

Algorithms for Omega-Regular Games with Incomplete Information

*Krishnendu Chatterjee
Laurent Doyen
Thomas A. Henzinger
Jean-Francois Raskin*

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2006-89

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-89.html>

June 26, 2006



Copyright © 2006, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

This research was supported in part by the NSF grants CCR-0225610 and CCR-0234690, and by the SNSF under the Indo-Swiss Joint Research Programme.

Algorithms for Omega-Regular Games with Incomplete Information ^{*}

Krishnendu Chatterjee ¹ Laurent Doyen ²
Thomas A. Henzinger ^{1,3}
Jean-François Raskin ²

¹ EECS, University of California at Berkeley, U.S.A.

² CS, Université Libre de Bruxelles, Belgium

³ CCS, École Polytechnique Fédérale de Lausanne, Switzerland

June 26, 2006

Abstract

We study observation-based strategies for two-player turn-based games on graphs with omega-regular objectives. An observation-based strategy relies on incomplete information about the history of a play, namely, on the past sequence of observations. Such games occur in the synthesis of a controller that does not see the private state of the plant. Our main results are twofold. First, we give a fixpoint algorithm for computing the set of states from which a player can win with a deterministic observation-based strategy for any omega-regular objective. The fixpoint is computed in the lattice of antichains of state sets. This algorithm has the advantages of being directed by the objective and of avoiding an explicit subset construction on the game graph. Second, we give an algorithm for computing the set of states from which a player can win with probability 1 with a randomized observation-based strategy for a Büchi objective. This set is of interest because in the absence of complete information, randomized strategies are more powerful than deterministic ones. We show that our algorithms are optimal by proving matching lower bounds.

^{*}This research was supported in part by the NSF grants CCR-0225610 and CCR-0234690, and by the SNSF under the Indo-Swiss Joint Research Programme.

1 Introduction

Two-player games on graphs play an important role in computer science. In particular, the *controller synthesis* problem asks, given a model for a plant, to construct a model for a controller such that the behaviors resulting from the parallel composition of the two models respects a given specification (e.g., are included in an ω -regular set). Controllers can be synthesized as winning strategies in a graph game whose vertices represent the plant states, and whose players represent the plant and the controller [17, 18]. Other applications of graph games include realizability and compatibility checking, where the players represent parallel processes of a system, or its environment [1, 11, 6].

Most results about two-player games played on graphs make the hypothesis of *complete information*. In this setting, the controller knows, during its interaction with the plant, the exact state of the plant. In practice, this hypothesis is often not reasonable. For example, in the context of hybrid systems, the controller acquires information about the state of the plant using sensors with finite precision, which return incomplete information about the state. Similarly, if the players represent individual processes, then a process has only access to the public variables of the other processes, not to their private variables [19, 2].

Two-player games of *incomplete information* are considerably more complicated than games of complete information. First, decision problems for incomplete-information games usually lie in higher complexity classes than their complete-information counterparts [19, 14, 2]. The algorithmic difference is often exponential, due to a subset construction that, similar to the determinization of finite automata, turns an incomplete-information game into an equivalent complete-information game. Second, because of the determinization, no symbolic algorithms are known to solve incomplete-information games. This is in contrast to the complete-information case, where (often) simple and elegant fixed-point algorithms exist [12, 8]. Third, in the context of incomplete information, deterministic strategies are sometimes insufficient. A game is *turn-based* if in every state one of the players chooses a successor state. While deterministic strategies suffice to win turn-based games of complete information, turn-based games of incomplete information require randomized strategies to win with probability 1. Fourth, winning strategies for incomplete-information games need memory even for simple objectives such as safety and reachability. This is again in contrast to the complete-information case, where turn-based safety and reachability games can be won with memoryless strategies. We will illustrate the needs

for randomization and memory with an example.

The contributions of this paper are two-fold. First, we provide a symbolic fixed-point algorithm to solve games of incomplete information for arbitrary ω -regular objectives. The novelty is that our algorithm is symbolic; it does not carry out an explicit subset construction. Instead, we compute fixed points on the lattice of antichains of state sets. Antichains of state sets can be seen as a symbolic and compact representation for \subseteq -downward-closed sets of sets of states.¹ This solution extends our recent result [10] from safety objectives to all ω -regular objectives. To justify the correctness of the algorithm, we transform games of incomplete information into games of complete information while preserving the existence of winning strategies for every Borel objective. The reduction is only part of the proof, not part of the algorithm. For the special case of parity objectives, we obtain a symbolic EXPTIME algorithm for solving parity games of incomplete information. This is optimal as the reachability problem for games of incomplete information is known to be EXPTIME-hard [19].

Second, we study randomized strategies and winning with probability 1 for incomplete-information games. To our knowledge, for these games no algorithms (symbolic or not) are present in the literature. Following [7], we refer to winning with probability 1 as *almost-sure* winning (*almost* winning, for short), in contrast to *sure* winning with deterministic strategies. We provide a symbolic EXPTIME algorithm to compute the set of almost-winning states for games of incomplete information with Büchi objectives (reachability objectives can be obtained as a special case and for safety objectives almost-winning and sure-winning coincide). Our solution is again justified by a reduction to games of complete information. However, for randomized strategies the reduction is different, and considerably more complicated. We prove our algorithm to be optimal, showing that computing the almost-winning states for reachability games of incomplete information is EXPTIME-hard.

The structure of the paper is simple: Section 2 presents the definitions; Section 3 gives the algorithm for the case of sure winning with deterministic strategies; Section 4, for the case of almost winning with randomized strategies; and Section 5 provides the lower bounds.

¹We have recently used this symbolic representation of \subseteq -downward-closed sets of state sets to propose a new algorithm to solve the universality problem for nondeterministic finite automata. First experiments show very promising performances, see [9] for details.

Related work. In [17], Pnueli and Rosner study the synthesis of reactive modules. In their framework, there is no game graph; instead, the environment and the objective are specified using an LTL formula. In [14], Kupferman and Vardi extend these results in two directions: they consider CTL* objectives and incomplete information. Again, no game graph, but a specification formula is given to the synthesis procedure. We believe that our setting, where a game graph is given explicitly, is more suited to fully and uniformly understand the role of incomplete information. Indeed, Kupferman and Vardi assert that incomplete information comes at no cost, because if the specification is given as a CTL (or CTL*) formula, then the synthesis problem is complete for EXPTIME (resp., 2EXPTIME), just as in the complete-information case. These hardness results, however, depend on the fact that the specification is given compactly as a formula. In our setting, with an explicit game graph, reachability games of complete information are known to be PTIME-complete, whereas we know EXPTIME-completeness for reachability games of incomplete information. In [19], Reif studies the complexity of sure winning for graph games of incomplete information with simple reachability objectives. The solution proposed by Reif comes in the form of an explicit subset construction. The hardness proof uses private and blind alternating turing machines. None of the above works provide symbolic solutions, and none of them consider randomized strategies.

It is known that for Partially Observable Markov Decision Processes (POMDPs) with boolean rewards and limit-average objectives the quantitative analysis (whether the value is greater than a specified value) is EXPTIME-complete, see [15]. However, the almost-winning is a qualitative question and the hardness result for almost-winning for incomplete information games does not follow from those results on POMDPs. We propose in section 5 an alternative proof of hardness for games of incomplete information with reachability objectives and we show that this result extends to almost winning conditions as well. To the best of our knowledge, this is the first hardness result that applies to qualitative analysis of almost-winning in incomplete information games.

A sub-class of semi-perfect (or semi-complete) information games (where only one player has incomplete information and the other player has complete information) is studied in [4]. The class of games in [4] is simpler than the class of games we study and can be solved in $\text{NP} \cap \text{coNP}$ for parity objectives.

2 Definitions

Games. A *game structure of incomplete information* is a tuple $G = \langle L, l_0, \Sigma, \Delta, \mathbf{Obs}, \gamma \rangle$ where L is a finite set of states, $l_0 \in L$ is the initial state, Σ is a finite alphabet, $\Delta \subseteq L \times \Sigma \times L$ is a set of labeled transitions, \mathbf{Obs} is a finite set of observations, and $\gamma : \mathbf{Obs} \rightarrow 2^L \setminus \emptyset$ maps each observation to the set of states that it represents. Furthermore, we require the following properties on G : (i) for all $\ell \in L$, for all $\sigma \in \Sigma$, there exists $\ell' \in L$: $(\ell, \sigma, \ell') \in \Delta$; (ii) the set of sets $\{\gamma(o) \mid o \in \mathbf{Obs}\}$ partitions L . The *size* of G is $|L|$. We say that $G = \langle L, l_0, \Sigma, \Delta, \mathbf{Obs}, \gamma \rangle$ is a game structure of *complete information* if $\mathbf{Obs} = L$ and $\gamma(\ell) = \{\ell\}$ for all $\ell \in L$. We often omit (\mathbf{Obs}, γ) in the description of games of complete information. For $\sigma \in \Sigma$ and $s \subseteq L$ let $\text{Post}_\sigma^G(s) = \{\ell' \in L \mid \exists \ell \in s : (\ell, \sigma, \ell') \in \Delta\}$.

Plays and observations. In a game structure, in each turn, Player 1 chooses a letter in Σ and Player 2 resolves nondeterminism by choosing the successor. A *play* in G is an infinite sequence $\pi = \ell_0 \sigma_0 \ell_1 \dots \sigma_{n-1} \ell_n \sigma_n \dots$ such that (i) $\ell_0 = l_0$, and (ii) for all $i \geq 0$: $(\ell_i, \sigma_i, \ell_{i+1}) \in \Delta$. The set of plays in G is noted $\text{Play}(G)$. Fix a play $\pi = \ell_0 \sigma_0 \ell_1 \dots \sigma_{n-1} \ell_n \sigma_n \dots$ in G . The *prefix up to ℓ_n* of π is denoted by $\pi(n)$, i.e., it is the finite sequence $\ell_0 \sigma_0 \ell_1 \dots \sigma_{n-1} \ell_n$. its *length* is $|\pi(n)| = n + 1$ and its *last* element is $\text{Last}(\pi(n)) = \ell_n$. The *observation* of π is the unique infinite sequence $\text{Obs}(\pi) = o_0 \sigma_0 o_1 \dots \sigma_{n-1} o_n \sigma_n \dots$ such that for all $i \geq 0$, $\ell_i \in \gamma(o_i)$. The *observation* of $\pi(n)$ is the prefix up to o_n of $\text{Obs}(\pi)$. We define similarly the *observation* of prefixes. The set of plays in G is noted $\text{Play}(G)$ and the set of prefixes in G is noted $\text{Pref}(G)$. A state $\ell \in L$ is *reachable* in G if there exists a prefix $\rho \in \text{Pref}(G)$ such that $\text{Last}(\rho) = \ell$. Let $\rho = \ell_0 \sigma_0 \ell_1 \dots \sigma_{n-1} \ell_n \in \text{Pref}(G)$, the *observation* of ρ , noted $\text{Obs}(\rho)$, is the unique sequence $o_0 \sigma_0 o_1 \dots \sigma_{n-1} o_n$ such that for all i , $0 \leq i \leq n$, $\ell_i \in \gamma(o_i)$. Given $\rho \in \text{Pref}(G)$, we define $[\rho]_{\text{Obs}}$ as the set of prefixes of plays in G that have the same observation as ρ : $[\rho]_{\text{Obs}} = \{\rho' \mid \rho' \in \text{Pref}(G) \text{ and } \text{Obs}(\rho) = \text{Obs}(\rho')\}$. The *knowledge* associated to a sequence of observations $\tau = o_0 \sigma_0 o_1 \sigma_1 \dots \sigma_{n-1} o_n$, noted $\text{K}(\tau)$, is the set of states in which the play can be after this sequence of observations: $\text{K}(\tau) = \{\text{Last}(\rho) \mid \rho \in \text{Pref}(G) \text{ and } \text{Obs}(\rho) = \tau\}$.

Given $\rho \in \text{Pref}(G)$, we define $[\rho]_{\text{K}}$ as the set of prefixes of plays in G that leads to the same knowledge: $[\rho]_{\text{K}} = \{\rho_1 \in \text{Pref}(G) \mid \text{K}(\text{Obs}(\rho_1)) = \text{K}(\text{Obs}(\rho))\}$.

Lemma 1 Let $G = \langle L, l_0, \Sigma, \Delta, \text{Obs}, \gamma \rangle$ be a game structure of incomplete information. For $\sigma \in \Sigma$, $\ell \in L$ and $\rho, \rho' \in \text{Pref}(G)$ with $\rho' = \rho \cdot \sigma \cdot \ell$, let $\text{obs}_\ell \in \text{Obs}$ be the unique observation such that $\ell \in \gamma(\text{obs}_\ell)$. Then we have $\text{K}(\text{Obs}(\rho')) = \text{Post}_\sigma^G(\text{K}(\text{Obs}(\rho))) \cap \gamma(\text{obs}_\ell)$.

Strategies. Let $G = \langle L, l_0, \Sigma, \Delta, \text{Obs}, \gamma \rangle$ be a game structure of incomplete information. A *deterministic strategy* in G for Player 1 is a function $\alpha : \text{Pref}(G) \rightarrow \Sigma$. For a finite set A , a *probability distribution* on A is a function $\kappa : A \rightarrow [0, 1]$ such that $\sum_{a \in A} \kappa(a) = 1$. We denote the set of probability distributions on A by $\mathcal{D}(A)$. Given a distribution $\kappa \in \mathcal{D}(A)$, let $\text{Supp}(\kappa) = \{a \in A \mid \kappa(a) > 0\}$ the *support* of κ . A *randomized strategy* in G for Player 1 is a function $\alpha : \text{Pref}(G) \rightarrow \mathcal{D}(\Sigma)$. A (deterministic or randomized) strategy α for Player 1 is said to be *observation-based* if for all prefixes $\rho, \rho' \in \text{Pref}(G)$, if $\text{Obs}(\rho) = \text{Obs}(\rho')$ then $\alpha(\rho) = \alpha(\rho')$. In the sequel, we are interested in the existence of observation-based strategies for Player 1. A *deterministic strategy* in G for Player 2 is a function $\beta : \text{Pref}(G) \times \Sigma \rightarrow L$ such that $\forall \rho \in \text{Pref}(G) \cdot \forall \sigma \in \Sigma : (\text{Last}(\rho), \sigma, \beta(\rho, \sigma)) \in \Delta$. A *randomized strategy* in G for Player 2 is a function $\beta : \text{Pref}(G) \times \Sigma \rightarrow \mathcal{D}(L)$ such that $\forall \rho \in \text{Pref}(G) \cdot \forall \sigma \in \Sigma \cdot \forall \ell \in \text{Supp}(\beta(\rho, \sigma)) : (\text{Last}(\rho), \sigma, \ell) \in \Delta$. We denote by $\mathcal{A}_G, \mathcal{A}_G^O$ and \mathcal{B}_G the set of all Player 1 strategies, the set of all observation-based Player 1 strategies, and the set of all Player 2 strategies, respectively. All the results of the paper can also be proved if strategies are mappings from sequence of states only (without the moves).

The *outcome* of two deterministic strategies α (for Player 1) and β (for Player 2) in G is the play $\pi = \ell_0 \sigma_0 \ell_1 \dots \sigma_{n-1} \ell_n \sigma_n \dots \in \text{Play}(G)$ such that for all $i \geq 0$, we have $\sigma_i = \alpha(\pi(i))$ and $\ell_{i+1} = \beta(\pi(i), \sigma_i)$. This play is noted $\text{outcome}(G, \alpha, \beta)$. The *outcome* of two randomized strategies α (for Player 1) and β (for Player 2) in G is the set of plays $\pi = \ell_0 \sigma_0 \ell_1 \dots \sigma_{n-1} \ell_n \sigma_n \dots \in \text{Play}(G)$ such that for all $i \geq 0$, we have $\alpha(\pi(i))(\sigma_i) > 0$ and $\beta(\pi(i), \sigma_i)(\ell_{i+1}) > 0$. The *outcome set* of the deterministic (resp. randomized) strategy α for Player 1 in G is the set $\text{Outcome}_1(G, \alpha)$ of plays $\text{outcome}(G, \alpha, \beta)$ such that β is a deterministic (resp. randomized) strategy for Player 2. Similarly, the *outcome set* of the deterministic (resp. randomized) strategy β for Player 2 in G is the set $\text{Outcome}_2(G, \beta)$ of plays $\text{outcome}(G, \alpha, \beta)$ such that α is a deterministic (resp. randomized) strategy for Player 1.

Objectives. An *objective* in G is a set ϕ of infinite sequences of observations and input letters, *i.e.* $\phi \subseteq (\text{Obs} \times \Sigma)^\omega$. A play $\pi =$

$\ell_0\sigma_0\ell_1\dots\sigma_{n-1}\ell_n\sigma_n\dots \in \text{Play}(G)$ satisfies an objective ϕ (noted $\pi \models \phi$) iff $\text{Obs}(\pi) \in \phi$. Objectives are generally Borel measurable; a Borel objective is a Borel set in the Cantor topology on $(\text{Obs} \times \Sigma)^\omega$ [13]. We also consider reachability, safety, (open and closed sets in the Cantor topology on $(\text{Obs} \times \Sigma)^\omega$, respectively) Büchi, coBüchi and parity objectives. The parity objectives are a canonical form to express all ω -regular objectives [20]. For a play $\pi = \ell_0\sigma_0\ell_1\dots$ we denote by $\text{Inf}(\pi)$ the set of observations that appear infinitely often in $\text{Obs}(\pi)$, i.e., $\text{Inf}(\pi) = \{o \in \text{Obs} \mid l_i \in \gamma(o) \text{ for infinitely many } i\}$.

- *Reachability and safety objectives.* Given a set of observations $O \subseteq \text{Obs}$, the *reachability* objective $\text{Reach}(O)$ requires that an observation in O is visited at least once, i.e. $\text{Reach}(O) = \{\pi = \ell_0\sigma_0\ell_1\sigma_1\dots \in \text{Play}(G) \mid \exists k \geq 0. \exists o \in O. \ell_k \in \gamma(o)\}$. Dually, the *safety* objective $\text{Safe}(O)$ requires that only observations in O be visited. Formally we have $\text{Safe}(O) = \{\pi = \ell_0\sigma_0\ell_1\sigma_1\dots \in \text{Play}(G) \mid \forall k \geq 0. \exists o \in O : \ell_k \in \gamma(o)\}$.
- *Büchi and coBüchi objectives.* The *Büchi* objective $\text{Buchi}(O)$ requires that an observation in O is visited infinitely often, i.e. $\text{Buchi}(O) = \{\pi \mid \text{Inf}(\pi) \cap O \neq \emptyset\}$. Dually, the *coBüchi* objective $\text{coBuchi}(O)$ requires that only observations in O be visited infinitely often. Formally we have $\text{coBuchi}(O) = \{\pi \mid \text{Inf}(\pi) \subseteq O\}$.
- *Parity objectives.* For $d \in \mathbb{N}$, let $p : \text{Obs} \rightarrow \{0, 1, \dots, d\}$ be a priority function that maps each observation to a non-negative integer. The *parity* objective $\text{Parity}(p)$ requires that the minimum priority that appears infinitely often is even. Formally we have $\text{Parity}(p) = \{\pi \mid \min\{p(O) \mid O \in \text{Inf}(\pi)\} \text{ is even}\}$.

Observe that by definition for all objectives ϕ , if $\pi \models \phi$ and $\text{Obs}(\pi) = \text{Obs}(\pi')$, then $\pi' \models \phi$.

Given a game structure of incomplete information G and a state ℓ of G , we write G_ℓ for the game structure of incomplete information that results from G by changing the initial state to ℓ : that is, if $G = \langle L, l_0, \Sigma, \Delta, \text{Obs}, \gamma \rangle$, then $G_\ell = \langle L, \ell, \Sigma, \Delta, \text{Obs}, \gamma \rangle$.

An *event* is a measurable set of paths (or a Borel objective), and given two strategies α and β , the probabilities of events are uniquely defined [21]. We denote by $\text{Pr}_\ell^{\alpha, \beta}(\phi)$ the probability that an objective ϕ is satisfied in the game G_ℓ given strategies α and β .

Sure winning and almost winning. A strategy λ_i for Player i in G is *sure-winning* for an objective ϕ iff for all $\pi \in \text{Outcome}_i(G, \lambda_i)$, we have $\pi \models \phi$. A strategy α for Player 1 is *almost-winning* for an objective ϕ , if for all randomized strategies β for Player 2 we have $\text{Pr}_{l_0}^{\alpha, \beta}(\phi) = 1$. The set of *sure-winning* (resp. *almost-winning*) *states* of a game structure of incomplete information G for the objective ϕ is the set of states ℓ such that Player 1 has a deterministic sure-winning (resp. randomized almost-winning) observation-based strategy in G_ℓ for the objective ϕ .

Theorem 1 (Determinacy [16]). *For all Borel objectives ϕ , for all complete-information game structures G either there exists a deterministic sure-winning strategy for Player 1 for objective ϕ , or there exists a deterministic sure-winning strategy for Player 2 for objective $\text{Play}(G) \setminus \phi$.*

Notice that deterministic strategies suffice for sure-winning a game: given a randomized strategy α for Player 1, let α^D be a deterministic strategy such that for all $\rho \in \text{Pref}(G)$, $\alpha^D(\rho)$ chooses an input letter from $\text{Supp}(\alpha(\rho))$. Clearly $\text{Outcome}_1(G, \alpha^D) \subseteq \text{Outcome}_1(G, \alpha)$ and hence if α is sure-winning, then so is α^D . The result also holds for observation-based strategies and for complete-information games. Hence deterministic strategies suffices for sure-winning and also for complete-information games randomized strategies are no more powerful than deterministic strategies. However, for almost-winning, randomized strategies are more powerful than deterministic strategies as shown by Example 1

Example 1 *Consider the game structure shown in Figure 1. The observations o_1, o_2, o_3, o_4 are such that $\gamma(o_1) = \{\ell_1\}$, $\gamma(o_2) = \{\ell_2, \ell'_2\}$, $\gamma(o_3) = \{\ell_3, \ell'_3\}$, $\gamma(o_4) = \{\ell_4\}$. The transitions are shown as labeled edges in the figure and the initial state is ℓ_1 . The objective of Player 1 is $\text{Reach}(\{\ell_4\})$, i.e. to reach state ℓ_4 . We argue that the game is not sure-winning for Player 1. Let α^D be any deterministic strategy for Player 1. Consider a strategy β for Player 2 as follows: for all $\rho \in \text{Pref}(G)$ such that $\text{Last}(\rho) \in \gamma(o_2)$, if $\alpha^D(\rho) = a$, then in the previous round β chooses the state ℓ_2 , and if $\alpha^D(\rho) = b$, then in the previous round β chooses the state ℓ'_2 . Given α^D and β , $\text{outcome}(G, \alpha^D, \beta)$ never reaches ℓ_4 . However we show that the game G is almost-winning for Player 1. Consider the randomized strategy that plays a and b uniformly at random at all states. Every time the game visits observation o_2 , for any strategy for Player 2, the game visits ℓ_3 and ℓ'_3 with probability $\frac{1}{2}$, and hence also reaches ℓ_4 with probability $\frac{1}{2}$. It easily follows that against all Player 2 strategies the play eventually reaches ℓ_4 with probability 1.*




Figure 1: Game structure G .

Spoiling strategies. To spoil a strategy of Player 1 (for sure-winning), Player 2 does not need the full memory of the history of the play, he only needs counting strategies. We say that a deterministic strategy $\beta : \text{Pref}(G) \times \Sigma \rightarrow L$ for Player 2 is *counting* if for all prefixes $\rho, \rho' \in \text{Pref}(G)$ such that $|\rho| = |\rho'|$ and $\text{Last}(\rho) = \text{Last}(\rho')$, and for all $\sigma \in \Sigma$, we have $\beta(\rho, \sigma) = \beta(\rho', \sigma)$. Let \mathcal{B}_G^c be the set of counting strategies for Player 2. The memory needed by a counting strategy is only the number of turns that have been played. This type of strategy is sufficient to spoil the non-winning strategies of Player 1.

Proposition 1 *Let G be a game structure of incomplete information and ϕ be an objective. There exists an observation-based strategy $\alpha^o \in \mathcal{A}_G^o$ such that for all $\beta \in \mathcal{B}_G$ we have $\text{outcome}(G, \alpha^o, \beta) \in \phi$ if and only if there exists an observation-based strategy $\alpha^o \in \mathcal{A}_G^o$ such that for all counting strategies $\beta^c \in \mathcal{B}_G^c$ we have $\text{outcome}(G, \alpha^o, \beta^c) \in \phi$.*

Proof. We prove the equivalent statement that: $\forall \alpha^o \in \mathcal{A}_G^o \cdot \exists \beta \in \mathcal{B}_G : \text{outcome}(G, \alpha^o, \beta) \notin \phi$ iff $\forall \alpha^o \in \mathcal{A}_G^o \cdot \exists \beta^c \in \mathcal{B}_G^c : \text{outcome}(G, \alpha^o, \beta^c) \notin \phi$. The right implication (\leftarrow) is trivial. For the left implication (\rightarrow), let $\alpha^o \in \mathcal{A}_G^o$ be an arbitrary observation-based strategy for Player 1 in G . Let $\beta \in \mathcal{B}_G$ be a strategy for Player 2 such that $\text{outcome}(G, \alpha^o, \beta) \notin \phi$. Let $\text{outcome}(G, \alpha^o, \beta) = l_0 \sigma_0 l_1 \dots \sigma_{n-1} l_n \sigma_n \dots$ and define a counting strategy β^c for Player 2 such that $\forall \rho \in \text{Pref}(G) \cdot \forall \sigma \in \Sigma : \text{if } \text{Last}(\rho) = l_{n-1} \text{ and } \sigma = \sigma_{n-1} \text{ for } n = |\rho|, \text{ then } \beta^c(\rho, \sigma) = l_n, \text{ and otherwise } \beta^c(\rho, \sigma) \text{ is fixed arbitrarily in the set } \text{Post}_\sigma^G(\text{Last}(\rho))$. Clearly, β^c is a counting strategy and we have $\text{outcome}(G, \alpha^o, \beta) = \text{outcome}(G, \alpha^o, \beta^c)$ and thus $\text{outcome}(G, \alpha^o, \beta^c) \notin \phi$. \blacksquare

3 Sure Winning

First, we show that a game structure of incomplete information G can be encoded in a game structure of complete information G^K such that for any Borel objective ϕ , there exists a deterministic observation-based sure-winning strategy for Player 1 in G for ϕ if and only if there exists a deterministic sure-winning strategy for Player 1 in G^K for ϕ . We obtain G^K using a subset construction. A state in G^K is a set of states of G that corresponds to the knowledge of Player 1. In the worst case, the size of G^K is exponentially larger than the size of G . Second, we present a fixed point algorithm based on antichains of set of states [10] whose correctness relies on our subset construction, but that avoids the explicit construction of G^K .

3.1 Subset construction for sure winning

Subset construction. Given a game structure of incomplete information $G = (L, l_0, \Sigma, \Delta, \text{Obs}, \gamma)$, we define the *knowledge-based subset construction* of G as the following game structure of complete information:

$$G^K = \langle \mathcal{L}, \{l_0\}, \Sigma, \Delta^K \rangle$$

where $\mathcal{L} = 2^L \setminus \{\emptyset\}$ and $(s_1, \sigma, s_2) \in \Delta^K$ iff $\exists \text{obs} \in \text{Obs}: s_2 = \text{Post}_\sigma^G(s_1) \cap \gamma(\text{obs})$ and $s_2 \neq \emptyset$. Notice that for all $s \in \mathcal{L}$, for all $\sigma \in \Sigma$, there exists $s' \in \mathcal{L}: (s, \sigma, s') \in \Delta^K$.

A (deterministic or randomized) strategy in G^K is called a *knowledge-based* strategy. To distinguish between a general strategy in G , an observation-based strategy in G and a knowledge-based strategy in G^K , we often use the notations α, α^o , and α^K respectively.

Lemma 2 *For all set $s \in \mathcal{L}$ that is reachable in G^K , for all observation $\text{obs} \in \text{Obs}$, we have either $s \subseteq \gamma(\text{obs})$ or $s \cap \gamma(\text{obs}) = \emptyset$.*

Proof. First, the property holds for $s = \{l_0\}$, the initial state in G^K as it is a singleton. Second, we show that the property holds for any successor s' of any state s in G^K . Assume that $(s, \sigma, s') \in \Delta^K$. Then we know that $s' = \text{Post}_\sigma^G(s) \cap \gamma(\text{obs})$ for some $\text{obs} \in \text{Obs}$. Hence, $s' \subseteq \gamma(\text{obs})$ and $s' \cap \gamma(\text{obs}') = \emptyset$ for any $\text{obs}' \neq \text{obs}$ since the set $\{\gamma(\text{obs}) \mid \text{obs} \in \text{Obs}\}$ partitions L . ■

By an abuse of notation, we define the *observation* of a play $\pi = s_0 \sigma_0 s_1 \dots \sigma_{n-1} s_n \sigma_n \dots \in \text{Play}(G^K)$ as the sequence $\text{Obs}(\pi) = o_0 \sigma_0 o_1 \dots \sigma_{n-1} o_n \sigma_n \dots$ such that for all $i \geq 0$, $s_i \subseteq \gamma(o_i)$. This sequence is

unique by Lemma 2. We also say that π *satisfies* an objective $\phi \subseteq (\text{Obs} \times \Sigma)^\omega$ iff $\text{Obs}(\pi) \in \phi$. Lemma 2 justifies the following definition:

Definition 1 For a play $\pi = s_0\sigma_0s_1 \dots \sigma_{n-1}s_n\sigma_n \dots \in \text{Play}(G^K)$, define the *observation* of π as the unique sequence $\text{Obs}(\pi) = o_0\sigma_0o_1 \dots \sigma_{n-1}o_n\sigma_n \dots$ such that for any $i \geq 0$, $s_i \subseteq \gamma(o_i)$. \square

As above, we say that a play $\pi = s_0\sigma_0s_1 \dots \sigma_{n-1}s_n\sigma_n \dots \in \text{Play}(G^K)$ *satisfies* an objective ϕ iff the sequence of observations $o_0o_1 \dots o_n \dots$ such that for all $i \geq 0$, $l_i \in \gamma(o_i)$ belongs to ϕ .

Lemma 3 *If Player 1 has a deterministic sure-winning strategy in G^K for an objective ϕ , then he has a deterministic observation-based sure-winning strategy in G for ϕ .*

Proof. Let α^K be a deterministic sure-winning strategy for Player 1 in G^K with the objective ϕ . Define α^o a strategy for Player 1 in G as follows: for every $\rho \in \text{Pref}(G)$, let $\alpha^o(\rho) = \alpha^K(\rho^K)$ where ρ^K is defined from $\rho = \ell_0\sigma_0\ell_1 \dots \sigma_{n-1}\ell_n$ by $\rho^K = s_0\sigma_0s_1 \dots \sigma_{n-1}s_n$ where $s_i = K(\text{Obs}(\ell_0\sigma_0\ell_1 \dots \sigma_{i-1}\ell_i))$ for each $0 \leq i \leq n$. Clearly, α^o is an deterministic observation-based strategy as $\text{Obs}(\rho) = \text{Obs}(\rho')$ implies $\rho^K = \rho'^K$.

By contradiction, assume that α^o is not a sure-winning strategy for Player 1 in G with the objective ϕ . Then there exists a play $\pi \in \text{Outcome}_1(G, \alpha^o)$ such that $\pi \not\models \phi$. Let $\pi = \ell_0\sigma_0\ell_1\sigma_1 \dots$ and consider the infinite sequence $\pi^K = s_0\sigma_0s_1\sigma_1 \dots$ where $s_i = K(\text{Obs}(\pi(i)))$ for each $i \geq 0$. We show that $\pi^K \in \text{Outcome}_1(G^K, \alpha^K)$. First, we have $s_0 = K(\text{Obs}(\pi(0))) = K(\text{Obs}(\ell_0)) = \{\ell_0\}$. Second, for any $i \geq 0$, we have $s_i = K(\text{Obs}(\pi(i)))$ and by Lemma 1 $s_{i+1} = \text{Post}_{\sigma_i}^G(s_i) \cap \gamma(\text{obs})$ where obs is such that $\ell_{i+1} \in \gamma(\text{obs})$ and so $(s_i, \sigma_i, s_{i+1}) \in \Delta^K$. Third, by definition of α^o , we have $\sigma_i = \alpha^o(\pi(i)) = \alpha^K(\pi^K(i))$. This entails that $\pi^K \in \text{Outcome}_1(G^K, \alpha^K)$.

Now, observe that trivially $\ell_i \in K(\text{Obs}(\pi(i)))$ for any $i \geq 0$, that is $\ell_i \in s_i$ and so $s_i \cap \gamma(\text{obs}_i) \neq \emptyset$ where obs_i is the unique observation such that $\ell_i \in \gamma(\text{obs}_i)$. Hence, by Lemma 2, we have $s_i \subseteq \gamma(\text{obs}_i)$. Consequently, $\text{Obs}(\pi^K) = \text{Obs}(\pi)$ and thus $\pi^K \not\models \phi$ which contradicts the fact that α^K is a sure-winning strategy for Player 1 in G^K with the objective ϕ . Therefore, α^o is a sure-winning strategy for Player 1 in G with the objective ϕ . \blacksquare

Lemma 4 *If Player 1 does not have a deterministic sure-winning strategy in G^K for a Borel objective ϕ , then Player 1 does not have a deterministic observation-based sure-winning strategy in G for ϕ .*

Proof. First, observe that if there exists no sure-winning strategy for Player 1 in G^K with objective ϕ , then by the Determinacy Theorem (Theorem 1), Player 2 has a spoiling strategy β^S such that for any play $\pi^K \in \text{Outcome}_2(G^K, \beta^S)$ we have $\pi^K \not\models \phi$. By contradiction, assume that Player 1 has an observation-based sure-winning strategy α^o in G for ϕ . We construct the tree $\mathsf{T}(\alpha^o, \beta^S) = \langle N, \Sigma, \Delta_T \rangle$ where $N = \bigcup_{i \geq 0} N_i$ and $\Delta_T = \bigcup_{i \geq 1} E_i$. The sets N_i ($i \geq 0$) and E_i ($i \geq 1$) are constructed inductively as follows: let $N_0 = \{\ell_0\}$ and $s_0 = \{\ell_0\}$. The following properties clearly hold for $i = 0$:

- (a) $s_i = \{\text{Last}(\rho) \mid \rho \in N_i\}$,
- (b) $N_i \subseteq \text{Pref}(G)$ and $s_0 \sigma_0 s_1 \sigma_1 \dots \sigma_{i-1} s_i \in \text{Pref}(G^K)$,
- (c) $\forall \rho, \rho' \in N_i : \text{Obs}(\rho) = \text{Obs}(\rho')$ and
- (d) $\forall \rho \in N_i : s_i = \mathsf{K}(\text{Obs}(\rho))$.

Assume that we have constructed the sets N_i and s_i for $i = 0, \dots, k$ and the labels σ_i for $i = 0, \dots, k-1$ such that (a), (b), (c), (d) hold and let us construct N_{k+1} , s_{k+1} and σ_k such that (a), (b), (c), (d) hold for $i = k+1$. Let

- $\rho \in N_k$ and $\tau = \text{Obs}(\rho)$,
- $\sigma_k = \alpha^o(\rho)$ (by (c), σ_k is independent of the choice of ρ since α^o is observation-based),
- $s_{k+1} = \beta^S(s_0 \sigma_0 s_1 \sigma_1 \dots \sigma_{k-1} s_k, \sigma_k)$,
- $N_{k+1} = \{\rho \sigma_k l' \mid \rho \in N_k \wedge l' \in s_{k+1} \wedge (\text{Last}(\rho), \sigma, l') \in \Delta\}$,
- $E_{k+1} = \{(\rho, \sigma_k, \rho \sigma_k l') \mid \rho \in N_k \wedge \rho \sigma_k l' \in N_{k+1}\}$.

Then, we have (i) $s_{k+1} \neq \emptyset$ (ii) $(s_k, \sigma_k, s_{k+1}) \in \Delta^K$ and thus $s_{k+1} \subseteq \text{Post}_{\sigma_k}^G(s_k)$ and (iii) $s_{k+1} \subseteq \gamma(\text{obs})$ for some $\text{obs} \in \text{Obs}$. We prove that:

- (a') $s_{k+1} = \{\text{Last}(\rho) \mid \rho \in N_{k+1}\}$. By definition of N_{k+1} , we have $\forall \rho \in N_{k+1} : \text{Last}(\rho) \in s_{k+1}$. By (ii), we have that $\forall l' \in s_{k+1} \cdot \exists l \in s_k : (l, \sigma_k, l') \in \Delta$ and by (a), this entails that $\forall l' \in s_{k+1} \cdot \exists \rho \in N_k : \rho \sigma_k l' \in N_{k+1}$, that is $s_{k+1} \subseteq \{\text{Last}(\rho) \mid \rho \in N_k\}$;
- (b') $N_{k+1} \subseteq \text{Pref}(G)$ and $s_0 \sigma_0 \dots \sigma_k s_{k+1} \in \text{Pref}(G^K)$. This follows from (b) and the definition of N_{k+1} and s_{k+1} ;

(c') $\forall \rho, \rho' \in N_{k+1} : \text{Obs}(\rho) = \text{Obs}(\rho')$. Let $\rho, \rho' \in N_{k+1}$. By definition of N_{k+1} and by (iii), we have $\text{Obs}(\rho) = \tau \sigma_k \text{obs}$ where τ is independent of ρ by (c). Therefore $\text{Obs}(\rho') = \tau \sigma_k \text{obs}$;

(d') $\forall \rho \in N_{k+1} : s_{k+1} = \mathsf{K}(\text{Obs}(\rho))$. This is established by Lemma 1.

Now, by (a) and (i), the tree $\mathsf{T}(\alpha^o, \beta^S)$ is infinite and by König's Lemma it contains an infinite sequence $\rho_0, \sigma_0, \rho_1, \sigma_1, \dots$ such that $\rho_0 = \{\ell_0\}$ and for all $i \geq 1$, $(\rho_{i-1}, \sigma_{i-1}, \rho_i) \in \Delta_T$. Correspondingly, let $\pi = \ell_0 \sigma_0 \ell_1 \sigma_1 \dots$ such that $\ell_i = \text{Last}(\rho_i)$ for each $i \geq 0$. By construction of $\mathsf{T}(\alpha^o, \beta^S)$, we have $\pi(i) \in \text{Pref}(G)$ and $\sigma_i = \alpha^o(\pi(i))$ for all $i \geq 0$. Hence $\pi \in \text{Outcome}_1(G, \alpha^o)$ and therefore $\pi \models \phi$. Consider the infinite sequence $\pi^K = s_0 \sigma_0 s_1 \sigma_1 \dots$ (the sets s_i have been constructed together with the sets N_i). By construction of $\mathsf{T}(\alpha^o, \beta^S)$, we have $s_{i+1} = \beta^S(s_0 \sigma_0 \dots \sigma_{i-1} s_i, \sigma_i)$ for all $i \geq 0$ and thus $\pi^K \in \text{Outcome}_2(G^K, \beta^S)$ and $\pi^K \not\models \phi$. However, by (a) and (d) we have $\text{Obs}(\pi) = \text{Obs}(\pi^K)$ which contradicts the fact that $\pi \models \phi$ and $\pi^K \not\models \phi$. Therefore, Player 1 cannot have a sure-winning observation-based strategy α^o in G for ϕ . ■

Lemma 3 and Lemma 4 yield Theorem 2.

Theorem 2 *Player 1 has a deterministic sure-winning strategy in G^K for a Borel objective ϕ if and only if Player 1 has a deterministic observation-based sure-winning strategy in G for ϕ .*

3.2 Two interpretations of the μ -calculus

From the results of Section 3.1, we could solve a game G of incomplete information with objective ϕ by constructing the knowledge-based subset construction G^K of G and solve the resulting game of complete information for the objective ϕ using standard methods. For the important class of ω -regular objectives, there exists a fixed point theory –the μ -calculus– for this purpose [8]. When run on G^K , those fixed point algorithms compute sets of sets of states of the game G . An important property of those sets is that they are *downward closed* for set inclusion: if Player 1 has a deterministic strategy to win the game G when her knowledge is the set s , then she also has a deterministic strategy to win the game when her knowledge is s' with $s' \subseteq s$. And thus, if s is a sure-winning state of G^K , then so is s' . Based on this property, we devise a new algorithm for solving games of incomplete information.

An *antichain* of nonempty sets of states is a set $q \subseteq 2^L \setminus \emptyset$ such that for all $s, s' \in q$, we have $s \not\subseteq s'$. Let \mathcal{A} be the set of antichains of nonempty subsets

of L and consider the following partial order on \mathcal{A} : for all $q, q' \in \mathcal{A}$, let $q \sqsubseteq q'$ iff $\forall s \in q : \exists s' \in q' : s \subseteq s'$. For $q \subseteq 2^L \setminus \emptyset$, define the set of *maximal* elements of q by $\lceil q \rceil = \{s \in q \mid s \neq \emptyset \wedge \nexists s' \in q : s \subset s'\}$. Clearly, $\lceil q \rceil$ is an antichain. The least upper bound of $q, q' \in \mathcal{A}$ is $q \sqcup q' = \lceil \{s \mid s \in q \vee s \in q'\} \rceil$ and their greatest lower bound is $q \sqcap q' = \lceil \{s \cap s' \mid s \in q \wedge s' \in q'\} \rceil$. The definition of those two operators extends naturally to sets of antichains, and the greatest element of \mathcal{A} is $\top = \{L\}$ and the least element is $\perp = \emptyset$. The partially ordered set $\langle \mathcal{A}, \sqsubseteq, \sqcup, \sqcap, \top, \perp \rangle$ forms a complete lattice. We see antichains of state sets as a symbolic representation of \sqsubseteq -downward closed sets of state sets.

A *game-lattice* is a complete lattice V together with a *predecessor operator* $\text{CPre} : V \rightarrow V$. Given a game structure $G = \langle L, l_0, \Sigma, \Delta, \text{Obs}, \gamma \rangle$ and its knowledge-based subset construction $G^K = \langle \mathcal{L}, \{l_0\}, \Sigma, \Delta^K \rangle$, we consider two game-lattices: the *lattice of subsets* $\langle \mathcal{S}, \sqsubseteq, \cup, \cap, \mathcal{L}, \emptyset \rangle$ (where $\mathcal{S} = 2^{\mathcal{L}}$) with the predecessor operator $\text{CPre} : \mathcal{S} \rightarrow \mathcal{S}$ defined by:

$$\text{CPre}(q) = \{s \in \mathcal{L} \mid \exists \sigma \in \Sigma \cdot \forall s' : \text{if } (s, \sigma, s') \in \Delta^K \text{ then } s' \in q\}$$

and the *lattice of antichains* $\langle \mathcal{A}, \sqsubseteq, \sqcup, \sqcap, \{L\}, \emptyset \rangle$ with the predecessor operator $\lceil \text{CPre} \rceil : \mathcal{A} \rightarrow \mathcal{A}$ defined by:

$$\lceil \text{CPre} \rceil(q) = \lceil \{s \in \mathcal{L} \mid \exists \sigma \in \Sigma \cdot \forall o \in \text{Obs} \cdot \exists s' \in q : \text{Post}_\sigma(s) \cap \gamma(o) \subseteq s'\} \rceil$$

The μ -calculus formulas are generated by the grammar:

$$\varphi ::= \text{obs} \mid x \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \text{pre}(\varphi) \mid \mu x. \varphi \mid \nu x. \varphi$$

for atomic propositions $\text{obs} \in \text{Obs}$ and variables x . We can define $\neg \text{obs}$ as a shortcut for $\bigvee_{o' \in \text{Obs} \setminus \{\text{obs}\}} o'$. A variable is *free* in a formula φ if it is not under the scope of a quantification μx or νx . A formula φ is *closed* if it contains no free variable. Given a game-lattice V , a *valuation* \mathcal{E} for the variables is a function that maps every variable x to an element in V . For $q \in V$, we write $\mathcal{E}[x \mapsto q]$ for the valuation that agrees with \mathcal{E} on all variables, except that x is mapped to q . Given a game-lattice V , and a valuation \mathcal{E} , each μ -calculus formula φ specifies an element $\llbracket \varphi \rrbracket_{\mathcal{E}}^V$ of V , which is defined inductively by the following equations:

Lattice of subsets
$\llbracket \text{obs} \rrbracket_{\mathcal{E}}^{\mathcal{S}} = \{s \in \mathcal{L} \mid s \subseteq \gamma(\text{obs})\}$
$\llbracket x \rrbracket_{\mathcal{E}}^{\mathcal{S}} = \mathcal{E}(x)$
$\llbracket \varphi_1 \left\{ \begin{smallmatrix} \vee \\ \wedge \end{smallmatrix} \right\} \varphi_2 \rrbracket_{\mathcal{E}}^{\mathcal{S}} = \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^{\mathcal{S}} \left\{ \begin{smallmatrix} \cup \\ \cap \end{smallmatrix} \right\} \llbracket \varphi_2 \rrbracket_{\mathcal{E}}^{\mathcal{S}}$
$\llbracket \text{pre}(\varphi) \rrbracket_{\mathcal{E}}^{\mathcal{S}} = \text{CPre}(\llbracket \varphi \rrbracket_{\mathcal{E}}^{\mathcal{S}})$
$\llbracket \left\{ \begin{smallmatrix} \mu \\ \nu \end{smallmatrix} \right\} x.\varphi \rrbracket_{\mathcal{E}}^{\mathcal{S}} = \left\{ \begin{smallmatrix} \cap \\ \cup \end{smallmatrix} \right\} \{q \mid q = \llbracket \varphi \rrbracket_{\mathcal{E}[x \mapsto q]}^{\mathcal{S}}\}$
Lattice of antichains
$\llbracket \text{obs} \rrbracket_{\mathcal{E}}^{\mathcal{A}} = \{\gamma(\text{obs})\}$
$\llbracket x \rrbracket_{\mathcal{E}}^{\mathcal{A}} = \mathcal{E}(x)$
$\llbracket \varphi_1 \left\{ \begin{smallmatrix} \vee \\ \wedge \end{smallmatrix} \right\} \varphi_2 \rrbracket_{\mathcal{E}}^{\mathcal{A}} = \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^{\mathcal{A}} \left\{ \begin{smallmatrix} \sqcup \\ \sqcap \end{smallmatrix} \right\} \llbracket \varphi_2 \rrbracket_{\mathcal{E}}^{\mathcal{A}}$
$\llbracket \text{pre}(\varphi) \rrbracket_{\mathcal{E}}^{\mathcal{A}} = \lceil \text{CPre} \rceil (\llbracket \varphi \rrbracket_{\mathcal{E}}^{\mathcal{A}})$
$\llbracket \left\{ \begin{smallmatrix} \mu \\ \nu \end{smallmatrix} \right\} x.\varphi \rrbracket_{\mathcal{E}}^{\mathcal{A}} = \left\{ \begin{smallmatrix} \sqcap \\ \sqcup \end{smallmatrix} \right\} \{q \mid q = \llbracket \varphi \rrbracket_{\mathcal{E}[x \mapsto q]}^{\mathcal{A}}\}$

If φ is a closed formula, then $\llbracket \varphi \rrbracket^V = \llbracket \varphi \rrbracket_{\mathcal{E}}^V$ for any valuation \mathcal{E} .

Theorem 3 ([8]) *For all ω -regular objectives ϕ there exists a closed μ -calculus formula φ_ϕ such that for all complete information game structures G , $\llbracket \varphi_\phi \rrbracket^{\mathcal{S}}$ is the set of sure-winning states of G for ϕ .*

Downward closure. Given a set $q \in \mathcal{S}$, the *downward closure* of q is the set $q \downarrow = \{s \in \mathcal{L} \mid \exists s' \in q : s \subseteq s'\}$. Observe that in particular, for all $q \in \mathcal{S}$ we have $\emptyset \notin q \downarrow$ and $\lceil q \rceil \downarrow = q \downarrow$. The sets $q \downarrow$ ($q \in \mathcal{S}$) are the *downward closed* sets. We say that a valuation \mathcal{E} for the variables is *downward closed* if every variable x is mapped to a downward closed set, that is $\mathcal{E}(x) = \mathcal{E}(x) \downarrow$.

Lemma 5 *For all downward closed sets $q, q' \in \mathcal{S}$, we have*

$$\lceil q \cap q' \rceil = \lceil q \rceil \sqcap \lceil q' \rceil \quad \text{and} \quad \lceil q \cup q' \rceil = \lceil q \rceil \sqcup \lceil q' \rceil.$$

Lemma 6 *For all μ -calculus formulas φ , for all downward closed valuations \mathcal{E} in the lattice of subsets, the set $\llbracket \varphi \rrbracket_{\mathcal{E}}^{\mathcal{S}}$ is downward closed.*

Proof. We prove this lemma by induction on the structure of φ .

- if $\varphi \equiv \text{obs}$ for $\text{obs} \in \text{Obs}$. It is immediate to show that $\llbracket \varphi \rrbracket_{\mathcal{E}}^{\mathcal{S}} = \llbracket \varphi \rrbracket_{\mathcal{E}}^{\mathcal{S}} \downarrow$.

- if $\varphi \equiv x$ for a variable x . We have $\llbracket \varphi \rrbracket_{\mathcal{E}}^{\mathcal{S}} = \mathcal{E}(x)$ which is downward closed by hypothesis.

- if $\varphi \equiv \varphi_1 \left\{ \begin{array}{c} \vee \\ \wedge \end{array} \right\} \varphi_2$ and both $\llbracket \varphi_1 \rrbracket_{\mathcal{E}}^{\mathcal{S}}$ and $\llbracket \varphi_2 \rrbracket_{\mathcal{E}}^{\mathcal{S}}$ are downward closed.

Then we have $\llbracket \varphi \rrbracket_{\mathcal{E}}^{\mathcal{S}} = \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^{\mathcal{S}} \left\{ \begin{array}{c} \cup \\ \cap \end{array} \right\} \llbracket \varphi_2 \rrbracket_{\mathcal{E}}^{\mathcal{S}}$ and the result follows from the fact that union and intersection of downward closed sets are downward closed.

- if $\varphi \equiv \text{pre}(\varphi_1)$ and $\llbracket \varphi_1 \rrbracket_{\mathcal{E}}^{\mathcal{S}}$ is downward closed. We show that $\llbracket \varphi \rrbracket_{\mathcal{E}}^{\mathcal{S}}$ is downward closed. Let $s_1 \in \llbracket \varphi \rrbracket_{\mathcal{E}}^{\mathcal{S}}$ and let $s_2 \in \mathcal{L}$ such that $s_2 \subseteq s_1$. Let us show that $s_2 \in \llbracket \varphi \rrbracket_{\mathcal{E}}^{\mathcal{S}}$. By definition of CPre , since $s_1 \in \llbracket \text{pre}(\varphi_1) \rrbracket_{\mathcal{E}}^{\mathcal{S}}$, there exists $\sigma \in \Sigma$ such that for any s'_1 , if $(s_1, \sigma, s'_1) \in \Delta^K$ then $s'_1 \in \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^{\mathcal{S}}$. Consider any s'_2 such that $(s_2, \sigma, s'_2) \in \Delta^K$. According to Definition 1, we have $s'_2 = \text{Post}_{\sigma}(s_2) \cap \gamma(\text{obs}) \neq \emptyset$ for some $\text{obs} \in \text{Obs}$. Now, let $s'_1 = \text{Post}_{\sigma}(s_1) \cap \gamma(\text{obs})$. Since $s_2 \subseteq s_1$, we have $s'_2 \subseteq s'_1$ and thus $s'_1 \neq \emptyset$. Therefore $(s_1, \sigma, s'_1) \in \Delta^K$ and so $s'_1 \in \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^{\mathcal{S}}$. As the latter set is downward closed, we also have $s'_2 \in \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^{\mathcal{S}}$ and thus $s_2 \in \llbracket \text{pre}(\varphi_1) \rrbracket_{\mathcal{E}}^{\mathcal{S}}$.

- if $\varphi \equiv \nu x. \varphi_1$ and $\llbracket \varphi_1 \rrbracket_{\mathcal{E}'}$ is downward closed for any downward closed valuation \mathcal{E}' . By Tarski's theorem, $\llbracket \varphi \rrbracket_{\mathcal{E}}^{\mathcal{S}}$ is one of the set in the infinite sequence q_0, q_1, \dots defined by $q_0 = \mathcal{L}$ and for every $i \geq 1$, $q_i = \llbracket \varphi \rrbracket_{\mathcal{E}[x \mapsto q_{i-1}]}^{\mathcal{S}}$. Since q_0 is downward closed, every q_i ($i \geq 1$) is also downward closed by the induction hypothesis.

- if $\varphi \equiv \mu x. \varphi_1$ and $\llbracket \varphi_1 \rrbracket_{\mathcal{E}'}$ is downward closed for any downward closed valuation \mathcal{E}' . The proof is similar to the previous case.

■

Lemma 7 *For all μ -calculus formulas φ , for all downward closed valuations \mathcal{E} in the lattice of subsets, we have $\lceil \llbracket \varphi \rrbracket_{\mathcal{E}}^{\mathcal{S}} \rceil = \llbracket \varphi \rrbracket_{\lceil \mathcal{E} \rceil}^{\mathcal{A}}$ where $\lceil \mathcal{E} \rceil$ is a valuation in the lattice of antichains defined by $\lceil \mathcal{E} \rceil(x) = \lceil \mathcal{E}(x) \rceil$ for all variables x .*

Proof. We prove this by induction on the structure of φ .

- if $\varphi \equiv \text{obs}$ for $\text{obs} \in \text{Obs}$. The claim is immediate.
- if $\varphi \equiv x$ for a variable x . We have $\lceil \llbracket \varphi \rrbracket_{\mathcal{E}}^{\mathcal{S}} \rceil = \lceil \mathcal{E}(x) \rceil = \llbracket \varphi \rrbracket_{\lceil \mathcal{E} \rceil}^{\mathcal{A}}$.

- if $\varphi \equiv \varphi_1 \left\{ \begin{array}{c} \vee \\ \wedge \end{array} \right\} \varphi_2$ and both $\lceil \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S \rceil = \llbracket \varphi_1 \rrbracket_{\mathcal{E}_1}^A$ and $\lceil \llbracket \varphi_2 \rrbracket_{\mathcal{E}}^S \rceil = \llbracket \varphi_2 \rrbracket_{\mathcal{E}_1}^A$.

Using Lemma 5 and Lemma 6, we have successively:

$$\begin{aligned} \lceil \llbracket \varphi \rrbracket_{\mathcal{E}}^S \rceil &= \lceil \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S \left\{ \begin{array}{c} \cup \\ \cap \end{array} \right\} \llbracket \varphi_2 \rrbracket_{\mathcal{E}}^S \rceil = \\ \lceil \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S \rceil \left\{ \begin{array}{c} \cup \\ \cap \end{array} \right\} \lceil \llbracket \varphi_2 \rrbracket_{\mathcal{E}}^S \rceil &= \llbracket \varphi_1 \rrbracket_{\mathcal{E}_1}^A \left\{ \begin{array}{c} \cup \\ \cap \end{array} \right\} \llbracket \varphi_2 \rrbracket_{\mathcal{E}_1}^A. \end{aligned}$$

- if $\varphi \equiv \text{pre}(\varphi_1)$ and $\lceil \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S \rceil = \llbracket \varphi_1 \rrbracket_{\mathcal{E}_1}^A$.

1. We prove the inclusion $\llbracket \varphi \rrbracket_{\mathcal{E}_1}^A \subseteq \lceil \llbracket \varphi \rrbracket_{\mathcal{E}}^S \rceil$. First, we show that for any state $s \in \lceil \text{CPre}(\llbracket \varphi_1 \rrbracket_{\mathcal{E}_1}^A) \rceil$, we have $s \in \text{CPre}(\llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S)$. We know that there exists $\sigma \in \Sigma$ such that $\forall o \in \text{Obs} \cdot \exists s'' \in \llbracket \varphi_1 \rrbracket_{\mathcal{E}_1}^A$: $\text{Post}_{\sigma}(s) \cap \gamma(o) \subseteq s''$. Since $\lceil \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S \rceil = \llbracket \varphi_1 \rrbracket_{\mathcal{E}_1}^A$ (induction hypothesis), it is clear that for such σ , if $(s, \sigma, s') \in \Delta^K$, then there exists $s'' \in \lceil \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S \rceil$ such that $s' \subseteq s''$. And since $\llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S$ is downward closed (by Lemma 6) we have $s' \in \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S$, so that $s \in \text{CPre}(\llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S)$ (and thus $s \in \llbracket \varphi \rrbracket_{\mathcal{E}}^S$).

Second, we show that s is maximal in $\llbracket \varphi \rrbracket_{\mathcal{E}}^S$. By contradiction, assume that there exists $s_1 \in \llbracket \varphi \rrbracket_{\mathcal{E}}^S$ with $s \subset s_1$. Then, by the same argument as in the first part of the proof of the inclusion $\lceil \llbracket \varphi \rrbracket_{\mathcal{E}}^S \rceil \subseteq \llbracket \varphi \rrbracket_{\mathcal{E}_1}^A$, we have that s_1 satisfies the definition of $\lceil \text{CPre}(\llbracket \varphi_1 \rrbracket_{\mathcal{E}_1}^A) \rceil$ up to the operator $\lceil \cdot \rceil$. This means that s is not maximal in $\lceil \text{CPre}(\llbracket \varphi_1 \rrbracket_{\mathcal{E}_1}^A) \rceil$, a contradiction.

2. We prove the inclusion $\lceil \llbracket \varphi \rrbracket_{\mathcal{E}}^S \rceil \subseteq \llbracket \varphi \rrbracket_{\mathcal{E}_1}^A$. This is trivial if $\lceil \llbracket \varphi \rrbracket_{\mathcal{E}}^S \rceil = \emptyset$. Otherwise, let us first show that $\llbracket \varphi_1 \rrbracket_{\mathcal{E}_1}^A \neq \emptyset$. Let $s \in \lceil \llbracket \varphi \rrbracket_{\mathcal{E}}^S \rceil$. Then, there exists $\sigma \in \Sigma$ such that for any s' , if $(s, \sigma, s') \in \Delta^K$ then $s' \in \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S$. Since the transition relation of G is total and the observations partition the state space, we have $\text{Post}_{\sigma}(s) \cap \gamma(\text{obs}) \neq \emptyset$ for some $\text{obs} \in \text{Obs}$. Therefore, $\llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S$ is nonempty and so is $\llbracket \varphi_1 \rrbracket_{\mathcal{E}_1}^A$.

Now, we proceed with the proof of inclusion. Let $s \in \lceil \llbracket \varphi \rrbracket_{\mathcal{E}}^S \rceil$, and let $\sigma \in \Sigma$ such that for any s' , if $(s, \sigma, s') \in \Delta^K$ then $s' \in \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S$. Let us show that $s \in \lceil \text{CPre}(\llbracket \varphi_1 \rrbracket_{\mathcal{E}_1}^A) \rceil$. First, consider an arbitrary observation $\text{obs} \in \text{Obs}$ and let $s' = \text{Post}_{\sigma}(s) \cap \gamma(\text{obs})$. We must show that there exists $s'' \in \llbracket \varphi_1 \rrbracket_{\mathcal{E}_1}^A$ such that $s' \subseteq s''$. This is obvious if $s' = \emptyset$ since $\llbracket \varphi_1 \rrbracket_{\mathcal{E}_1}^A$ is nonempty. Otherwise, by Definition 1 we have $(s, \sigma, s') \in \Delta^K$ and therefore $s' \in \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S$. Since

$\llbracket \varphi_1 \rrbracket_{\mathcal{E}}^A = \lceil \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^S \rceil$ (induction hypothesis), there exists $s'' \in \llbracket \varphi_1 \rrbracket_{\mathcal{E}}^A$ such that $s' \subseteq s''$.

Second, let us show that s is maximal in $\lceil \text{CPre} \rceil(\llbracket \varphi_1 \rrbracket_{\mathcal{E}}^A)$. By contradiction, assume that there exists $s_1 \in \lceil \text{CPre} \rceil(\llbracket \varphi_1 \rrbracket_{\mathcal{E}}^A)$ with $s \subset s_1$. Then, by the same argument as in the first part of the proof of the inclusion $\llbracket \varphi \rrbracket_{\mathcal{E}}^A \subseteq \lceil \llbracket \varphi \rrbracket_{\mathcal{E}}^S \rceil$, we have $s_1 \in \llbracket \varphi \rrbracket_{\mathcal{E}}^S$. This implies that $s \notin \lceil \llbracket \varphi \rrbracket_{\mathcal{E}}^S \rceil$, a contradiction.

- if $\varphi \equiv \nu x. \varphi_1$ and $\lceil \llbracket \varphi_1 \rrbracket_{\mathcal{E}'}^S \rceil = \llbracket \varphi_1 \rrbracket_{\mathcal{E}'}^A$ for any downward closed valuation \mathcal{E}' . By Tarski's theorem, $\llbracket \varphi \rrbracket_{\mathcal{E}}^S$ is one of the set in the infinite sequence q_0, q_1, \dots defined by $q_0 = \mathcal{L}$ and for every $i \geq 1$, $q_i = \llbracket \varphi \rrbracket_{\mathcal{E}[x \mapsto q_{i-1}]}^S$; and similarly, $\llbracket \varphi \rrbracket_{\mathcal{E}}^A$ is one of the set in the infinite sequence q'_0, q'_1, \dots defined by $q'_0 = \{L\}$ and for all $i \geq 1$, $q'_i = \llbracket \varphi \rrbracket_{\mathcal{E}[x \mapsto q'_{i-1}]}^A$. Observe that $q'_0 = \lceil q_0 \rceil$. By induction, assume that $q'_{i-1} = \lceil q_{i-1} \rceil$ for some $i \geq 1$. Then clearly $q'_i = \lceil q_i \rceil$ as $\mathcal{E}[x \mapsto q'_{i-1}] = \lceil \mathcal{E}[x \mapsto q_{i-1}] \rceil$.
- if $\varphi \equiv \mu x. \varphi_1$ and $\lceil \llbracket \varphi_1 \rrbracket_{\mathcal{E}'}^S \rceil = \llbracket \varphi_1 \rrbracket_{\mathcal{E}'}^A$ for any downward closed valuation \mathcal{E}' . The proof is similar to the previous case.

■

From Theorems 2 and 3 and Lemma 7, we can decide the existence of a deterministic observation-based sure-winning strategy for Player 1 without constructing G^K .

Theorem 4 *Let G be a game structure of incomplete information with initial state l_0 . For all ω -regular objectives ϕ , Player 1 has a deterministic observation-based strategy in G for ϕ if and only if $\{l_0\} \sqsubseteq \llbracket \varphi_\phi \rrbracket^A$.*

Corollary 1 *Let G be a game structure of incomplete information. Whether a state ℓ of G is a sure-winning state for a parity objective $\text{Parity}(p)$ can be decided in EXPTIME.*

The argument to prove Corollary 1 is that the size and the number of fixed-point quantifier alternations of a μ -calculus formula φ for a parity objective ϕ are polynomial in the number of priorities, and the size of the parity objective ϕ . The evaluation of φ is exponential in the size of ϕ .

4 Almost Winning

Given a game structure G of incomplete information we construct game structure H by subset construction and establish certain equivalences of randomized strategies in G and H . We then show how the reduction can be used to obtain a symbolic EXPTIME algorithm to compute almost-winning states in G for Büchi objectives.

4.1 Subset construction for almost winning

Reduction. Given a game structure of incomplete information $G = (L, l_0, \Sigma, \Delta, \text{Obs}, \gamma)$, we construct a game structure of complete information $H = (Q, q_{l_0}, \Sigma, \Delta_H)$ as follows: $Q = \{ (s, \ell) \mid \exists o \in \text{Obs}. s \subseteq \gamma(o), \ell \in s \}$; when H is in a state (s, ℓ) , it corresponds in G to the state ℓ and the knowledge of Player 1 is s ; the initial state is $q_{l_0} = (\{ l_0 \}, l_0)$; the transition relation $\Delta_H \subseteq Q \times \Sigma \times Q$ is defined as follows: $((s, \ell), \sigma, (s', \ell')) \in \Delta_H$ iff $\exists o \in \text{Obs} : s' = \text{Post}_\sigma^G(s) \cap \gamma(o)$ and $(\ell, \sigma, \ell') \in \Delta$.

Plays, prefixes and cones of prefixes A play in H is an infinite alternating sequence $q_0 \sigma_0 q_1 \sigma_1 \dots$ of states in Q and input letters from Σ . We denote by $\text{Play}(H)$ the set of plays of H . A prefix $\rho_H = q_0 \sigma_0 q_1 \sigma_1 \dots \sigma_{n-1} q_n$ is a prefix of a play in H . For a prefix $\rho_H = q_0 \sigma_0 q_1 \sigma_1 \dots \sigma_{n-1} q_n$ we denote by $\text{Last}(\rho_H) = q_n$ the last state in ρ_H and by $|\rho_H| = n + 1$ the length of ρ_H . We denote by $\text{Pref}(H)$ the set of prefixes in H . For a prefix $\rho_H \in \text{Pref}(H)$, let $\text{Cone}(\rho_H) = \{ \pi_H \in \text{Play}(H) \mid \rho_H \text{ is a prefix of } \pi_H \}$. The definition of $\text{Cone}(\rho_G)$ for prefixes $\rho_G \in \text{Pref}(G)$ is similar.

Equivalence of states, prefixes and plays Two states $q = (s, \ell)$ and $q' = (s', \ell')$ of H are *equivalent*, noted $q \approx q'$, iff $s = s'$. Two prefixes $\rho_H = q_0 \sigma_0 q_1 \dots \sigma_{n-1} q_n$ and $\rho'_H = q'_0 \sigma'_0 q'_1 \dots \sigma'_{n-1} q'_n$, are *equivalent*, noted $\rho_H \approx \rho'_H$, iff for all $0 \leq i \leq n$ we have $q_i \approx q'_i$ and for all $0 \leq i \leq n - 1$ we have $\sigma_i = \sigma'_i$. Two plays $\pi_H, \pi'_H \in \text{Play}(H)$ are *equivalent*, noted $\pi_H \approx \pi'_H$, iff for all $i \geq 0$ we have $\pi_H(i) \approx \pi'_H(i)$. For a state $q \in Q$ we denote by $[q]_{\approx} = \{ q' \in Q \mid q \approx q' \}$ the \approx -equivalence class of q . We define a similar notation for prefixes and plays.

Positional and equivalence preserving strategies. A strategy for Player 1 is a function $\alpha_H : \text{Pref}(H) \rightarrow \mathcal{D}(\Sigma)$. A strategy α_H for Player 1 in H is *positional* if it is independent of the prefix of plays and depends only on the last state, *i.e.* for all $\rho_H, \rho'_H \in \text{Pref}(H)$ such that $\text{Last}(\rho_H) = \text{Last}(\rho'_H)$




Figure 2: Game structure H (for G of Figure 1).

we have $\alpha_H(\rho_H) = \alpha_H(\rho'_H)$. A positional strategy α_H can be viewed as a function $\alpha_H : Q \rightarrow \mathcal{D}(\Sigma)$. A strategy α_H for Player 1 is *equivalence preserving* if for all $\rho_H, \rho'_H \in \text{Pref}(H)$ such that $\rho_H \approx \rho'_H$ we have $\alpha_H(\rho_H) = \alpha_H(\rho'_H)$. We denote by $\mathcal{A}_H, \mathcal{A}_H^P, \mathcal{A}_H^{\approx}$ the set of all strategies, the set of all positional strategies, and the set of all equivalence preserving strategies for Player 1, respectively. We denote by $\mathcal{A}_H^{\approx(P)} = \mathcal{A}_H^{\approx} \cap \mathcal{A}_H^P$ the set of equivalence preserving positional strategies. A strategy for Player 2 is a function $\beta_H : \text{Pref}(H) \times \Sigma \rightarrow \mathcal{D}(Q)$ such that for all $\rho_H \in \text{Pref}(H)$, for all $\sigma \in \Sigma$, $\text{Supp}(\beta_H(\rho_H, \sigma)) \subseteq \text{Post}_{\text{Last}}^H(\rho_H, \{\sigma\})$. We denote by \mathcal{B}_H the set of all Player 2 strategies.

Objectives. Objectives Φ_H in H are subsets of $(Q \times \Sigma)^\omega$. The definition of objectives are similar to as in Section 2. An objective Φ_H is equivalence preserving if for all $\pi_H \in \Phi_H$, $[\pi_H]_{\approx} \subseteq \Phi_H$.

Example 2 Given the game structure G of incomplete information G of Figure 1 the game structure H produced by the reduction of Section 4 is shown in Figure 2. All the states are almost-winning for the Büchi objective $\text{Buchi}(\{(\{l_4\}, l_4)\})$. For a state q we say $\text{Rank}(q) = j$ if $q \in \text{Rank}(j) \setminus \text{Rank}(j-1)$ (see notations after Lemma 11 for definition of Rank). The rank of the states as obtained from Lemma 12 are shown along side the states. The positional strategy that plays both a and b with equal probability is almost-winning for all states. For states q of Rank(1), Rank(3) and Rank(4) we have if $\text{Rank}(q) = j$, then $\text{Post}_a^H(q) \subseteq \text{Rank}(j-1)$ and $\text{Post}_b^H(q) \subseteq \text{Rank}(j-1)$. For states with Rank(2) we have if $q = (\{l_2, l'_2\}, l_2)$, then $\text{Post}_b^H(q) \subseteq \text{Rank}(1)$; and if $q = (\{l_2, l'_2\}, l'_2)$, then $\text{Post}_a^H(q) \subseteq \text{Rank}(1)$.

Mapping of prefixes and plays. We define a mapping $h : \text{Pref}(G) \rightarrow \text{Pref}(H)$ that maps prefixes in G to prefixes in H as follows: given $\rho_G =$

$\ell_0\sigma_0\ell_1\sigma_1\dots\sigma_{n-1}\ell_n$ we have $h(\rho_G) = q_0\sigma_0q_1\sigma_1\dots\sigma_{n-1}q_n$, where $q_i = (s_i, \ell_i)$ and $s_i = \mathsf{K}(\mathsf{Obs}(\rho(i)))$. The mapping h for plays is defined similarly.

Proposition 2 *The mapping h is a bijection.*

Mapping of strategies for Player 1. Given a strategy α_H in H we construct a strategy $\alpha_G = h(\alpha_H)$ as follows: given $\rho_G \in \mathsf{Pref}(G)$ we have $\alpha_G(\rho_G) = \alpha_H(h(\rho_G))$. Similarly given a strategy α_G in G we construct a strategy $\alpha_H = g(\alpha_G)$ as follows: given $\rho_H \in \mathsf{Pref}(H)$ we have $\alpha_H(\rho_H) = \alpha_G(h^{-1}(\rho_H))$.

Proposition 3 *The following properties hold.*

1. *For the strategy mapping $h : \mathcal{A}_H \rightarrow \mathcal{A}_G$ we have*
 - (a) *if $\mathsf{Obs}(\rho_G) = \mathsf{Obs}(\rho'_G)$, then $h(\rho_G) \approx h(\rho'_G)$; and*
 - (b) *if α_H is equivalence preserving, then $h(\alpha_H)$ is observation-based.*
2. *For the strategy mapping $g : \mathcal{A}_G \rightarrow \mathcal{A}_H$ we have*
 - (a) *if $\rho_H \approx \rho'_H$, then $\mathsf{Obs}(h^{-1}(\rho_H)) = \mathsf{Obs}(h^{-1}(\rho'_H))$; and*
 - (b) *if α_G is observation-based, then $g(\alpha_G)$ is equivalence preserving.*

Mapping of strategies for Player 2. Observe that for all $q \in Q$, $\sigma \in \Sigma$ and $\ell \in L$ we have $|\{q' = (s', \ell) \mid (q, \sigma, q') \in \Delta_H\}| \leq 1$. Given a strategy β_H in H we construct a strategy $\beta_G = h(\beta_H)$ as follows: for $\rho_G \in \mathsf{Pref}(G)$, $\sigma \in \Sigma$ and $\ell \in L$ we have $\beta_G(\rho_G, \sigma)(\ell) = \beta_H(h(\rho_G), \sigma)(s, \ell)$, where $(s, \ell) \in \mathsf{Post}_\sigma^H(\mathsf{Last}(h(\rho_G)))$. Similarly given a strategy β_G in G we construct a strategy $\beta_H = g(\beta_G)$ as follows: for $\rho_H \in \mathsf{Pref}(H)$, $\sigma \in \Sigma$, $q = (s, \ell) \in Q$ we have $\beta_H(\rho_H, \sigma)(q) = \beta_G(h^{-1}(\rho_H), \sigma)(\ell)$.

Lemma 8 *The following assertions hold.*

1. *For all $\rho_H \in \mathsf{Pref}(H)$, for every equivalence preserving strategy α_H , for every strategy β_H we have*

$$\Pr_{q_0}^{\alpha_H, \beta_H}(\mathsf{Cone}(\rho_H)) = \Pr_{l_0}^{h(\alpha_H), h(\beta_H)}(h^{-1}(\mathsf{Cone}(\rho_H))).$$

2. *For all $\rho_G \in \mathsf{Pref}(G)$, for every observational strategy α_G , for every strategy β_G we have*

$$\Pr_{l_0}^{\alpha_G, \beta_G}(\mathsf{Cone}(\rho_G)) = \Pr_{q_0}^{g(\alpha_G), g(\beta_G)}(h(\mathsf{Cone}(\rho_G))).$$

Proof. The following properties follow from the construction of strategies in Subsection 4.1.

1. For all $\rho_H \in \text{Pref}(H)$, for every equivalence preserving strategy α_H , for every strategy β_H and for all ρ'_H such that $|\rho'_H| = |\rho_H| + 1$, we have $\Pr_{q_0}^{\alpha_H, \beta_H}(\text{Cone}(\rho'_H) \mid \text{Cone}(\rho_H)) = \Pr_{l_0}^{h(\alpha_H), h(\beta_H)}(h^{-1}(\text{Cone}(\rho'_H)) \mid h^{-1}(\text{Cone}(\rho_H)))$.
2. For all $\rho_G \in \text{Pref}(G)$, for every observation-based strategy α_G , for every strategy β_G and for all ρ'_G such that $|\rho'_G| = |\rho_G| + 1$, we have $\Pr_{l_0}^{\alpha_G, \beta_G}(\text{Cone}(\rho'_G) \mid \text{Cone}(\rho_G)) = \Pr_{q_0}^{g(\alpha_G), g(\beta_G)}(h(\text{Cone}(\rho'_G)) \mid h(\text{Cone}(\rho_G)))$.

The proof for the first part is as follows: observe that

$$\begin{aligned} \Pr_{q_0}^{\alpha_H, \beta_H}(\text{Cone}(q_{l_0})) &= \Pr_{l_0}^{g(\alpha_H), g(\beta_H)}(h^{-1}\text{Cone}(q_0)) \\ &= \Pr_{l_0}^{g(\alpha_H), g(\beta_H)}(\text{Cone}(l_0)) = 1. \end{aligned}$$

The rest follows from the above properties and induction. The proof for the second part is similar. \blacksquare

Theorem 5 *For all incomplete information game structures G , all Borel objectives Φ_G and for all equivalence preserving Borel objectives Φ_H , for all observation-based player 1 strategies α_G , for all equivalence preserving player 1 strategies α_H , for all player 2 strategies β_G and β_H we have*

$$\begin{aligned} \Pr_{l_0}^{\alpha_G, \beta_G}(\Phi_G) &= \Pr_{q_0}^{g(\alpha_G), g(\beta_G)}(h(\Phi_G)); \\ \Pr_{q_0}^{\alpha_H, \beta_H}(\Phi_H) &= \Pr_{l_0}^{h(\alpha_G), h(\beta_G)}(h^{-1}(\Phi_G)). \end{aligned}$$

Proof. The proof is as follows: by the Caratheodary unique extension theorem a probability measure defined on Cone's has an unique extension to all Borel objectives. The result then follows from Lemma 8 and the unique extension theorem. \blacksquare

Corollary 2 follows from Theorem 5.

Corollary 2 *For all Borel objectives Φ_G we have*

$$\begin{aligned} \sup_{\alpha_G \in \mathcal{A}_G^O} \inf_{\beta_G \in \mathcal{B}_G} \Pr_{l_0}^{\alpha_G, \beta_G}(\Phi_G) &= \sup_{\alpha_H \in \mathcal{A}_H^\approx} \inf_{\beta_H \in \mathcal{B}_H} \Pr_{q_0}^{\alpha_H, \beta_H}(h(\Phi_G)); \\ \exists \alpha_G \in \mathcal{A}_G^O. \forall \beta_G \in \mathcal{B}_G. \Pr_{l_0}^{\alpha_G, \beta_G}(\Phi_G) &= 1 \\ \text{iff } \exists \alpha_H \in \mathcal{A}_H^\approx. \forall \beta_H \in \mathcal{B}_H. \Pr_{q_0}^{\alpha_H, \beta_H}(h(\Phi_G)) &= 1. \end{aligned}$$

4.2 Almost winning for Büchi objectives

We first illustrate the need of memory and randomization for almost-winning in incomplete information games with Büchi objectives.

Example 3 (Memory is needed to almost-win) Consider the example of Figure 3. The objective of Player 1 is to reach a state with observation o_4 . We show that Player 1 has no observation based sure-winning strategy




Figure 3: The observations are as follows: $o_1 = \{q_0\}$, $o_2 = \{q_1, q_2\}$, $o_3 = \{q_3, q_4\}$, $o_4 = \{q_5\}$, $o_5 = \{q_6\}$.

in this game. This is because when we fix an observation-based strategy for Player 1, Player 2 has a spoiling strategy to maintain the game into the states $\{q_0, q_1, q_2\}$. Indeed, at q_0 , the only reasonable choice for Player 1 is to play a . Then Player 2 can choose to go either in q_1 or q_2 . In both cases, the observation will be the same for Player 1. After seeing $o_1 a o_2$, if the strategy of Player 1 is to play a then Player 2 chooses q_2 , otherwise, if Player 1 strategy is to play b then Player 2 chooses q_1 . This can be repeated and so Player 2 has a spoiling strategy against any observation based strategy of Player 1.

We now show that almost-winning strategies exist for Player 1. Consider that Player 1 plays an observation-based randomized strategy α as follows: after a sequence of observations τ ,

- in any observation sequence τ such that $\text{Last}(\tau) = o_1$, $\alpha(\tau)(a) = 1$ and $\alpha(\tau)(b) = 0$,
- in any observation sequence τ such that $\text{Last}(\tau) = o_2$, $\alpha(\tau)(a) = 0.5$ and $\alpha(\tau)(b) = 0.5$,
- in any observation sequence $\tau = \tau' \cdot o_1 \cdot \sigma \cdot o_3$, $\alpha(\tau)(a) = 1$ and $\alpha(\tau)(b) = 0$,
- in any observation sequence $\tau = \tau' \cdot o_2 \cdot \sigma \cdot o_3$, $\alpha(\tau)(a) = 0$ and $\alpha(\tau)(b) = 1$,
- otherwise take, arbitrarily, $\alpha(\tau)(a) = 1$ and $\alpha(\tau)(b) = 0$.

The strategy α is almost-winning against any randomized strategy of Player 2. Note that the strategy α uses memory and this is necessary because when receiving observation o_3 , Player 1 has to play a if the game was just before in a state satisfying observation o_1 and b if the game was just before in state satisfying o_2 .

Given a game structure G let H be the corresponding game structure constructed by the reduction. Given $O \subseteq \text{Obs}$ let $B_O = \{ q = (s, l) \in Q \mid \exists o \in O. s \subseteq \gamma(o) \}$. We have $h(\text{Buchi}(O)) = \text{Buchi}(B_O) = \{ \pi_H \in \text{Play}(H) \mid \text{Inf}(\pi_H) \cap B_O \neq \emptyset \}$. We first show that almost-winning for Büchi objectives wrt to equivalence preserving strategies is equivalent to almost-winning wrt to equivalence preserving positional strategies. Formally, for $B_O \subseteq Q$, let

$$Q_{\text{AS}}^{\approx} = \{ q \in Q \mid \exists \alpha_H \in \mathcal{A}_H^{\approx}. \forall \beta_H \in \mathcal{B}_H. \forall q' \in [q]_{\approx}. \text{Pr}_{q'}^{\beta_H, \alpha_H}(\text{Buchi}(B_O)) = 1 \};$$

$$Q_{\text{AS}}^{\approx(P)} = \{ q \in Q \mid \exists \alpha_H \in \mathcal{A}_H^{\approx(P)}. \forall \beta_H \in \mathcal{B}_H. \forall q' \in [q]_{\approx}. \text{Pr}_{q'}^{\beta_H, \alpha_H}(\text{Buchi}(B_O)) = 1 \}.$$

We will prove that $Q_{\text{AS}}^{\approx} = Q_{\text{AS}}^{\approx(P)}$. Proposition 4 follows from the construction of H from G ; and yields Lemma 9.

Proposition 4 For all $q \in Q$, for all $q_1 \in [q]_{\approx}$, for all $\sigma \in \Sigma$, if $(q_1, \sigma, q'_1) \in \Delta_H$, then for all $q'_2 \in [q'_1]_{\approx}$, there exists $q_2 \in [q_1]_{\approx}$ such that $(q_2, \sigma, q'_2) \in \Delta_H$.

Lemma 9 Given an equivalence preserving strategy α_H , a prefix $\rho_H \in \text{Pref}(H)$ and a state q , if there exists a strategy β_H such that $\text{Pr}_q^{\alpha_H, \beta_H}(\text{Cone}(\rho_H)) > 0$, then for every ρ'_H such that $\rho'_H \approx \rho_H$ there exists a strategy β'_H and a state $q' \in [q]_{\approx}$ such that $\text{Pr}_{q'}^{\alpha_H, \beta'_H}(\text{Cone}(\rho'_H)) > 0$.

Observe that

$$Q \setminus Q_{\text{AS}}^{\approx} = \{q \in Q \mid \forall \alpha_H \in \mathcal{A}_H^{\approx}. \exists \beta_H \in \mathcal{B}_H. \exists q' \in [q]_{\approx}. \Pr_{q'}^{\alpha_H, \beta_H}(\text{Buchi}(B_O)) < 1\}.$$

It follows from Lemma 9 that if a play starts in Q_{AS}^{\approx} and reaches $(Q \setminus Q_{\text{AS}}^{\approx})$ with positive probability, then for all equivalence preserving strategies for Player 1, there is a Player 2 strategy that ensures that the Büchi objective $\text{Buchi}(B_O)$ is not satisfied with probability 1.

Notations. For a state $q \in Q$ and $Y \subseteq Q$ we denote by $\text{Allow}(q, Y) = \{\sigma \in \Sigma \mid \text{Post}_{\sigma}^H(q) \subseteq Y\}$. For a state $q \in Q$ and $Y \subseteq Q$ we denote by $\text{Allow}([q], Y) = \bigcap_{q' \in [q]_{\approx}} \text{Allow}(q', Y)$.

Lemma 10 *For all $q \in Q_{\text{AS}}^{\approx}$ we have $\text{Allow}([q], Q_{\text{AS}}^{\approx}) \neq \emptyset$.*

Proof. Assume towards contradiction that there exists $q \in Q_{\text{AS}}^{\approx}$ such that $\text{Allow}([q], Q_{\text{AS}}^{\approx}) = \emptyset$. Then for all $\sigma \in \Sigma$ there exists $q' \in [q]_{\approx}$ such that $\text{Post}_{\sigma}^H(q') \cap (Q \setminus Q_{\text{AS}}^{\approx}) \neq \emptyset$. Hence for all equivalence preserving strategy α_H there exists $q' \in [q]_{\approx}$ such that $\alpha_H(q')(\sigma) > 0$ and $\text{Post}_{\sigma}^H(q') \cap (Q \setminus Q_{\text{AS}}^{\approx}) \neq \emptyset$. Hence for every equivalence strategy α_H there is a state $q' \in [q]_{\approx}$ and a strategy β_H for Player 2 such that $Q \setminus Q_{\text{AS}}^{\approx}$ is reached with positive probability. This contradicts that $[q]_{\approx} \subseteq Q_{\text{AS}}^{\approx}$. ■

Lemma 11 *Given $q \in Q_{\text{AS}}^{\approx}$ let α_H be an equivalence preserving strategy such that for all β_H and for all $q' \in [q]_{\approx}$ we have $\Pr_{q'}^{\alpha_H, \beta_H}(\text{Buchi}(B_O)) = 1$. For all $\rho_H = q_0 \sigma_0 q_1 \dots \sigma_{n-1} q_n \in \text{Pref}(H)$ such that for all $0 \leq i \leq n$, $q_i \in Q_{\text{AS}}^{\approx}$, if there is a strategy β_H such that $\Pr_{q'}^{\alpha_H, \beta_H}(\text{Cone}(\rho_H)) > 0$, then $\text{Supp}(\alpha_H(\rho_H)) \subseteq \text{Allow}([q], Q_{\text{AS}}^{\approx})$.*

Proof. Assume towards contradiction for a history ρ_H satisfying the conditions of the lemma there exists $\sigma \in \text{Supp}(\alpha_H(\rho_H)) \setminus \text{Allow}([q], Q_{\text{AS}}^{\approx})$. Then there exists $q'_n \in [q]_{Q_{\text{AS}}^{\approx}}$ such that $\text{Post}_{\sigma}^H(q'_n) \cap (Q \setminus Q_{\text{AS}}^{\approx}) \neq \emptyset$. Then there exists ρ'_H such that $\rho_H \approx \rho'_H$ and $\text{Last}(\rho'_H) = q'_n$. Then by Lemma 9 there exists a strategy β'_H and $q' \in [q]_{\approx}$ such that $\Pr_{q'}^{\alpha_H, \beta'_H}(\text{Cone}(\rho'_H)) > 0$. Then given ρ'_H and the strategy α_H there exists a Player 2 strategy such that $Q \setminus Q_{\text{AS}}^{\approx}$ is reached with positive probability. This contradicts that α_H is an almost-winning strategy. ■

Notations. We inductively define *ranks* of the states in Q_{AS}^{\approx} as follows: let $\text{Rank}(0) = B_O \cap Q_{\text{AS}}^{\approx}$, and

$$\text{Rank}(j+1) = \text{Rank}(j) \cup \{ q \in Q_{\text{AS}}^{\approx} \mid \exists \sigma \in \text{Allow}([q], Q_{\text{AS}}^{\approx}). \text{Post}_{\sigma}^H(q) \subseteq \text{Rank}(j) \}.$$

Let $j^* = \min\{i \mid \text{Rank}(i) = \text{Rank}(i+1)\}$, and $Q^* = \text{Rank}(j^*)$.

Lemma 12 $Q^* = Q_{\text{AS}}^{\approx}$.

Proof. Clearly by definition $Q^* \subseteq Q_{\text{AS}}^{\approx}$. We now prove that $Q_{\text{AS}}^{\approx} \subseteq Q^*$. Assume towards contradiction $X = Q_{\text{AS}}^{\approx} \setminus Q^* \neq \emptyset$. For all states $q \in X$, for all $\sigma \in \text{Allow}([q], Q_{\text{AS}}^{\approx})$, we have $\text{Post}_{\sigma}^H(q) \cap X \neq \emptyset$; as otherwise q would have been in Q^* . Hence for every state $q \in X$ and $\sigma \in \text{Allow}([q], Q_{\text{AS}}^{\approx})$ there exists $q' \in X$ and $(q, \sigma, q') \in \Delta_H$. Fix a strategy β_H for Player 2 as follows: for a state $q \in X$ and the input letter $\sigma \in \text{Allow}([q], Q_{\text{AS}}^{\approx})$ choose a successor $q' \in X$ such that $(q, \sigma, q') \in \Delta_H$. Consider a state $q \in X$ and an equivalence preserving almost-winning strategy for Player 1 from q . By Lemma 11, the strategy α_H satisfies that for histories $\rho_H = q_0\sigma_0q_1 \dots \sigma_{n-1}q_n$ (satisfying the conditions of Lemma 11) we have $\text{Supp}(\alpha_H(\rho_H)) \subseteq \text{Allow}([q_n], Q_{\text{AS}}^{\approx})$. Hence it follows that $\Pr_q^{\alpha_H, \beta_H}(\text{Safe}(X)) = 1$. Since $B_O \cap Q_{\text{AS}}^{\approx} \subseteq Q^*$, it follows that $B_O \cap X = \emptyset$. Hence $\Pr_q^{\alpha_H, \beta_H}(\text{Reach}(B_O)) = 0$, and $\Pr_q^{\alpha_H, \beta_H}(\text{Buchi}(B_O)) = 0$. This contradicts that α_H is an equivalence preserving almost-winning strategy. \blacksquare

An equivalence-preserving positional strategy. Consider the equivalence preserving positional strategy α_H^p defined as follows: for a state $q \in Q_{\text{AS}}^{\approx}$ play all the moves in $\text{Allow}([q], Q_{\text{AS}}^{\approx})$ uniformly at random.

Lemma 13 For all states $q \in Q_{\text{AS}}^{\approx}$, for all strategies β_H we have:

- $\Pr_q^{\alpha_H^p, \beta_H}(\text{Safe}(Q_{\text{AS}}^{\approx})) = 1$; and
- $\Pr_q^{\alpha_H^p, \beta_H}(\text{Reach}(B_O \cap Q_{\text{AS}}^{\approx})) = 1$.

Proof. By Lemma 12 we have $Q^* = Q_{\text{AS}}^{\approx}$. Let $z = |Q^*|$.

- For a state $q \in Q_{\text{AS}}^{\approx}$ we have $\text{Post}_{\sigma}^H(q) \subseteq Q_{\text{AS}}^{\approx}$ for all $\sigma \in \text{Allow}([q], Q_{\text{AS}}^{\approx})$. Hence it follows for all states $q \in Q_{\text{AS}}^{\approx}$, for all strategies β_H for Player 2 we have $\Pr_q^{\alpha_H^p, \beta_H}(\text{Safe}(Q_{\text{AS}}^{\approx})) = 1$.
- For a state $q \in (\text{Rank}(j+1) \setminus \text{Rank}(j))$, there exists $\sigma \in \text{Allow}([q], Q_{\text{AS}}^{\approx})$ such that $\text{Post}_{\sigma}^H(q) \subseteq \text{Rank}(j)$. For a set $Y \subseteq Q$ let $\diamond^j(Y)$ denote the set of paths that reaches Y within j steps. It follows that for all states $q \in \text{Rank}(j+1)$, for all strategies β_H we have

$$\Pr_q^{\alpha_H^p, \beta_H}(\diamond^1(\text{Rank}(j))) \geq \frac{1}{|\Sigma|}.$$

Let $B = B_O \cap Q_{AS}^\approx$. By induction on the ranks it follows that for all states $q \in Q^*$, for all strategies β_H we have

$$\Pr_q^{\alpha_H^p, \beta_H}(\diamond^z(\text{Rank}(0))) = \Pr_q^{\alpha_H^p, \beta_H}(\diamond^z(B)) \geq \left(\frac{1}{|\Sigma|}\right)^z = r > 0.$$

For $m > 0$ we have $\Pr_q^{\alpha_H^p, \beta_H}(\diamond^{m \cdot z}(B)) \geq 1 - (1 - r)^m$. Thus we have

$$\Pr_q^{\alpha_H^p, \beta_H}(\text{Reach}(B)) = \lim_{m \rightarrow \infty} \Pr_q^{\alpha_H^p, \beta_H}(\diamond^{m \cdot z}(B)) \geq \lim_{m \rightarrow \infty} 1 - (1 - r)^m = 1.$$

The result follows. ■

It follows from Lemma 13 that given the strategy α_H^p , the set Q_{AS}^\approx is never left and the states in $B_O \cap Q_{AS}^\approx$ are reached with probability 1. Since this happens for every state in Q_{AS}^\approx it follows that the set $B_O \cap Q_{AS}^\approx$ is visited infinitely often with probability 1, i.e., the Büchi objective $\text{Buchi}(B_O)$ is satisfied with probability 1. The above analysis along with the fact that $[q_{l_0}]_\approx$ is a singleton and Corollary 2 proves that $Q_{AS}^\approx = Q_{AS}^{\approx(P)}$. Theorem 6 follows and the result also easily specializes for the case of reachability objectives.

Theorem 6 *For all incomplete information game structures G , for a set O of observations, there exists an observation-based almost-winning strategy in G for the objective $\text{Buchi}(O)$ iff there exists an equivalence preserving positional almost-winning strategy in H for the objective $\text{Buchi}(B_O)$.*

Symbolic algorithm We present a symbolic quadratic time (in the size of H) algorithm to compute the set Q_{AS}^\approx . For $Y \subseteq Q$ and $X \subseteq Y$, let

$$\text{Apre}(Y, X) = \{ q \in Y \mid \exists \sigma \in \text{Allow}([q], Y). \text{Post}_\sigma^H(q) \subseteq X \};$$

$$\text{Spre}(Y) = \{ q \in Y \mid \text{Allow}([q], Y) \neq \emptyset \} = \text{Apre}(Y, Y).$$

Let $\phi = \nu Y. \mu X. (\text{Apre}(Y, X) \vee (B_O \wedge \text{Spre}(Y)))$ and $Z = \llbracket \phi \rrbracket$.

Lemma 14 $Z = Q_{AS}^\approx$.

Proof of Lemma 14. We prove $Z = Q_{AS}^\approx$ by proving inclusion in both directions. We have $Z = \llbracket \phi \rrbracket$ and $\phi = \nu Y. \mu X. (\text{Apre}(Y, X) \vee (B_O \wedge \text{Spre}(Y)))$

1. We first show that $Z \subseteq Q_{\text{AS}}^{\approx}$. Since Z is a fixed-point of ϕ we have

$$Z = \llbracket \mu X. (\text{Apre}(Z, X) \vee (B_O \wedge \text{Spre}(Z))) \rrbracket .$$

We analyze the evaluation of Z as the fixed-point as follows: let $X_0 = \emptyset$ and $X_{i+1} = \text{Apre}(Z, X_i) \vee (B_O \wedge \text{Spre}(Z))$. Observe that since $X_0 = \emptyset$ we have $\text{Apre}(Z, X_0) = \emptyset$ and hence $X_1 = B_O \cap \text{Spre}(Z) \subseteq B_O$. Let $j^* = \min\{i \mid X_{i+1} = X_i\}$ and we have $Z = X_{j^*}$. Consider the equivalence preserving strategy α_H^p for Player 1 that at a state $q \in Z$ plays all moves in $\text{Allow}([q], Z)$ uniformly at random. For all strategies β_H for Player 2 and for all states $q \in Z$ we have $\text{Pr}_q^{\alpha_H^p, \beta_H}(\text{Safe}(Z)) = 1$. Also for a state $q \in (X_{i+1} \setminus X_i) \setminus B_O$ we have there exists $\sigma \in \text{Allow}([q], Z)$ such that $\text{Post}_\sigma^H(q) \subseteq X_i$, i.e., for a state $q \in (X_{i+1} \setminus X_i) \setminus B_O$, given α_H^p against all strategies β_H the next state is in X_i with probability at least $\frac{1}{|\Sigma|}$. Arguments similar to Lemma 13 establishes that α_H^p is an almost-winning strategy for all states $q \in Z$. Hence we have $Z \subseteq Q_{\text{AS}}^{\approx}$.

2. We now show that $Q_{\text{AS}}^{\approx} \subseteq Z$. We first show that Q_{AS}^{\approx} satisfies that

$$Q_{\text{AS}}^{\approx} = \llbracket \mu X. (\text{Apre}(Q_{\text{AS}}^{\approx}, X) \vee (B_O \wedge \text{Spre}(Q_{\text{AS}}^{\approx}))) \rrbracket .$$

Observe that $Q_{\text{AS}}^{\approx} = \text{Spre}(Q_{\text{AS}}^{\approx})$. We now analyze the evaluation of Q_{AS}^{\approx} as the fixed-point Q^* as shown in Lemma 12. Let $X_0 = \emptyset$, then $\text{Apre}(Q_{\text{AS}}^{\approx}, X_0) = \emptyset$. Hence $X_1 = \text{Apre}(Q_{\text{AS}}^{\approx}, X_0) \vee (B_O \wedge \text{Spre}(Q_{\text{AS}}^{\approx})) = B_O \wedge Q_{\text{AS}}^{\approx} = \text{Rank}(0)$ (as defined before Lemma 12). By the definition of $\text{Rank}(j+1)$ from $\text{Rank}(j)$ and the definition of $\text{Apre}(\cdot, \cdot)$ and $\text{Spre}(\cdot)$ it follows that for all $i > 0$, given $\text{Rank}(i-1) = X_i$, we have $\text{Rank}(i+1) = X_i = \text{Apre}(Q_{\text{AS}}^{\approx}, X_i) \vee (B_O \wedge \text{Spre}(Q_{\text{AS}}^{\approx}))$. By induction we have $Q^* = \llbracket \mu X. (\text{Apre}(Q_{\text{AS}}^{\approx}, X) \vee (B_O \wedge \text{Spre}(Q_{\text{AS}}^{\approx}))) \rrbracket$. Since $Q^* = Q_{\text{AS}}^{\approx}$ (by Lemma 12) we have

$$Q_{\text{AS}}^{\approx} = \llbracket \mu X. (\text{Apre}(Q_{\text{AS}}^{\approx}, X) \vee (B_O \wedge \text{Spre}(Q_{\text{AS}}^{\approx}))) \rrbracket .$$

Since Z is the greatest fixed-point we have $Q_{\text{AS}}^{\approx} \subseteq Z$.

The result follows. ■

Theorem 7 *Let G be a game structure of incomplete information and a set O of observations. Whether a state ℓ of G is an almost-winning state for a Büchi objective $\text{Buchi}(O)$ can be decided in EXPTIME.*

Lemma 14 and the fact that H is exponential in size of G yield Theorem 7. The arguments for the proof of Theorem 6 and Theorem 7 do not directly extend to coBüchi or parity objectives in general. In fact Theorem 6 does not hold for parity objectives in general for the following reason: in concurrent games with parity objectives with more than two priorities almost-winning strategies require infinite memory in general; for an example see [5]. Such concurrent games are reducible to semi-perfect information games [4] and semi-perfect information games are reducible to the incomplete information games we study. Hence reduction to a finite game structure and obtaining randomized memoryless strategy is not possible in general for parity objectives. Theorem 6 and Theorem 7 may hold for coBüchi objectives, but there does not seem to be a simple extension of arguments of Büchi objectives to the coBüchi case. The results similar to Theorem 6 and Theorem 7 for coBüchi objectives is open.

Direct symbolic algorithm As in Section 3.2, the structure H has not to be constructed explicitly and we can obtain a direct fixed point algorithm on a well chosen lattice. The fixed point formula to compute the set Q_{AS}^{\approx} is evaluated on the lattice $\langle 2^Q, \subseteq, \cup, \cap, Q, \emptyset \rangle$. It is easy to show that the sets computed by the fixed point algorithm are downward closed for the following order on Q : for $(s, \ell), (s', \ell') \in Q$ let $(s, \ell) \preceq (s', \ell')$ iff $\ell = \ell'$ and $s \subseteq s'$. Then, we can define an antichain over Q as a set of pairwise \preceq -incomparable elements of Q , and compute the almost-sure winning states in the lattice of antichains over Q , without explicitly constructing H .

5 Lower Bounds

We show that deciding the existence of a deterministic (resp. randomized) observation-based sure-winning (resp. almost-winning) strategy for Player 1 in games of incomplete information is EXPTIME-hard already for reachability objectives. Note that the result for sure-winning follows from results in [19] but our new proof extends to almost-winning as well.

Sure winning. To show the lower bound result, we use a reduction of the membership problem for polynomial space Alternating Turing Machine. An *alternating Turing machine* (ATM) is a tuple $M = \langle Q, q_0, g, \Sigma_i, \Sigma_t, \delta, F \rangle$ where:

- Q is a finite set of control states;

- $q_0 \in Q$ is the initial state;
- $g : Q \rightarrow \{\wedge, \vee\}$;
- $\Sigma_i = \{0, 1\}$ is the input alphabet;
- $\Sigma_t = \{0, 1, 2\}$ is the tape alphabet and 2 is the *blank* symbol;
- $\delta \subseteq Q \times \Sigma_t \times Q \times \Sigma_t \times \{-1, 1\}$ is a transition relation; and
- $F \subseteq Q$ is the set of accepting states.

We say that M is a *polynomial space* ATM if for some polynomial $p(\cdot)$, the space used by M on any input word w is bounded by $p(|w|)$.

Without loss of generality, we make the hypothesis that the initial control state of the machine is a \vee -state and that transitions link \vee -state to \wedge -state and vice versa. A word w is *accepted* by an ATM M if there exists a run tree of M on w whose all leaf nodes are accepting configurations. The AND-OR graph of the polynomial space ATM (M, p) on the input word $w \in \Sigma^*$ is $G(M, p) = \langle S_\vee, S_\wedge, s_0, \Rightarrow, R \rangle$ where

- $S_\vee = \{(q, h, t) \mid q \in Q, g(q) = \vee, 1 \leq h \leq p(|w|) \text{ and } t \in \Gamma^{p(|w|)}\}$;
- $S_\wedge = \{(q, h, t) \mid q \in Q, g(q) = \wedge, 1 \leq h \leq p(|w|) \text{ and } t \in \Gamma^{p(|w|)}\}$;
- $s_0 = (q_0, 1, t)$ where $t = w.\Gamma^{p(|w|)-|w|}$;
- $((q_1, h_1, t_1), (q_2, h_2, t_2)) \in \Rightarrow$ iff there exists $(q, \gamma, d) \in \delta$ such that $q_2 = q$, $h_2 = h_1 + d$, $t_2(h_1) = \gamma$ and $t_2(i) = t_1(i)$ for all $i \neq h_1$;
- $R = \{(q, h, t) \in S_\vee \cup S_\wedge \mid q \in F\}$.

A word w is *accepted* by (M, p) iff R is reachable in $G(M, p)$. The *membership problem* is to decide if a given word w is accepted by a given polynomial space ATM (M, p) . This problem is known to be EXPTIME-HARD [3].

Idea of the reduction. Given a polynomial space ATM M and a word w , we construct a game of size polynomial in the size of (M, w) to simulate the execution of M on w . Player 1 makes choices in \vee -states and Player 2 makes choices in \wedge -states. Furthermore, Player 1 is responsible for maintaining the symbol under the tape head. The objective is to reach an accepting configuration of the ATM.

Each turn proceeds as follows. In an \vee -state, by choosing a letter (t, a) in the alphabet of the game, Player 1 reveals (i) the transition t of the ATM

that he has chosen (this way he also reveals the symbol that is currently under the tape head) and (ii) the symbol a under the next position of the tape head. If Player 1 lies about the current or the next symbol under the tape head, he should lose the game, otherwise the game proceeds. The machine is now in an \wedge -state and Player 1 has no choice: he announces a special symbol ϵ and Player 2, by resolving nondeterminism on ϵ , chooses a transition of the Turing machine which is compatible with the current symbol under the tape head revealed by Player 1 at the previous turn. The state of the ATM is updated and the game proceeds. The transition chosen by Player 2 is visible in the next state of the game and so Player 1 can update his knowledge about the configuration of the ATM. Player 1 wins whenever an accepting configuration of the ATM is reached, that is w is accepted.

The difficulty is to ensure that Player 1 loses when he announces a wrong content for the cell under the tape head. As the number of configurations of the polynomial ATM is exponential, we can not directly encode the full configuration of the ATM in the states of the game. To overcome this difficulty, we use the power of incomplete information as follows. Initially, Player 2 chooses a position k , $1 \leq k \leq p(|w|)$, on the tape: this number as well as the symbol $\sigma \in \{0, 1, 2\}$ that lies in the tape cell number k is maintained all along the game in the non-observable portion of the game states. The pair (σ, k) is thus private to Player 2 and invisible to Player 1. Thus, at any point in the game, Player 2 can check whether Player 1 is lying when announcing the content of cell number k , and go to a sink state if Player 1 cheats (no other states can be reached from there). Since Player 1 does not know which cell is monitored by Player 2 (k is private), to avoid losing, he should not lie about any of the tape cells and thus he should faithfully simulate the machine. Then, he wins the game if and only if the ATM accepts the words w .

Almost winning. To establish lower bound for almost-winning, we can use the same reduction. Randomization can not help Player I in this game. Indeed, at any point of the game, if Player I takes a chance in either not faithfully simulating the ATM or lying about the symbol under the tape head, the sink state is reached. In those case, the probability to reach the sink state is positive and so the probability to win the game is strictly less than one. We now present the details of the reduction of the hardness proof.

Reduction Given a polynomial space ATM (M, p) , with $M = \langle Q, q_0, g, \Sigma_i, \Sigma_t, \delta, F \rangle$ and a word w , we construct the following game structure $G_{M,p,w} = \langle L, l_0, \Sigma, \Delta, \text{Obs}, \gamma \rangle$, where:

- The set of positions $L = \{\text{init}\} \cup \{\text{sink}\} \cup L_1 \cup L_2$ where: $L_1 = (\delta \cup \{-\}) \times Q \times \{1, \dots, p(|w|)\} \times \{1, \dots, p(|w|)\} \times \Sigma_t$. A state (t, q, h, k, σ) consists of a transition t of the ATM chosen by Player 2 at the previous round or $-$ if this is the first round where Player 1 plays, the current control state q of M , the position h of the tape head, the pair (k, σ) such that the k -th symbol of the tape is σ , this pair (k, σ) will be kept invisible for Player 1. $L_2 = Q \times \{1, \dots, p(|w|)\} \times \Sigma_t \times \{1, \dots, p(|w|)\} \times \Sigma_t$. A state $(q, h, \gamma, k, \sigma)$ consists of q, h, k, σ as in L_1 and γ is the symbol that Player 1 claims to be under the tape head. The objective for Player 1 will be to reach a state $l \in L$ associated with an accepting control state of M .
- $l_0 = \text{init}$.
- $\Sigma = \{\epsilon\} \cup (\delta \times \Sigma_t)$.
- The transition relation Δ contains the following sets of transitions:
 - I_1 that contains transitions $(\text{init}, \epsilon, (-, q_0, 1, k, \sigma))$ where (i) $1 \leq k \leq p(|w|)$ and (ii) $\sigma = w(k)$ if $1 \leq k \leq |w|$ and $\sigma = 2$ otherwise.
 - I_2 that contains transitions $(\text{init}, (t, \gamma), \text{sink})$ where $(t, \gamma) \in \delta \times \Sigma_t$. So, at the initial state init Player 1 has to play ϵ in order to avoid entering sink . By resolving nondeterminism on ϵ , Player 2 chooses a tape cell to monitor.
 - S that contains transitions $(\text{sink}, \sigma, \text{sink})$ for all $\sigma \in \Sigma$. When the sink state is entered, the game stays there forever.
 - $L_{1.1}$ that contains transitions $((t, q, h, k, \sigma), \epsilon, \text{sink})$. $L_{1.2}$ that contains transitions $((t, q, h, k, \sigma), ((q_1, \gamma_1, q_2, \gamma_2, d), \gamma_3), \text{sink})$ where $q_1 \neq q$ or $\neg(1 \leq h + d \leq p(|w|))$. $L_{1.3}$ contains the transitions $((t, q, h, k, \sigma), ((q_1, \gamma_1, q_2, \gamma_2, d), \gamma_3), \text{sink})$ where $h = k \wedge \gamma_1 \neq \sigma$ or $h + d = k \wedge \gamma_3 \neq \sigma$. $L_{1.4}$ contains the transitions $((t, q, h, k, \sigma), ((q_1, \gamma_1, q_2, \gamma_2, d), \gamma_3), (q_2, h + d, \gamma_3, k, \sigma'))$ such that $q = q_1$, $1 \leq h + d \leq p(|w|)$, $h = k \rightarrow (\gamma_1 = \sigma \wedge \sigma' = \gamma_2)$, and $h \neq k \rightarrow \sigma' = \sigma$. Those transitions are associated with states of the game where Player 1 chooses a transition of the ATM to execute (if he proposes ϵ , the game evolves to the sink state, see $L_{1.1}$). The transition proposed by Player 1 should be valid for the

current control state of the ATM and the head should not exit the bounded tape after execution of the transition by the ATM, otherwise the game evolves to the sink state, see L.1.2. When choosing a letter, Player 1 also reveals the current letter under the tape head (given by the transition) as well as the letter under the next position of the tape head. If one of those positions is the one that is monitored by Player 2, the game evolves to the sink state in case Player 1 lies, see $L_{1.3}, L_{1.4}$.

- $L_{2.1}$ contains the transitions $((q, h, \gamma_1, k, \sigma), \epsilon, ((q_1, \gamma_2, q_2, \gamma_3, d), q_3, h + d, k, \sigma'))$ such that $q = q_1, q_2 = q_3, \gamma_1 = \gamma_2, 1 \leq h + d \leq p(|w|), h = k \rightarrow \sigma' = \gamma_3$, and $h \neq k \rightarrow \sigma' = \sigma$. $L_{2.2}$ contains the transitions $((q, h, \gamma_1, k, \sigma), \epsilon, \text{sink})$ such that there does not exist a transition $(q, \gamma_1, q_1, \gamma_2, d) \in \delta$ with $1 \leq h + d \leq p(|w|)$. $L_{2.3}$ contains the transitions $((q, h, \gamma_1, k, \sigma), (t, \gamma), \text{sink})$ where $(t, \gamma) \in \Sigma \setminus \{\epsilon\}$. Those transitions are associated with states of the game where Player 2 chooses the next transition of the ATM to execute. Player 1 should play ϵ otherwise the game goes to the sink state (see $L_{2.3}$, also the game goes to the sink state if there is no valid transition to execute in the ATM (see $L_{2.2}$). In the other cases, Player 1 proposes ϵ and Player 2 chooses a valid transition by resolving nondeterminism. The copy of the monitored cell is updated if necessary.

- $\text{Obs} = \{\text{init}, \text{sink}\} \cup \text{Obs}_1 \cup \text{Obs}_2$ where $\text{Obs}_1 = \{(t, q, h) \mid \exists (t, q, h, k, \sigma) \in L_1\}$ and $\text{Obs}_2 = \{(q, h, \gamma) \mid \exists (q, h, \gamma, k, \sigma) \in L_2\}$.
- γ is defined as follows: $\gamma(\text{init}) = \{\text{init}\}, \gamma(\text{sink}) = \{\text{sink}\}$, for all $(t, q, h) \in \text{Obs}_1, \gamma(t, q, h) = \{(t, q, h, k, \sigma) \in L_1\}$, for all $(q, h, \gamma) \in \text{Obs}_2, \gamma(q, h, \gamma) = \{(q, h, \gamma, k, \sigma) \in L_2\}$.

Finally, the objective ϕ of this game for Player 1 is to reach a state where the associated control state of the ATM is accepting, i.e. $\phi = \{o_1 o_2 \dots o_n \dots \in \text{Obs}^\omega \mid \exists i \geq 0 : (o_i = (t, q, h) \in \text{Obs}_1 \vee o_i = (q, h, \gamma) \in \text{Obs}_2) \wedge q \in F\}$.

It follows that Player 1 has an observation-based sure-winning (or almost-winning) strategy in the game $G_{M,p,w}$ for the objective ϕ iff the word w is accepted by the polynomial space ATM (M, p) .

Theorem 8 *Player 1 has a deterministic (resp. randomized) observation-based sure-winning (resp. almost-winning) strategy in the game $G_{M,p,w}$ for*

the objective ϕ iff the word w is accepted by the polynomial space ATM (M, p) .

Theorem 9 *Given a game structure of incomplete information G with reachability objective ϕ and a state ℓ , deciding (a) whether ℓ is a sure-winning state for ϕ in G is EXPTIME-HARD; and (b) whether ℓ is an almost-winning state for ϕ in G is EXPTIME-HARD.*

References

- [1] M. Abadi, L. Lamport, and P. Wolper. Realizable and unrealizable specifications of reactive systems. In *ICALP'89*, pages 1–17. Springer, 1989.
- [2] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *JACM*, 49:672–713, 2002.
- [3] A. K. Chandra, D. Kozen, and L. J. Stockmeyer. Alternation. *JACM*, 28(1):114–133, 1981.
- [4] K. Chatterjee and T.A. Henzinger. Semiperfect-information games. In *FSTTCS'05*, pages 1–18. Springer, 2005.
- [5] L. de Alfaro and T.A. Henzinger. Concurrent omega-regular games. In *LICS 00*, pages 141–154. IEEE Computer Society Press, 2000.
- [6] L. de Alfaro and T.A. Henzinger. Interface automata. In *FSE'01*, pages 109–120. ACM Press, 2001.
- [7] L. de Alfaro, T.A. Henzinger, and O. Kupferman. Concurrent reachability games. In *FOCS'98*, pages 564–575. IEEE Computer Society Press, 1998.
- [8] L. de Alfaro, T.A. Henzinger, and R. Majumdar. From verification to control: dynamic programs for omega-regular objectives. In *LICS'01*, pages 279–290. IEEE Computer Society Press, 2001.
- [9] M. De Wulf, L. Doyen, T.A. Henzinger and J.-F. Raskin. Antichains: a New Algorithm to Solve Universality of FA In *CAV*, Springer, To appear, 2006.
- [10] M. De Wulf, L. Doyen, and J.-F. Raskin. A lattice theory for solving games of imperfect information. In *HSCC*, pages 153–168, Springer, 2006.
- [11] D. Dill. *Trace Theory for Automatic Hierarchical Verification of Speed-independent Circuits*. The MIT Press, 1989.
- [12] E.A. Emerson and C.S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *FOCS'91*, pages 368–377. IEEE Computer Society Press, 1991.

- [13] A. Kechris. *Classical Descriptive Set Theory*. Springer, 1995.
- [14] O. Kupferman and M. Vardi. Synthesis with incomplete informatio. *ICTL '97*, 1997.
- [15] M.L. Littman. *Algorithms for sequential decision making*. PhD Thesis, Brown University, 1996.
- [16] D. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.
- [17] A. Pnueli and R. Rosner. On the synthesis of an asynchronous reactive module. In *ICALP'89*, pages 652–671, Springer, 1989.
- [18] P. Ramadge and W. Wonham. Supervisory control of a class of discrete-event processes. *SIAM Journal of Control and Optimization*, 25(1):206–230, 1987.
- [19] J. H. Reif. The complexity of two-player games of incomplete information. *JCSS*, 29(2):274–301, 1984.
- [20] W. Thomas. Languages, automata, and logic. In *Handbook of Formal Languages*, volume 3, Beyond Words, chapter 7, pages 389–455. Springer, 1997.
- [21] M. Vardi. Automatic verification of probabilistic concurrent finite-state systems. In *FOCS'85*, pages 327–338. IEEE Computer Society Press, 1985.