Algorithms for Omega-Regular Games with Imperfect Information ⋆,⋆⋆

Krishnendu Chatterjee[1], Laurent Doyen[2⋆⋆⋆], Thomas A. Henzinger[1,3], and Jean-François Raskin[2]

[1] EECS, University of California at Berkeley, U.S.A.
[2] CS, Université Libre de Bruxelles, Belgium
[3] I&C, École Polytechnique Fédérale de Lausanne, Switzerland

Abstract. We study observation-based strategies for two-player turn-based games on graphs with omega-regular objectives. An observation-based strategy relies on imperfect 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 fixed-point 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 fixed point 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 perfect information, randomized strategies are more powerful than deterministic ones. We show that our algorithms are optimal by proving matching lower bounds.

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 [18, 17]. Other applications of graph games

⋆ This research was supported in part by the NSF grants CCR-0225610 and CCR-0234690, by the SNSF under the Indo-Swiss Joint Research Programme, and by the FRFC project "Centre Fédéré en Vérification" funded by the FNRS under grant 2.4530.02.

⋆⋆ A fuller version with proofs is available as UC Berkeley Tech. Rep. EECS-2006-89.

⋆⋆⋆ Research fellow supported by the Belgian National Science Foundation (FNRS).

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 *perfect 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 imperfect 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 *imperfect information* are considerably more complicated than games of perfect information. First, decision problems for imperfect-information games usually lie in higher complexity classes than their perfect-information counter-parts [19, 14, 2]. The algorithmic difference is often exponential, due to a subset construction that, similar to the determinization of finite automata, turns an imperfect-information game into an equivalent perfect-information game. Second, because of the determinization, no symbolic algorithms are known to solve imperfect-information games. This is in contrast to the perfect-information case, where (often) simple and elegant fixed-point algorithms exist [12, 8]. Third, in the context of imperfect 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 perfect information, turn-based games of imperfect information require randomized strategies to win with probability 1 (see Example 1). Fourth, winning strategies for imperfect-information games need memory even for simple objectives such as safety and reachability (for an example see the technical-report version of this paper). This is again in contrast to the perfect-information case, where turn-based safety and reachability games can be won with memoryless strategies.

The contributions of this paper are twofold. First, we provide a symbolic fixed-point algorithm to solve games of imperfect 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.[4] 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 imperfect information into games of perfect 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

[4] We recently used this symbolic representation of $\subseteq$-downward-closed sets of state sets to propose a new algorithm for solving the universality problem of nondeterministic finite automata. First experiments show a very promising performance; (see [9]).

imperfect information. This is optimal, as the reachability problem for games of imperfect information is known to be EXPTIME-hard [19].

Second, we study randomized strategies and winning with probability 1 for imperfect-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 imperfect 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 perfect 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 imperfect information is EXPTIME-hard. The problem of computing the almost-winning states for coBüchi objectives under imperfect information in EXPTIME remains an open problem.

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 imperfect 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 imperfect information. For example, Kupferman and Vardi assert that imperfect 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 perfect-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 perfect information are PTIME-complete, whereas reachability games of imperfect information are EXPTIME-complete [19]. None of the above papers 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 threshold) is EXPTIME-complete [15]. However, almost winning is a qualitative question, and our hardness result for almost winning of imperfect-information games does not follow from the known results on POMDPs. We propose in Section 5 a new proof of the hardness for sure winning of imperfect-information games with reachability objectives, and we extend the proof to almost winning as well. To the best of our knowledge, this is the first hardness result that applies to the qualitative analysis of almost winning in imperfect-information games. A class of *semiperfect*-information games, where one player has imperfect information and the other player has perfect information, is studied in [4]. That class is simpler than the games studied here; it can be solved in NP $\cap$ coNP for parity objectives.

2 Definitions

A *game structure* (*of imperfect information*) is a tuple $G = \langle L, l_0, \Sigma, \Delta, \mathcal{O}, \gamma \rangle$, 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, $\mathcal{O}$ is a finite set of observations, and $\gamma : \mathcal{O} \to 2^L \setminus \emptyset$ maps each observation to the set of states that it represents. We require the following two properties on G: (i) for all $\ell \in L$ and all $\sigma \in \Sigma$, there exists $\ell' \in L$ such that $(\ell, \sigma, \ell') \in \Delta$; and (ii) the set $\{\gamma(o) \mid o \in \mathcal{O}\}$ partitions L. We say that G is a game structure of *perfect information* if $\mathcal{O} = L$ and $\gamma(\ell) = \{\ell\}$ for all $\ell \in L$. We often omit $(\mathcal{O}, \gamma)$ in the description of games of perfect information. For $\sigma \in \Sigma$ and $s \subseteq L$, let $\mathsf{Post}^G_\sigma(s) = \{\ell' \in L \mid \exists \ell \in s : (\ell, \sigma, \ell') \in \Delta\}$.

Plays. In a game structure, in each turn, Player 1 chooses a letter in Σ, and Player 2 resolves nondeterminism by choosing the successor state. 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 $i \geq 0$, we have $(\ell_i, \sigma_i, \ell_{i+1}) \in \Delta$. The *prefix up to* ℓ_n of the play π is denoted by $\pi(n)$; its length is $|\pi(n)| = n + 1$; and its last element is $\mathsf{Last}(\pi(n)) = \ell_n$. The *observation sequence* of π is the unique infinite sequence $\gamma^{-1}(\pi) = o_0 \sigma_0 o_1 \dots \sigma_{n-1} o_n \sigma_n \dots$ such that for all $i \geq 0$, we have $\ell_i \in \gamma(o_i)$. Similarly, the *observation sequence* of $\pi(n)$ is the prefix up to o_n of $\gamma^{-1}(\pi)$. The set of infinite plays in G is denoted $\mathsf{Plays}(G)$, and the set of corresponding finite prefixes is denoted $\mathsf{Prefs}(G)$. A state $\ell \in L$ is *reachable* in G if there exists a prefix $\rho \in \mathsf{Prefs}(G)$ such that $\mathsf{Last}(\rho) = \ell$. For a prefix $\rho \in \mathsf{Prefs}(G)$, the *cone* $\mathsf{Cone}(\rho) = \{\, \pi \in \mathsf{Plays}(G) \mid \rho \text{ is a prefix of } \pi \,\}$ is the set of plays that extend ρ. The *knowledge* associated with a finite observation sequence $\tau = o_0 \sigma_0 o_1 \sigma_1 \dots \sigma_{n-1} o_n$ is the set $\mathsf{K}(\tau)$ of states in which a play can be after this sequence of observations, that is, $\mathsf{K}(\tau) = \{\mathsf{Last}(\rho) \mid \rho \in \mathsf{Prefs}(G) \text{ and } \gamma^{-1}(\rho) = \tau\}$. For $\sigma \in \Sigma$, $\ell \in L$, and $\rho, \rho' \in \mathsf{Prefs}(G)$ with $\rho' = \rho \cdot \sigma \cdot \ell$, let $o_\ell \in \mathcal{O}$ be the unique observation such that $\ell \in \gamma(o_\ell)$. Then $\mathsf{K}(\gamma^{-1}(\rho')) = \mathsf{Post}^G_\sigma(\mathsf{K}(\gamma^{-1}(\rho))) \cap \gamma(o_\ell)$.

Strategies. A *deterministic strategy* in G for Player 1 is a function $\alpha : \mathsf{Prefs}(G) \to \Sigma$. For a finite set A, a probability distribution on A is a function $\kappa : A \to [0, 1]$ 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 $\mathsf{Supp}(\kappa) = \{a \in A \mid \kappa(a) > 0\}$ be the *support* of κ. A *randomized strategy* in G for Player 1 is a function $\alpha : \mathsf{Prefs}(G) \to \mathcal{D}(\Sigma)$. A (deterministic or randomized) strategy α for Player 1 is *observation-based* if for all prefixes $\rho, \rho' \in \mathsf{Prefs}(G)$, if $\gamma^{-1}(\rho) = \gamma^{-1}(\rho')$, 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 $\beta : \mathsf{Prefs}(G) \times \Sigma \to L$ such that for all $\rho \in \mathsf{Prefs}(G)$ and all $\sigma \in \Sigma$, we have $(\mathsf{Last}(\rho), \sigma, \beta(\rho, \sigma)) \in \Delta$. A *randomized strategy* in G for Player 2 is a function $\beta : \mathsf{Prefs}(G) \times \Sigma \to \mathcal{D}(L)$ such that for all $\rho \in \mathsf{Prefs}(G)$, all $\sigma \in \Sigma$, and all $\ell \in \mathsf{Supp}(\beta(\rho, \sigma))$, we have $(\mathsf{Last}(\rho), \sigma, \ell) \in \Delta$. We denote by $\mathcal{A}_G$, $\mathcal{A}^O_G$, 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 in G, respectively. All results of

this paper can be proved also if strategies depend on state sequences only, and not on the past moves of a play.

The *outcome* of two deterministic strategies α (for Player 1) and β (for Player 2) in G is the play $\pi = \ell_0\sigma_0\ell_1 \dots \sigma_{n-1}\ell_n\sigma_n \dots \in \mathsf{Plays}(G)$ 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 denoted $\mathsf{outcome}(G, \alpha, \beta)$. The *outcome* of two randomized strategies α (for Player 1) and β (for Player 2) in G is the set of plays $\pi = \ell_0\sigma_0\ell_1 \dots \sigma_{n-1}\ell_n\sigma_n \dots \in \mathsf{Plays}(G)$ such that for all $i \geq 0$, we have $\alpha(\pi(i))(\sigma_i) > 0$ and $\beta(\pi(i), \sigma_i)(\ell_{i+1}) > 0$. This set is denoted $\mathsf{outcome}(G, \alpha, \beta)$. The *outcome set* of the deterministic (resp. randomized) strategy α for Player 1 in G is the set $\mathsf{Outcome}_i(G, \alpha)$ of plays π such that there exists a deterministic (resp. randomized) strategy β for Player 2 with $\pi = \mathsf{outcome}(G, \alpha, \beta)$ (resp. $\pi \in \mathsf{outcome}(G, \alpha, \beta)$). The outcome sets for Player 2 are defined symmetrically.

Objectives. An *objective* for G is a set ϕ of infinite sequences of observations and input letters, that is, $\phi \subseteq (\mathcal{O} \times \Sigma)^\omega$. A play $\pi = \ell_0\sigma_0\ell_1 \dots \sigma_{n-1}\ell_n\sigma_n \dots \in \mathsf{Plays}(G)$ *satisfies* the objective ϕ, denoted $\pi \models \phi$, if $\gamma^{-1}(\pi) \in \phi$. Objectives are generally Borel measurable: a Borel objective is a Borel set in the Cantor topology on $(\mathcal{O} \times \Sigma)^\omega$ [13]. We specifically consider reachability, safety, Büchi, coBüchi, and parity objectives, all of them Borel measurable. 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 write $\mathsf{Inf}(\pi)$ for the set of observations that appear infinitely often in $\gamma^{-1}(\pi)$, that is, $\mathsf{Inf}(\pi) = \{o \in \mathcal{O} \mid \ell_i \in \gamma(o) \text{ for infinitely many } i\text{'s}\}$.

Given a set $\mathcal{T} \subseteq \mathcal{O}$ of target observations, the *reachability* objective $\mathsf{Reach}(\mathcal{T})$ requires that an observation in $\mathcal{T}$ be visited at least once, that is, $\mathsf{Reach}(\mathcal{T}) = \{\, \ell_0\sigma_0\ell_1\sigma_1 \dots \in \mathsf{Plays}(G) \mid \exists k \geq 0 \cdot \exists o \in \mathcal{T} : \ell_k \in \gamma(o) \,\}$. Dually, the *safety* objective $\mathsf{Safe}(\mathcal{T})$ requires that only observations in $\mathcal{T}$ be visited. Formally, $\mathsf{Safe}(\mathcal{T}) = \{\, \ell_0\sigma_0\ell_1\sigma_1 \dots \in \mathsf{Plays}(G) \mid \forall k \geq 0 \cdot \exists o \in \mathcal{T} : \ell_k \in \gamma(o) \,\}$. The *Büchi* objective $\mathsf{Buchi}(\mathcal{T})$ requires that an observation in $\mathcal{T}$ be visited infinitely often, that is, $\mathsf{Buchi}(\mathcal{T}) = \{\, \pi \mid \mathsf{Inf}(\pi) \cap \mathcal{T} \neq \emptyset \,\}$. Dually, the *coBüchi* objective $\mathsf{coBuchi}(\mathcal{T})$ requires that only observations in $\mathcal{T}$ be visited infinitely often. Formally, $\mathsf{coBuchi}(\mathcal{T}) = \{\, \pi \mid \mathsf{Inf}(\pi) \subseteq \mathcal{T} \,\}$. For $d \in \mathbb{N}$, let $p : \mathcal{O} \to \{\, 0, 1, \dots, d \,\}$ be a *priority function*, which maps each observation to a nonnegative integer priority. The *parity* objective $\mathsf{Parity}(p)$ requires that the minimum priority that appears infinitely often be even. Formally, $\mathsf{Parity}(p) = \{\, \pi \mid \min\{\, p(o) \mid o \in \mathsf{Inf}(\pi) \,\} \text{ is even} \,\}$. Observe that by definition, for all objectives ϕ, if $\pi \models \phi$ and $\gamma^{-1}(\pi) = \gamma^{-1}(\pi')$, then $\pi' \models \phi$.

Sure winning and almost winning. A strategy λ_i for Player i in G is *sure winning* for an objective ϕ if for all $\pi \in \mathsf{Outcome}_i(G, \lambda_i)$, we have $\pi \models \phi$. Given a game structure G and a state ℓ of G, we write G_ℓ for the game structure that results from G by changing the initial state to ℓ, that is, if $G = \langle L, l_0, \Sigma, \Delta, \mathcal{O}, \gamma \rangle$, then $G_\ell = \langle L, \ell, \Sigma, \Delta, \mathcal{O}, \gamma \rangle$. An *event* is a measurable set of plays, and given strategies α and β for the two players, the probabilities of events are uniquely defined [21]. For a Borel objective ϕ, we denote by $\Pr_\ell^{\alpha,\beta}(\phi)$ the probability ϕ is satisfied in the game G_ℓ given the strategies α and β. A strategy α for Player 1 in G is *almost winning* for the objective ϕ if for all randomized strategies β for

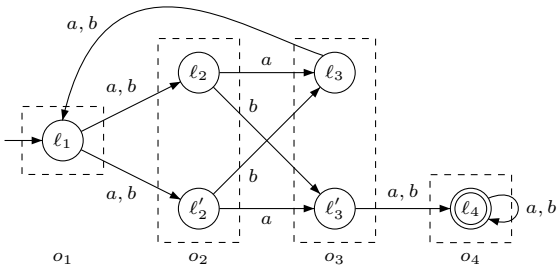

Fig. 1. Game structure G.

Player 2, we have $\Pr_{l_0}^{\alpha,\beta}(\phi) = 1$. The set of *sure-winning* (resp. *almost-winning*) *states* of a game structure 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 perfect-information game structures G and all Borel objectives ϕ, either there exists a deterministic sure-winning strategy for Player* 1 *for the objective ϕ, or there exists a deterministic sure-winning strategy for Player* 2 *for the complementary objective* $\mathsf{Plays}(G) \setminus \phi$.

Notice that deterministic strategies suffice for sure winning a game: given a randomized strategy α for Player 1, let α^D be the deterministic strategy such that for all $\rho \in \mathsf{Prefs}(G)$, the strategy $\alpha^D(\rho)$ chooses an input letter from $\mathsf{Supp}(\alpha(\rho))$. Then $\mathsf{Outcome}_1(G, \alpha^D) \subseteq \mathsf{Outcome}_1(G, \alpha)$, and thus, if α is sure winning, then so is α^D. The result also holds for observation-based strategies and for perfect-information games. 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 Fig. 1. The observations o_1, o_2, o_3, o_4 are such that $\gamma(o_1) = \{\ell_1\}$, $\gamma(o_2) = \{\ell_2, \ell'_2\}$, $\gamma(o_3) = \{\ell_3, \ell'_3\}$, and $\gamma(o_4) = \{\ell_4\}$. The transitions are shown as labeled edges in the figure, and the initial state is ℓ_1. The objective of Player 1 is $\mathsf{Reach}(\{o_4\})$, to reach state ℓ_4. We argue that the game is not sure winning for Player 1. Let α be any deterministic strategy for Player 1. Consider the deterministic strategy β for Player 2 as follows: for all $\rho \in \mathsf{Prefs}(G)$ such that $\mathsf{Last}(\rho) \in \gamma(o_2)$, if $\alpha(\rho) = a$, then in the previous round β chooses the state ℓ_2, and if $\alpha(\rho) = b$, then in the previous round β chooses the state ℓ'_2. Given α and β, the play $\mathsf{outcome}(G, \alpha, \beta)$ never reaches ℓ_4. However, 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 follows that against all Player 2 strategies the play eventually reaches ℓ_4 with probability 1.

Remarks. First, the hypothesis that the observations form a partition of the state space can be weakened to a covering of the state space, where observations can overlap. In that case, Player 2 chooses both the next state of the game ℓ and the next observation o such that $\ell \in \gamma(o)$. The definitions related to plays, strategies, and objectives are adapted accordingly. Such a game structure G with overlapping observations can be encoded by an equivalent game structure G' of imperfect information, whose state space is the set of pairs (ℓ, o) such that $\ell \in \gamma(o)$. The set of labeled transitions Δ' of G' is defined by $\Delta' = \{((\ell, o), \sigma, (\ell', o')) \mid (\ell, \sigma, \ell') \in \Delta\}$ and $\gamma'^{-1}(\ell, o) = o$. The games G and G' are equivalent in the sense that for every Borel objective ϕ, there exists a sure (resp. almost) winning strategy for Player i in G forϕ if and only if there exists such a winning strategy for Player i in G' for ϕ. Second, it is essential that the objective is expressed in terms of the observations. Indeed, the games of imperfect information with a nonobservable winning condition are more complicated to solve. For instance, the universality problem for Büchi automata can be reduced to such games, but the construction that we propose in Section 3 cannot be used.

3 Sure Winning

We show that a game structure G of imperfect information can be encoded by a game structure G^{K} of perfect information such that for all Borel objectives 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} by a subset construction. Each state in G^{K} is a set of states of G which represents 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 the subset construction, but avoids the explicit construction of G^{K}.

3.1 Subset construction for sure winning

Given a game structure of imperfect information $G = \langle L, l_0, \Sigma, \Delta, \mathcal{O}, \gamma \rangle$, we define the *knowledge-based subset construction* of G as the following game structure of perfect information: $G^{\mathsf{K}} = \langle \mathcal{L}, \{l_0\}, \Sigma, \Delta^{\mathsf{K}} \rangle$, where $\mathcal{L} = 2^L \setminus \{\emptyset\}$, and $(s_1, \sigma, s_2) \in \Delta^{\mathsf{K}}$ iff there exists an observation $o \in \mathcal{O}$ such that $s_2 = \mathsf{Post}^G_\sigma(s_1) \cap \gamma(o)$ and $s_2 \neq \emptyset$. Notice that for all $s \in \mathcal{L}$ and all $\sigma \in \Sigma$, there exists a set $s' \in \mathcal{L}$ such that $(s, \sigma, s') \in \Delta^{\mathsf{K}}$.

A (deterministic or randomized) strategy in G^{K} is called a *knowledge-based* strategy. For all sets $s \in \mathcal{L}$ that are reachable in G^{K}, and all observations $o \in \mathcal{O}$, either $s \subseteq \gamma(o)$ or $s \cap \gamma(o) = \emptyset$. By an abuse of notation, we define the *observation sequence* of a play $\pi = s_0\sigma_0 s_1 \dots \sigma_{n-1} s_n \sigma_n \dots \in \mathsf{Plays}(G^{\mathsf{K}})$ as the infinite sequence $\gamma^{-1}(\pi) = o_0\sigma_0 o_1 \dots \sigma_{n-1} o_n \sigma_n \dots$ of observations such that for all $i \geq 0$, we have $s_i \subseteq \gamma(o_i)$; this sequence is unique. The play π *satisfies* an objective $\phi \subseteq (\mathcal{O} \times \Sigma)^\omega$ if $\gamma^{-1}(\pi) \in \phi$. The proof of the following theorem can be found in the technical-report version for this paper.

Theorem 2 (Sure-winning reduction). *Player* 1 *has a deterministic observation-based sure-winning strategy in a game structure G of imperfect information for a Borel objective ϕ if and only if Player* 1 *has a deterministic sure-winning strategy in the game structure G^{K} of perfect information for ϕ.*

3.2 Two interpretations of the μ-calculus

Form the results of Section 3.1, we can solve a game G of imperfect information with objective ϕ by constructing the knowledge-based subset construction G^{K} and solving the resulting game of perfect 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}, these 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* with respect to set inclusion: if Player 1 has a deterministic strategy to win the game G when her knowledge is a 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 perfect 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 $s \not\subset s'$. Let $\mathcal{C}$ be the set of antichains of nonempty subsets of L, and consider the following partial order on $\mathcal{C}$: for all $q, q' \in \mathcal{C}$, let $q \sqsubseteq q'$ iff $\forall s \in q \cdot \exists s' \in q' : s \subseteq s'$. For $q \subseteq 2^L \setminus \emptyset$, define the set of *maximal* elements of q by $\lceil q \rceil = \{s \in q \mid s \neq \emptyset \text{ and } \forall s' \in q : s \not\subset s'\}$. Clearly, $\lceil q \rceil$ is an antichain. The least upper bound of $q, q' \in \mathcal{C}$ is $q \sqcup q' = \lceil \{s \mid s \in q \text{ or } s \in q'\} \rceil$, and their greatest lower bound is $q \sqcap q' = \lceil \{s \cap s' \mid s \in q \text{ and } s' \in q'\} \rceil$. The definition of these two operators extends naturally to sets of antichains, and the greatest element of $\mathcal{C}$ is $\top = \{L\}$ and the least element is $\bot = \emptyset$. The partially ordered set $\langle \mathcal{C}, \sqsubseteq, \sqcup, \sqcap, \top, \bot \rangle$ forms a complete lattice. We view antichains of state sets as a symbolic representation of $\subseteq$-downward-closed sets of state sets.

A *game lattice* is a complete lattice V together with a *predecessor operator* $\mathsf{CPre} : V \to V$. Given a game structure $G = \langle L, l_0, \Sigma, \Delta, \mathcal{O}, \gamma \rangle$ of imperfect information, and its knowledge-based subset construction $G^{\mathsf{K}} = \langle \mathcal{L}, \{l_0\}, \Sigma, \Delta^{\mathsf{K}} \rangle$, we consider two game lattices: the *lattice of subsets* $\langle \mathcal{S}, \subseteq, \cup, \cap, \mathcal{L}, \emptyset \rangle$, where $\mathcal{S} = 2^{\mathcal{L}}$ and $\mathsf{CPre} : \mathcal{S} \to \mathcal{S}$ is defined by $\mathsf{CPre}(q) = \{s \in \mathcal{L} \mid \exists \sigma \in \Sigma \cdot \forall s' \in \mathcal{L} : \text{ if } (s, \sigma, s') \in \Delta^{\mathsf{K}}, \text{ then } s' \in q\}$; and the *lattice of antichains* $\langle \mathcal{C}, \sqsubseteq, \sqcup, \sqcap, \{L\}, \emptyset \rangle$, with the operator $\lceil \mathsf{CPre} \rceil : \mathcal{C} \to \mathcal{C}$ defined by $\lceil \mathsf{CPre} \rceil(q) = \lceil \{s \in \mathcal{L} \mid \exists \sigma \in \Sigma \cdot \forall o \in \mathcal{O} \cdot \exists s' \in q : \mathsf{Post}_\sigma(s) \cap \gamma(o) \subseteq s'\} \rceil$.

The *μ-calculus formulas* are generated by the grammar

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

for atomic propositions $o \in \mathcal{O}$ and variables x. We can define $\neg o$ as a shortcut for $\bigvee_{o' \in \mathcal{O} \setminus \{o\}} o'$. A variable is *free* in a formula φ if it is not in the scope of a quantifier 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

Lattice of subsets
$[\![o]\!]^{\mathcal{S}}_{\mathcal{E}} = \{s \in \mathcal{L} \mid s \subseteq \gamma(o)\}$
$[\![x]\!]^{\mathcal{S}}_{\mathcal{E}} = \mathcal{E}(x)$
$[\![\varphi_1 \left\{\begin{smallmatrix}\vee\\ \wedge\end{smallmatrix}\right\} \varphi_2]\!]^{\mathcal{S}}_{\mathcal{E}} = [\![\varphi_1]\!]^{\mathcal{S}}_{\mathcal{E}} \left\{\begin{smallmatrix}\cup\\ \cap\end{smallmatrix}\right\} [\![\varphi_2]\!]^{\mathcal{S}}_{\mathcal{E}}$
$[\![\mathsf{pre}(\varphi)]\!]^{\mathcal{S}}_{\mathcal{E}} = \mathsf{CPre}([\![\varphi]\!]^{\mathcal{S}}_{\mathcal{E}})$
$[\![\left\{\begin{smallmatrix}\mu\\ \nu\end{smallmatrix}\right\} x.\varphi]\!]^{\mathcal{S}}_{\mathcal{E}} = \left\{\begin{smallmatrix}\cap\\ \cup\end{smallmatrix}\right\} \{q \mid q = [\![\varphi]\!]^{\mathcal{S}}_{\mathcal{E}[x \mapsto q]}\}$

Lattice of antichains
$[\![o]\!]^{\mathcal{C}}_{\mathcal{E}} = \{\gamma(o)\}$
$[\![x]\!]^{\mathcal{C}}_{\mathcal{E}} = \mathcal{E}(x)$
$[\![\varphi_1 \left\{\begin{smallmatrix}\vee\\ \wedge\end{smallmatrix}\right\} \varphi_2]\!]^{\mathcal{C}}_{\mathcal{E}} = [\![\varphi_1]\!]^{\mathcal{C}}_{\mathcal{E}} \left\{\begin{smallmatrix}\sqcup\\ \sqcap\end{smallmatrix}\right\} [\![\varphi_2]\!]^{\mathcal{C}}_{\mathcal{E}}$
$[\![\mathsf{pre}(\varphi)]\!]^{\mathcal{C}}_{\mathcal{E}} = \lceil \mathsf{CPre} \rceil([\![\varphi]\!]^{\mathcal{C}}_{\mathcal{E}})$
$[\![\left\{\begin{smallmatrix}\mu\\ \nu\end{smallmatrix}\right\} x.\varphi]\!]^{\mathcal{C}}_{\mathcal{E}} = \left\{\begin{smallmatrix}\sqcap\\ \sqcup\end{smallmatrix}\right\} \{q \mid q = [\![\varphi]\!]^{\mathcal{C}}_{\mathcal{E}[x \mapsto q]}\}$

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 $[\![\varphi]\!]^V_{\mathcal{E}}$ of V, which is defined inductively by the equations shown in the two tables. If φ is a closed formula, then $[\![\varphi]\!]^V = [\![\varphi]\!]^V_{\mathcal{E}}$ for any valuation $\mathcal{E}$. The following theorem recalls that perfect-information games can be solved by evaluating fixed-point formulas in the lattice of subsets.

Theorem 3 (Symbolic solution of perfect-information games). [8] *For every ω-regular objective ϕ, there exists a closed μ-calculus formula $\mu\mathsf{Form}(\phi)$, called the* characteristic formula *of ϕ, such that for all game structures G of perfect information, the set of sure-winning states of G for ϕ is $[\![\mu\mathsf{Form}(\phi)]\!]^{\mathcal{S}}$.*

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 $q{\downarrow}$, for $q \in \mathcal{S}$, are the *downward-closed* sets. A valuation $\mathcal{E}$ for the variables in the lattice $\mathcal{S}$ of subsets is *downward closed* if every variable x is mapped to a downward-closed set, that is, $\mathcal{E}(x) = \mathcal{E}(x){\downarrow}$.

Lemma 1. *For all downward-closed sets $q, q' \in \mathcal{S}$, we have $\lceil q \cap q' \rceil = \lceil q \rceil \sqcap \lceil q' \rceil$ and $\lceil q \cup q' \rceil = \lceil q \rceil \sqcup \lceil q' \rceil$.*

Lemma 2. *For all μ-calculus formulas φ and all downward-closed valuations $\mathcal{E}$ in the lattice of subsets, the set $[\![\varphi]\!]^{\mathcal{S}}_{\mathcal{E}}$ is downward closed.*

Lemma 3. *For all μ-calculus formulas φ, and all downward-closed valuations $\mathcal{E}$ in the lattice of subsets, we have $\lceil [\![\varphi]\!]^{\mathcal{S}}_{\mathcal{E}} \rceil = [\![\varphi]\!]^{\mathcal{C}}_{\lceil \mathcal{E} \rceil}$, 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.*

Consider a game structure G of imperfect information and a parity objective ϕ. From Theorems 2 and 3 and Lemma 3, we can decide the existence of a deterministic observation-based sure-winning strategy for Player 1 in G for ϕ without explicitly constructing the knowledge-based subset construction G^{K}, by instead evaluating a fixed-point formula in the lattice of antichains.

Theorem 4 (Symbolic solution of imperfect-information games). *Let G be a game structure of imperfect information with initial state l_0. For every*

ω-regular objective ϕ, Player 1 has a deterministic observation-based strategy in G for ϕ if and only if $\{l_0\} \sqsubseteq [\![\mu \mathsf{Form}(\phi)]\!]^{\mathcal{C}}$.

Corollary 1. *Let G be a game structure of imperfect information, let p be a priority function, and let ℓ be a state of G. Whether ℓ is a sure-winning state in G for the parity objective $\mathsf{Parity}(p)$ can be decided in* EXPTIME.

Corollary 1 is proved as follows: for a parity objective ϕ, an equivalent μ-calculus formula φ can be obtained, where the size and the fixed-point quantifier alternations of φ is polynomial in ϕ. Thus given G and ϕ, we can evaluate φ in G^{K} in EXPTIME.

4 Almost Winning

Given a game structure G of imperfect information, we first construct a game structure H of perfect information by a subset construction (different from the one used for sure winning), and then establish certain equivalences between randomized strategies in G and H. Finally, we show how the reduction can be used to obtain a symbolic EXPTIME algorithm for computing almost-winning states in G for Büchi objectives. An EXPTIME algorithm for almost winning for coBüchi objectives under imperfect information remains unknown.

4.1 Subset construction for almost winning

Given a game structure of imperfect information $G = \langle L, l_0, \Sigma, \Delta, \mathcal{O}, \gamma \rangle$, we construct game structure of perfect information $H = \mathsf{Pft}(G) = \langle Q, q_0, \Sigma, \Delta_H \rangle$ as follows: $Q = \{ (s, \ell) \mid \exists o \in \mathcal{O} : s \subseteq \gamma(o) \text{ and } \ell \in s \}$; the initial state is $q_0 = (\{l_0\}, l_0)$; the transition relation $\Delta_H \subseteq Q \times \Sigma \times Q$ is defined by $((s, \ell), \sigma, (s', \ell')) \in \Delta_H$ iff there is an observation $o \in \mathcal{O}$ such that $s' = \mathsf{Post}^G_\sigma(s) \cap \gamma(o)$ and $(\ell, \sigma, \ell') \in \Delta$. Intuitively, when H is in state (s, ℓ), it corresponds to G being in state ℓ and the knowledge of Player 1 being s. Two states $q = (s, \ell)$ and $q' = (s', \ell')$ of H are *equivalent*, written $q \approx q'$, if $s = s'$. Two prefixes $\rho = q_0 \sigma_0 q_1 \dots \sigma_{n-1} q_n$ and $\rho' = q'_0 \sigma'_0 q'_1 \dots \sigma'_{n-1} q'_n$ of H are *equivalent*, written $\rho \approx \rho$, if 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 $\pi, \pi' \in \mathsf{Plays}(H)$ are *equivalent*, written $\pi_H \approx \pi'_H$, if for all $i \geq 0$, we have $\pi(i) \approx \pi'(i)$. 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 equivalence classes for prefixes and plays similarly.

Equivalence-preserving strategies and objectives. A strategy α for Player 1 in H is *positional* if it is independent of the prefix of plays and depends only on the last state, that is, for all $\rho, \rho' \in \mathsf{Prefs}(H)$ with $\mathsf{Last}(\rho) = \mathsf{Last}(\rho')$, we have $\alpha(\rho) = \alpha(\rho')$. A positional strategy α can be viewed as a function $\alpha : Q \to \mathcal{D}(\Sigma)$. A strategy α for Player 1 in H is *equivalence-preserving* if for all $\rho, \rho' \in \mathsf{Prefs}(H)$ with $\rho \approx \rho'$, we have $\alpha(\rho) = \alpha(\rho')$. We denote by $\mathcal{A}_H$, $\mathcal{A}^P_H$, and $\mathcal{A}^{\approx}_H$ the set of all Player-1 strategies, the set of all positional Player-1 strategies, and the set

of all equivalence-preserving Player-1 strategies in H, respectively. We write $\mathcal{A}_H^{\approx(P)} = \mathcal{A}_H^{\approx} \cap \mathcal{A}_H^P$ for the set of equivalence-preserving positional strategies.

An objective ϕ for H is a subset of $(Q \times \Sigma)^\omega$, that is, the objective ϕ is a set of plays. The objective ϕ is *equivalence-preserving* if for all plays $\pi \in \phi$, we have $[\pi]_\approx \subseteq \phi$.

Relating prefixes and plays. We define a mapping $h : \mathsf{Prefs}(G) \to \mathsf{Prefs}(H)$ that maps prefixes in G to prefixes in H as follows: given $\rho = \ell_0\sigma_0\ell_1\sigma_1 \dots \sigma_{n-1}\ell_n$, let $h(\rho) = q_0\sigma_0 q_1\sigma_1 \dots \sigma_{n-1}q_n$, where for all $0 \leq i \leq n$, we have $q_i = (s_i, \ell_i)$, and for all $0 \leq i \leq n-1$, we have $s_i = \mathsf{K}(\gamma^{-1}(\rho(i)))$. The following properties hold: (i) for all $\rho, \rho' \in \mathsf{Prefs}(G)$, if $\gamma^{-1}(\rho) = \gamma^{-1}(\rho')$, then $h(\rho) \approx h(\rho')$; and (ii) for all $\rho, \rho' \in \mathsf{Prefs}(H)$, if $\rho \approx \rho'$, then $\gamma^{-1}(h^{-1}(\rho)) = \gamma^{-1}(h^{-1}(\rho'))$. The mapping $h : \mathsf{Plays}(G) \to \mathsf{Plays}(H)$ for plays is defined similarly, and has similar properties.

Relating strategies for Player 1. We define two strategy mappings $h : \mathcal{A}_H \to \mathcal{A}_G$ and $g : \mathcal{A}_G \to \mathcal{A}_H$. Given a Player-1 strategy α_H in H, we construct a Player-1 strategy $\alpha_G = h(\alpha_H)$ in G as follows: for all $\rho \in \mathsf{Prefs}(G)$, let $\alpha_G(\rho) = \alpha_H(h(\rho))$. Similarly, given a Player-1 strategy α_G in G, we construct a Player-1 strategy $\alpha_H = g(\alpha_G)$ in H as follows: for all $\rho \in \mathsf{Prefs}(H)$, let $\alpha_H(\rho) = \alpha_G(h^{-1}(\rho))$. The following properties hold: (i) for all strategies $\alpha_H \in \mathcal{A}_H$, if α_H is equivalence-preserving, then $h(\alpha_H)$ is observation-based; and (ii) for all strategies $\alpha_G \in \mathcal{A}_G$, if α_G is observation-based, then $g(\alpha_G)$ is equivalence-preserving.

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

Lemma 4. *For all $\rho \in \mathsf{Prefs}(H)$, for every equivalence-preserving strategy α of Player 1 in H, and for every strategy β of Player 2 in H, we have $\Pr_{q_0}^{\alpha,\beta}(\mathsf{Cone}(\rho)) = \Pr_{l_0}^{h(\alpha),h(\beta)}(h^{-1}(\mathsf{Cone}(\rho)))$.*

Lemma 5. *For all $\rho \in \mathsf{Prefs}(G)$, for every observational strategy α of Player 1 in G, and for every strategy β of Player 2 in G, we have $\Pr_{l_0}^{\alpha,\beta}(\mathsf{Cone}(\rho)) = \Pr_{q_0}^{g(\alpha),g(\beta)}(h(\mathsf{Cone}(\rho)))$.*

Theorem 5 (Almost-winning reduction). *Let G be a game structure of imperfect information, and let $H = \mathsf{Pft}(G)$ be the game structure of perfect information. For all Borel objectives ϕ for G, all observation-based Player-1 strategies α in G, and all Player-2 strategies β in G, we have $\Pr_{l_0}^{\alpha,\beta}(\phi) = \Pr_{q_0}^{g(\alpha),g(\beta)}(h(\phi))$. Dually, for all equivalence-preserving Borel objectives ϕ for H, all equivalence-preserving Player-1 strategies α in H, and all Player-2 strategies β in H, we have $\Pr_{q_0}^{\alpha,\beta}(\phi) = \Pr_{l_0}^{h(\alpha),h(\beta)}(h^{-1}(\phi))$.*

The proof is as follows: by the Caratheódary unique-extension theorem, a probability measure defined on cones has a unique extension to all Borel objectives. The theorem then follows from Lemma 4.

Corollary 2. *For every Borel objective ϕ for G, we have $\exists \alpha_G \in \mathcal{A}_G^O \cdot \forall \beta_G \in \mathcal{B}_G : \Pr_{\ell_0}^{\alpha_G,\beta_G}(\phi) = 1$ if and only if $\exists \alpha_H \in \mathcal{A}_H^{\approx} \cdot \forall \beta_H \in \mathcal{B}_H : \Pr_{q_0}^{\alpha_H,\beta_H}(h(\phi)) = 1$.*

4.2 Almost winning for Büchi objectives

Given a game structure G of imperfect information, let $H = \mathsf{Pft}(G)$ be the game structure of perfect information. Given a set $\mathcal{T} \subseteq \mathcal{O}$ of target observations, let $B_{\mathcal{T}} = \{(s,l) \in Q \mid \exists o \in \mathcal{T} : s \subseteq \gamma(o)\}$. Then $h(\mathsf{Buchi}(\mathcal{T})) = \mathsf{Buchi}(B_{\mathcal{T}}) = \{\pi_H \in \mathsf{Plays}(H) \mid \mathsf{Inf}(\pi_H) \cap B_{\mathcal{T}} \neq \emptyset\}$. We first show that almost winning in H for the Büchi objective $\mathsf{Buchi}(B_{\mathcal{T}})$ with respect to equivalence-preserving strategies is equivalent to almost winning with respect to equivalence-preserving positional strategies. Formally, for $B_{\mathcal{T}} \subseteq Q$, let $Q_{\mathsf{AS}}^{\approx} = \{ q \in Q \mid \exists \alpha \in \mathcal{A}_H^{\approx} \cdot \forall \beta \in \mathcal{B}_H \cdot \forall q' \in [q]_{\approx} : \Pr_{q'}^{\alpha,\beta}(\mathsf{Buchi}(B_{\mathcal{T}})) = 1 \}$, and $Q_{\mathsf{AS}}^{\approx(P)} = \{ q \in Q \mid \exists \alpha \in \mathcal{A}_H^{\approx(P)} \cdot \forall \beta \in \mathcal{B}_H \cdot \forall q' \in [q]_{\approx} : \Pr_{q'}^{\alpha,\beta}(\mathsf{Buchi}(B_{\mathcal{T}})) = 1 \}$. We will prove that $Q_{\mathsf{AS}}^{\approx} = Q_{\mathsf{AS}}^{\approx(P)}$. Lemma 6 follows from the construction of H from G.

Lemma 6. *Given an equivalence-preserving Player-1 strategy $\alpha \in \mathcal{A}_H$, a prefix $\rho \in \mathsf{Prefs}(H)$, and a state $q \in Q$, if there exists a Player-2 strategy $\beta \in \mathcal{B}_H$ such that $\Pr_q^{\alpha,\beta}(\mathsf{Cone}(\rho)) > 0$, then for every prefix $\rho' \in \mathsf{Prefs}(H)$ with $\rho \approx \rho'$, there exist a Player-2 strategy $\beta' \in \mathcal{B}_H$ and a state $q' \in [q]_{\approx}$ such that $\Pr_{q'}^{\alpha,\beta'}(\mathsf{Cone}(\rho')) > 0$.*

Observe that $Q \setminus Q_{\mathsf{AS}}^{\approx} = \{ q \in Q \mid \forall \alpha \in \mathcal{A}_H^{\approx} \cdot \exists \beta \in \mathcal{B}_H \cdot \exists q' \in [q]_{\approx} : \Pr_{q'}^{\alpha,\beta}(\mathsf{Buchi}(B_{\mathcal{T}})) < 1 \}$. It follows from Lemma 6 that if a play starts in $Q_{\mathsf{AS}}^{\approx}$ and reaches $Q \setminus Q_{\mathsf{AS}}^{\approx}$ with positive probability, then for all equivalence-preserving strategies for Player 1, there is a Player 2 strategy that ensures that the Büchi objective $\mathsf{Buchi}(B_{\mathcal{T}})$ is not satisfied with probability 1.

Notation. For a state $q \in Q$ and $Y \subseteq Q$, let $\mathsf{Allow}(q, Y) = \{ \sigma \in \Sigma \mid \mathsf{Post}_\sigma^H(q) \subseteq Y \}$. For a state $q \in Q$ and $Y \subseteq Q$, let $\mathsf{Allow}([q]_{\approx}, Y) = \bigcap_{q' \in [q]_{\approx}} \mathsf{Allow}(q', Y)$.

Lemma 7. *For all $q \in Q_{\mathsf{AS}}^{\approx}$, we have $\mathsf{Allow}([q]_{\approx}, Q_{\mathsf{AS}}^{\approx}) \neq \emptyset$.*

Lemma 8. *Given a state $q \in Q_{\mathsf{AS}}^{\approx}$, let $\alpha \in \mathcal{A}_H$ be an equivalence-preserving Player-1 strategy such that for all Player-2 strategies $\beta \in \mathcal{B}_H$ and all states $q' \in [q]_{\approx}$, we have $\Pr_{q'}^{\alpha,\beta}(\mathsf{Buchi}(B_{\mathcal{T}})) = 1$. Let $\rho = q_0\sigma_0 q_1 \dots \sigma_{n-1} q_n$ be a prefix in $\mathsf{Prefs}(H)$ such that for all $0 \leq i \leq n$, we have $q_i \in Q_{\mathsf{AS}}^{\approx}$. If there is a Player-2 strategy $\beta \in \mathcal{B}_H$ and a state $q' \in [q]_{\approx}$ such that $\Pr_{q'}^{\alpha,\beta}(\mathsf{Cone}(\rho)) > 0$, then $\mathsf{Supp}(\alpha(\rho)) \subseteq \mathsf{Allow}([q]_{\approx}, Q_{\mathsf{AS}}^{\approx})$.*

Notation. We inductively define the *ranks* of states in $Q_{\mathsf{AS}}^{\approx}$ as follows: let $\mathsf{Rank}(0) = B_{\mathcal{T}} \cap Q_{\mathsf{AS}}^{\approx}$, and for all $j \geq 0$, let $\mathsf{Rank}(j+1) = \mathsf{Rank}(j) \cup \{ q \in Q_{\mathsf{AS}}^{\approx} \mid \exists \sigma \in \mathsf{Allow}([q]_{\approx}, Q_{\mathsf{AS}}^{\approx}) : \mathsf{Post}_\sigma^H(q) \subseteq \mathsf{Rank}(j) \}$. Let $j^* = \min\{ j \geq 0 \mid \mathsf{Rank}(j) = \mathsf{Rank}(j+1) \}$, and let $Q^* = \mathsf{Rank}(j^*)$. We say that the set $\mathsf{Rank}(j+1) \setminus \mathsf{Rank}(j)$ contains the *states of rank* $j+1$, for all $j \geq 0$.

Lemma 9. $Q^* = Q^{\approx}_{\mathsf{AS}}$.

Equivalence-preserving positional strategy. Consider the equivalence-preserving positional strategy α^p for Player 1 in H, which is defined as follows: for a state $q \in Q^{\approx}_{\mathsf{AS}}$, choose all moves in $\mathsf{Allow}([q]_{\equiv}, Q^{\approx}_{\mathsf{AS}})$ uniformly at random.

Lemma 10. *For all states $q \in Q^{\approx}_{\mathsf{AS}}$ and all Player-2 strategies β in H, we have $\Pr_q^{\alpha^p,\beta}(\mathsf{Safe}(Q^{\approx}_{\mathsf{AS}})) = 1$ and $\Pr_q^{\alpha^p,\beta}(\mathsf{Reach}(B_{\mathcal{T}} \cap Q^{\approx}_{\mathsf{AS}})) = 1$.*

Proof. By Lemma 9, we have $Q^* = Q^{\approx}_{\mathsf{AS}}$. Let $z = |Q^*|$. For a state $q \in Q^{\approx}_{\mathsf{AS}}$, we have $\mathsf{Post}^H_\sigma(q) \subseteq Q^{\approx}_{\mathsf{AS}}$ for all $\sigma \in \mathsf{Allow}([q]_{\approx}, Q^{\approx}_{\mathsf{AS}})$. It follows for all states $q \in Q^{\approx}_{\mathsf{AS}}$ and all strategies β for Player 2, we have $\Pr_q^{\alpha^p,\beta}(\mathsf{Safe}(Q^{\approx}_{\mathsf{AS}})) = 1$.

For a state $q \in (\mathsf{Rank}(j+1) \setminus \mathsf{Rank}(j))$, there exists $\sigma \in \mathsf{Allow}([q]_{\approx}, Q^{\approx}_{\mathsf{AS}})$ such that $\mathsf{Post}^H_\sigma(q) \subseteq \mathsf{Rank}(j)$. For a set $Y \subseteq Q$, let $\Diamond^j(Y)$ denote the set of prefixes that reach Y within j steps. It follows that for all states $q \in \mathsf{Rank}(j+1)$ and all strategies β for Player 2, we have $\Pr_q^{\alpha^p,\beta}(\Diamond^1(\mathsf{Rank}(j))) \geq \frac{1}{|\Sigma|}$. Let $B = B_{\mathcal{T}} \cap Q^{\approx}_{\mathsf{AS}}$. By induction on the ranks it follows that for all states $q \in Q^*$ and all strategies β for Player 2: $\Pr_q^{\alpha^p,\beta}(\Diamond^z(\mathsf{Rank}(0))) = \Pr_q^{\alpha^p,\beta}(\Diamond^z(B)) \geq \left(\frac{1}{|\Sigma|}\right)^z = r > 0$. For $m > 0$, we have $\Pr_q^{\alpha^p,\beta}(\Diamond^{m \cdot z}(B)) \geq 1 - (1-r)^m$. Thus:

$$\Pr_q^{\alpha^p,\beta}(\mathsf{Reach}(B)) = \lim_{m\to\infty} \Pr_q^{\alpha^p,\beta}(\Diamond^{m\cdot z}(B)) \geq \lim_{m\to\infty} 1 - (1-r)^m = 1. \qquad \blacksquare$$

Lemma 10 implies that, given the Player-1 strategy α^p, the set $Q^{\approx}_{\mathsf{AS}}$ is never left, and the states in $B_{\mathcal{T}} \cap Q^{\approx}_{\mathsf{AS}}$ are reached with probability 1. Since this happens for every state in $Q^{\approx}_{\mathsf{AS}}$, it follows that the set $B_{\mathcal{T}} \cap Q^{\approx}_{\mathsf{AS}}$ is visited infinitely often with probability 1, that is, the Büchi objective $\mathsf{Buchi}(B_{\mathcal{T}})$ is satisfied with probability 1. This analysis, together with the fact that $[q_0]_{\approx}$ is a singleton and Corollary 2, proves that $Q^{\approx}_{\mathsf{AS}} = Q^{\approx(P)}_{\mathsf{AS}}$. Theorem 6 follows.

Theorem 6 (Positional almost winning for Büchi objectives under imperfect information). *Let G be a game structure of imperfect information, and let $H = \mathsf{Pft}(G)$ be the game structure of perfect information. For all sets $\mathcal{T}$ of observations, there exists an observation-based almost-winning strategy for Player 1 in G for the objective $\mathsf{Buchi}(\mathcal{T})$ iff there exists an equivalence-preserving positional almost-winning strategy for Player 1 in H for the objective $\mathsf{Buchi}(B_{\mathcal{T}})$.*

Symbolic algorithm. We present a symbolic quadratic-time (in the size of H) algorithm to compute the set $Q^{\approx}_{\mathsf{AS}}$. For $Y \subseteq Q$ and $X \subseteq Y$, let $\mathsf{Apre}(Y, X) = \{ q \in Y \mid \exists \sigma \in \mathsf{Allow}([q]_{\approx}, Y) : \mathsf{Post}^H_\sigma(q) \subseteq X \}$ and $\mathsf{Spre}(Y) = \{ q \in Y \mid \mathsf{Allow}([q]_{\approx}, Y) \neq \emptyset \}$. Note that $\mathsf{Spre}(Y) = \mathsf{Apre}(Y, Y)$. Let $\phi = \nu Y.\mu X.\big(\mathsf{Apre}(Y, X) \vee (B_{\mathcal{T}} \wedge \mathsf{Spre}(Y)\big)$ and let $Z = [\![\phi]\!]$. It can be shown that $Z = Q^{\approx}_{\mathsf{AS}}$.

Theorem 7 (Complexity of almost winning for Büchi objectives under imperfect information). *Let G be a game structure of imperfect information, let $\mathcal{T}$ be a set of observations, and let ℓ be a state of G. Whether ℓ is an almost-winning state in G for the Büchi objective $\mathsf{Buchi}(\mathcal{T})$ can be decided in* EXPTIME.

The facts that $Z = Q^{\approx}_{\mathsf{AS}}$ and that H is exponential in the size of G yield Theorem 7. The arguments for the proofs of Theorem 6 and 7 do not directly extend to coBüchi or parity objectives. 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 may require infinite memory; for an example, see [5]. Such concurrent games are reducible to semiperfect-information games [4], and semiperfect-information games are reducible to the imperfect-information games we study. Hence a reduction to finite game structures of perfect information in order to obtain randomized positional strategies is not possible with respect to almost winning for general parity objectives. Theorem 6 and Theorem 7 may hold for coBüchi objectives, but there does not seem to be a simple extension of our arguments for Büchi objectives to the coBüchi case. The results that correspond to Theorems 6 and 7 for coBüchi objectives are open.

Direct symbolic algorithm. As in Section 3.2, the subset structure H does not have to be constructed explicitly. Instead, we can evaluate a fixed-point formula on a well-chosen lattice. The fixed-point formula to compute the set $Q^{\approx}_{\mathsf{AS}}$ 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 the exponential game structure 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 imperfect information is EXPTIME-hard already for reachability objectives. The result for sure winning follows from [19], but our new proof extends to almost winning as well.

Sure winning. To show the lower bound, we use a reduction from the membership problem for polynomial-space alternating Turing machines. 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; $g : Q \to \{\wedge, \vee\}$; $\Sigma_i = \{0, 1\}$ 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 the transition relation; and $F \subseteq Q$ is the set of accepting states. We say that M is a *polynomial-space* ATM if there exists a polynomial $p(\cdot)$ such that for every word w, the tape space used by M on input w is bounded by $p(|w|)$. Without loss of generality, we assume that the initial control state of the machine is a $\vee$-state, and that transitions connect $\vee$-states to $\wedge$-states, and vice versa. A word w is *accepted* by the ATM M if there exists a run tree of M on w all of whose leaf nodes are accepting configurations (i.e., configurations containing an accepting state); see [3] for details. The *membership*

problem is to decide if a given word w is accepted by a given polynomial-space ATM (M, p). This problem is EXPTIME-hard [3].

Sketch of the reduction. Given a polynomial-space ATM M and a word w, we construct a game structure of imperfect information, 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. Player 1 is responsible for maintaining the symbol under the tape head. His 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, then 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 the nondeterminism on ϵ, chooses a transition of the ATM that 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 thus Player 1 can update his knowledge about the configuration of the ATM. Player 1 wins whenever an accepting configuration of the ATM is reached.

The difficulty is to ensure that Player 1 loses when he announces a wrong symbol under the tape head. As the number of configurations of the polynomial-space ATM is exponential, we cannot directly encode the full configuration of the ATM in the states of the game. To overcome this difficulty, we use the power of imperfect information as follows. Initially, Player 2 chooses a position k, where $1 \leq k \leq p(|w|)$, on the tape. The chosen number k, as well as the symbol $\sigma \in \{0, 1, 2\}$ that lies in the tape cell with number k, are maintained all along the game in the nonobservable 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 (since k is private), to avoid losing, he must not lie about any of the tape cells, and thus he must faithfully simulate the machine. Then, he wins the game if and only if the ATM accepts the words w.

Almost winning. To establish the lower bound for almost winning, we can use the same reduction. Randomization cannot help Player 1 in this game. Indeed, at any point in the game, if Player 1 takes a chance in either not faithfully simulating the ATM or lying about the symbol under the tape head, then the sink state is reached. In these cases, the probability to reach the sink state is positive, and so the probability to win the game is strictly less than one.

Theorem 8 (Lower bounds). *Let G be a game structure of imperfect information, let $\mathcal{T}$ be a set of observations, and let ℓ be a state of G. Deciding whether ℓ*

is a sure-winning state in G for the reachability objective $\mathsf{Reach}(\mathcal{T})$ *is* EXPTIME*-hard. Deciding whether ℓ is an almost-winning state in G for* $\mathsf{Reach}(\mathcal{T})$ *is also* EXPTIME*-hard.*

References

1. M. Abadi, L. Lamport, and P. Wolper. Realizable and unrealizable specifications of reactive systems. In *ICALP*, LNCS 372, pages 1–17. Springer, 1989.
2. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49:672–713, 2002.
3. A.K. Chandra, D. Kozen, and L.J. Stockmeyer. Alternation. *J. ACM*, 28:114–133, 1981.
4. K. Chatterjee and T.A. Henzinger. Semiperfect-information games. In *FSTTCS*, LNCS 3821, pages 1–18. Springer, 2005.
5. L. de Alfaro and T.A. Henzinger. Concurrent ω-regular games. In *Proc. LICS*, pages 141–154. IEEE Computer Society, 2000.
6. L. de Alfaro and T.A. Henzinger. Interface automata. In *Proc. FSE*, pages 109–120. ACM, 2001.
7. L. de Alfaro, T.A. Henzinger, and O. Kupferman. Concurrent reachability games. In *Proc. FOCS*, pages 564–575. IEEE Computer Society, 1998.
8. L. de Alfaro, T.A. Henzinger, and R. Majumdar. From verification to control: Dynamic programs for ω-regular objectives. In *Proc. LICS*, pages 279–290. IEEE Computer Society, 2001.
9. M. De Wulf, L. Doyen, T.A. Henzinger and J.-F. Raskin. Antichains: A new algorithm for checking universality of finite automata. In *CAV*, LNCS (to appear). Springer, 2006.
10. M. De Wulf, L. Doyen, and J.-F. Raskin. A lattice theory for solving games of imperfect information. In *HSCC*, LNCS 3927, pages 153–168. Springer, 2006.
11. D. Dill. *Trace Theory for Automatic Hierarchical Verification of Speed-independent Circuits*. MIT Press, 1989.
12. E.A. Emerson and C.S. Jutla. Tree automata, μ-calculus, and determinacy. In *Proc. FOCS*, pages 368–377. IEEE Computer Society, 1991.
13. A. Kechris. *Classical Descriptive Set Theory*. Springer, 1995.
14. O. Kupferman and M.Y. Vardi. Synthesis with incomplete informatio. In *Advances in Temporal Logic* (H. Barringer et al., eds.), pages 109–127. Kluwer, 1999.
15. M.L. Littman. *Algorithms for Sequential Decision Making*. PhD Thesis, Brown Univ., 1996.
16. D. Martin. Borel determinacy. *Annals of Mathematics*, 102:363–371, 1975.
17. A. Pnueli and R. Rosner. On the synthesis of an asynchronous reactive module. In *ICALP*, LNCS 372, pages 652–671. Springer, 1989.
18. P.J. Ramadge and W.M. Wonham. Supervisory control of a class of discrete-event processes. *SIAM J. Control and Optimization*, 25:206–230, 1987.
19. J.H. Reif. The complexity of two-player games of incomplete information. *J. Computer and System Sciences*, 29:274–301, 1984.
20. W. Thomas. Languages, automata, and logic. In *Handbook of Formal Languages* (G. Rozenberg et al., eds.), volume 3, pages 389–455. Springer, 1997.
21. M.Y. Vardi. Automatic verification of probabilistic concurrent finite-state systems. In *Proc. FOCS*, pages 327–338. IEEE Computer Society, 1985.