell\}$. If a node n of the tree is reached by a run of the APT that ends up in state (v, ℓ) then we must verify that the strategy σ encoded in the tree cannot win no matter what Player 2 chooses. So, the automaton follows the choices prescribed by the strategy σ for nodes annotated by $i \in \{0, 1\}$ and branches universally on the choices of Player 2 for nodes annotated by $*$. The difference with the first case (w) is handled by the acceptance condition. Accordingly:
 - for $i \in \{0, 1\}$: $\Delta((v, \ell), i) = ((s(v, i), \ell), \mathcal{O}^1(s(v, i)))$
 - for $*$: $\Delta((v, \ell), *) = ((s(v, 0), \ell), \mathcal{O}^1(s(v, 0))) \wedge ((s(v, 1), \ell), \mathcal{O}^1(s(v, 1)))$

It remains now to define the acceptance condition. First, we remark that a run of the APT directly jump from its initial state q_0 to one of the four parts whose transitions have been described above. When entering one of those parts, the run will stay in that part of the state space for ever. The runs that have entered the part associated to w and to c_w must simulate paths in the game graph that are in W , those that are in the part associated to c_ℓ and to ℓ must simulate paths that are outside W . W is defined by a parity condition pr . We thus define the parity condition pr' of our automaton as follows: for all $v \in V$, $pr'(v, w) = pr'(v, c_w) = pr(v)$ (the parity condition is preserved) and $pr'(v, c_\ell) = pr'(v, \ell) = pr(v) + 1$ (the parity condition is inverted). ◀

As the emptiness problem for APT is solvable in EXPTIME, we deduce the following corollary:

► **Corollary 30.** *Given a partition $M = M_w \uplus M_c \uplus M_\ell$ of a monochromatic set, we can decide in EXPTIME if there exists an observation-based strategy $\sigma : (\mathcal{O}^1)^* \cdot \mathcal{O}_{\text{turn}}^1 \rightarrow \{0, 1\}$ that induces this partition.*

► **Lemma 31.** *We can construct in EXPTIME an APT P that accepts the $\{0, 1, *\}$ -labelled \mathcal{O}^1 -trees that are the tree encodings of the observation-based strategies $\sigma \in \Sigma_1^{\omega, 1}$ that are dominated in $R^U(v_0)$ where Player 2 is perfectly informed.*

Proof. Let $R(v_0)$ be a regular game. First, we note that Lemma 29 allows us to compute, for all monochromatic sets K , all the partitions $K_w \uplus K_c \uplus K_\ell$ that are witnessed by observation-based strategies of Player 1. Among those, we can extract all the partitions $K_w \uplus K_c \uplus K_\ell$ that correspond to strategies that are dominated following Theorem 20. There are at most an exponential number of them in the size of the regular game. We note \mathcal{B} all those partitions.

Our APT will first guess a finite observation history ρ and compute the knowledge associated to ρ while following the strategy σ encoded in the tree. Let K be the resulting monochromatic set of vertices. This set is stored in the state of the APT and then the APT nondeterministically chooses a partition $K_w \uplus K_c \uplus K_\ell$ of K in \mathcal{B} . Then the APT verifies that the strategy in the tree induces that partition. This is done as in the proof of Lemma 29. ◀

As APT are closed under all Boolean operations, see e.g. [19], the set of admissible observation-based strategies of Player 1 is effectively omega-regular:

► **Theorem 32.** *When Player 2 is perfectly informed, the set of admissible observation-based strategies of Player 1 in $R^U(v_0)$ is effectively ω -regular and the emptiness of this set is decidable in 2EXPTIME.*

7.3 Player 2 imperfectly informed

The hierarchical case

We start by considering the case in which both players have imperfect information but Player 2 is more informed than Player 1. In this case, the informedness of players is hierarchical in the sense of [7]. For the hierarchical case, we can decide if the set of admissible observation-based strategies of Player 1 is empty or not. We obtain this result by a reduction to the model-checking problem of *Strategy Logic with Imperfect Information* [5], SL_{ii} for short.

We only recall here informally the syntax and semantics of SL_{ii} formulas and refer the interested reader to [5] for formal definitions. We start with the case where the players are perfectly informed. Strategy Logic (SL) extends the linear temporal logic (LTL [3]) and treats strategies x as first-order objects that can be quantified: $\langle\langle x \rangle\rangle$ reads "there exists a strategy x " while $\llbracket x \rrbracket$ reads "for all strategies x ". In SL, strategies can be bound to players: $(x, 1)$ reads "Player 1 uses strategy x ". As an example, let ϕ be an LTL formula, and consider the following SL formula $\langle\langle x \rangle\rangle \llbracket y \rrbracket (x, 1)(y, 2)\phi$. This formula expresses that there exists a strategy x for Player 1 such that for all strategies y of Player 2, when the two players play their respective strategies x and y then the outcome from the initial vertex of the game arena is a path that satisfies ϕ . So, this formula expresses the existence of a *winning strategy* for Player 1 in a two-player zero sum game with objective ϕ for Player 1 and $\neg\phi$ for Player 2. SL can express many interesting game properties such as the existence of Nash equilibria, the existence of dominating strategies, etc., see [14, 22] for more examples.

When strategy logic is interpreted over a game in which players have imperfect information, then the strategy quantifier explicitly limits the quantification to observation-based strategies: $\langle\langle x \rangle\rangle^{\mathcal{O}}$ reads "there exists an \mathcal{O} -observation-based strategy x ". However, to obtain the decidability of the model-checking problem for SL_{ii} , quantifiers must respect constraints that ensure "hierarchical instances".

Intuitively, a formula of SL_{ii} is *hierarchical* if, as one goes down the syntactic tree of the formula, the observations annotating strategy quantifications can only become finer. So, for example, if in the syntactic tree, quantification $\langle\langle x \rangle\rangle^{\mathcal{O}_1}$ is followed by $\langle\langle y \rangle\rangle^{\mathcal{O}_2}$ then it must be the case that \mathcal{O}_2 is finer than \mathcal{O}_1 meaning that for all $v_1, v_2 \in V$, if \mathcal{O}_2 does not distinguish v_1 and v_2 , i.e. $\mathcal{O}_2(v_1) = \mathcal{O}_2(v_2)$, then it is also the case for \mathcal{O}_1 , i.e. $\mathcal{O}_1(v_1) = \mathcal{O}_1(v_2)$.

Let us now consider a regular game $R(v_0) = (V, E, V_1, V_2, \mathcal{O}^1, \mathcal{O}^2)$ such that \mathcal{O}_2 is finer than \mathcal{O}_1 that induces the infinite sequential game with imperfect information $R^U(v_0) = (d, W, \cong_1, \cong_2)$. Let ϕ_1 be a LTL specification for the objective of Player 1 then the following SL_{ii} is hierarchical and evaluates to true in v_0 , if and only if, Player 1 has admissible \mathcal{O}_1 -observation-based strategies when Player 2 plays \mathcal{O}_2 -observation-based strategies:

$$\langle\langle x \rangle\rangle^{\mathcal{O}_1} \llbracket x' \rrbracket^{\mathcal{O}_1} \langle\langle y \rangle\rangle^{\mathcal{O}_2} (x, 1)(y, 2)\phi_1 \wedge (x', 1)(y, 2)\neg\phi_1 \\ \vee \llbracket y \rrbracket^{\mathcal{O}_2} (x', 1)(y, 2)\phi_1 \rightarrow (x, 1)(y, 2)\phi_1$$

► **Theorem 33.** *Let $R(v_0) = (V, E, V_1, V_2, \mathcal{O}^1, \mathcal{O}^2)$ be a regular game such that \mathcal{O}_2 is finer than \mathcal{O}_1 and that induces $R^U(v_0) = (d, W, \cong_1, \cong_2)$. The existence of admissible strategies for Player 1 in $R^U(v_0)$ is decidable.*

The general case

This case is more involved and so far we did not succeed to obtain a general characterization of dominated strategies in the form of Theorem 20. Nevertheless, we are able to characterize all the observation-based strategies of Player 1 that dominate a regular observation-based

strategy given as a finite automaton \mathcal{A} (as in Definition 6). Our characterization is effective: we can construct from a regular game R and a finite automaton \mathcal{A} , an APT P that accepts the tree encodings of all the observation-based strategies σ that dominate the regular strategy $\sigma_{\mathcal{A}}$ defined \mathcal{A} . The construction of the APT is given in the proof of the following lemma:

► **Lemma 34.** *Let $R(v_0)$ be a regular game that induces $R^U(v_0) = (d, W, \cong_1, \cong_2)$, and let \mathcal{A} be a finite automaton that encodes a regular observation-based strategy $\sigma_{\mathcal{A}} \in \Sigma_1^{\cong_1}$. We can construct in PTIME an AMT P that accepts all the tree encodings of observation-based strategies $\sigma \in \Sigma_1^{\cong_1}$ that dominates $\sigma_{\mathcal{A}}$ when Player 2 plays strategies in $\Sigma_2^{\cong_2}$.*

Proof. Remember that the strategy $\sigma_{\mathcal{A}}$ is dominated by a strategy σ when Player 2 plays observation-based strategies if and only if:

1. $\forall \pi \in \Sigma_2^{\cong_2} : \text{Out}(\sigma_{\mathcal{A}}, \pi) \in W \rightarrow \text{Out}(\sigma, \pi) \in W$
2. $\exists \pi \in \Sigma_2^{\cong_2} : \text{Out}(\sigma_{\mathcal{A}}, \pi) \notin W \wedge \text{Out}(\sigma, \pi) \in W$

Here the strategy $\sigma_{\mathcal{A}}$ is encoded by the automaton $\mathcal{A} = (Q^{\mathcal{A}}, q_0^{\mathcal{A}}, \mathcal{O}^1, \delta^{\mathcal{A}}, F^{\mathcal{A}})$, and the strategy σ is the strategy encoded in the tree. We use a AMT to check those two properties: its state space is structured in two disjunct parts plus an initial state q_0 : $Q = \{q_0\} \cup Q^1 \cup Q^2$. States in Q^1 are used to check the first property, and in Q^2 the second property. We now give a detailed description of those sets of states: $Q^i = \{i\} \times V \times V \times \{0, 1\} \times Q^{\mathcal{A}}$, $i \in \{1, 2\}$. Intuitively, a state $(i, v_1, v_2, \text{dist}, q)$ is reached when:

- if $i = 1$, the ATP is checking the first condition, and if $i = 2$, the AMT is checking the second condition;
- the interaction of the strategy π played by Player 2 against $\sigma_{\mathcal{A}}$ leads to vertex v_1 ;
- the interaction of the strategy π played by Player 2 against σ leads to vertex v_2 ;
- dist is true if and only if Player 2 has been able to distinguish, based on its observation \mathcal{O}^2 , between the history generated by strategy σ against π and the history generated by the strategy $\sigma_{\mathcal{A}}$ against π , and false otherwise;
- q is the state reached by the automaton \mathcal{A} , that encodes $\sigma_{\mathcal{A}}$ on the current sequence of \mathcal{O}^1 -observations.

In the part Q^1 of the AMT, the automaton branches universally on the choices of Player 2 in order to consider all possible strategies π that Player 2 can play and it verifies that, for all of them, if the outcome against $\sigma_{\mathcal{A}}$ is winning then the outcome against σ is also winning. In the part Q^2 of the AMT, the automaton branches existentially on the choices of Player 2 in order to guess a strategy of Player 2 that forces an outcome in the complement of W against $\sigma_{\mathcal{A}}$ and in W against σ . This is realized by the following transition relation and acceptance condition:

- **Transition function.** Let $(i, v_1, v_2, \text{dist}, q^{\mathcal{A}})$, and we distinguish between the universal ($i = 1$) and existential ($i = 2$) cases:
 - in the universal part:
 - * if the label of the current node in the tree is $*$ then it is the turn of Player 2 to play. Then either Player 2 has already distinguished the histories that ends up in v_1 and v_2 respectively, i.e. $\text{dist} = 1$ and so he can possibly play differently in the two cases:

$$\begin{aligned} & \Delta((1, v_1, v_2, 1, q^{\mathcal{A}}), *) \\ & = \\ & \bigwedge_{j_1, j_2 \in \{0, 1\}} ((1, s(v_1, j_1), s(v_2, j_2), 1, \delta^{\mathcal{A}}(q^{\mathcal{A}}, \mathcal{O}^1(s(v_1, j_1))))), \mathcal{O}^1(s(v_2, j_2))) \end{aligned}$$

2:20 Admissibility in Games with Imperfect Information

or Player 2 has not yet distinguished the histories that ends up in v_1 and v_2 respectively, i.e. $\text{dist} = 0$, and so Player 2 makes the same choices in the two branches:

$$\begin{aligned} & \Delta((1, v_1, v_2, 0, q^A), *) \\ &= \\ & \bigwedge_{j \in \{0,1\}} ((1, s(v_1, j), s(v_2, j), \text{dist}'(j), \delta^A(q^A, \mathcal{O}^1(s(v_1, j))))), \mathcal{O}^1(s(v_2, j))) \end{aligned}$$

where $\text{dist}'(j) = (\mathcal{O}^2(s(v_1, j)) \neq \mathcal{O}^2(s(v_2, j)))$, i.e. the new value of dist is set to 1 only if Player 2 received two different observations for $s(v_1, j)$ and $s(v_2, j)$, in that case, he can now distinguish between the two branches.

- * If the label of the current node in the tree is $c \in \{0, 1\}$ then it is the turn of Player 1 to play and c is the choice made by σ , while the choice d made by σ_A is equal to 0 if $q^A \in F^A$ otherwise it is equal to 1:

$$\begin{aligned} & \Delta((1, v_1, v_2, \text{dist}, q^A), c) \\ &= \\ & ((1, s(v_1, d), s(v_2, c), \text{dist}'(j), \delta^A(q^A, \mathcal{O}^1(s(v_1, d))))), \mathcal{O}^1(s(v_2, c))) \end{aligned}$$

where $\text{dist}'(j) = \text{dist} \vee (\mathcal{O}^2(s(v_1, d)) \neq \mathcal{O}^2(s(v_2, c)))$, i.e. the two histories are distinguishable if they were before or if Player 2 gets two different observations in this round.

- in the existential part:

- * If the label of the current node in the tree is $*$ then it is the turn of Player 2 to play. Then either Player 2 has already distinguished the histories that ends up in v_1 and v_2 respectively, i.e. $\text{dist} = 1$:

$$\begin{aligned} & \Delta((1, v_1, v_2, 1, q^A), *) \\ &= \\ & \bigvee_{j_1, j_2 \in \{0,1\}} ((1, s(v_1, j_1), s(v_2, j_2), 1, \delta^A(q^A, \mathcal{O}^1(s(v_1, j_1))))), \mathcal{O}^1(s(v_2, j_2))) \end{aligned}$$

or Player 2 has not already distinguished the histories that ends up in v_1 and v_2 respectively, i.e. $\text{dist} = 0$, so Player 2 makes the same choices in the two branches:

$$\begin{aligned} & \Delta((1, v_1, v_2, 0, q^A), *) \\ &= \\ & \bigvee_{j \in \{0,1\}} ((1, s(v_1, j), s(v_2, j), \text{dist}'(j), \delta^A(q^A, \mathcal{O}^1(s(v_1, j))))), \mathcal{O}^1(s(v_2, j))) \end{aligned}$$

where $\text{dist}'(j) = (\mathcal{O}^2(s(v_1, j)) \neq \mathcal{O}^2(s(v_2, j)))$, i.e. the new value of dist becomes equal to 1 only if Player 2 received two different observations for $s(v_1, j)$ and $s(v_2, j)$, if so, he can now distinguish between the two branches. Note that those formulas are the same as the one for the universal part but with the conjunction replaced by a disjunction.

- * If the label of the current node in the tree is $d \in \{0, 1\}$ then it is the turn of Player 1 to play and d is the choice made by σ , while the choice c made by σ_A is equal to 0 if $q^A \in F^A$ otherwise it is equal to 1, and then we have that:

$$\begin{aligned} & \Delta((1, v_1, v_2, 1, q^A), c) \\ &= \\ & ((1, s(v_1, c), s(v_2, d), \text{dist}'(j), \delta^A(q^A, \mathcal{O}^1(s(v_1, c))))), \mathcal{O}^1(s(v_2, d))) \end{aligned}$$

where $\text{dist}'(j) = \text{dist} \vee (\mathcal{O}^2(s(v_1, c)) \neq \mathcal{O}^2(s(v_2, d)))$, i.e. the two histories are distinguishable if they were before or if Player 2 gets two different observations now.

- **Acceptance condition.** Again, we distinguish between the universal and existential parts:
 - Let r be a run of the AMT in the universal part and let $\text{inf}(r)$ be the set of states that repeats infinitely often along r . Let P_r be the set of pairs of vertices (v_1, v_2) that appear in states of the automaton in $\text{inf}(r)$, i.e. pairs of vertices that appear infinitely often along the run r . We declare the set of states $\text{inf}(r)$ to be accepting when we verify that: if $\min_{(v_1, v_2) \in P_r} p(v_1)$ is even then $\min_{(v_1, v_2) \in P_r} p(v_2)$ is even, i.e. a run is good if the run simulates an execution in which Player 2 plays a strategy π , and if the outcome of σ_A and π is winning for the parity condition pr , then it is also the case for the outcome of σ and π . This is a valid Muller condition.
 - Let r be a run of the AMT in the existential part and let $\text{inf}(r)$ be the set of states that repeats infinitely often along r . Let P_r be the set of pairs of vertices (v_1, v_2) that are in $\text{inf}(r)$, i.e. the pairs of vertices that appear infinitely often along the run r . We declare the set of state $\text{inf}(r)$ to be accepting when we verify that: $\min_{(v_1, v_2) \in P_r} p(v_1)$ is odd and $\min_{(v_1, v_2) \in P_r} p(v_2)$ is even, i.e. a run is good if the run simulates an execution in which Player 2 plays a strategy π , and the outcome of σ_A and π is losing for the parity condition pr , and the outcome of σ and π is winning for the parity condition pr . This is a valid Muller condition. ◀

As a consequence, we obtain:

► **Theorem 35.** *Let $R(v_0)$ be a regular game that induces $R^U(v_0) = (d, W, \cong_1, \cong_2)$, and let \mathcal{A} be a finite automaton that encodes a regular observation-based strategy $\sigma_A \in \Sigma_1^{\cong_1}$. The problem to decide if σ_A is dominated when Player 2 plays strategies in $\Sigma_2^{\cong_2}$ is EXPTIME-C.*

Proof. We can solve the problem in EXPTIME thanks to the AMT construction of polynomial size given in lemma 34 and the fact that emptiness of AMT is solvable in EXPTIME. For hardness, it is easy to reduce the problem of deciding the winner in a zero-sum reachability game with imperfect information to our problem, and this problem is complete for EXPTIME [13]. ◀

Acknowledgements. We would like to thank Marie Van den Bogaard for carefully reading and commenting on a previous version of this paper.

References

- 1 Martín Abadi, Leslie Lamport, and Pierre Wolper. Realizable and unrealizable specifications of reactive systems. In *ICALP'89*, volume 372 of *LNCS*. Springer, 1989.
- 2 Brandenburger Adam, Friedenberg Amanda, H Jerome, et al. Admissibility in games. *Econometrica*, 2008.
- 3 Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
- 4 Nicolas Basset, Gilles Geeraerts, Jean-François Raskin, and Ocan Sankur. Admissibility in concurrent games. *CoRR*, abs/1702.06439, 2017. URL: <http://arxiv.org/abs/1702.06439>.
- 5 Raphael Berthon, Bastien Maubert, Aniello Murano, Sasha Rubin, and Moshe Vardi. Strategy logic with imperfect information. In *LICS'14*. IEEE, 2014.
- 6 Dietmar Berwanger. Admissibility in infinite games. In *STACS 2007, 24th Annual Symposium on Theoretical Aspects of Computer Science, Aachen, Germany, February 22-24, 2007, Proceedings*, volume 4393 of *Lecture Notes in Computer Science*, pages 188–199. Springer, 2007.

- 7 Dietmar Berwanger, Anup Basil Mathew, and Marie van den Bogaard. Hierarchical information patterns and distributed strategy synthesis. In *Automated Technology for Verification and Analysis - 13th International Symposium, ATVA 2015, Shanghai, China, October 12-15, 2015, Proceedings*, volume 9364 of *Lecture Notes in Computer Science*, pages 378–393. Springer, 2015.
- 8 Felix Brandt, Markus Brill, Felix Fischer, and Paul Harrenstein. On the complexity of iterated weak dominance in constant-sum games. *Theory of Computing Systems*, 49(1):162–181, 2011. doi:10.1007/s00224-010-9282-7.
- 9 Romain Brenguier, Lorenzo Clemente, Paul Hunter, Guillermo A. Pérez, Mickael Randour, Jean-François Raskin, Ocan Sankur, and Mathieu Sassolas. Non-zero sum games for reactive synthesis. In *Language and Automata Theory and Applications - 10th International Conference, LATA 2016, Prague, Czech Republic, March 14-18, 2016, Proceedings*, volume 9618 of *Lecture Notes in Computer Science*, pages 3–23. Springer, 2016.
- 10 Romain Brenguier, Guillermo A. Pérez, Jean-François Raskin, and Ocan Sankur. Admissibility in quantitative graph games. In *36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2016, December 13-15, 2016, Chennai, India*, volume 65 of *LIPICs*, pages 42:1–42:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
- 11 Romain Brenguier, Jean-François Raskin, and Ocan Sankur. Assume-admissible synthesis. *Acta Inf.*, 54(1):41–83, 2017. doi:10.1007/s00236-016-0273-2.
- 12 Romain Brenguier, Jean-François Raskin, and Mathieu Sassolas. The complexity of admissibility in omega-regular games. In *CSL-LICS '14, 2014*. ACM, 2014. doi:10.1145/2603088.2603143.
- 13 Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, and Jean-François Raskin. Algorithms for omega-regular games with imperfect information. *Logical Methods in Computer Science*, 3(4), 2007.
- 14 Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010.
- 15 Vincent Conitzer and Tuomas Sandholm. Complexity of (iterated) dominance. In *EC '05: Proceedings of the 6th ACM conference on Electronic commerce*, pages 88–97, New York, NY, USA, 2005. ACM. doi:10.1145/1064009.1064019.
- 16 Werner Damm and Bernd Finkbeiner. Automatic compositional synthesis of distributed systems. In *FM 2014*, volume 8442 of *LNCS*, pages 179–193. Springer, 2014.
- 17 Luca de Alfaro, Thomas A. Henzinger, and Orna Kupferman. Concurrent reachability games. In *39th Annual Symposium on Foundations of Computer Science, FOCS '98, November 8-11, 1998, Palo Alto, California, USA*, pages 564–575. IEEE Computer Society, 1998.
- 18 Marco Faella. Admissible strategies in infinite games over graphs. In *MFCS 2009*, volume 5734 of *Lecture Notes in Computer Science*, pages 307–318. Springer, 2009.
- 19 Daniel Kirsten. Alternating tree automata and parity games. In *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume 2500 of *Lecture Notes in Computer Science*, pages 153–167. Springer, 2001.
- 20 Donald Knuth, Christos Papadimitriou, and John Tsitsiklis. A note on strategy elimination in bimatrix games. *Operations Research Letters*, 7(3):103–107, 1988. doi:10.1016/0167-6377(88)90075-2.
- 21 Orna Kupferman and Moshe Y. Vardi. Synthesis with incomplete informatio. In Howard Barringer, Michael Fisher, Dov Gabbay, and Graham Gough, editors, *Advances in Temporal Logic*, pages 109–127, Dordrecht, 2000. Springer Netherlands. doi:10.1007/978-94-015-9586-5_6.

- 22 Fabio Mogavero, Aniello Murano, and Moshe Y. Vardi. Reasoning about strategies. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2010, December 15-18, 2010, Chennai, India*, volume 8 of *LIPICs*, pages 133–144. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2010.
- 23 Arno Pauly. The computational complexity of iterated elimination of dominated strategies. *Theory of Computing Systems*, pages 52–75, 2016. doi:10.1007/s00224-015-9637-1.
- 24 Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190, 1989.
- 25 John H. Reif. The complexity of two-player games of incomplete information. *J. Comput. Syst. Sci.*, 29(2):274–301, 1984. doi:10.1016/0022-0000(84)90034-5.