

Pure Strategies in Imperfect Information Stochastic Games

Arnaud Carayol^{*1}, Christof Löding^{†2}, and Olivier Serre^{‡3}

¹LIGM (CNRS & Université Paris Est)

²Informatik 7, RWTH Aachen, Germany

³IRIF (CNRS & Université Paris Diderot – Paris 7)

Abstract

We consider imperfect information stochastic games where we require the players to use pure (*i.e.* non randomised) strategies. We consider reachability, safety, Büchi and co-Büchi objectives, and investigate the existence of almost-sure/positively winning strategies for the first player when the second player is perfectly informed or more informed than the first player. We obtain decidability results for positive reachability and almost-sure Büchi with optimal algorithms to decide existence of a pure winning strategy and to compute one if it exists. We complete the picture by showing that positive safety is undecidable when restricting to pure strategies even if the second player is perfectly informed.

1 Introduction

The study of two-player games has received a lot of attention in the last decade, mainly motivated by applications to the verification of reactive open systems. Those systems are composed of a program (represented by the first player, Eve) and some (possibly hostile) environment (represented by the second player, Adam). The verification problem consists in deciding whether the program can be restricted so that the system meets some given specification *whatever* the environment does. Here, restricting the program means synthesizing a controller [15], which, in terms of games, is equivalent to designing a strategy for Eve that is winning against *any* strategy of Adam.

Of course, the class of games to consider depends on the class of systems that one intends to model. This may lead to consider various features such as concurrency (the players *independently* and *simultaneously* choose their action, whose *parallel* execution determines the next state), stochastic transitions (the next state is chosen according to a probability distribution depending on the current state and on the actions chosen by the players) or imperfect information (the players do not observe the exact state). Note that imperfect information is necessary if one wants for instance to model a system where the program and the environment share some public variables while having also their own private variables [16].

Recently in [13, 2, 3] two (mainly equivalent) models of concurrent stochastic games with imperfect information have been introduced. They permit to capture several known models (as those from [9, 6, 7]) while preserving the main decidability results.

In this paper we consider the games as introduced in [13, 2, 3] (we use the formalism of [13]). These are finite state games in which, at each round, the two players choose concurrently an action and based on these actions the successor state is chosen according to some fixed probability distribution. The resulting infinite play is won by Eve if it satisfies a given *objective*. The

*Arnaud.Carayol@univ-mlv.fr

†loeding@informatik.rwth-aachen.de

‡Olivier.Serre@cnrs.fr

objectives we consider here are reachability (*i.e.* a final state is eventually visited), safety (*i.e.* no forbidden state is visited), Büchi (*i.e.* some final state is visited infinitely often) and co-Büchi (*i.e.* no forbidden states is visited infinitely often). Imperfect information is modelled as follows: both players have an equivalence relation over states and, instead of observing the exact state, they only observe its equivalence class. Intuitively, two equivalent states are indistinguishable by the corresponding player.

In [13, 2, 3] the authors were considering general strategies where a player is allowed to use *randomisation* when choosing her/his next action. It was then shown, for Büchi objectives, that one can decide whether Eve has such a strategy φ that is almost-surely winning against any strategy ψ of Adam (meaning that an infinite play played according to φ and ψ is won by Eve with probability 1). It was also established in [2, 3] that one can decide for co-Büchi objectives whether Eve has a positively winning strategy.

In the present work we restrict our attention to *pure* strategies, *i.e.* we forbid the players to randomise when choosing their actions. Our initial motivation for this work comes from automata theory. The emptiness problem for automata on infinite trees can be described as the problem of computing a winning strategy in a two-player game of infinite duration. The required game model depends on the class of automata that is considered. In particular, [11] proposes a reduction of the emptiness problem for *alternating* tree automata to the existence of a pure winning strategy for Eve in an imperfect information game. For capturing the automaton model with a qualitative acceptance condition as introduced in [5], one furthermore needs stochastic games (and up to now this is the only known method for checking emptiness of such automata). So one of our aims is to obtain a toolbox and to understand the limits of this method for checking emptiness of tree automata.

Another motivation for studying pure strategies comes from controller synthesis. Indeed, a classical way to consider an open system (*i.e.* a controllable program interacting with an uncontrollable environment) is as a two-player game, and in this setting synthesising a controller for the program boils down to compute a winning strategy in the game for the player standing for the program. In the setting of open systems, imperfect information naturally arises, for instance when the program and its environment use private variables. A desirable property of a controller (in addition to its optimality) is its implementability which could be limited by its size but also by the required features. In the imperfect information setting, the main needed feature might be randomisation which is well-known to be non-trivial to implement in a non-biased fashion: hence, existence of an optimal non-randomised controller (equivalently a pure strategy) is a natural question.

Our main results are the following. On the negative side, by a reduction of the value 1 problem for probabilistic word automata [12], we prove that even if Adam is fully informed and Eve is totally blind (*i.e.* all states are indistinguishable for her), it is undecidable whether Eve can positively win a safety game (Section 3). Under the same restrictions, positive winning in Büchi games and almost-sure winning in co-Büchi games are proved to be undecidable by reduction from the emptiness problem for probabilistic ω -word automata [1].

To obtain positive results, we have to impose restrictions on how Adam is informed. We consider the case where he has perfect information and the case where he is more informed than Eve¹. In both situations we show that it is decidable whether Eve has a positively winning pure strategy in a reachability game (Section 4). Using this result in a fixpoint computation, we prove that one can decide whether Eve has an almost-surely winning pure strategy in a Büchi game (Section 5). Moreover, if it exists, such a strategy can be constructed and requires *finite* memory. In both cases, we obtain matching upper and lower complexity bounds.

The decidability results for the special case where Adam is perfectly informed were also obtained in [8]. However, the technique we develop here is different and in particular uses the positive winning case as a toolbox, which later permits us to handle the more general case where Adam is more informed than Eve. And while [8] focuses on reachability conditions and studies the memory required for winning strategies depending on how the players are informed, we focus

¹We say that Adam is more informed than Eve when his equivalence relation on the states of the games refines that of Eve. In particular, this is the case when Adam is perfectly informed.

on the case in which Adam is better informed than Eve (or even perfectly informed), and study different winning conditions.

In our setting (restricting to pure strategies) the algorithmic complexity increases to 2-EXPTIME while it is only EXPTIME in the general setting (allowing randomised strategies) when Adam is more informed than Eve. Also note that in the latter setting one has decidability results in the case where no assumption is made on how Adam and Eve are informed while in our setting this question is left open.

The resulting complete picture is summarised in the table at the end of this paper.

2 Definitions

A **probability distribution** over a finite set X is a mapping $d : X \rightarrow [0, 1]$ such that $\sum_{x \in X} d(x) = 1$. In the sequel we denote by $\mathcal{D}(X)$ the set of probability distributions over X . Given some set X and some equivalence relation \sim over X , $[x]_{\sim}$ stands for the equivalence class of x for \sim and $X/\sim = \{[x]_{\sim} \mid x \in X\}$ denotes the set of equivalence classes of \sim . As usual we write A^* (resp. A^ω) for the set of finite (resp. infinite) words over some finite alphabet A . For $k \geq 0$ we denote by $A^{\geq k}$ (resp. $A^{\leq k}$) the set of words of length at least (resp. at most) k .

A **concurrent arena with imperfect information** (or simply an **arena**) is defined as a tuple $\mathcal{A} = (S, \Sigma_E, \Sigma_A, \delta, \sim_E, \sim_A)$ where S is a finite set of **states**; Σ_E (resp. Σ_A) is the (finite) set of **actions** for Eve (resp. Adam); $\delta : S \times \Sigma_E \times \Sigma_A \rightarrow \mathcal{D}(S)$ is the (total) transition function; and \sim_E and \sim_A are equivalence relations over states.

A play in such an arena proceeds as follows. First it starts in some initial state s . Then the first player, Eve, picks an action $\sigma_E \in \Sigma_E$ and, *simultaneously* and *independently*, the second player, Adam, chooses an action $\sigma_A \in \Sigma_A$. Then a successor state is chosen according to the probability distribution $\delta(s, \sigma_E, \sigma_A)$, and the process restarts: the players choose a new pair of actions that induces, together with the current state, a new state and so on forever. Hence, a **play** is an infinite sequence $s_0(\sigma_E^0, \sigma_A^0)s_1(\sigma_E^1, \sigma_A^1)s_2 \cdots$ in $(S \cdot (\Sigma_E \times \Sigma_A))^\omega$ such that for every $i \geq 0$, $\delta(s_i, \sigma_E^i, \sigma_A^i)(s_{i+1}) > 0$. In the sequel we refer to a prefix of a play ending by a state as a **partial play**.

The intuitive meaning of \sim_E (resp. \sim_A) is that two states s_1 and s_2 such that $s_1 \sim_E s_2$ (resp. $s_1 \sim_A s_2$) cannot be distinguished by Eve (resp. by Adam). We easily extend relation \sim_X , with $X \in \{E, A\}$, to partial plays as follows. First, for any partial play $\lambda = s_0(\sigma_E^0, \sigma_A^0)s_1(\sigma_E^1, \sigma_A^1) \cdots s_k$ denote $[\lambda]_{\sim_X} = [s_0]_{\sim_X}[s_1]_{\sim_X} \cdots [s_k]_{\sim_X}$; then define $\lambda \sim_X \lambda'$ if and only if $[\lambda]_{\sim_X} = [\lambda']_{\sim_X}$.

We say that Adam is **more informed** than Eve if $\sim_A \subseteq \sim_E$, and Adam is **perfectly informed** if \sim_A is the equality relation.

Example 1. Consider the concurrent game with imperfect information depicted in Figure 1. Let $\Sigma_E = \Sigma_A = \{a, b\}$. The initial state is s_0 and from s_0 if Adam plays the action a then any action played by Eve leads with probability $\frac{1}{2}$ either to s_1 or to s_2 . Similarly if Adam plays b then any action played by Eve leads with probability $\frac{1}{2}$ either to s_3 or to s_4 . In the states s_1, s_2, s_3 and s_4 , which are indistinguishable by Eve, the action of Adam has no impact. If Eve plays a from s_1 or s_4 or b from s_2 or s_3 the play goes to the final state f which is a sink state. Any other action by Eve from one of those states leave the current state unchanged.

In order to choose their moves the players respect strategies, and, for this, they may use all the information they have about what was played so far. However, if two partial plays are equivalent for \sim_E (resp. \sim_A), then Eve (resp. Adam) cannot distinguish between them, and should behave the same. This leads to the following notion.

An **observation-based pure strategy** (simply called a strategy in the following) for Eve is a function $\varphi : (S/\sim_E)^* \rightarrow \Sigma_E$, i.e., to choose her next action, Eve considers the sequence of observations she has seen so far. We overload φ by writing $\varphi(\lambda)$ instead of $\varphi([\lambda]_{\sim_E})$: in particular, a strategy φ for Eve is such that $\varphi(\lambda) = \varphi(\lambda')$ whenever $\lambda \sim_E \lambda'$ (and similarly for Adam).

A **finite-memory strategy** for Eve is a strategy that can be computed by a finite automaton with output that reads the observation sequence of the partial play and outputs the next action




Figure 1: A concurrent arena where Adam is perfectly informed while Eve cannot distinguish states s_1, s_2, s_3 and s_4 .

of Eve. We do not give a precise technical definition because it is not needed in this work. The size of such a strategy corresponds to the number of states of the automaton.

Strategies for Adam are defined in a similar way by replacing \sim_E by \sim_A .

Remark 1. *In our definition of a strategy we implicitly assume that the players only observe the sequence of states and not the corresponding sequence of actions. While the fact that a player does not observe what his adversary has played is reasonable (otherwise imperfect information on states would make less sense) one could object that the player should observe the actions she has played so far. However, as the players do not use randomisation in their strategies, they can always retrieve the actions they played so far.*

Moreover, in the special case where Adam is more informed than Eve (as later studied in Section 4.4), we can also note that, when playing against a fixed strategy of Eve, he can always retrieve the actions she played so far (as he knows the strategy of Eve and also $[\lambda]_{\sim_E}$ for any partial play λ).

Let $\mathcal{A} = (S, \Sigma_E, \Sigma_A, \delta, \sim_E, \sim_A)$ be an arena, let $s_0 \in S$ be an initial state, φ_E be a strategy for Eve and φ_A be a strategy for Adam. First we let $\text{Outcomes}(s_0, \varphi_E, \varphi_A)$ to be the set of all possible plays when the game starts in s_0 and when Eve and Adam respectively follows φ_E and φ_A : $\lambda = s_0(\sigma_E^0, \sigma_A^0)s_1(\sigma_E^1, \sigma_A^1) \cdots$ belongs to $\text{Outcomes}(s_0, \varphi_E, \varphi_A)$ iff

$$\delta(s_i, \varphi_E([s_0]_{\sim_E}[s_0]_{\sim_E} \cdots [s_i]_{\sim_E}), \varphi_A([s_0]_{\sim_A}[s_0]_{\sim_A} \cdots [s_i]_{\sim_A}))(s_{i+1}) > 0$$

for every $i \geq 0$. Then we are interested in defining the probability of a (measurable) set of plays, knowing that Eve (resp. Adam) uses φ_E (resp. φ_A). This is done in the usual way (see e.g. [6]): once a pair (φ_E, φ_A) of strategies for both players is fixed, one is left with a (possibly infinite) Markov chain that naturally induces a probability space over the Borel σ -field generated by the cones, where for any partial play λ starting in s_0 the cone for λ is the set $\text{cone}(\lambda) = \lambda \cdot ((\Sigma_E \times \Sigma_A) \cdot S)^\omega$ of all infinite plays with prefix λ . We let $\text{Pr}_{s_0}^{\varphi_E, \varphi_A}$ denote the corresponding probability measure over this space.

An **objective** for Eve is a (measurable) set \mathcal{O} of plays: a play is won by Eve if it belongs to \mathcal{O} ; otherwise it is won by Adam. A **concurrent game with imperfect information** (simply called a game in the following) is a triple $\mathbb{G} = (\mathcal{A}, s_0, \mathcal{O})$ where \mathcal{A} is an arena, s_0 is an initial state and \mathcal{O} is an objective. In the sequel we focus on the following special classes of ω -regular objectives (note that all of them are Borel sets hence, measurable) that we define using a subset $F \subseteq S$ of **final** states.

A **reachability objective** (resp. **safety**) is of the form $(S \cdot (\Sigma_E \times \Sigma_A))^* F ((\Sigma_E \times \Sigma_A) \cdot S)^\omega$ (resp. of the form $((S \setminus F) \cdot (\Sigma_E \times \Sigma_A))^\omega$): a play is winning if it contains (resp. does not contain) a final state.


A **Büchi objective** (resp. **co-Büchi objective**) is of the form $\bigcap_{k \geq 0} (S \cdot (\Sigma_E \times \Sigma_A))^{\geq k} F ((\Sigma_E \times \Sigma_A) \cdot S)^\omega$ (resp. of the form $(S \cdot (\Sigma_E \times \Sigma_A))^* ((S \setminus F) \cdot (\Sigma_E \times \Sigma_A))^\omega$) : a play is winning if it goes infinitely often (resp. finitely often) through final states.

A reachability (resp. safety, Büchi, co-Büchi) game is a game equipped with a reachability (resp. safety, Büchi, co-Büchi) objective. In the sequel we may replace \mathcal{O} by F when it is clear from the context which objective we consider.

Fix a game $\mathbb{G} = (\mathcal{A}, s_0, \mathcal{O})$. A strategy φ_E for Eve is **surely winning** if, for any counter-strategy φ_A for Adam, $\text{Outcomes}(s_0, \varphi_E, \varphi_A) \subseteq \mathcal{O}$. If such a strategy exists, we say that Eve **surely wins** \mathbb{G} . A strategy φ_E for Eve is **almost-surely winning** (resp. **positively winning**) if, for any counter-strategy φ_A for Adam, $\Pr_{s_0}^{\varphi_E, \varphi_A}(\mathcal{O}) = 1$ (resp. > 0). If such a strategy exists, we say that Eve **almost-surely wins** (resp. **positively wins**) \mathbb{G} .

In this paper, we are interested in deciding existence of almost-surely/positively winning strategies for Eve for safety/reachability/Büchi/co-Büchi games.

Example 2. Consider the (perfect information) concurrent reachability game depicted below with q_f as the unique final state. In state q_w , if both players choose the same action then they stay in state q_w and otherwise they move to state q_f . In state q_f , all choices of actions stay in state q_f . Eve does not have any almost-surely winning strategy.



Indeed, given any strategy φ_E for Eve, the counter-strategy φ_A for Adam mirroring the strategy of Eve (i.e. $\varphi_A = \varphi_E$) only allows for the play q_w^ω and hence, $\Pr_{s_0}^{\varphi_E, \varphi_A}(\mathcal{O}) = 0$. Similarly Adam does not have an almost-surely winning strategy. For any fixed strategy φ_A of Adam, any counter-strategy φ_E for Eve that satisfies $\varphi_E(q_w) \neq \varphi_A(q_w)$ is such that $\Pr_{s_0}^{\varphi_E, \varphi_A}(\mathcal{O}) = 1$.

3 Undecidability Results

In this section we provide undecidability results for certain combinations of types of winning strategies and objectives. An easy consequence of undecidability results for probabilistic ω -automata from [1] is stated in the following theorem. In these reductions, Eve plays alone and cannot distinguish any states of the game. The states and transitions of the game are those of the ω -automaton and the strategy of Eve corresponds to the input word.

Theorem 1. *The decision problems whether Eve almost-surely wins a given co-Büchi game or positively wins a given Büchi game are undecidable (even if the set of actions of Adam is a singleton).*

Proof. Consider a probabilistic automaton \mathcal{A} on ω -words as in [1]. Now consider a concurrent game with imperfect information where Adam plays no role and where Eve's actions are the letters from the input alphabet A of \mathcal{A} and whose states are the ones of the automaton. Moreover all states are \sim_E -equivalent. Now the transition function of the game mimics the one of the automaton. As Eve does not observe anything, a (pure) strategy φ of Eve can be described as an infinite word u_φ in A (the i -th letter being the i -th action played by Eve), and φ is almost-surely (resp. positively) winning iff the probability of a run of \mathcal{A} over u_φ to be accepting is 1 (resp. strictly positive). The undecidability results follow from the undecidability of the emptiness problem for co-Büchi (resp. Büchi) probabilistic automaton with the almost-sure (resp. positive) semantics [1]. \square

In the following we prove that the existence of positively winning strategies for safety objectives is undecidable. Our result is based on the undecidability of the value 1 problem for probabilistic automata on finite words [12]. For simpler use in our reduction we reformulate this problem in terms of games.




Figure 2: The Hide-or-Run game

Consider the class of concurrent reachability games \mathbb{G} with imperfect information with the following properties. Eve is blind (i.e. \sim_E consists of a unique equivalence class), and Adam has no impact on the game (i.e. his set of actions is a singleton). Furthermore, there is a special action \sharp that Eve can play at any time, and that leads (depending on the current state) either to a final sink state or to a non-final sink state. The final sink state is the only final state. Intuitively, one can think of such a game as one where Eve plays a sequence of actions and then declares by \sharp that she stops (and she wins if she stopped in a state that leads to the winning sink).

We refer to this type of game as *probabilistic automaton game* (PA game) because it corresponds to probabilistic automaton on finite words (see [14] for an introduction to probabilistic automata): a strategy of Eve corresponds to a finite word followed by \sharp (without playing \sharp Eve surely loses), and the probability that it is winning is the probability of the word to be accepted in the automaton. Then we have the following result, which directly follows from the undecidability of the value 1 problem for probabilistic automata [12].

Lemma 1. *For a given a PA game, it is undecidable whether Eve has for each $0 < \varepsilon < 1$ a strategy that is winning with probability $(1 - \varepsilon) < p < 1$.*

Our reduction that uses Lemma 1 starts from an example of a concurrent safety game \mathbb{G}_{HR} known as *Hide-or-Run* [10] (see Figure 2). In this game, Adam can choose between hiding (h) and running (r), and Eve can choose between waiting (w) and throwing (t) her only snowball. If Adam hides and Eve waits, the game stays in state s_{hide} . If Adam runs and Eve throws the snowball, then Adam is hit, and the game proceeds to sink state s_{wet} . In all other cases, Adam gets home (either he runs without being hit or he can safely run after Eve has thrown her snowball) and the game proceeds to sink state s_{home} . This is a safety game where Eve wants to avoid visiting s_{home} .

In [10] it is shown that Eve can only win by using a randomised strategy that plays action w in round i with probability p_i such that $0 < p_i < 1$ for every i and $\prod_i p_i > 0$ (for this, Eve does not have to distinguish the states).

Now the idea is to incorporate a gadget in \mathbb{G}_{HR} that permits Eve to simulate random choices while playing deterministically.

Theorem 2. *It is undecidable whether Eve positively wins in a safety game (resp. co-Büchi game), even if \sim_E consists of a single equivalence class.*

Proof. Consider a probabilistic automaton game \mathbb{G} with a set of actions disjoint from the one in the game \mathbb{G}_{HR} . Let \mathbb{G}_r and \mathbb{G}_h be two disjoint copies of \mathbb{G} where we removed the two states reachable by Eve playing \sharp (the \sharp -edges are redirected as described below).

In the game \mathbb{G}'_{HR} (see Figure 3), the concurrent choices of the actions in \mathbb{G}_{HR} are simulated by the imperfect information. All states are indistinguishable by Eve. First Adam makes his choice r or h from s_{hide} (Eve's action has no impact). The game then moves to the initial state of \mathbb{G}_r or \mathbb{G}_h , depending on the choice of Adam (ignore the action *cheat* for the moment, which is explained later). Because of the imperfect information Eve does not observe Adam's choice.

In \mathbb{G}_r and \mathbb{G}_h we removed the target states of \sharp but Eve can still play \sharp : if in \mathbb{G} it was leading to the final state it now behaves as Eve playing w from s_{hide} , and otherwise it behaves as Eve playing t from s_{hide} (see Figure 3).




Figure 3: The modified version of *Hide-or-Run*: \mathbb{G}'_{HR} . Black states in $\mathbb{G}_r/\mathbb{G}_h$ correspond to states from which $\#$ led to the final state in \mathbb{G} , and x denotes any letter different from $\#$.

Finally, in order to prevent Eve from playing an infinite sequence of actions without $\#$, we add an extra small gadget where Adam is allowed to declare that Eve will cheat. If he plays *cheat* from s_{hide} this leads to a new state s_c where the following may happen depending on the next move of Eve (the action of Adam has no impact): if she plays $\#$ from s_c then the play goes to a sink state s_w (that is not final); if she does not play $\#$ from s_c then with probability $1/2$ the play stays in s_c and with probability $1/2$ the play goes to a sink final state s_l . Hence, from s_c if she never plays $\#$, then the play almost-surely ends in s_l .

Let \mathbb{G}'_{HR} be this new game, where we recall that all states are indistinguishable for Eve, s_{hide} is the initial state and $\{s_{home}, s_l\}$ are the final states. We claim that Eve positively wins game \mathbb{G}'_{HR} iff Eve in \mathbb{G} has strategies winning with probability arbitrarily close to 1. Indeed, consider a strategy φ for Eve in \mathbb{G}'_{HR} . As Eve cannot distinguish any state in \mathbb{G}'_{HR} , and does not observe the actions played by Adam, φ is independent of Adam's choices.

If the strategy of Eve consists in playing $\#$ only finitely often, it cannot be positively winning as it suffices for Adam to wait for the last $\#$ and then play *cheat*. More precisely, the strategy of Adam consists in playing (in state s_{hide}) the action h whenever Eve's strategy will still play $\#$ in the future, and *cheat* if Eve will never play $\#$ in the future. It can be shown that following this strategy Adam wins against the strategy of Eve with probability 1.

Thus, in the following we only consider strategies φ of Eve that play $\#$ infinitely often. An equivalent description of such strategies φ is by a sequence $(\varphi_i)_{i \geq 1}$ of strategies for Eve in \mathbb{G} : φ consists in playing an arbitrary letter then playing as φ_1 until playing $\#$, then playing an arbitrary letter, then playing as φ_2 until playing $\#$ and so on (the arbitrary letter is used here when Adam chooses to move to \mathbb{G}_r , \mathbb{G}_h or s_c).

For one direction, assume that φ is positively winning in \mathbb{G}'_{HR} . Let p_i be the probability that Eve wins in \mathbb{G} when playing according to φ_i . Then, from the properties of \mathbb{G}_{HR} , it follows that φ is winning iff $0 < p_i < 1$ for every $i \geq 1$ and $\prod_i p_i > 0$. This implies that the sequence $(p_i)_{i \geq 1}$ converges to 1 and hence the φ_i are strategies as in Lemma 1.

Conversely, if Eve has strategies winning with probabilities arbitrarily close to 1 as in Lemma 1, then one can choose the φ_i such that $1 > p_i \geq 1 - \frac{1}{(i+1)^2}$ which ensures $0 < p_i < 1$ for every $i \geq 1$ and $\prod_i p_i > 0$. Indeed,

$$\prod_{i \geq 1} 1 - \frac{1}{(i+1)^2} = \lim_{m \rightarrow \infty} \prod_{i=1}^m 1 - \frac{1}{(i+1)^2} = \lim_{m \rightarrow \infty} \frac{m+2}{2m+2} = \frac{1}{2}$$

This family φ_i defines a strategy for Eve in \mathbb{G}'_{HR} . Again using the properties of \mathbb{G}_{HR} , this implies that Eve positively wins against all strategies of Adam: either no outcome ever reaches s_c , in

which case \mathbb{G}_{HR} is simulated, or if an outcome reaches s_c , then it does with positive probability, and then it also reaches s_w with positive probability. \square

4 Positive Winning in Reachability Games

We now address the decidability of whether Eve positively wins in a reachability game, and we show decidability (and matching lower bounds) for the case where (i) Adam is perfectly informed and (ii) Adam is more informed than Eve.

For the rest of this section fix an arena $\mathcal{A} = (S, \Sigma_E, \Sigma_A, \delta, \sim_E, \sim_A)$ and a set of final states $F \subseteq S$.

To later address almost-sure winning (Section 5) we need to consider games that may start in different states, and we are interested in strategies that are winning from all of these states. For this reason, we define for any subset B of states a game $(\mathcal{A}, B, \mathcal{O})$ that is played as follows: there is a new initial step where Adam picks a state s_0 in B and then the play proceeds as in $(\mathcal{A}, s_0, \mathcal{O})$. Hence, a strategy φ for Eve in such a game is almost-surely (*resp.* positively) winning iff φ is almost-surely (*resp.* positively) winning in $(\mathcal{A}, s_0, \mathcal{O})$ for *every* state $s_0 \in B$.

4.1 Winning in a Finite Number of Moves.

We start with a general result that does not depend on how the players are informed. It states that if Eve can positively win in a reachability game then she can do so within a bounded number of moves.

Proposition 1. *Let $B \subseteq S$ be a subset of states and assume that Eve has a positively winning strategy φ in the reachability game (\mathcal{A}, B, F) . Then, there is a bound N and some $0 < \varepsilon_B \leq 1$ such that whenever Eve respects φ in the game (\mathcal{A}, B, F) , the probability that the resulting play visits a final state within the N first moves is at least ε_B .*

Proof. For any $N > 0$, any $s \in B$ and any strategy ψ_N for Adam, call $p_N^{\psi_N, s}$ the probability of the event "a play in (\mathcal{A}, s, F) , where Eve respects φ and Adam respects ψ_N visits a final state within the N first moves".

Let $x_N^{\psi_N} = \min\{p_N^{\psi_N, s} \mid s \in B\}$. We aim to show that there exists some $N > 0$ such that for each strategy ψ_N for Adam, $x_N^{\psi_N} > 0$.

For this, we reason by contradiction, assuming that for any bound $N > 0$, Adam has a counter strategy ψ_N such that $x_N^{\psi_N} = 0$. In particular, there is a state $s \in B$ such that $p_N^{\psi_N, s} = 0$ for infinitely many N . Hence, we can assume that the ψ_N are such that $p_N^{\psi_N, s} = 0$ for every $N \geq 0$ (as to get the property for some N Adam can always use the strategy for some $N' > N$).

Using $(\psi_N)_{N \geq 0}$ we define a strategy ψ for Adam as follows. We first let $I_0 = \mathbb{N}$ be the set of naturals. Next we define ψ and $(I_k)_{k \geq 0}$, a decreasing sequence (for inclusion) of *infinite* subsets of the naturals. First we sort partial plays by increasing length. We assume that ψ is defined on all partial plays of length smaller than k (hence initialization for $k = 0$ comes for free) and for plays of length $k + 1$ we do the following. As there are finitely many plays of length $k + 1$ while I_k is infinite there exists an infinite subset $I_{k+1} \subseteq I_k$ such that, for all $j_1, j_2 \in I_{k+1}$, both strategies ψ_{j_1} and ψ_{j_2} agree on plays of length $k + 1$; we define ψ to behave accordingly on plays of length $k + 1$.

Then, the following is a direct consequence of the definitions of ψ and $(I_k)_{k \geq 0}$: for every $k \geq 0$, the set I_k is infinite; and for every $j \in I_k$ and every partial play λ of length smaller than k , both ψ and ψ_j agree on λ .

In particular it implies that $x_N^{\psi} = 0$ for every $N \geq 0$: indeed, $x_N^{\psi_M} = 0$ for any $M \geq N$ and ψ agrees with all ψ_M with $M \in I_N$ (and as I_N is infinite such an M exists). Finally, as $0 \leq \Pr_s^{\varphi, \psi}(\mathcal{O}) \leq \sum_{N \geq 0} x_N^{\psi} = 0$ (here \mathcal{O} denotes the reachability objective defined by F), we conclude that $\Pr_s^{\varphi, \psi}(\mathcal{O}) = 0$ which contradicts our initial assumption of φ being positively winning in (\mathcal{A}, s, F) .

The fact that there is some $\varepsilon_B > 0$ such that φ ensures to reach a final state in less than N moves with a probability greater than ε_B is a direct consequence of the fact that one bounds the number of moves by N . \square

Remark 2. Proposition 1 implies that finite memory suffices for Eve to positively win in a reachability game. Indeed, it suffices to follow φ for the N first moves and then play the same action forever.

Another important consequence is that the values of the probabilities do not have any influence on whether Eve positively wins in a reachability game. More precisely consider another arena \mathcal{A}' that is exactly as \mathcal{A} except that its transition function δ' is such that for every state s and every pair of actions (σ_E, σ_A) one has $\delta(s, \sigma_E, \sigma_A) = 0$ iff $\delta'(s, \sigma_E, \sigma_A) = 0$. Then Eve positively wins in the reachability game (\mathcal{A}, B, F) iff she positively wins in the reachability game (\mathcal{A}', B, F) .

4.2 Positively Winning When Adam Is Perfectly Informed

We now assume that Adam is perfectly informed.

Consider for every $n \geq 0$, the objective $\text{Reach}^{\leq n}(F) = (S \cdot (\Sigma_E \times \Sigma_A))^{\leq n} F ((\Sigma_E \times \Sigma_A) \cdot S)^\omega$ where a final state has to be visited within the first n steps. The following inductively characterises the sets B for which Eve can win $(\mathcal{A}, B, \text{Reach}^{\leq n}(F))$.

Proposition 2. Let $B \subseteq S$ be a set of pairwise \sim_E -equivalent states and let $n > 0$. Eve positively wins $(\mathcal{A}, B, \text{Reach}^{\leq n}(F))$ if and only if there exists an action $\sigma_E \in \Sigma_E$ and a set $B' \subseteq S$ such that

- Eve positively wins $(\mathcal{A}, B', \text{Reach}^{\leq n-1}(F))$,
- for every $s \in B \setminus F$ and for every $\sigma_A \in \Sigma_A$, there exists $s' \in B'$ such that $\delta(s, \sigma_E, \sigma_A)(s') > 0$.

Proof. Fix a set $B \subseteq S$ of pairwise equivalent states and an integer $n > 0$. For the direct implication assume that Eve has a positively winning strategy φ in $(\mathcal{A}, B, \text{Reach}^{\leq n}(F))$. Let $\sigma_E = \varphi([B]_{\sim_E})$ be the first action played by Eve and let φ' be the strategy followed by Eve after this first step (i.e. $\varphi'(\lambda) = \varphi([B]_{\sim_E} \cdot \lambda)$ for every partial play λ). Let B' be the set of states s' such that φ' is positively winning in $(\mathcal{A}, s', \text{Reach}^{\leq n-1}(F))$.

We claim σ_E and B' satisfy the property of the statement. First and by definition φ' is positively winning in $(\mathcal{A}, B', \text{Reach}^{\leq n-1}(F))$. For the second property assume toward a contradiction that there exist some $\sigma_A \in \Sigma_A$ and some state $s_0 \in B$ such that for every $s \in B'$, $\delta(s_0, \sigma_E, \sigma_A, s) = 0$ then φ' is not positively winning in $(\mathcal{A}, s, \text{Reach}^{\leq n-1}(F))$ (i.e. there exists a strategy ψ_s of Adam such that $\text{Pr}_s^{\varphi', \psi_s}(\text{Reach}^{\leq n-1}(F)) = 0$). Consider the strategy ψ for Adam consisting in playing first σ_A and then the ψ_s corresponding to the observed state s (i.e. $\psi(s_0) = \sigma_A$ and $\psi(s_0(\sigma_E, \sigma_A)\lambda) = \psi_s(\lambda)$ for any partial play λ starting with $s \in B'$). We have the following contradiction:

$$\text{Pr}_{s_0}^{\varphi, \psi}(\text{Reach}^{\leq n}(F)) = \sum_{s \in B'} \delta(s_0, \sigma_E, \sigma_A)(s) \cdot \underbrace{\text{Pr}_s^{\varphi', \psi_s}(\text{Reach}^{\leq n-1}(F))}_{=0} = 0.$$

For the converse implication assume that there exists an action $\sigma_E \in \Sigma_E$ and a set $B' \subseteq S$ satisfying the properties of the statement. Fix a positively winning strategy φ' for Eve in $(\mathcal{A}, B', \text{Reach}^{\leq n-1}(F))$.

Consider the strategy φ for Eve in $(\mathcal{A}, B, \text{Reach}^{\leq n}(F))$ consisting of first playing σ_E and then following φ' (i.e. $\varphi([B]_{\sim_E}) = \sigma_E$ and $\varphi([B]_{\sim_E} \cdot [\lambda]_{\sim_E}) = \varphi'([\lambda]_{\sim_E})$ for every partial play λ). We claim that this strategy is positively winning in $(\mathcal{A}, B, \text{Reach}^{\leq n}(F))$.

Indeed, let ψ be a strategy for Adam in $(\mathcal{A}, B, \text{Reach}^{\leq n}(F))$ and let $s_0 \in B$. Let σ_A be the first action played by Adam when using ψ and let ψ' be the strategy followed by Adam after this first step (i.e. $\psi'(\lambda) = \psi(s_0(\sigma_E, \sigma_A) \cdot \lambda)$ for every partial play λ). By definition of B' , there exists $s' \in B'$ such that $\delta(s_0, \sigma_E, \sigma_A)(s') > 0$. We have:

$$\text{Pr}_{s_0}^{\varphi, \psi}(\text{Reach}^{\leq n}(F)) \geq \delta(s_0, \sigma_E, \sigma_A)(s') \cdot \text{Pr}_{s'}^{\varphi', \psi'}(\text{Reach}^{\leq n-1}(F)) > 0.$$

□

Now, consider the increasing family of sets $(\mathcal{W}_i)_{i \geq 0}$ defined by:

- $\mathcal{W}_0 = \{B \mid B \subseteq F\}$

- $\mathcal{W}_{i+1} = \{B \subseteq S \mid \forall r \in B, \exists \sigma_E \exists B' \in \mathcal{W}_i \text{ s.t. } \forall s \sim_E r \text{ with } s \in B \setminus F, \forall \sigma_A \exists s' \in B' \text{ s.t. } \delta(s, \sigma_E, \sigma_A)(s') > 0\}$

and call \mathcal{W} its limit. Then the following is a simple consequence of Proposition 2.

Theorem 3. *Let $B \subseteq S$ be a non-empty set. Eve has a positively winning strategy φ in the game (A, B, F) if and only if $B \in \mathcal{W}$. In particular it can be decided in time exponential in $|S|$ whether Eve has a positively winning strategy. If such a strategy exists, one can construct one that uses the set 2^S as memory, and this strategy guarantees to positively reach a final state within the $2^{|S|}$ first moves.*

Proof. Using Proposition 2, by a direct induction on n one gets that $B \neq \emptyset$ belongs to \mathcal{W}_n if and only if Eve positively wins $(A, B, \text{Reach}^{\leq n}(F))$.

For any $B \in \mathcal{W}$, we denote by $\text{rk}(B)$ the smallest n such that $B \in \mathcal{W}_n$. Now, for any $B \in \mathcal{W}$, we define a strategy for Eve denoted φ_B that uses \mathcal{W} as a finite memory. Initially the memory is B . For a partial play λ ending in a state in some equivalence class $[s]_{\sim_E}$ and assuming that the memory is B' , we define the strategy as follows:

- If $\text{rk}(B') > 0$ and if there exists $r \in [s]_{\sim_E} \cap B'$, then by definition of $(\mathcal{W}_i)_{i \geq 0}$ there exists some action σ_E and some set B'' such that the following holds:
 - $\text{rk}(B'') = \text{rk}(B') - 1$,
 - $\forall r' \sim_E r$ with $r' \in B \setminus F, \forall \sigma_A, \exists s' \in B'', \text{ s.t. } \delta(s, \sigma_E, \sigma_A)(s') > 0$.

Then we let $\varphi_B(\lambda) = \sigma_E$ and update the memory to B'' .

- In all other cases, we take $\varphi_B(\lambda)$ to be an arbitrary action and update the memory to \emptyset .

By induction on n , we show that for every non-empty $B \in \mathcal{W}_n$, the strategy φ_B is positively winning in $(A, B, \text{Reach}^{\leq n}(F))$. The base case is immediate. Assume that the property is established for $n - 1 \geq 0$. Let B be a non-empty element of \mathcal{W}_n . Let $s_0 \in B$, $\sigma_E = \varphi_B([s_0]_{\sim_E})$ and $B' \in \mathcal{W}_{n-1}$ be the memory of φ_B after the first move. Let ψ be a strategy for Adam. Let σ_A be the first action played by Adam when using ψ and let ψ' be the strategy followed by Adam after this first step, *i.e.* $\psi'(\lambda) = \psi(s_0 \cdot (\sigma_E, \sigma_A) \cdot \lambda)$ for every partial play λ . By definition of B' , there exists $s' \in B'$ such that $\delta(s_0, \sigma_E, \sigma_A)(s') > 0$. Hence, we have:

$$\Pr_{s_0}^{\varphi_B, \psi}(\text{Reach}^{\leq n}(F)) \geq \delta(s_0, \sigma_E, \sigma_A)(s') \cdot \Pr_{s'}^{\varphi_{B'}, \psi'}(\text{Reach}^{\leq n-1}(F)) > 0.$$

which concludes the proof. \square

The following is a restatement of the end of Theorem 3.

Corollary 1. *In Proposition 1, when Adam is perfectly informed, one can always choose φ such that $N \leq 2^{|S|}$.*

4.3 Automaton-Compatible Strategies

The aim of this section is to refine Theorem 3 to positively winning strategies that satisfy further constraints. The motivation is that in Section 5 we compute almost-sure winning strategies for Büchi conditions using a fixpoint computation. In one iteration of this computation, we compute positively winning strategies for reachability that satisfy an extra constraint (roughly, that Eve can positively win the reachability game while ensuring that she can win another round of the reachability game once the target set is reached). This further constraint is expressible by finite automata that read partial plays and restrict the set of admissible next actions for Eve. Thus, below we develop the notion of a strategy that is compatible with such an automaton and then later apply it to the specific setting that we need.

Let $\mathcal{T} = (Q, \Sigma_E \times S_{/\sim_E}, q_0, q_s, \delta_{\mathcal{T}}, Act)$ be a deterministic finite automaton with input alphabet $\Sigma_E \times S_{/\sim_E}$, a finite set of states Q , an initial state q_0 , a sink state q_s , a transition function $\delta_{\mathcal{T}} : Q \times (\Sigma_E \times S_{/\sim_E}) \rightarrow Q$ and a function $Act : Q \rightarrow 2^{\Sigma_E}$ associating with any state of \mathcal{T} a subset of actions for Eve. Moreover, we require that the following holds

- $Act(q) = \emptyset$ if and only if $q = q_s$.
- For every state q and for every $(\sigma, x) \in \Sigma_E \times S_{/\sim_E}$ one has $\delta_{\mathcal{T}}(q, (\sigma, x)) = q_s$ if and only if $\sigma \notin Act(q)$.

Such a machine associates with any partial play λ a unique state q_λ defined by $q_{s_0} = q_0$ and $q_{\lambda \cdot (\sigma_E, \sigma_A) \cdot s} = \delta_{\mathcal{T}}(q_\lambda, (\sigma_E, [s]_{\sim_E}))$; it also permits to associate with any partial play a subset of actions by letting $Act_{\mathcal{T}}(\lambda) = Act(q_\lambda)$.

A strategy φ of Eve is \mathcal{T} -**compatible** if for every partial play λ where Eve respects φ one has $\varphi(\lambda) \in Act_{\mathcal{T}}(\lambda)$. Note that it implies that $q_\lambda \neq q_s$.

Remark 3. Consider the special case of the automaton \mathcal{T}_0 defined as follows: Q consists only of two states, the initial state and the sink state; $\delta_{\mathcal{T}}(q_0, (\sigma, x)) = q_0$ and $\delta_{\mathcal{T}}(q_s, (\sigma, x)) = q_s$ for any $(\sigma, x) \in \Sigma_E \times S_{/\sim_E}$ (i.e. all transitions are looping); and Act equals all actions Σ_E in the initial state. Then it follows that any strategy is \mathcal{T}_0 -compatible.

Hence, by considering the special case of \mathcal{T}_0 , any result we obtain later will also hold if we drop the \mathcal{T} -compatibility constraint.

In Section 5 and for the proof of Theorem 4, we work with automata that compute the belief of Eve along a play, as explained below. For an initial belief set $B_0 \subseteq S$ of pairwise \sim_E -equivalent states, the belief (also known as knowledge) $Belief_E^{B_0}(\lambda)$ of Eve after a partial play λ starting in a state of B_0 , intuitively corresponds to the set of possible states that can have been reached in a play \sim_E -equivalent to λ .

Formally, the value of $Belief_E^{B_0}(\lambda)$ can be inductively defined as follows: $Belief_E^{B_0}(s_0) = B_0$ and $Belief_E^{B_0}(\lambda \cdot (\sigma_E, \sigma_A) \cdot s) = \text{UpBelief}_E(Belief_E^{B_0}(\lambda), \sigma_E, [s]_{\sim_E})$ where the function $\text{UpBelief}_E : 2^S \times \Sigma_E \times [S]_{/\sim_E} \rightarrow 2^S$ is defined by:

$$\text{UpBelief}_E(B, \sigma_E, [s]_{\sim_E}) = \{t \in [s]_{\sim_E} \mid \exists r \in B, \exists \sigma_A \in \Sigma_A \text{ s.t. } \delta(r, \sigma_E, \sigma_A)(t) > 0\}.$$




Figure 4: Arenas and beliefs

Remark 4. The belief is in general smaller than the currently observed equivalence class. For instance, consider the reachability game depicted in Figure 4a in which all states are equivalent. If the strategy of Eve is to play $(abb)^\omega$, then her observation is always the same (as all states are equivalent). Her initial belief is $\{q_0, q_1\}$ but after playing a it becomes $\{q_1\}$ and after a it becomes $\{q_0\}$ and after another b it becomes $\{q_0, q_1\}$.

Remark 5. Given a family $\mathcal{B} \subseteq 2^S$ of beliefs for Eve (in the sense that each $B \in \mathcal{B}$ is a subset of a \sim_E -class), one can construct an automaton $\mathcal{T}_{\mathcal{B}}$ such that the $\mathcal{T}_{\mathcal{B}}$ -compatible strategies are precisely those such that Eve's belief always remains inside \mathcal{B} . The states of $\mathcal{T}_{\mathcal{B}}$ are the elements of \mathcal{B} , the transition function is defined by UpBelief_E , and the actions $Act(B)$ enabled at a state B are those that ensure that the belief remains inside \mathcal{B} .

Remark 6. In [13, 2, 3], it is shown that if Eve can almost-surely win (using randomised strategies) a Büchi game, she can do so using a strategy φ that only depends on the belief, i.e. $\varphi(\lambda) = \varphi(\lambda')$ whenever $\text{Belief}^\varphi(\lambda) = \text{Belief}^\varphi(\lambda')$. However, even if Eve is playing alone, this is no longer true² (even for reachability games) in our setting where we restrict to pure (i.e. non-randomised) strategies. Consider the reachability game in Figure 4 where Eve is playing alone. The equivalence relation is given by $s_1 \sim_E s_2$, $t_1 \sim_E f_2$ and $t_2 \sim_E f_1$.

If the game starts in s_0 then whatever strategy Eve uses, her belief always coincides with her observation. Eve can surely win (she can simply play the sequence $aaab$). But if her strategy only depends on her belief then she necessarily plays a sequence of actions of the form xu^ω where $x \in \{a, b\}$ and u is a two-letter word, and thus she has a probability $\frac{1}{2}$ to win using such a strategy.

We now return to the strengthening of Theorem 3. We assume that Adam is perfectly informed and we fix an automaton $\mathcal{T} = (Q, \Sigma_E \times S_{/\sim_E}, q_0, q_s, \delta_{\mathcal{T}}, \text{Act})$ as in Section 4.3. We are interested in checking whether Eve has a \mathcal{T} -compatible strategy that is positively winning in the reachability game (\mathcal{A}, B, F) .

Our main result is the following and its proof is by two successive reductions and an application of Theorem 3.

Theorem 4. *When Adam is perfectly informed, one can decide in exponential time in $|S|$ and polynomial in $|Q|$ whether Eve has a \mathcal{T} -compatible strategy that is positively winning in the reachability game (\mathcal{A}, B, F) . If such a strategy exists, one can construct one that uses memory of size polynomial in $|Q|$ and exponential in $|S|$.*

Proof. Note that adding the condition on the strategy being \mathcal{T} -compatible somehow means that once a final state is reached the play is not yet won by Eve because she needs to keep playing in accordance with \mathcal{T} (i.e. she must avoid to produce a partial play λ with $q_\lambda = q_s$). Hence, it is natural to consider an enriched arena $\mathcal{A}_{\mathcal{T}}$ that embeds \mathcal{T} . For this let $\mathcal{A}_{\mathcal{T}} = (S \times Q, \Sigma_E, \Sigma_A, \delta', \approx_E, \approx_A)$ where

- $\delta'((s, q), \sigma_E, \sigma_A)(s', q')$ equals $\delta(s, \sigma_E, \sigma_A)(s')$ if $q' = \delta_{\mathcal{T}}(q, (\sigma_E, [s']_{\sim_E}))$ and otherwise it equals 0;
- $(s_1, q_1) \approx_E (s_2, q_2)$ if and only if $s_1 \sim_E s_2$ and $q_1 = q_2$; and
- \approx_A is the equality relation, i.e. Adam is perfectly informed.

Of special interest is the safety game $(\mathcal{A}_{\mathcal{T}}, B \times \{q_0\}, S \times \{q_s\})$ and we are interested in sure winning for Eve because of the following straightforward lemma

Lemma 2. *Eve has a (possibly losing) \mathcal{T} -compatible strategy in the reachability game (\mathcal{A}, B, F) if and only if she has a surely winning strategy in the safety game $(\mathcal{A}_{\mathcal{T}}, B \times \{q_0\}, S \times \{q_s\})$.*

It is a known result [4] that when one considers sure winning for Eve in a safety game, winning strategies only depend on the belief of Eve (in the sense of Section 4.3). More precisely consider the (unique) largest subset \mathcal{B} of beliefs and the (unique) mapping $\text{Aut} : \mathcal{B} \rightarrow 2^{\Sigma_E}$ such that the following holds.

- No belief $B \in \mathcal{B}$ contains a forbidden state.
- For every $B \in \mathcal{B}$, the set $\text{Aut}(B)$ which consists of all those actions $\sigma_E \in \text{Aut}(B)$ such that for every action $\sigma_A \in \Sigma_A$ one has $\text{UpBelief}_E(B, \sigma_E, [s]_{\sim_E}) \in \mathcal{B}$, is not empty; i.e. actions in $\text{Aut}(B)$ are those that ensure that the updated belief will still be in \mathcal{B} regardless of the action of Adam.

Then Eve surely wins the safety game from configurations where her belief B is in \mathcal{B} and a strategy consists in choosing any action in $\text{Aut}(B)$.

Note that in the safety game $(\mathcal{A}_{\mathcal{T}}, B \times \{q_0\}, S \times \{q_s\})$, Eve's beliefs are elements in $2^S \times Q$ (as we have that $(s_1, q_1) \approx_E (s_2, q_2)$ implies $q_1 = q_2$).

²This fact is also observed in [8].

Now consider an automaton $\mathcal{T}' = (Q', \Sigma_E \times S_{/\sim_E}, q'_0, q'_s, \delta_{\mathcal{T}'}, Act')$ that *computes* Eve's belief (as explained in Remark 5). Hence, \mathcal{T}' is the same as \mathcal{T}_B in the previous safety game and uses function $Aut = Act'$ to define those authorised actions. To fit the definition, merge all beliefs not in B in a sink state and define Aut to be equal to \emptyset on it. The states Q' of \mathcal{T}' are elements of B (plus the sink state) and one takes as the initial state $q'_0 = B \times \{q_0\}$ (which possibly is the sink state). In particular the number of states of \mathcal{T}' is exponential in $|S|$ and linear in $|Q|$.

Now one can go back to the original arena and consider the enriched arena $\mathcal{A}_{\mathcal{T}'}$. Then we have the following easy lemma.

Lemma 3. *Eve has a \mathcal{T} -compatible positively winning strategy in the reachability game (\mathcal{A}, B, F) if and only if she has a positively winning strategy in the reachability game $(\mathcal{A}_{\mathcal{T}'}, B \times \{q'_0\}, F \times (Q' \setminus \{q'_s\}))$.*

Moreover, from a positively winning strategy in the second game using memory of size N one can effectively construct a \mathcal{T} -compatible positively winning strategy in the reachability game (\mathcal{A}, B, F) that uses a memory of size $\mathcal{O}(N \times 2^{|\mathcal{S}|} \times |Q|)$.

Proof. If Eve positively wins in $(\mathcal{A}_{\mathcal{T}'}, B \times \{q'_0\}, F \times (Q' \setminus \{q'_s\}))$ then we can safely assume that she necessarily always plays authorised (according to Act') actions (otherwise the play goes directly to $S \times \{q'_s\}$ and gets trap in it forever, hence cannot reach $F \times (Q' \setminus \{q'_s\})$, hence is \mathcal{T} -compatible thanks to Lemma 2. Such a strategy can be mimicked in the original game and it requires to simulate automaton \mathcal{T}' hence, costs an extra memory of size the one of \mathcal{T}' . Conversely, if it she has a positively winning \mathcal{T} -compatible strategy in the original game, the same strategy can be mimicked in the reduced game and is still positively winning. \square

Now combining Lemma 3 together with Theorem 3 concludes the proof of Theorem 4. \square

4.4 The Case Where Adam Is More Informed Than Eve

We now assume that Adam is more informed than Eve and we fix an automaton $\mathcal{T} = (Q, \Sigma_E \times S_{/\sim_E}, q_0, q_s, \delta_{\mathcal{T}}, Act)$ as in Section 4.3. Again, we are interested in checking whether Eve has a \mathcal{T} -compatible strategy that is positively winning in the reachability game (\mathcal{A}, B, F) . The idea here is to reduce this question to one on a game where Adam is perfectly informed and therefore conclude thanks to Theorem 4.

Recall that in this setting, as noted in Remark 1, we can safely assume that, against a fixed strategy of Eve, Adam observes the actions played by both players.

For this let \mathcal{H} be all those subsets of S that consist of \sim_A -equivalent states. For such a subset H and for any pair of actions $(\sigma_E, \sigma_A) \in (\Sigma_E \times \Sigma_A)$ define the set $Up(H, \sigma_E, \sigma_A) \in \mathcal{H}$ as follows. First, define $M = \{s' \in S \mid \exists s \in H \text{ s.t. } \delta(s, \sigma_E, \sigma_A)(s') > 0\}$ as the set of all possible successors of states in H when playing the pair of actions (σ_E, σ_A) and let $Up(H, \sigma_E, \sigma_A)$ consist of all those non-empty subsets H' that can be written as $H' = M \cap [s]_{\sim_A}$, *i.e.* all possible indistinguishable (for Adam) subsets of M .

Define now a new arena $\mathcal{A}' = (\mathcal{H}, \Sigma_E, \Sigma_A, \delta', \approx_E, \approx_A)$ by letting

- $\delta'(H, \sigma_E, \sigma_A)(H') = 1/|Up(H, \sigma_E, \sigma_A)|$ if $H' \in Up(H, \sigma_E, \sigma_A)$ and 0 otherwise;
- $H_1 \approx_E H_2$ if $s_1 \sim_E s_2$ for every $s_1 \in H_1$ and $s_2 \in H_2$; and
- \approx_A is the equality relation, *i.e.* Adam is perfectly informed.

Define the set of final states F' as those elements H in \mathcal{H} such that $H \cap F \neq \emptyset$.

Note that the equivalence classes of \approx_E can be identified with the equivalence classes of \sim_E (because $\sim_A \subseteq \sim_E$) and therefore one can define \mathcal{T} -compatible strategies for Eve also in a play in \mathcal{A}' . More generally, any Eve's strategy in one game can be used in the other one.

For a set $B \subseteq S$ define $\nu(B) \in \mathcal{H}$ as $\nu(B) = \{\{s\} \mid s \in B\}$. The following proposition relates game (\mathcal{A}, B, F) and game $(\mathcal{A}', \nu(B), F')$.

Proposition 3. *A strategy of Eve is a positively winning \mathcal{T} -compatible strategy in $\mathbb{G} = (\mathcal{A}, B, F)$ if and only if it is a positively winning \mathcal{T} -compatible strategy in $\mathbb{G}' = (\mathcal{A}', \nu(B), F')$.*

Proof. Let φ be a positively winning \mathcal{T} -compatible strategy in \mathbb{G} . Now use φ in \mathbb{G}' : obviously it is still \mathcal{T} -compatible and we only have to prove that it is positively winning. Consider a strategy ψ' of Adam in \mathbb{G}' . Then, assuming Eve respects φ , strategy ψ' can be mimicked in game \mathbb{G} : indeed, Adam simply has to update a state H in \mathcal{A}' which is done by computing $Up(H, \sigma_E, \sigma_A)$ and observing the equivalence class for \sim_A relation; assuming Eve respects φ it means that Adam always knows what action σ_E she will play and therefore can compute $Up(H, \sigma_E, \sigma_A)$. Call ψ the strategy in \mathbb{G} mimicking ψ' .

Now let N be some integer and consider all those partial plays of length N in \mathbb{G} where Eve respects φ and Adam respects ψ . Consider the \sim_A -equivalent classes among these partial plays and for every class consider the set H of possible last states. Then those such H are exactly those states that can be reached in \mathbb{G}' in a partial play of length N when Eve respects φ and Adam respects ψ' . As φ is positively winning in \mathbb{G} , thanks to Proposition 1 there is some N such that Eve positively wins within the N first moves and therefore for the same N we conclude that Eve positively wins within the N first moves in \mathbb{G}' using φ against ψ' . As this property does not depend on ψ' we conclude that φ is positively winning in \mathbb{G}' .

Conversely, assume she has a positively winning \mathcal{T} -compatible strategy in \mathbb{G}' . Now use φ in \mathbb{G} : obviously it is still \mathcal{T} -compatible and we only have to prove that it is positively winning. By contradiction, assume Adam has a strategy ψ that ensures, provided Eve uses φ in \mathbb{G} , that no final state is reached. Then, from ψ one can define a strategy in ψ' that consists in a partial play $H_0(\sigma_E^0, \sigma_A^0)H_1(\sigma_E^1, \sigma_A^1) \cdots H_k$ to play action $\psi([s_0]_{\sim_A} \cdots [s_k]_{\sim_A})$ where s_i is any (they are all \sim_A -equivalent) element in H_i for every i . Using the same argument as in the direct implication relating plays in \mathbb{G} when using strategies (φ, ψ) and plays in \mathbb{G}' when using strategies (φ', ψ') , one concludes that playing ψ' against φ in \mathbb{G}' ensures that no final state is visited hence, leading a contradiction with φ being positively winning in \mathbb{G}' . □

Combining Proposition 3 with Theorem 4 directly leads the following result.

Theorem 5. *When Adam is more informed than Eve, one can decide in double exponential time in $|S|$ and polynomial in $|Q|$ whether Eve has a \mathcal{T} -compatible strategy that is positively winning in the reachability game (\mathcal{A}, B, F) . If such a strategy exists, one can construct one that uses memory of size polynomial in $|Q|$ and doubly exponential in $|S|$.*

5 Almost-Surely Winning for Büchi Conditions

For the rest of this section fix an arena $\mathcal{A} = (S, \Sigma_E, \Sigma_A, \delta, \sim_E, \sim_A)$ and a set of final states $F \subseteq S$. We are interested in almost-sure winning strategies, and we focus on Büchi conditions, as a solution for this case permits to obtain a solution for reachability condition by a simple reduction (change the arena so that whenever a final state is reached then the play stays in it forever). For the moment we do not make any assumption on how Adam is informed.

We show how to compute the set of **almost-surely winning beliefs** of Eve, denoted \mathcal{B}^{AS} , which is the set of subsets $B \subseteq S$ such that $B \subseteq [s]_{\sim_E}$ for some $s \in S$ and for which Eve has an almost-surely winning strategy in the Büchi game $\mathbb{G}_B = (\mathcal{A}, B, F)$. For some $B \in \mathcal{B}^{AS}$ we let $[B]_{\sim_E} = [s]_{\sim_E}$ for $s \in S$ such that $B \subseteq [s]_{\sim_E}$.

5.1 Fixpoint Characterisation

Lemma 4 below states that the set \mathcal{B}^{AS} can be expressed as the greatest fix-point of a (monotone) mapping $\Xi : 2^{2^S} \rightarrow 2^{2^S}$ defined as follows. Let $\mathcal{B} \subseteq 2^S$ and let $B \in \mathcal{B}$. We say that B belongs to $\Xi(\mathcal{B})$ if Eve has a strategy in the reachability game (\mathcal{A}, B, F) which is positively winning and guarantees that her belief always stays in \mathcal{B} .

Lemma 4. \mathcal{B}^{AS} is the greatest fixpoint of Ξ .

Proof. We first argue that \mathcal{B}^{AS} is a fixpoint for Ξ . For this we consider any $B \in \mathcal{B}^{AS}$ and prove that $B \in \Xi(\mathcal{B}^{AS})$. We denote by \mathbb{G}_B the Büchi game (\mathcal{A}, B, F) and we start with a simple lemma.

Lemma 5. Let $B \in \mathcal{B}^{AS}$. Let φ be any strategy for Eve that is almost-surely winning for her in \mathbb{G}_B and let $\sigma_E = \varphi([B]_{\sim_E})$. Then, for any $\sigma_A \in \Sigma_A$, for any t such that $\exists s \in B$ with $\delta(s, \sigma_E, \sigma_A)(t) > 0$, $\text{UpBelief}_E(B, \sigma_E, [t]_{\sim_E}) \in \mathcal{B}^{AS}$.

Proof. Consider some action σ_A and some t such that $\delta(s, \sigma_E, \sigma_A)(t) > 0$ and let $B' = \text{UpBelief}_E(B, \sigma_E, [t]_{\sim_E})$. By definition of UpBelief_E , for every $t' \in B'$, there is some $s' \in B$ and some action $\sigma_A^{t'}$ such that $\delta(s', \sigma_E, \sigma_A^{t'})(t') > 0$. Now, define the strategy φ' of Eve by letting $\varphi'(\lambda) = \varphi([s]_{\sim_E} \cdot \lambda)$ for any partial play λ . We claim that φ' is almost-surely winning for Eve in $\mathbb{G}_{t'}$ for any $t' \in B'$, hence implying that $B' \in \mathcal{B}^{AS}$. By contradiction, assume that φ' is not almost-surely winning for some $\mathbb{G}_{t'}$ with $t' \in B'$ and let ψ' be a counter-strategy for Adam in $\mathbb{G}_{t'}$, i.e. $\Pr_{\psi'}^{\varphi', \psi'}(\mathcal{O}) < 1$ (recall that \mathcal{O} denotes here the Büchi objective). Now, pick $s' \in B$ such that $\delta(s', \sigma_E, \sigma_A^{t'})(t') > 0$ and define a strategy ψ of Adam by letting $\psi(s') = \sigma_A^{t'}$ and $\psi(s' \cdot \lambda) = \psi'(\lambda)$. Then as $\Pr_{\psi'}^{\varphi', \psi'}(\mathcal{O}) < 1$ one also has that $\Pr_{\psi}^{\varphi, \psi}(\mathcal{O}) < 1$ which leads to a contradiction. \square

Fix a strategy φ_B as in Lemma 5: a play λ in \mathbb{G}_B where Eve respects φ_B is such that $\text{Belief}_E^B(\lambda) \in \mathcal{B}^{AS}$. Moreover, as φ_B is almost-surely winning for the Büchi game \mathbb{G}_B , it is in particular positively winning in the *reachability* game (\mathcal{A}, B, F) . Hence, using Proposition 1, one gets a bound N_B and some ε_B , meaning that the probability of a play λ in \mathbb{G}_B where Eve respects φ_B to visit a final state within its first N_B moves is $\geq \varepsilon_B$. Hence, $B \in \Xi(\mathcal{B}^{AS})$, implying that \mathcal{B}^{AS} is a fixpoint for Ξ .

Now we show that any fixpoint of Ξ is included in \mathcal{B}^{AS} . For this assume that $\Xi(\mathcal{B}) = \mathcal{B}$ for some \mathcal{B} . As any $B \in \mathcal{B}$ is such that $B \in \Xi(\mathcal{B})$ it comes with some φ_B , N_B and ε_B . We let $N = \max\{N_B \mid B \in \mathcal{B}\}$ and $\varepsilon = \min\{\varepsilon_B \mid B \in \mathcal{B}\}$.

Now we define a strategy φ that consists in playing in rounds of length N : at the beginning of some round, Eve considers her current belief H and plays according to φ_H in the next N moves; then she restarts with the updated belief, and so on forever.

Now consider some $B \in \mathcal{B}$. We claim that φ is almost-surely winning for Eve in any in \mathbb{G}_B . Indeed, from the properties of the φ_H , it follows that any play in \mathbb{G}_B where Eve respects φ is such that the belief is in \mathcal{B} . Now, as the φ_H ensure to visit a final state with probability $\geq \varepsilon$ in less than N moves the Borel-Cantelli Lemma implies that φ is almost-surely winning. Hence, $B \in \mathcal{B}^{AS}$ and this concludes the proof. \square

5.2 Decidability Issues

As Ξ is monotone for set inclusion, it suffices to compute \mathcal{B}^{AS} by successive applications (starting with the set of all subsets) of the operator Ξ until reaching the fixpoint. Since $\mathcal{B}^{AS} \subseteq 2^S$, the fixpoint is reached in at most $2^{|S|}$ steps.

Now, as noted in Remark 5 the property for a strategy to guarantee that Eve's belief remains in a set \mathcal{B} can be expressed as the strategy being $\mathcal{T}_{\mathcal{B}}$ -compatible (and the number of states of $\mathcal{T}_{\mathcal{B}}$ is at most exponential in $|S|$). Therefore, thanks to Theorem 4 (resp. Theorem 5) every step in the fixpoint computation can be achieved in time exponential (resp. doubly exponential) in $|S|$ if Adam is perfectly informed (resp. more informed than Eve). This leads the following result.

Theorem 6. Let \mathbb{G} be a Büchi (or reachability) game with n states.

- If Adam is perfectly informed, one can decide whether Eve has an almost-surely winning strategy in time exponential in n . If such a strategy exists, it can be effectively constructed and requires memory at most exponential in n .
- If Adam is more informed than Eve, one can decide whether Eve has an almost-surely winning strategy in time doubly exponential in n . If such a strategy exists, it can be effectively constructed and requires memory at most doubly exponential in n .

Proof. Decidability follows from Theorem 4/Theorem 5 and the fixpoint characterisation given in Lemma 4. The result on the strategies is also a consequence of Theorem 4/Theorem 5 combined with Corollary 1 which permits to bound the size of N in the proof of Lemma 4. \square

6 Lower Bound

We now give a matching lower bound to the upper bounds in Theorem 5 and in Theorem 6 for the case where Adam is more informed than Eve. Note that in the case where Adam is perfectly informed one can get a matching lower bound (ExpTime-hardness) as in the case where randomised strategies are allowed [7]. Also note that in the case where Adam is more informed than Eve similar lower bounds, when randomised strategies are allowed, were obtained in [13, 2, 3] for almost-sure winning³; however, even if the ideas of the proof below are similar to the ones in [13, 2, 3], namely the players simulate a run of an exponential space alternating Turing machine while gadgets prevent cheating, there is a slight but crucial difference. Indeed, Eve is in charge of describing the successive configuration while in previous proofs Adam was; we actually believe that this is needed for the proof to work (mainly because we are not only interested in almost-sure winning but also in positively winning).

Theorem 7. *Deciding whether Eve has a positively winning (resp. an almost-surely winning) strategy in a reachability game where Adam is more informed than her is a 2-ExpTime-hard problem.*

Proof. The idea is to simulate a computation of an alternating Turing machine that uses a space of exponential size and to reduce termination to almost-surely winning for Eve. As alternating Turing machines of exponential space are equivalent to deterministic Turing machines working in doubly exponential time it permits to obtain the desired lower bound. We can safely assume that the tape alphabet A contains a blank symbol as well as a special symbol \sharp , and that initially the input tape is made of n successive \sharp symbols followed by $2^n - n$ blank symbols. A configuration of the machine can be described by a word of length 2^n in A^*QA^* where Q is the set of states of the machine (including some final states): the meaning of a configuration $a_1 \cdots a_\ell q a_{\ell+1} \cdots a_{2^n}$ is that the tape content is $a_1 \cdots a_\ell a_{2^n}$, the state is q and the reading/writing head is on the ℓ -th cell. A run of the machine is a sequence of successive configurations separated by transitions of the machine; it is accepting if it contains a final configuration (and in that case the run is of finite length; otherwise it is of infinite length).

A classical way of thinking of an alternating Turing machine is as a game where Eve is in charge of the choice of transitions when the machine is in an existential state while Adam takes care of the universal states. The machine accepts if and only if Eve has a winning strategy to eventually reach a configuration with a final control state.

Consider now the following (informal) game. Eve is in charge of describing the run of the Turing machine (her actions' alphabet contains all the necessary symbols for that *i.e.* $A \cup Q$ that permits the game to go in some associated states). After she described a configuration either she (in case the state is existential) or Adam (in case the state is universal) describes a valid transition of the machine (again by playing some special actions), and then Eve describes the successive configuration and so on until possibly a final configuration is reached (in which case she wins the game). Hence, in the game's state one stores the state (of the Turing machine) when described as well as the adjacent symbols; this information is used when Eve/Adam has to describe the next transition of the machine (that should be consistent).

Of course the problem is that Eve could cheat by not describing a valid run. For this, Adam can, in every configuration, secretly (*i.e.* Eve does not observe it) mark a cell of the tape, and in the next configuration he can indicate a cell (supposedly of same index than the previously marked one) and it is checked whether it has been wrongly updated: this is easily done as the cell before and after the marked cell have been stored in the arena (and Eve does not observe it of course) and together with the transition one can compute the correct update of the cell. Now in

³Actually the lower bound in [13] uses a game where none of the player is more informed than the other but it is easily seen how to modify it to obtain a game where Adam is more informed than Eve

case there is indeed a wrong update of the cell content, the play restarts (*i.e.* the players restart from the initial configuration of the Turing machine); otherwise the play goes to a final state and Eve wins. Hence, in the game's state one also stores (and hides to Eve) information of a cell marked by Adam and of the adjacent symbols for a later check. This marking by Adam as well as the checking later is done by him playing a distinguished action.

One problem in the previous simulation is that Adam could cheat by indicating two cells that are not with the same index. If the space used by the machine was of linear size, one could of course store the actual index and formally check it. Here, we use an extra coding to circumvent this problem. When describing the configuration, after every symbol Eve produces a sequence of n bits whose meaning is to describe, in binary counting, the index of the last symbol. When she describes such a binary number, Adam can secretly mark a bit that he claims will be not correctly updated when describing the index of the next symbol (for this he just plays an action that stands for a number between 1 and n) and this is checked next: if she made an incorrect update, the play restarts (*i.e.* the players restart from the initial configuration of the Turing machine); otherwise the play goes to a final state where she wins. Hence, in the game's state one also stores (and hides to Eve) the index (between 1 and n) of a bit marked by Adam as well as the value corresponding to the bit of the same index in the incremented version of the described number (this can be computed on the fly). This marking by Adam is done by him playing a distinguished action; the checking is done deterministically (thanks to a counter). One also uses this binary encoding of the index of the cell in the following way: whenever Adam marks a symbol that he claims will be incorrectly updated in the next configuration, a bit of its binary encoding is guessed (*i.e.* randomly chosen) and its index is stored and not observed by none of the players. Later, when Adam indicates the supposed corresponding symbol in the next configuration, the guessed bit is checked and should match: if not the play goes to a final state and Eve wins; otherwise one does as previously explained (*i.e.* one checks whether the symbol is correct: if not the play restarts otherwise the play goes to a final state and Eve wins). Hence, in the game's state one also stores (and hides to both players) the value and index of the randomly chosen bit.

We claim that Eve positively wins (equivalently almost-surely wins) this game if and only if the Turing machine accepts. Once this is established, the proof will be over, as one can easily notice that the previous informal game can be encoded formally as a two-player game with imperfect information of polynomial size in the one of the Turing machine.

First, assume that the Turing machine accepts. Hence, it means that the existential player Eve has a winning strategy in the acceptance game of the machine. Now, mimic this strategy in the above described game: Eve always makes a correct description of a run and, when she has to choose a transition of the machine, she does as in her winning strategy in the acceptance game of the machine. We claim that this strategy is almost-surely winning (hence, also positively winning). Indeed, any strategy of Adam that does not infinitely often claim that a cell is incorrectly updated is surely losing for him because either he makes a wrong claim (actually his claims are always wrong but here we mean he get discovered because of the hidden bit), or after some point the simulation goes to the end and finishes by a final configuration of the Turing machine. Now, against this strategy of Eve, when Adam infinitely often claims that a cell is incorrectly updated, he almost-surely gets caught because at every claim there is a (fixed positive) probability (at least $1/n$) that the secret bit does not match, and therefore, by Borel-Cantelli Lemma, the probability that he gets caught eventually is 1. Of course, if Adam claims at some point that a bit is incorrectly updated by Eve he also loses (because she describes a valid run). Hence, Eve's strategy almost-surely defeats any strategy of Adam.

Conversely, assume that the Turing machine does not accept. Hence, it means that the existential player Eve has no winning strategy in the acceptance game of the machine. Now, consider a strategy of Eve. There are two possibilities.

- Either there is a strategy⁴ of Adam against which Eve's strategy eventually cheats. Then, consider the strategy of Adam that plays the same except that he points the moment where she cheats: then, Eve must behave the same and therefore the play restarts. Now, consider

⁴In fact a set of indistinguishable strategies from Eve's point of view, including the ones where Adam claims she cheats.

how she behaves in the restarted play and do the same reasoning. If we are always in the same situation, by iteratively playing a strategy pointing where she cheats in every simulation of the Turing machine ensures that no final configuration is reached and therefore that she surely loses.

- Or, against any strategy of Adam, Eve’s strategy never cheats (*i.e.* describes a valid run). Hence, Eve’s strategy can be seen as a strategy in the acceptance game of the machine and therefore, one can consider the strategy of the universal player that beats it in the acceptance game and let Adam mimic it in the simulation game (and he never claims that she cheats). Then, this strategy leads to an infinite play that corresponds to the description of an infinite run of the alternating Turing machine that never visits a final configuration: hence, it surely defeats Eve’s strategy

In conclusion, for any strategy of Eve in the above described game there is a strategy of Adam that surely beats this strategy, which implies that there is no positively winning (hence almost-surely winning) strategy for Eve. This terminates the proof. \square

7 Summary

In this paper we considered finite state games in which, at each round, the two players (called Eve and Adam) choose concurrently an action and based on these actions the successor state is chosen according to some fixed probability distribution. We considered several classical winning conditions: safety, reachability, Büchi and co-Büchi. Moreover, the players are imperfectly informed: each player has an equivalence relation over states and, instead of observing the exact state, he observes its equivalence class. Finally, we restricted our attention to pure strategies, *i.e.* we forbid the players to randomise when choosing their actions.

We studied the decidability and complexity status of the problem of deciding whether Eve has a positively (resp. almost-surely) winning strategy. To obtain positive results, we imposed restrictions on how Adam is informed: we considered the case where he has perfect information and the case where he is more informed than Eve.

The landscape of decidability and undecidability results with pointers to the literature and to the results in our paper are shown in the Table 1. The entries of the form “1/2-Exptime-complete” refer to the two cases of Adam being perfectly informed and being better informed than Eve, respectively (the results from [8] are for the case of Adam being perfectly informed). The implication \Rightarrow means that our result is an easy consequence of a result from the literature. The undecidability results already hold for the case in which Adam is perfectly informed.

	Safety	Reachability	Büchi	co-Büchi
Positively	Undecidable Th. 2	1/2-Exptime-complete [8], Th. 3/Th. 5 + Th. 7	Undecidable [1] \Rightarrow Th. 1	Undecidable Th. 2
Almost Sure	ExpTime-comp. [4]	1/2-Exptime-complete [8], Th. 6 + Th. 7	1/2-Exptime-complete Th. 6 + Th. 7	Undecidable [1] \Rightarrow Th. 1

Table 1: Landscape of decidability and undecidability results

References

- [1] Christel Baier, Nathalie Bertrand, and Marcus Größer. On decision problems for probabilistic büchi automata. In *Proceedings of the 11th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS 2008)*, volume 4962 of *Lecture Notes in Computer Science*, pages 287–301. Springer-Verlag, 2008.

- [2] N. Bertrand, B. Genest, and H. Gimbert. Qualitative determinacy and decidability of stochastic games with signals. In *Proceedings of Logic in Computer Science (LiCS'09)*, pages 319–328. IEEE Computer Society, 2009.
- [3] Nathalie Bertrand, Blaise Genest, and Hugo Gimbert. Qualitative determinacy and decidability of stochastic games with signals. *Journal of the Association for Computing Machinery (ACM)*, 64(5):33:1–33:48, 2017.
- [4] Dietmar Berwanger and Laurent Doyen. On the power of imperfect information. In *Proceedings of the 28th International Conference on Foundations of Software Technology and Theoretical Computer Science (FST&TCS 2008)*, volume 2 of *LIPICs*, pages 73–82. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2008.
- [5] Arnaud Carayol, Axel Haddad, and Olivier Serre. Randomisation in automata on infinite trees. *ACM Transactions on Computational Logic*, 15(3), 2014.
- [6] K. Chatterjee. *Stochastic ω -Regular Games*. PhD thesis, University of California, 2007.
- [7] K. Chatterjee, L. Doyen, T.A. Henzinger, and J.-F. Raskin. Algorithms for omega-regular games with imperfect information. *LMCS*, 3(3), 2007.
- [8] Krishnendu Chatterjee and Laurent Doyen. Partial-observation stochastic games: How to win when belief fails. *ACM Transactions on Computational Logic*, 15(2):16, 2014.
- [9] L. de Alfaro and T.A. Henzinger. Concurrent omega-regular games. In *Proceedings of Logic in Computer Science (LiCS'00)*, pages 141–154. IEEE Computer Society, 2000.
- [10] Luca de Alfaro, Thomas A. Henzinger, and Orna Kupferman. Concurrent reachability games. *Theoretical Computer Science*, 386(3):188–217, 2007.
- [11] Nathanaël Fijalkow, Sophie Pinchinat, and Olivier Serre. Emptiness of alternating tree automata using games with imperfect information. In *Proceedings of the 33rd International Conference on Foundations of Software Technology and Theoretical Computer Science (FST&TCS 2013)*, volume 24 of *LIPICs*, pages 299–311. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2013.
- [12] Hugo Gimbert and Youssouf Oualhadj. Probabilistic automata on finite words: Decidable and undecidable problems. In *Proceedings of the 37th International Colloquium on Automata, Languages, and Programming (ICALP 2010)*, volume 6199 of *Lecture Notes in Computer Science*, pages 527–538. Springer-Verlag, 2010.
- [13] V. Gripon and O. Serre. Qualitative concurrent stochastic games with imperfect information. In *Proceedings of the 36th International Colloquium on Automata, Languages, and Programming (ICALP 2009)*, volume 5556 of *Lecture Notes in Computer Science*, pages 200–211. Springer-Verlag, 2009.
- [14] A. Paz. *Introduction to probabilistic automata*. Academic Press New York, 1971.
- [15] P.J. Ramadge and W.M. Wonham. Supervisory Control of a Class of Discrete Event Processes. *SIAM Journal on Control and Optimization*, 25:206, 1987.
- [16] J.H. Reif. The complexity of two-player games of incomplete information. *Journal of Computer and System Sciences*, 29(2):274–301, 1984.