

ON THE EXISTENCE OF REACTIVE STRATEGIES RESILIENT TO DELAY

MARTIN FRÄNZLE ^a, PAUL KRÖGER ^a, SARAH WINTER ^b, AND MARTIN ZIMMERMANN ^c

^a Carl von Ossietzky Universität, Oldenburg, Germany
e-mail address: martin.fraenzle@uol.de, p.kroeger@uol.de

^b IRIF & Université Paris Cité, Paris, France
e-mail address: sarah.winter@irif.fr

^c Aalborg University, Aalborg, Denmark
e-mail address: mzi@cs.aau.dk

ABSTRACT. We compare games under delayed control and delay games, two types of infinite games modelling asynchronicity in reactive synthesis. In games under delayed control both players suffer from partial informedness due to symmetrically delayed communication, while in delay games, the protagonist has to grant lookahead to the alter player.

Our first main result, the interreducibility of the existence of sure winning strategies for the protagonist, allows to transfer known complexity results and bounds on the delay from delay games to games under delayed control, for which no such results had been known. We furthermore analyse existence of randomized strategies that win almost surely, where this correspondence between the two types of games breaks down. In this setting, some games surely won by the alter player in delay games can now be won almost surely by the protagonist in the corresponding game under delayed control, showing that it indeed makes a difference whether the protagonist has to grant lookahead or both players suffer from partial informedness. These results get even more pronounced when we finally address the quantitative goal of winning with a probability in $[0, 1]$. We show that for any rational threshold $\theta \in [0, 1]$ there is a game that can be won by the protagonist with exactly probability θ under delayed control, while being surely won by alter in the delay game setting. All these findings refine our original result that games under delayed control are not determined.

1. INTRODUCTION

Two-player zero-sum games of infinite duration are a standard model for the synthesis of reactive controllers, i.e., correct-by-construction controllers that satisfy their specification even in the presence of a malicious environment. In such games, the interaction between the controller and the environment is captured by the rules of the game and the specification on the controller induces the winning condition of the game. Then, computing a correct controller boils down to computing a winning strategy.

Key words and phrases: Games under delayed control, delay games, synthesis.

* Full version of [FWZ23], contains all proofs omitted in the conference version as well as a new section on winning games under delayed control with mixed strategies with respect to a fixed threshold (Section 6).

Often, it is convenient to express the rules in terms of a graph capturing the state-space such that moves correspond to transitions between these states. The interaction between the controller and the environment then corresponds to a path through the graph and the winning condition is a language of such paths, containing those that correspond to interactions that satisfy the specification on the controller.

In other settings, it is more convenient to consider a slightly more abstract setting without game graphs, so-called Gale-Stewart games [GS53]. In such games, the players alternately pick a sequence of letters, thereby constructing an infinite word. The winning condition is a language over infinite words, containing the winning words for one player. To capture the synthesis problem, the winning condition has to encode both the specification on the controller as well as the rules of interaction. It is straightforward to transform a graph-based game into a Gale-Stewart game and a Gale-Stewart game into a graph-based game such that the existence of winning strategies for both players is preserved.

In the most basic setting of synthesis, both the controller and the environment are fully informed about the current state of the game (complete information). However, this scenario is not always realistic. Thus, much effort has been poured into studying games under incomplete information where the players are only partially informed about the current state of the game. Here, we are concerned with a special type of partial information designed to capture delays in perception and action. Such delays either render the most recent moves of the opponent invisible to a player or induce a time lag between the selection and the implementation of an own move, respectively.

As a motivating example, consider the domain of cooperative driving: Here, the exchange of information between cars is limited (and therefore delayed) by communication protocols that have to manage the available bandwidth to transfer information between cars. Other delaying factors include, e.g., complex signal processing chains based on computer vision to detect the locations of obstacles. Thus, decisions have to be made based on incomplete information, which only arrives after some delay.

1.1. Games under Delayed Control. Chen et al. [CFL⁺21] introduced (graph) games under delayed control to capture this type of incomplete information. Intuitively, assume the players so far have constructed a finite path $v_0 \cdots v_n$ through the graph. Then, the controller has to base their decision on a visible proper prefix $v_0 \cdots v_{n-\delta}$, where δ is the amount of delay. Hence, the suffix $v_{n-\delta+1} \cdots v_n$ is not yet available to base the decision on, although the decision to be made is to be applied at the last state v_n in the sequence.

Chen et al. showed that solving games under delayed control with safety conditions and with respect to a given delay is decidable: They presented two algorithms, an exponential one based on a reduction to delay-free safety games using a queue of length δ , and a more practical incremental algorithm synthesizing a series of controllers handling increasing delays and reducing game-graph size in between. They showed that even a naïve implementation of this algorithm outperforms the reduction-based one, even when the latter is used with state-of-the-art solvers for delay-free games. However, the exact complexity of the incremental algorithm and that of solving games under delayed control remained open.

Note that asking whether there is some delay δ that allows controller to win reduces to solving standard, i.e., delay-free games, as they correspond to the case $\delta = 0$. The reason is monotonicity in the delay: if the controller can win for delay δ then also for any $\delta' < \delta$. More interesting is the question whether controller wins with respect to every possible delay. Chen

et al. conjectured that there is some exponential δ such that if the controller wins under delay δ , then also under every δ' . We call a δ with this property decisive.

1.2. Delay Games. There is also a variant of Gale-Stewart games modelling delayed interaction between the players [HL72]. Here, the player representing the environment (often called Player I for input) has to provide a lookahead on their moves, i.e., the player representing the controller (accordingly called Player O for output) has access to the first $n + k$ letters picked by Player I when picking their n -th letter. So, k is the amount of lookahead that Player I has to grant Player O . Note that the lookahead benefits Player O (representing the controller) while the delay in a game under delayed control disadvantages the controller.

Only three years after the seminal Büchi-Landweber theorem showing that delay-free games with ω -regular winning conditions are decidable [BL69], Hosch and Landweber showed that it is decidable whether there is a k such that Player O wins a given Gale-Stewart game with lookahead k [HL72]. Forty years later, Holtmann, Kaiser, and Thomas [HKT12] revisited these games (and dubbed them delay games). They proved that if Player O wins a delay game then Player O wins it already with at most doubly-exponential lookahead (in the size of a given deterministic parity automaton recognizing the winning condition). Thus, unbounded lookahead does not offer any advantage over doubly-exponential lookahead in games with ω -regular winning conditions. Furthermore, they presented an algorithm with doubly-exponential running time solving delay games with ω -regular winning conditions, i.e., determining whether there exists a k such that Player O wins a given delay game (with its winning condition again given by a deterministic parity automaton) with lookahead k .

Both upper bounds were improved and matching lower bounds were later proven by Klein and Zimmermann [KZ16a]: Solving delay games is EXPTIME-complete and exponential lookahead is both necessary to win some games and sufficient to win all games that can be won. Both lower bounds already hold for winning conditions specified by deterministic safety automata while the upper bounds hold for deterministic parity automata. The special case of solving delay games with conditions given as reachability automata is PSPACE-complete, but exponential lookahead is still necessary and sufficient. Thus, there are tight complexity results for delay games, unlike for games under delayed control.

1.3. On Player Names. Note that games under delayed control and delay games are in a sense asymmetric: In a game under delayed control, one models situations in which the controller is disadvantaged by the delayed access to environment's moves while in a delay game, one models situations in which the controller (Player O) benefits from a lookahead on environment's (Player I) moves. Said differently, in a delay game, Player I is disadvantaged by having to provide a lookahead on their moves to Player O . Thus, it is natural to ask whether there is a connection between the two disadvantaged players (i.e., controller in a game under delayed control and Player I in a delay game) and between the other two players (environment and Player O).

For the sake of consistency with existing literature, we prefer to stick to the player names “controller” and “environment” as well as “Player I ” and “Player O ”, even though in a delay game Player O is typically understood to represent the controller.

Also, let us remark that winning conditions in both types of games are always formulated from the perspective of controller, i.e., for controller in games under delayed control and for

winning conditon	complexity	decisive delay
safety	PSPACE-complete	exponential
reachability	EXPTIME-complete	exponential
parity	EXPTIME-complete	exponential
LTL	3EXPTIME-complete	triply-exponential

Table 1: Our results for games under delayed control. “Complexity” refers to the problem of determining whether controller wins the game for every possible delay. All bounds for the decisive delays are tight.

Player O in delay games. Hence, as we will relate controller and Player I , we always have to complement winning conditions.

1.4. Our Contributions. In this work, we prove that there is indeed a tight relation between controller in a game under delayed control and Player I in a delay game. More precisely, we show that one can transform a safety game under delayed control in polynomial time into a delay game with a reachability condition for Player O (i.e., with a safety condition for Player I) such that controller wins the game under delayed control with even delay δ if and only if Player I wins the resulting delay game with lookahead of size $\frac{\delta}{2}$. Dually, we show that one can transform a delay game with safety condition for Player I in polynomial time into a reachability game under delayed control such that Player I wins the delay game with lookahead of size δ if and only if controller wins the resulting game under delayed control with delay 2δ . Thus, we can transfer both upper and lower bound results on complexity and on (necessary and sufficient) lookahead from delay games to delayed control. In particular, determining whether controller wins a given safety game under delayed control for every possible delay is PSPACE-complete. Our reductions also prove the conjecture by Chen et al. on the delays that allow controller to win such games: There is an exponential decisive delay in games under delayed control with safety conditions.

Furthermore, we generalize our translation from games with safety conditions to games with reachability conditions, games with parity conditions, and games with winning conditions given by formulas of Linear Temporal Logic (LTL) [Pnu77], again allowing us to transfer known results for delay games to games under delayed control. Table 1 lists our results.

Note that we have only claimed that the existence of winning strategies for the controller in the game under delayed control and Player I in the delay game coincides. This is no accident! In fact, the analogous result for relating environment and Player O fails. This follows immediately from the fact that delay games are determined while games under delayed control are undetermined, even with safety conditions. The reason is that the latter games are truly incomplete information games (which are typically undetermined) while delay games are complete information games.

We furthermore refine these findings by a detailed comparison between environment and Player O both in the setting of pure (deterministic) and in the setting of mixed (randomized) strategies. The latter setting increases power for both the controller and the environment, making them win (almost surely) games under delayed control that remain undetermined in the deterministic setting, but it also breaks the correspondence between controller and Player I observed in the deterministic setting: there are games that controller wins almost

surely while the corresponding Player I surely loses them — and thus Player O , generally corresponding to the environment, wins.

We finally provide results concerning existence of strategies that do not have to win surely or almost surely, but guarantee a win with a given probability $\theta \in]0, 1[$ at least. Again, this relaxation of the winning condition to at least θ -sure winning does not change anything for delay games, as they are determined, meaning that one of the two parties Player I and Player O has a sure (and thus also an almost sure and an at least θ -sure) winning strategy. For games under delayed control, the relaxation however makes a tremendous difference: we show that for every rational $\theta \in [0, 1]$ there exists an ω -regular game that the controller, taking its optimal strategy, wins with probability θ exactly. By symmetry, the same applies to the environment.

These findings confirm our original intuition that delay games and games under delayed control are indeed different, with the latter in relation giving more power to controller by symmetrically subjecting both parties, i.e., the environment also, to incomplete information. In the asymmetric setting of delay games, delay comes as a penalty for Player I only, while it is a benefit for Player O , who is granted a lookahead.

2. PRELIMINARIES

We denote the non-negative integers by \mathbb{N} . An alphabet Σ is a non-empty finite set of letters. A word over Σ is a finite or infinite sequence of letters of Σ : The set of finite words (non-empty finite words, infinite words) over Σ is denoted by Σ^* (Σ^+ , Σ^ω). The empty word is denoted by ε , the length of a finite word w is denoted by $|w|$. Given two infinite words $\alpha \in (\Sigma_0)^\omega$ and $\beta \in (\Sigma_1)^\omega$, we define $\binom{\alpha}{\beta} = \binom{\alpha(0)}{\beta(0)} \binom{\alpha(1)}{\beta(1)} \binom{\alpha(2)}{\beta(2)} \cdots \in (\Sigma_0 \times \Sigma_1)^\omega$.

2.1. Games under Delayed Control. Games under delayed control are played between two players, controller and environment. For pronomial convenience [McN00], we refer to controller as she and environment as he.

A game $\mathcal{G} = (S, s_0, S_c, S_e, \Sigma_c, \Sigma_e, \rightarrow, \text{Win})$ consists of a finite set S of states partitioned into the states $S_c \subseteq S$ of the controller and the states $S_e \subseteq S$ of the environment, an initial state $s_0 \in S_c$, the sets of actions Σ_c for the controller and Σ_e for the environment, and a transition function $\rightarrow: (S_c \times \Sigma_c) \cup (S_e \times \Sigma_e) \rightarrow S$ such that $s \in S_c$, $\sigma \in \Sigma_c$ implies $\rightarrow(s, \sigma) \in S_e$, and $s \in S_e$, $\sigma \in \Sigma_e$ implies $\rightarrow(s, \sigma) \in S_c$. Finally, $\text{Win} \subseteq S^\omega$ is the winning condition of the game. We write $s \xrightarrow{\sigma} s'$ as shorthand for $s' = \rightarrow(s, \sigma)$.

A play in \mathcal{G} is an infinite sequence $\pi = \pi_0 \sigma_0 \pi_1 \sigma_1 \pi_2 \sigma_2 \cdots$ satisfying $\pi_0 = s_0$ and $\pi_n \xrightarrow{\sigma_n} \pi_{n+1}$ for all $n \geq 0$. We say that controller wins π if $\pi_0 \pi_1 \pi_2 \cdots \in \text{Win}$; otherwise, we say that environment wins π . The play prefix of π of length n is defined as $\pi[n] = \pi_0 \sigma_0 \cdots \sigma_{n-1} \pi_n$, i.e., n is the number of actions (equivalently, the number of transitions). We denote by $\text{Pref}(\mathcal{G})$ the set of play prefixes of all plays in \mathcal{G} , which is partitioned into the sets $\text{Pref}_c(\mathcal{G})$ and $\text{Pref}_e(\mathcal{G})$ of play prefixes ending in S_c and S_e , respectively. Due to our alternation assumption, play prefixes of even (odd) length are in $\text{Pref}_c(\mathcal{G})$ ($\text{Pref}_e(\mathcal{G})$).

Fix some even $\delta \geq 0$. A strategy for the controller in \mathcal{G} under delay δ is a pair (α, τ_c) where $\alpha \in (\Sigma_c)^{\frac{\delta}{2}}$ and $\tau_c: \text{Pref}_c(\mathcal{G}) \rightarrow \Sigma_c$ maps play prefixes ending in S_c to actions of the controller. A play $\pi_0 \sigma_0 \pi_1 \sigma_1 \pi_2 \sigma_2 \cdots$ is consistent with (α, τ_c) if $\sigma_0 \sigma_2 \cdots \sigma_{\delta-4} \sigma_{\delta-2} = \alpha$ and $\sigma_{2n} = \tau_c(\pi[2n - \delta])$ for all $2n > \delta - 2$, i.e., controller has access to environment's actions with a delay of δ . In particular, her first $\frac{\delta}{2} + 1$ actions are independent of environment's

actions and, in general, her n -th action σ_{2n} only depends on the actions $\sigma_1, \sigma_3, \dots, \sigma_{(2n-\delta)-1}$ picked by environment, but not on the actions $\sigma_{(2n-\delta)+1}, \sigma_{(2n-\delta)+3}, \dots, \sigma_{2n-1}$ picked by environment. The strategy (α, τ_e) is winning under delay δ if every play that is consistent with it is winning for controller. Controller wins \mathcal{G} under delay δ if she has a winning strategy under delay δ for \mathcal{G} .

Remark 2.1.

- (1) The notion of winning strategy for controller under delay 0 is the classical one for delay-free games (cf. [GTW02]).
- (2) If controller wins \mathcal{G} under delay δ , then also under every delay $\delta' < \delta$ [CFL⁺21].

A strategy for environment is a mapping $\tau_e: \text{Pref}_e(\mathcal{G}) \rightarrow \Sigma_e$. A play $\pi_0\sigma_0\pi_1\sigma_1\pi_2\sigma_2\cdots$ is consistent with τ_e if $\sigma_{2n+1} = \tau_e(\pi_0\sigma_0\cdots\sigma_{2n-1}\pi_{2n+1})$ for all $n \geq 0$, i.e., environment has access to the full play prefix when picking his next action. The strategy τ_e is winning, if every play that is consistent with it is winning for the environment (i.e., the sequence of states is not in Win). Further, we say that environment wins \mathcal{G} , if he has a winning strategy for \mathcal{G} . Note that the two definitions of strategies are in general not dual, e.g., the one for environment is not defined with respect to a delay δ .

Remark 2.2. The notion of winning strategy for environment is the classical one for delay-free games (cf. [GTW02]).

We say that a game under delayed control \mathcal{G} is determined under delay δ , if either controller wins \mathcal{G} under delay δ or environment wins \mathcal{G} . Let us stress that determinacy is defined with respect to some fixed δ and that \mathcal{G} may be determined for some δ , but undetermined for some other δ' (due to the non-dual definition of strategies). Remark 4.11 shows an undetermined safety game under delayed control.

Example 2.3. Consider the game $\mathcal{G} = (S, s_0, S_c, S_e, \Sigma_c, \Sigma_e, \rightarrow, \text{Win})$ depicted in Figure 1 where Win contains all plays that do not visit the black vertex. Note that this is a safety condition. In particular, if controller does not pick action b at c_2 and does not pick action a at c_3 , then the vertex e_3 is never reached. This is straightforward without delay, but we claim that controller can also win \mathcal{G} under delay 2.

To gain some intuition, consider a play prefix $\pi_0\sigma_0\pi_1\cdots\pi_{n-1}\sigma_{n-1}\pi_n$ with $n \geq 4$ and $\pi_n \in S_c$. Then, controller has to pick an action σ_n to continue the prefix. However, due to the delayed control, she has to do so based on the prefix $\pi_0\sigma_0\pi_1\cdots\pi_{n-3}\sigma_{n-3}\pi_{n-2}$.

If π_{n-2} is c_2 , then π_n is either c_3 or c_1 . Hence, picking $\sigma_n = b$ is the only safe choice. Dually, if π_{n-2} is c_3 , then π_n is either c_2 or c_1 . Hence, picking $\sigma_n = a$ is the only safe choice.

Finally, assume π_{n-2} is c_1 . Then, π_n is either c_2 or c_3 . In the former case, picking $\sigma_n = a$ is the only safe choice, in the latter case, picking $\sigma_n = b$ is the only safe choice. So, controller needs to distinguish these two cases, although she has no access to π_n .

But she can do so by inspecting π_{n-3} (which she has access to): As a predecessor of $\pi_{n-2} = c_1$, it can either be e_4 , e_5 , or e_3 . In the latter case, the play is already losing. Thus, we disregard this case, as we construct a winning strategy. So, assume we have $\pi_{n-3} = e_4$ (the case $\pi_{n-3} = e_5$ is dual). Then, we must have $\pi_{n-4} = c_2$ (the only predecessor of e_4) and, by our analysis of the safe moves above, controller must have picked $\sigma_{n-2} = b$ (based, due to delay, on the prefix ending in $\sigma_{n-4} = c_2$). From this we can conclude $\pi_{n-1} = e_2$ and thus $\pi_n = c_3$ (the only successor of e_2). Thus, she can safely pick $\sigma_n = b$.

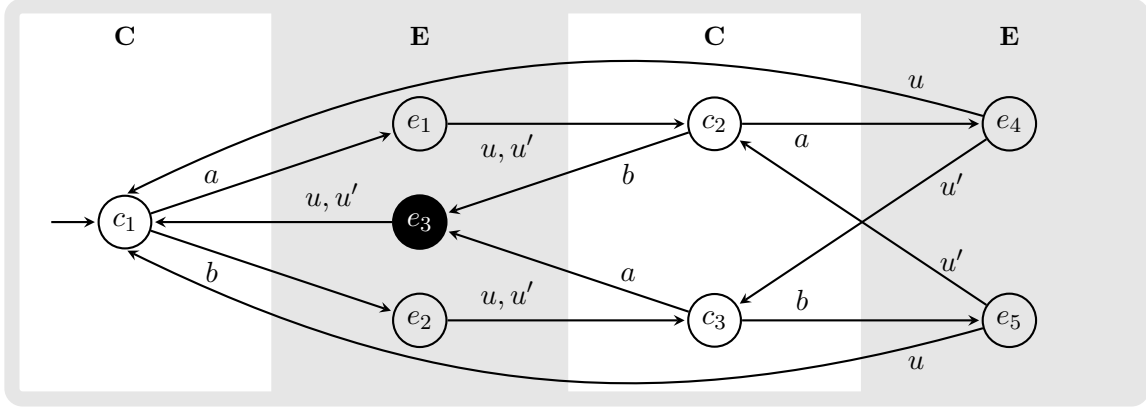


Figure 1: The game for Example 2.3. Controller wins all plays that never visit the black vertex. Note that we have $\Sigma_c = \{a, b\}$ and $\Sigma_e = \{u, u'\}$.

This intuition, and the necessary initialization, is implemented by the strategy (α, τ_c) with $\alpha = a$ and

$$\tau_c(\pi_0 \sigma_0 \pi_1 \cdots \pi_{n-3} \sigma_{n-3} \pi_{n-2}) = \begin{cases} a & n = 2 \text{ and } \pi_0 = c_1, \\ b & n > 2, \pi_{n-2} = c_1, \text{ and } \pi_{n-3} = e_4, \\ a & n > 2, \pi_{n-2} = c_1, \text{ and } \pi_{n-3} = e_5, \\ b & \pi_{n-2} = c_2, \\ a & \pi_{n-2} = c_3. \end{cases}$$

An induction over the play length shows that (α, τ_c) is winning for controller under delay 2.

Remark 2.4. Our definition of games under delayed control differs in three aspects from the original definition of Chen et al. [CFL⁺21].

- We allow arbitrary winning conditions while Chen et al. focused on safety conditions.
- The original definition allows nondeterministic strategies (a strategy that returns a nonempty set of actions, each one of which can be taken), while we restrict ourselves here to deterministic strategies (a strategy that returns a single action to be taken). The motivation for their use of nondeterministic strategies is the fact that they can be refined if additional constraints are imposed, which Chen et al.'s algorithm computing a winning strategy relies on.

Here, on the other hand, we are just interested in the existence of winning strategies. In this context, it is sufficient to consider deterministic strategies, as controller has a nondeterministic winning strategy if and only if she has a deterministic winning strategy. Also, strategies in delay games are deterministic, so the transformation between games under delayed control and delay games can be formulated more naturally for deterministic strategies.

- The original definition also allowed odd delays δ while we only allow even delays. As we will see in Section 3, the transformation of games under delayed control to delay games is naturally formulated for even delays. This choice also simplifies definitions, as accounting for odd delays imposes an additional notational burden.

2.2. Delay Games. Delay games are played between two players, Player I (she) and Player O (he). A delay game $\Gamma_k(L)$ (with constant lookahead) consists of a lookahead $k \in \mathbb{N}$ and a winning condition $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ for some alphabets Σ_I and Σ_O . Such a game is played in rounds $n = 0, 1, 2, \dots$ as follows: in round 0, first Player I picks a word $x_0 \in \Sigma_I^{k+1}$, then Player O picks a letter $y_0 \in \Sigma_O$. In round $n > 0$, Player I picks a letter $x_n \in \Sigma_I$, then Player O picks a letter $y_n \in \Sigma_O$. So, a play of $\Gamma_k(L)$ has the form $(x_0, y_0)(x_1, y_1)(x_2, y_2) \cdots$, where x_0 is a word of length $k + 1$ and all other x_n (those with $n > 0$) and all y_n are letters. Player O wins such a play $(x_0, y_0)(x_1, y_1)(x_2, y_2) \cdots$ if the outcome $\binom{x_0 x_1 x_2 \cdots}{y_0 y_1 y_2 \cdots}$ is in L ; otherwise, Player I wins.

A strategy for Player I in $\Gamma_k(L)$ is a mapping $\tau_I: \Sigma_O^* \rightarrow \Sigma_I^*$ satisfying $|\tau_I(\varepsilon)| = k + 1$ and $|\tau_I(w)| = 1$ for all $w \in \Sigma_O^+$. A strategy for Player O is a mapping $\tau_O: \Sigma_I^+ \rightarrow \Sigma_O$. A play $(x_0, y_0)(x_1, y_1)(x_2, y_2) \cdots$ is consistent with τ_I if $x_n = \tau_I(y_0 \cdots y_{n-1})$ for all $n \geq 0$, and it is consistent with τ_O if $y_n = \tau_O(x_0 \cdots x_n)$ for all $n \geq 0$. So, strategies are dual in delay games, i.e., Player I has to grant some lookahead on her moves that Player O has access to. A strategy for Player $P \in \{I, O\}$ is winning, if every play that is consistent with the strategy is won by Player P . We say that Player $P \in \{I, O\}$ wins a game $\Gamma_k(L)$ if Player P has a winning strategy in $\Gamma_k(L)$.

Remark 2.5.

- If Player O wins $\Gamma_k(L)$, then he also wins $\Gamma_{k'}(L)$ for every $k' > k$.
- If Player I wins $\Gamma_k(L)$, then she also wins $\Gamma_{k'}(L)$ for every $k' < k$.

Unlike games under delayed control, delay games with Borel winning conditions are determined [KZ16a], i.e., each delay game $\Gamma_k(L)$ with Borel L and fixed k is won by one of the players.

Example 2.6. Consider

$$L = \left\{ \binom{a_0}{b_0} \binom{a_1}{b_1} \binom{a_2}{b_2} \cdots \mid b_0 \notin \{a_0, a_1, a_2\} \right\}$$

over the alphabets $\Sigma_I = \Sigma_O = \{1, 2, 3, 4\}$.

Player I wins $\Gamma_k(L)$ for $k = 1$ with the following strategy τ_I : $\tau_I(\varepsilon) = 12$ and $\tau_I(b_0) = b_0$, and $\tau_I(w)$ arbitrary for all $w \in \Sigma_O^+$ with $|w| > 1$: In round 0, after Player I has picked $a_0 a_1 = 12$, Player O has to pick some b_0 . In order to not lose immediately, he has to pick $b_0 \notin \{1, 2\}$. Then, in round 1, Player I picks $a_2 = b_0$ and thereby ensures $b_0 \in \{a_0, a_1, a_2\}$. Hence, the play is not won by Player O (it's outcome is not in L), therefore it is winning for Player I .

However, Player O wins $\Gamma_k(L)$ for $k = 2$ with the following strategy τ_O : $\tau_O(a_0 a_1 a_2)$ is a letter in the nonempty set $\Sigma_O \setminus \{a_0, a_1, a_2\}$ and $\tau_O(w)$ arbitrary for all $w \in \Sigma_I^+$ with $|w| \neq 3$. In round 0, after Player I has picked $a_0 a_1 a_2$, Player O picks $b_0 \notin \{a_0, a_1, a_2\}$ and thus ensures that the outcome is in L .

Remark 2.7. We restrict ourselves here to the setting of constant lookahead, i.e., in a delay game $\Gamma_k(L)$ in round n when Player O picks her n -th letter, Player I has already picked $k + n + 1$ letters (note that we start in round 0 with the zeroth letter). Delay games have also been studied with respect to growing lookahead, i.e., the lookahead increases during a play [HKT12]. However, it is known that constant lookahead is sufficient for all ω -regular winning conditions: if Player O wins for any lookahead (no matter how fast it is growing), then she also wins with respect to constant lookahead, which can even be bounded exponentially in the size of a deterministic parity automaton recognizing the winning condition [KZ16a].

Stated differently, growing lookahead does not allow to win any more games than constant lookahead. Finally, the setting of constant lookahead in delay games considered here is the natural counterpart to games under delayed control, where the delay is fixed during a play.

2.3. Games under Delayed Control vs. Delay Games. Let us illustrate the differences between the strategic abilities of the players in a game \mathcal{G} under delayed control (say with delay δ) and the players in a delay game $\Gamma_k(L)$.

Let us start with the delay game $\Gamma_k(L)$. Here, Player I has to grant a lookahead of k moves to Player O , who can base her moves on this additional information. Formally, consider a play prefix $(x_0, y_0) \cdots (x_n, y_n)$, keeping in mind that x_0 is a word of length $k + 1$ and each $x_{n'}$ with $n' > 0$ and each $y_{n'}$ is a single letter. Hence, Player I has picked $x_0x_1 \cdots x_n$ of length $n + 1 + k$ while Player O has picked $y_0y_1 \cdots y_n$ of length $n + 1$. Now, Player I 's next move x_{n+1} may depend on $y_0y_1 \cdots y_n$ and Player O 's subsequent move depends on $x_0x_1 \cdots x_nx_{n+1}$. Thus, a delay game is a game of complete information, both players have full access to the moves made by their respective opponent.¹

On the other hand, in the game \mathcal{G} under delayed control, controller has to base her moves on a proper prefix of the play prefix constructed thus far while environment does not gain a symmetric advantage from controller's disadvantage. To understand this, consider a play prefix $\pi_0\sigma_0\pi_1 \cdots \pi_{n-1}\sigma_{n-1}\pi_n$ with (w.l.o.g.) $n > \delta$. If $\pi_n \in S_c$, i.e., it is controller's turn, then she has to base her decision on the prefix $\pi_0\sigma_0\pi_1 \cdots \pi_{n-1-\delta}\sigma_{n-1-\delta}\pi_{n-\delta}$, picking some action σ_n leading to a state $\pi_{n+1} \in S_e$. Thus, the actions $\sigma_{n+1-\delta}, \sigma_{n+3-\delta}, \dots, \sigma_{n-1}$ have already been picked by environment, but are not available to controller to base her decision on. Thus, she has incomplete information about the evolution of the play. Crucially, environment does not benefit from controller's disadvantage: Environment has to pick action σ_n only based on the actions $\sigma_0, \dots, \sigma_{n-1}$ while controller has to pick the action $\sigma_{n+\delta}$ based on the actions $\sigma_0, \dots, \sigma_{n-1}$.

We will study this difference in informedness of the players in a game under delayed control and in a delay game in the setting of deterministic and randomized strategies.

2.4. ω -Automata. A deterministic reachability automaton $\mathcal{A} = (Q, \Sigma, q_I, \delta_{\mathcal{A}}, F)$ consists of a finite set Q of states containing the initial state $q_I \in Q$ and the set of accepting states $F \subseteq Q$, an alphabet Σ , and a transition function $\delta_{\mathcal{A}}: Q \times \Sigma \rightarrow Q$. The size of \mathcal{A} is defined as $|\mathcal{A}| = |Q|$. Let $w = w_0w_1w_2 \cdots \in \Sigma^\omega$. The run of \mathcal{A} on w is the sequence $q_0q_1q_2 \cdots$ such that $q_0 = q_I$ and $q_{n+1} = \delta_{\mathcal{A}}(q_n, w_n)$ for all $n \geq 0$. A run $q_0q_1q_2 \cdots$ is (reachability) accepting if $q_n \in F$ for some $n \geq 0$. The language (reachability) recognized by \mathcal{A} , denoted by $L(\mathcal{A})$, is the set of infinite words over Σ such that the run of \mathcal{A} on w is (reachability) accepting.

A deterministic safety automaton has the form $\mathcal{A} = (Q, \Sigma, q_I, \delta_{\mathcal{A}}, U)$ where $Q, \Sigma, q_I, \delta_{\mathcal{A}}$ are as in a deterministic reachability automaton and where $U \subseteq Q$ is a set of unsafe states. The notions of size and runs are defined as for reachability automata, too. A run $q_0q_1q_2 \cdots$ is (safety) accepting if $q_n \notin U$ for all $n \geq 0$. The language (safety) recognized by \mathcal{A} , again denoted by $L(\mathcal{A})$, is the set of infinite words over Σ such that the run of \mathcal{A} on w is (safety) accepting.

Reachability and safety automata accept each a fragment of the ω -regular languages, while parity automata accept exactly the ω -regular languages [GTW02]. A deterministic parity

¹As usual, a player can always reconstruct their own previous moves, if necessary.

automaton has the form $\mathcal{A} = (Q, \Sigma, q_I, \delta_{\mathcal{A}}, \Omega)$ where $Q, \Sigma, q_I, \delta_{\mathcal{A}}$ are as in a deterministic reachability automaton and where $\Omega: Q \rightarrow \mathbb{N}$ is a coloring of the states. The notions of size and runs are defined as for reachability automata, too. A run $q_0q_1q_2 \cdots$ is (parity) accepting if the maximal color appearing infinitely often in the sequence $\Omega(q_0)\Omega(q_1)\Omega(q_2) \cdots$ is even. The language (parity) recognized by \mathcal{A} , again denoted by $L(\mathcal{A})$, is the set of infinite words over Σ such that the run of \mathcal{A} on w is (parity) accepting.

Reachability and safety automata are dual while parity automata are self-dual.

Remark 2.8.

- (1) Let $\mathcal{A} = (Q, \Sigma, q_I, \delta_{\mathcal{A}}, F)$ be a deterministic reachability automaton and let $\bar{\mathcal{A}}$ be the deterministic safety automaton $(Q, \Sigma, q_I, \delta_{\mathcal{A}}, F)$. Then, $L(\bar{\mathcal{A}}) = \overline{L(\mathcal{A})}$.
- (2) Let $\mathcal{A} = (Q, \Sigma, q_I, \delta_{\mathcal{A}}, F)$ be a deterministic safety automaton and let $\bar{\mathcal{A}}$ be the deterministic reachability automaton $(Q, \Sigma, q_I, \delta_{\mathcal{A}}, F)$. Then, $L(\bar{\mathcal{A}}) = \overline{L(\mathcal{A})}$.
- (3) Let $\mathcal{A} = (Q, \Sigma, q_I, \delta_{\mathcal{A}}, \Omega)$ be a deterministic parity automaton and let $\bar{\mathcal{A}}$ be the deterministic parity automaton $(Q, \Sigma, q_I, \delta_{\mathcal{A}}, q \mapsto \Omega(q) + 1)$. Then, $L(\bar{\mathcal{A}}) = \overline{L(\mathcal{A})}$.

3. FROM GAMES UNDER DELAYED CONTROL TO DELAY GAMES AND BACK

In this section, we exhibit a tight correspondence between controller in games under delayed control and Player I in delay games. Recall that in a game under delayed control, it is the controller whose control is delayed, i.e., she is at a disadvantage as she only gets delayed access to the action picked by environment. In a delay game, it is Player I who is at a disadvantage as she has to grant a lookahead on her moves to Player O . Thus, when simulating a game under delayed control by a delay game, it is natural to let Player I take the role of controller and let Player O take the role of environment. Also recall that the winning condition Win in a game under delayed control is formulated from controller's point-of-view: the winning condition requires her to enforce a play in Win . On the other hand, the winning condition L of a delay game is formulated from the point-of-view of Player O : Player O has to enforce a play whose outcome is in L . Thus, as Player I takes the role of controller, we need to complement the winning condition to reflect this change in perspective: The set of winning outcomes for Player I in the simulating delay game is the complement of Win .

In the remainder of this section, we show how to simulate a game under delayed control by a delay game and then the converse, i.e., we show how to simulate a delay game by a game under delayed control.

Transformation 3.1. *First, we transform a game under delayed control into a winning condition for a delay game, i.e., a language. In a delay game with this winning condition, the players simulate a play in the game under delayed control by picking actions, which uniquely induce such a play. To formalize this, we need to introduce some notation. Fix a game $\mathcal{G} = (S, s_0, S_c, S_e, \Sigma_c, \Sigma_e, \rightarrow, \text{Win})$. Note that a sequence $\sigma_0\sigma_1\sigma_2 \cdots \in (\Sigma_c\Sigma_e)^\omega$ induces a unique play $\text{play}(\sigma_0\sigma_1\sigma_2 \cdots) = \pi_0\sigma_0\pi_1\sigma_1\pi_2\sigma_2 \cdots$ in \mathcal{G} which is defined as follows: $\pi_0 = s_0$ and $\pi_{n+1} = \rightarrow(\pi_n, \sigma_n)$ for all $n \geq 0$. Likewise, a finite sequence $\sigma_0\sigma_1 \cdots \sigma_n \in (\Sigma_c\Sigma_e)^*(\Sigma_c + \varepsilon)$ induces a unique play prefix $\text{play}(\sigma_0\sigma_1 \cdots \sigma_n)$ which is defined analogously.*

Now, we define the language $L(\mathcal{G}) \subseteq (\Sigma_c \times \Sigma_e)^\omega$ such that $\binom{\sigma_0}{\sigma_1} \binom{\sigma_2}{\sigma_3} \binom{\sigma_4}{\sigma_5} \cdots \in L(\mathcal{G})$ if and only if $\text{play}(\sigma_0\sigma_1\sigma_2 \cdots)$ is winning for controller.

Example 3.2. We continue Example 2.3 to illustrate the transformation: consider the game under delayed control $\mathcal{G} = (S, s_0, S_c, S_e, \Sigma_c, \Sigma_e, \rightarrow, \text{Win})$ depicted in Figure 1 on Page 6 where




Figure 2: The safety automaton accepting the winning condition $L(\mathcal{G})$ obtained by applying Transformation 3.1 to the game from Example 2.3. A run is accepting if it does not visit the black state.

Win contains all plays that do not visit the black vertex. The language $L(\mathcal{G})$ is accepted by the safety automaton depicted in Figure 2, which accepts an infinite sequence of actions (arranged in tuples) if and only if the unique induced play in \mathcal{G} is winning for controller. In a delay game with winning condition $\overline{L(\mathcal{G})}$ (recall we need to complement the winning condition), Player I picks actions from Σ_c , Player O picks actions from Σ_e , and Player I wins if and only if the induced play in \mathcal{G} is in Win.

Now, we prove the correspondence between \mathcal{G} and $\Gamma_k(\overline{L(\mathcal{G})})$. The winning condition of the delay game is the complement of $L(\mathcal{G})$, which implements the switch of perspective described above.

Lemma 3.3. *Let \mathcal{G} be a game and $\delta \geq 0$ even. Controller wins \mathcal{G} under delay δ if and only if Player I wins $\Gamma_k(\overline{L(\mathcal{G})})$ for $k = \frac{\delta}{2}$.*

Proof. To simplify our notation in the following, we denote $\Gamma_k(\overline{L(\mathcal{G})})$ by Γ .

First, let (α, τ_c) be a winning strategy for controller under delay δ in \mathcal{G} . We inductively define a strategy τ_I for Player I in Γ .

For the induction start, we define $\tau_I(\varepsilon) = \alpha \cdot \tau_c(s_0)$, where s_0 is the initial vertex of \mathcal{G} . Note that $|\alpha| = k$, i.e., $\tau_I(\varepsilon)$ has the required length. For the induction step, we define $\tau_I(\sigma_1\sigma_3 \cdots \sigma_{2n+1}) = \tau_c(\text{play}(\sigma_0\sigma_1 \cdots \sigma_{2n+1}))$ for all $n \geq 0$, where $\sigma_0\sigma_2 \cdots \sigma_{\delta-2} = \alpha$ and $\sigma_{2i} = \tau_I(\sigma_1\sigma_3 \cdots \sigma_{2i-1})$ for all $2i$ in the range $\delta - 2 \leq 2i \leq 2n$.

We show that τ_I is a winning strategy for Player I in Γ . Consider a play in Γ that is consistent with τ_I , say with outcome $(\sigma_0\sigma_2\sigma_4 \cdots)$. We show that the play is winning for Player I by showing that its outcome is in $L(\mathcal{G})$, i.e., $\text{play}(\sigma_0\sigma_1\sigma_2 \cdots)$ is winning for controller in \mathcal{G} . To this end, consider the unique play π in \mathcal{G} under delay δ that is consistent with τ_c and where environment plays such that his moves spell $\sigma_1\sigma_3 \cdots$. By definition of τ_I , we have $\pi = \text{play}(\sigma_0\sigma_1\sigma_2 \cdots)$. Since τ_c is a winning strategy for controller, we indeed obtain that π is winning for controller in \mathcal{G} .

For the other direction, let τ_I be a winning strategy for Player I in Γ . We define a strategy (α, τ_c) for controller in \mathcal{G} under delay δ . Let $\tau_I(\varepsilon) = a_0 \cdots a_k$. We define $\alpha = a_0 \cdots a_{k-1}$ and $\tau_c(\pi_0) = a_k$. Note that $|\alpha| = \frac{\delta}{2}$, i.e., α has the required length. Furthermore, we define $\tau_c(\pi_0\sigma_0\pi_1 \cdots \sigma_{2n+1}\pi_{2n+2}) = \tau_I(\sigma_1\sigma_3 \cdots \sigma_{2n+1})$ for all play prefixes $\pi_0\sigma_0\pi_1 \cdots \sigma_{2n+1}\pi_{2n+2}$ of length $2n + 2$ for $n \geq 0$.

We show that (α, τ_c) is a winning strategy for controller in \mathcal{G} under delay δ . Consider a play $\pi = \pi_0\sigma_0\pi_1\sigma_1\pi_2\sigma_2 \cdots$ consistent with (α, τ_c) . We show that π is winning for controller in \mathcal{G} . To this end, consider the unique play in Γ that is consistent with τ_I and where Player O plays such that his moves spell $\sigma_1\sigma_3 \cdots$. By definition of τ_c , the outcome of this play is $\binom{\sigma_0\sigma_2\sigma_4 \cdots}{\sigma_1\sigma_3\sigma_5 \cdots}$. Since τ_I is a winning strategy for Player I in Γ , we obtain $\binom{\sigma_0\sigma_2\sigma_4 \cdots}{\sigma_1\sigma_3\sigma_5 \cdots} \notin \overline{L(\mathcal{G})}$, i.e., $\binom{\sigma_0\sigma_2\sigma_4 \cdots}{\sigma_1\sigma_3\sigma_5 \cdots} \in L(\mathcal{G})$. Hence, $\text{play}(\sigma_0\sigma_1\sigma_2 \cdots) = \pi$ is indeed winning for controller in \mathcal{G} . \square

Now, we consider the converse and transform a delay game into a game under delayed control.

Transformation 3.4. Fix a delay game $\Gamma_k(L)$ with $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$. We construct a game under delayed control to simulate $\Gamma_k(L)$ as follows: The actions of controller are the letters in Σ_I , and the actions of environment are the letters in Σ_O . Thus, by picking actions, controller and environment construct the outcome of a play of $\Gamma_k(L)$. As winning conditions of games under delayed control only refer to states visited by a play, but not the actions picked by the players, we reflect the action picked by a player in the state reached by picking that action. Here, we have to require without loss of generality that Σ_I and Σ_O are disjoint.

Formally, we define $\mathcal{G}(L) = (S, s_I, S_c, S_e, \Sigma_c, \Sigma_e, \rightarrow, \text{Win})$ with $S = S_c \cup S_e$, $S_c = \{s_I\} \cup \Sigma_O$, $S_e = \Sigma_I$, $\Sigma_c = \Sigma_I$, $\Sigma_e = \Sigma_O$, $\rightarrow(s, a) = a$ for all $s \in S_c$ and $a \in \Sigma_I$, and $\rightarrow(s, b) = b$ for all $s \in S_e$ and $b \in \Sigma_O$. Finally, we define $\text{Win} = \{s_I s_0 s_1 s_2 \cdots \mid \binom{s_0}{s_1} \binom{s_2}{s_3} \binom{s_4}{s_5} \cdots \in L\}$.

Example 3.5. We continue Example 2.6 to illustrate the second transformation: consider the delay game $\Gamma_k(L)$ with

$$L = \left\{ \binom{a_0}{b_0} \binom{a_1}{b_1} \binom{a_2}{b_2} \cdots \mid b_0 \notin \{a_0, a_1, a_2\} \right\}$$

over the alphabets $\Sigma_I = \Sigma_O = \{1, 2, 3, 4\}$. Note that this game does not satisfy our assumption that Σ_I and Σ_O are disjoint. Hence, in the following we decorate all letters in Σ_O by a prime and consider the winning condition

$$L = \left\{ \binom{a_0}{b_0} \binom{a_1}{b_1} \binom{a_2}{b_2} \cdots \mid b_0 \notin \{a'_0, a'_1, a'_2\} \right\}$$

over the alphabets $\Sigma_I = \{1, 2, 3, 4\}$ and $\Sigma_O = \{1', 2', 3', 4'\}$.

Now, applying Transformation 3.4 yields the game under delayed control depicted in Figure 3. In this game, controller picks letters from Σ_I (by moving to the corresponding state) and environment picks letters from Σ_O (again by moving to the corresponding state). Thus, controller and environment construct the outcome of a play of $\Gamma_k(L)$. Now, the winning condition of $\mathcal{G}(L)$ is defined such that controller wins $\mathcal{G}(\overline{L})$ (recall that we need to complement the winning condition) if and only if the resulting outcome is in L .

The following remark simplifies the proof of correctness of the second transformation. It follows by a careful inspection of the definitions.

Remark 3.6. Let $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$. Then, $L = L(\mathcal{G}(L))$.

Now, we show that the second transformation is correct, again using complementation to implement the perspective switch.

Lemma 3.7. Let $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ and $k \geq 0$. Player I wins $\Gamma_k(L)$ if and only if controller wins $\mathcal{G}(\overline{L})$ under delay $2k$.




Figure 3: The game under delayed control obtained by applying Transformation 3.4 to the delay game from Example 2.6. All incoming edges to vertex j are labeled by action j , all incoming edges to vertex j' are labeled by action j' .

Proof. Consider the delay game $\Gamma_k(\overline{L(\mathcal{G}(L))})$. By Lemma 3.3, controller wins $\mathcal{G}(\overline{L})$ under delay $2k$ if and only if Player I wins $\Gamma_k(L(\mathcal{G}(\overline{L})))$. Applying Remark 3.6 yields $\overline{L} = L(\mathcal{G}(\overline{L}))$ and thus $L = \overline{L(\mathcal{G}(\overline{L}))}$. Hence, as $\Gamma_k(L)$ and $\Gamma_k(L(\mathcal{G}(\overline{L})))$ have the same winning condition, Player I wins the former if and only if she wins the latter. Thus, the statement of the lemma directly follows. \square

4. RESULTS

Lemma 3.3 and Lemma 3.7 allow us to transfer results from delay games to games under delayed control. Due to the definitions of strategies in games under delayed control not being dual, we first consider controller and then environment.

Recall that delay that allows controller to win satisfies a monotonicity property (see Remark 2.1(2)): if controller wins a game under delay δ , then also under every delay $\delta' < \delta$. Thus, the set of delays for which controller wins is downward-closed, i.e., it is either a finite set $\{0, 2, 4, \dots, \delta_{\max}\}$ or it is equal to the set $2\mathbb{N}$ of even numbers. In the following, we study the complexity of determining whether controller wins under all possible delays, whether she wins under a given delay, and determine bounds on δ_{\max} .

Note that winning for environment is independent of delay and boils down to the classical notion of winning delay-free games [GTW02], which is a well-studied problem. Hence, we disregard this problem. However, we do discuss the relation between environment in a game under delayed control and Player O in the simulating delay game constructed in the previous section.

Before we present our results, we need to specify how to measure the size of games and delay games, especially how winning conditions are represented (recall that, so far, they are just ω -languages). In the following, we only consider ω -regular winning conditions specified by ω -automata (see Subsection 2.4) or formulas of Linear Temporal Logic (LTL) [Pnu77], which subsume the typical specification languages for winning conditions. Hence, the size of

a game $(S, s_0, S_c, S_e, \Sigma_c, \Sigma_e, \rightarrow, \text{Win})$ under delayed control is given by the sum $|S| + |\Sigma_c| + |\Sigma_e| + |\text{Win}|$, where $|\text{Win}|$ is the size of an automaton or LTL formula (measured in the number of distinct subformulas) representing Win , depending on whether Win is specified by an automaton or a formula. Analogously, for a delay game $\Gamma_k(L)$, we define the size of L as the size of an automaton or LTL formula (measured in the number of distinct subformulas) representing L . The bound k is encoded in binary, if necessary.

4.1. Safety. A game $\mathcal{G} = (S, s_0, S_c, S_e, \Sigma_c, \Sigma_e, \rightarrow, \text{Win})$ with winning condition Win is a safety game if Win is accepted by a deterministic safety automaton.

Remark 4.1. When Chen et al. introduced safety games under delayed control, they did not use automata to specify their winning plays, but instead equipped the game with a set of unsafe states and declared all those plays winning for controller that never visit an unsafe state. It is straightforward to see that our definition is equivalent, as their definition is captured by a deterministic safety automaton with two states. Conversely, taking the product of a game and a deterministic safety automaton yields an equivalent game with a state-based safety condition.

Our results rely on the following two bounds on the transformations presented in Section 3, which are obtained by applying Remark 2.8:

- (1) If the winning condition Win for a game \mathcal{G} under delayed control is given by a deterministic safety automaton with n states, then the winning condition $\overline{L(\mathcal{G})}$ is recognized by a deterministic reachability automaton with n states.
- (2) Dually, if the winning condition $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ of a delay game is given by a deterministic reachability automaton with n states, then the winning condition of the game $\mathcal{G}(\overline{L})$ under delayed control is recognized by a deterministic safety automaton with $\mathcal{O}(n \cdot |\Sigma_I|)$ states. Here, the factor $|\Sigma_I|$ stems from the fact that an automaton \mathcal{A} for L processes pairs $(a, b) \in \Sigma_I \times \Sigma_O$ of letters while an automaton for the winning condition of $\mathcal{G}(\overline{L})$ needs to simulate \mathcal{A} while first receiving a and then b , i.e., it needs to store $a \in \Sigma_I$ in its state.

We begin by settling the complexity of determining whether controller wins a given safety game under every delay, which follows from the PSPACE-completeness of determining whether there is a lookahead that allows Player O to win a given delay game with reachability winning condition [KZ16a].

Theorem 4.2. *The following problem is PSPACE-complete: Given a safety game \mathcal{G} , does controller win \mathcal{G} under every delay δ ?*

Proof. The following problem is PSPACE-complete [KZ16a]: Given a deterministic reachability automaton \mathcal{A} , does Player O win $\Gamma_k(L(\mathcal{A}))$ for some k ? Hence, due to determinacy of delay games with ω -regular winning conditions [KZ15] and due to closure of PSPACE under complementation, the following problem is also PSPACE-complete: Given a deterministic reachability automaton \mathcal{A} , does Player I win $\Gamma_k(L(\mathcal{A}))$ for all k ? Thus, the PSPACE lower bound follows from Lemma 3.7 and the corresponding upper bound from Lemma 3.3. \square

Next, we give a lower bound on the complexity of determining whether controller wins a given safety game under a given delay, which is derived from a lower bound for delay games with reachability winning conditions.

Theorem 4.3. *The following problem is PSPACE-hard: Given a safety game \mathcal{G} and δ (encoded in binary), does controller win \mathcal{G} under delay δ .*

Proof. The following problem is PSPACE-hard [KZ16a]²: Given a deterministic reachability automaton \mathcal{A} , does Player O win $\Gamma_k(L(\mathcal{A}))$ for $k = 2^{|\mathcal{A}|}$? Hence, due to determinacy of delay games with ω -regular winning conditions [KZ15] and due to closure of PSPACE under complementation, the following problem is also PSPACE-complete: Given a deterministic reachability automaton \mathcal{A} , does Player I win $\Gamma_k(L(\mathcal{A}))$ for $k = 2^{|\mathcal{A}|}$? Thus, the PSPACE lower bound follows from Lemma 3.7 using $\delta = 2k$. \square

Note that we do not claim any upper bound on the problem considered in Theorem 4.3. There is a trivial 2EXPTIME upper bound obtained by hardcoding the delay into the graph of the safety game, thereby obtaining a classical delay-free safety game. It is open whether the complexity can be improved. Let us remark though that, via the correspondence to delay games presented in Section 3, improvements here would also yield improvements on the analogous problem for delay games, which is open too [Zim22].

Next, we turn our attention to bounds on the delay for which controller wins. Recall that due to monotonicity, the set of delays for which controller wins is downward-closed, i.e., it is either a finite set $\{0, 2, 4, \dots, \delta_{\max}\}$ or it is equal to $2\mathbb{N}$. In the following, we present tight bounds on the value δ_{\max} .

As a consequence, we settle a conjecture by Chen et al.: They conjectured that there is some delay δ_t (exponential in $|\mathcal{G}|$), such that if controller wins \mathcal{G} under delay δ_t , then she wins under every delay. Note that this conjecture implies that δ_{\max} is at most exponential.

The following theorem proves Chen et al.’s conjecture, while Theorem 4.5 shows that δ_t must necessarily be exponential. For δ_{\max} this means it is at most exponential for every game, and can be exponential for some games.

The following two results are again obtained from similar bounds for delay games with reachability winning conditions.

Theorem 4.4. *Let \mathcal{G} be a safety game. There is a $\delta_t \in \mathcal{O}(2^{|\mathcal{G}|})$ such that if controller wins \mathcal{G} under delay δ_t , then she wins \mathcal{G} under every δ .*

Proof. Exponential lookahead is sufficient for Player O to win delay games with reachability conditions: Let \mathcal{A} be a deterministic reachability automaton with n states recognizing a language L . If Player O wins $\Gamma_k(L)$ for some k , then he also wins $\Gamma_{k'}(L)$ for $k' = 2^n$ [KZ16a]. So, dually, if Player O does not win $\Gamma_{k'}(L)$ for $k' = 2^n$, then he does not win $\Gamma_k(L)$ for any k . Due to determinacy [KZ15], this is equivalent to the following statement: if Player I wins $\Gamma_{k'}(L)$ for $k' = 2^n$, then she wins $\Gamma_k(L)$ for every k . We transfer this result to safety games under delayed control.

Given \mathcal{G} , let $\delta_t = 2 \cdot 2^{|\mathcal{G}|}$. Now, assume controller wins \mathcal{G} under delay δ_t . Then, Player I wins $\Gamma_k(\overline{L(\mathcal{G})})$ for $k = \frac{\delta_t}{2} = 2^{|\mathcal{G}|}$ due to Lemma 3.3. As argued above, $\overline{L(\mathcal{G})}$ is recognized by a deterministic reachability automaton of size at most $|\mathcal{G}|$. Thus, Player I wins $\Gamma_k(\overline{L(\mathcal{G})})$ for every k . Hence, due to Lemma 3.7, controller wins \mathcal{G} under every δ . \square

Finally, we show that the exponential upper bound on δ_{\max} is tight.

²The result is not stated as such, but follows by inspecting the proof of Theorem 4.1 in [KZ16a], which shows that Player O wins a delay game with reachability condition with respect to any lookahead if and only if he wins it with respect to constant lookahead $2^{|\mathcal{A}|}$.

Theorem 4.5. *For every $n > 1$, there is a safety game \mathcal{G}_n of size $\mathcal{O}(n)$ such that controller wins \mathcal{G} under delay 2^n , but not under delay $2^n + 2$.*

Proof. For every $n > 1$, there is a deterministic reachability automaton \mathcal{A}_n of size $\mathcal{O}(n)$ and where Σ_I is of size n such that Player O wins $\Gamma_k(L(\mathcal{A}_n))$ for $k = 2^n$, but Player I wins $\Gamma_{k'}(L(\mathcal{A}_n))$ for every $k' < 2^n$ [KZ16a]. Now, we consider the safety games $\mathcal{G}(L(\mathcal{A}_n))$, which have size $\mathcal{O}(n)$ as argued above. Applying Lemma 3.7 shows that controller wins $\mathcal{G}(L(\mathcal{A}_n))$ under delay $2 \cdot 2^n$. Now, towards a contradiction, assume controller wins $\mathcal{G}(L(\mathcal{A}_n))$ under delay $2^n + 2$. Then, Lemma 3.3 yields that Player I wins $\Gamma_k(L(\mathcal{G}(L(\mathcal{A}_n)))) = \Gamma_k(L(\mathcal{A}_n))$ (see Remark 3.6) for $k = 2^n$, yielding the desired contradiction. \square

4.2. Reachability. In this subsection, we consider the case of winning conditions given by reachability automata. A game $\mathcal{G} = (S, s_0, S_c, S_e, \Sigma_c, \Sigma_e, \rightarrow, \text{Win})$ with winning condition Win is a reachability game if Win is accepted by a deterministic reachability automaton.

Applying Remark 2.8 yields the following two bounds on the transformations from Section 3:

- (1) If the winning condition Win for a game \mathcal{G} under delayed control is given by a deterministic reachability automaton with n states, then the winning condition $\overline{L(\mathcal{G})}$ is recognized by a deterministic safety automaton with n states.
- (2) Dually, if the winning condition $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ of a delay game is given by a deterministic safety automaton with n states, then the winning condition of the game $\mathcal{G}(\overline{L})$ under delayed control is recognized by a deterministic reachability automaton with $\mathcal{O}(n \cdot |\Sigma_I|)$ states. Here, the additional factor $|\Sigma_I|$ is again required to store an input letter, as in the case of safety automata (see Page 14).

Exponential lookahead is both sufficient to win all delay games with safety conditions that can be won and required to win some of these games [KZ16a]. Furthermore, determining whether there is some lookahead that allows Player O to win a given delay game with safety condition is EXPTIME-complete [KZ16a]. As in the case of safety games, we can transfer these results to games under delayed control with ω -regular winning conditions.

Theorem 4.6.

- (1) *The following problem is EXPTIME-complete: Given a reachability game \mathcal{G} , does controller win \mathcal{G} under every delay δ ?*
- (2) *Let \mathcal{G} be a reachability game with winning condition specified by a deterministic reachability automaton with n states. There is a $\delta_t \in \mathcal{O}(2^{n^2})$ such that if controller wins \mathcal{G} under delay δ_t , then she wins \mathcal{G} under every δ .*
- (3) *For every $n > 1$, there is a reachability game \mathcal{G}_n of size $\mathcal{O}(n^2)$ with a winning condition specified by a two-state deterministic reachability automaton \mathcal{A}_n such that controller wins \mathcal{G} under delay 2^n , but not under delay $2^n + 2$.*

4.3. Parity. Next, we consider the case of ω -regular winning conditions, given by deterministic parity automata. Applying Remark 2.8 yields the following two bounds on the transformations from Section 3:

- (1) If the winning condition Win for a game \mathcal{G} under delayed control is given by a deterministic parity automaton with n states, then the winning condition $\overline{L(\mathcal{G})}$ is recognized by a deterministic parity automaton with n states.

- (2) Dually, if the winning condition $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ of a delay game is given by a deterministic parity automaton with n states, then the winning condition of the game $\mathcal{G}(\bar{L})$ under delayed control is recognized by a deterministic parity automaton with $\mathcal{O}(n \cdot |\Sigma_I|)$ states. Here, the additional factor $|\Sigma_I|$ is again required to store an input letter, as in the case of safety automata (see Page 14)

Exponential lookahead is both sufficient to win all ω -regular delay games that can be won and required to win some of these games [KZ16a]. Furthermore, determining whether there is some lookahead that allows Player O to win a given ω -regular delay game is EXPTIME-complete [KZ16a]. As in the case of safety games, we can transfer these results to games under delayed control with ω -regular winning conditions.

Theorem 4.7.

- (1) *The following problem is EXPTIME-complete: Given a game \mathcal{G} with ω -regular winning condition specified by a deterministic parity automaton, does controller win \mathcal{G} under every delay δ ?*
- (2) *Let \mathcal{G} be a game with ω -regular winning condition specified by a deterministic parity automaton with n states. There is a $\delta_t \in \mathcal{O}(2^{n^3})$ such that if controller wins \mathcal{G} under delay δ_t , then she wins \mathcal{G} under every δ .*
- (3) *For every $n > 1$, there is a game \mathcal{G}_n of size $\mathcal{O}(n^2)$ with ω -regular winning condition specified by a two-state deterministic parity automaton \mathcal{A}_n such that controller wins \mathcal{G} under delay 2^n , but not under delay $2^n + 2$.*

Note that the lower bound on δ_t is just a restatement of Theorem 4.5, as safety games have ω -regular winning conditions.

4.4. Linear Temporal Logic. Finally, one can also transfer the triply-exponential upper and lower bounds on the necessary lookahead in delay games with LTL winning conditions as well as the 3EXPTIME-completeness of determining whether Player O wins such a delay game with respect to some lookahead [KZ16b] to games under delayed control with LTL winning conditions. Here, we exploit the following facts:

- (1) If the winning condition Win for a game \mathcal{G} under delayed control is given by an LTL formula φ , then the winning condition $\bar{L}(\mathcal{G})$ is given by an LTL formula of size $\mathcal{O}(|\varphi|)$.
- (2) Dually, if the winning condition $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ of a delay game is given by an LTL formula φ , then the winning condition of the game $\mathcal{G}(\bar{L})$ under given action is given by an LTL formula of size $\mathcal{O}(|\varphi|)$.

Theorem 4.8.

- (1) *The following problem is 3EXPTIME-complete: Given a game \mathcal{G} with winning condition specified by an LTL formula φ , does controller win \mathcal{G} under every delay δ ?*
- (2) *Let \mathcal{G} be a game with ω -regular winning condition specified by an LTL formula φ . There is a $\delta_t \in \mathcal{O}(2^{2^{|\varphi|+|\mathcal{G}|}})$ such that if controller wins \mathcal{G} under delay δ_t , then she wins \mathcal{G} under every δ .*
- (3) *For every $n > 1$, there is a game \mathcal{G}_n of size $\mathcal{O}(n^2)$ with winning condition specified by an LTL formula φ_n of size $\mathcal{O}(n^2)$ such that controller wins \mathcal{G} under delay 2^{2^n} , but not under delay $2^{2^n} + 2$.*

4.5. Environment's View. In Section 3, we proved a tight correspondence between controller in a game under delayed control and Player I in a delay game. Thus, it is natural to ask whether environment and Player O also share such a tight correspondence. A first indication that this is not the case can be obtained by considering the determinacy of these games: While delay games with Borel winning conditions are determined [KZ15], even safety games under delayed control are not necessarily determined [CFL⁺21].

Upon closer inspection, the lack of a correspondence between environment and Player O is not surprising, as the strategies in games under delayed control are not dual between the players: controller is at a disadvantage as she only gets delayed access to the actions picked by environment while environment does not benefit from this disadvantage. He does not get access to the actions picked by controller in advance. In a delay game however, the strategy definitions are completely dual: Player I has to grant lookahead on her moves which Player O gets access to. Thus, environment is in a weaker position than Player O .³

In this section, we study the correspondence between environment and Player O in detail by formally proving that environment is weaker than Player O .⁴

Lemma 4.9. *Let \mathcal{G} be a game. If environment wins \mathcal{G} then Player O wins $\Gamma_k(\overline{L(\mathcal{G})})$ for every k .*

Proof. Fix some $k \geq 0$. To simplify our notation in the following, we denote $\Gamma_k(\overline{L(\mathcal{G})})$ by Γ .

Let $\tau_e: \text{Pref}_e(\mathcal{G}) \rightarrow \Sigma_c$ be a winning strategy for environment in \mathcal{G} . We inductively define a strategy $\tau_O: \Sigma_I^+ \rightarrow \Sigma_O$ for Player O in Γ , which is a straightforward simulation of τ_e ignoring the granted lookahead of k additional letters.

For the induction start, we define $\tau_O(\sigma_0\sigma_2 \cdots \sigma_{2k}) = \tau_e(\text{play}(\sigma_0))$. For the induction step, we define $\tau_O(\sigma_0\sigma_2 \cdots \sigma_{2n}) = \tau_e(\text{play}(\sigma_0\sigma_1 \cdots \sigma_{2(n-k)}))$ for all $n \geq k$, where $\sigma_{2i+1} = \tau_O(\sigma_0\sigma_2 \cdots \sigma_{2(i+k)})$ for all $2i$ in the range $0 \leq 2i \leq 2(n-k) - 2$.

We show that τ_O is a winning strategy for Player O in Γ . Consider a play in Γ that is consistent with τ_O , say with outcome $(\begin{smallmatrix} \sigma_0\sigma_2\sigma_4\cdots \\ \sigma_1\sigma_3\sigma_5\cdots \end{smallmatrix})$. We show that the play is winning for Player O by showing that its outcome is not in $L(\mathcal{G})$, i.e., $\text{play}(\sigma_0\sigma_1\sigma_2 \cdots)$ is winning for environment in \mathcal{G} . To this end, consider the unique play π in \mathcal{G} that is consistent with τ_e and where controller plays such that her moves spell $\sigma_0\sigma_2 \cdots$. By definition of τ_O , we have $\pi = \text{play}(\sigma_0\sigma_1\sigma_2 \cdots)$. Since τ_e is a winning strategy for environment, we indeed obtain that π is winning for environment in \mathcal{G} . \square

Now, we show that the converse direction of Lemma 4.9 fails.

Lemma 4.10. *There is a safety game \mathcal{G} such that Player O wins $\Gamma_k(\overline{L(\mathcal{G})})$ for some k , but environment does not win \mathcal{G} under any delay.*

Proof. Let \mathcal{G} be the safety game depicted in Figure 4. With each move, the players place a coin (by either picking heads or tails) and the environment wins a play by correctly predicting the second action of controller with his first action. Note that environment has only one nontrivial choice to make, i.e., the choice of head or tails for the first action picked by him.

³The difference can be formalized in terms of the information the players have access to: games under delayed control are incomplete-information games while delay games are complete-information games. Although interesting, we do not pursue this angle any further.

⁴Note that the following lemma can easily be proven using determinacy of delay games and the correspondence between controller and Player I exhibited by Lemma 3.3 and Lemma 3.7, if the winning condition is Borel. However, we state (and prove) the lemma for arbitrary winning conditions.




Figure 4: A safety game that environment does not win under any delay, but Player O wins the associated delay game with $k \geq 1$. The initial state is marked by an arrow and the unsafe vertices are black. Note that both players have the actions h and t available.

Clearly, environment has no winning strategy in \mathcal{G} (under any delay) because he has no access to future moves of controller. Stated differently, if environment picks h (t) in his first move, then the play in which the second action of controller is t (h) is winning for controller.⁵

Now, we consider the delay game $\Gamma_k(\overline{L(\mathcal{G})})$ for $k = 1$. Recall that the winning condition $\overline{L(\mathcal{G})}$ contains the winning plays for Player O , i.e., we have $(\sigma_0\sigma_2\sigma_4\cdots) \in \overline{L(\mathcal{G})}$ if and only if $\sigma_1 \neq \sigma_2$. It is easy to see that Player O has a winning strategy in $\Gamma_k(\overline{L(\mathcal{G})})$ by simply flipping the second letter picked by Player I . This is possible since Player I has to provide two letters during the first round. \square

Remark 4.11. The safety game \mathcal{G} depicted in Figure 4 is in fact undetermined under every delay $\delta > 0$. In the proof of Lemma 4.10, we have already established that environment does not win \mathcal{G} . Now, under every delay $\delta > 0$, controller has to fix at least two actions before getting access to the first action picked by environment. This implies that there is, for every strategy for controller under delay δ , at least one consistent play that is losing for her, i.e., a play in which environment picks h (t) if the second move fixed by controller is t (h). Thus, no strategy is winning for controller under delay δ .

Let us remark that, according to our definition of environment strategies, he is not able to enforce a losing play for controller (the game is undetermined after all), as he does not get access to the second action fixed by controller. Also, this is again the difference to delay games: Player O has access to these first two actions when making his first move, and is thereby able to win.

The full relation between games under delayed control and delay games is depicted in Figure 5, restricted to Borel winning conditions (note that both transformations described in Section 3 preserve Borelness). The equivalence between controller winning the game under delayed control and Player I winning the corresponding delay game has been shown in Lemma 3.3 and Lemma 3.7. Also, Lemma 3.7 and Remark 3.6 imply that undetermined safety games under delayed control and those won by environment get transformed into delay games that are won by Player O . Lemma 4.9 shows that games under delayed control won by environment are transformed into delay games won by Player O while chaining Lemma 3.3 and Lemma 3.7 shows that undetermined games under delayed control are also transformed

⁵Note that under any delay $\delta > 0$, controller cannot do this strategically, as she has to fix her first two actions in advance. But as environment has no access to these fixed actions, he cannot react to them strategically. Also, see Remark 4.11.




Figure 5: The relation between games under delayed control and delay games with Borel winning conditions. The upper ellipsis contains pairs (\mathcal{G}, δ) consisting of a game \mathcal{G} under delayed control and a fixed delay δ ; the lower one contains delay games $\Gamma_k(L)$ for some fixed k . The arrows represent the two transformations described in Section 3.

into delay games won by Player O . Finally, Lemma 3.3 and Remark 3.6 imply that delay games won by Player O get transformed into undetermined safety games under delayed control or to ones that are won by environment (and it straightforward to construct delay games that realize both cases).

5. REFINING THE CORRESPONDENCE: SURE WINNING AND ALMOST SURE WINNING

It should be noted that the above transformations of games under delayed control into delay games and vice versa hinge on the fact that environment in the game under delayed control could, though lacking recent state information to do so strategically, by mere chance play the very same actions that the informed Player O in the delay game plays in his optimal adversarial strategy. That this constitutes a fundamental difference becomes apparent if we consider almost sure winning instead of sure winning. Almost sure winning calls for the existence of a mixed strategy that wins with probability 1, i.e., may fail on a set of plays with measure 0. This is different from sure winning in the sense of the definition of winning strategies for games under delayed control in Subsection 2.1, which calls for a strategy that never fails.

Remark 5.1. We introduce mixed strategies for games under delayed control only, as delay games (with Borel winning conditions) are determined, which means that mixed strategies do not offer any advantage over pure strategies as introduced in Subsection 2.2.

Given an even $\delta \geq 0$, a mixed strategy for controller in \mathcal{G} under delay δ is a pair (α, τ_c) where $\alpha \in \mathcal{P}\left((\Sigma_c)^{\frac{\delta}{2}}\right)$ is a probability distribution over $(\Sigma_c)^{\frac{\delta}{2}}$ and $\tau_c: \text{Pref}_c(\mathcal{G}) \rightarrow \mathcal{P}(\Sigma_c)$ maps play prefixes ending in S_c to probability distributions over actions of controller. A mixed strategy for environment is a mapping $\tau_e: \text{Pref}_e(\mathcal{G}) \rightarrow \mathcal{P}(\Sigma_e)$.

The notion of consistency of a play with a strategy simply carries over, now inducing a Markov chain due to the probabilistic nature of the strategies. We say that a mixed strategy for controller (environment) *wins almost surely* if and only if it wins against any strategy of its opponent environment (controller) with probability 1, i.e., if and only if the winning condition is satisfied with probability 1 over the Markov chain induced by the game and




Figure 6: A reachability game that, under any positive delay, is won by controller almost surely via the simple randomized strategy of coin tossing (thus randomly generating head and tail events h and t), but won by player O surely if interpreted as a delay game due to the lookahead on Player I 's actions granted to Player O . The initial state is marked by an arrow and controller wins if and only if the black vertex is visited at least once.

the particular strategy combination. In this section, we write sure winning for winning as defined in Section 2, as is usual for games with randomized strategies.

The notion of almost sure winning alters chances for the players substantially by excluding the possibility of reliably playing an optimal strategy though lacking the information for doing so due to delayed observation. This can be seen from the following lemma, stating a fundamental difference between controller's power in games under delayed control and Player I 's power in the corresponding delay games.

Lemma 5.2. *There is a game \mathcal{G} under delayed control and some (even) delay δ such that controller wins \mathcal{G} under delay δ almost surely while Player O (not Player I , which is the player corresponding to controller) wins the corresponding delay game $\Gamma_k(\overline{L(\mathcal{G})})$ for $k = \frac{\delta}{2}$, and surely so.*

Proof. Consider the reachability game in Figure 6 under delay 2 (or any larger delay). Intuitively, the players place a coin in each round (by picking either heads to tails with each move) and controller wins a play if the black state is visited, which happens if she selects a different coin placement than chosen by environment in the previous move.

Under any even (by definition) positive delay, controller wins this game with probability 1, i.e., almost surely, by a simple randomized strategy of coin tossing: in each step randomly selecting action h or t with positive probability each, an eventual visit of the black state is guaranteed with probability 1, irrespective of being uninformed about environment's preceding move due to the delay.

The corresponding delay game $\Gamma_k(\overline{L(\mathcal{G})})$ for $k = \frac{\delta}{2}$, however, is easily won by Player O , because in delay games, the delayed Player I grants a lookahead to Player O . Hence, Player O can, due to the delay, already see the next move of Player I such that he can simply copy the next coin placement by Player I , safely staying in the non-black states and thereby win. \square

Note that Lemma 5.2 implies that the previously observed correspondence between Player I and controller breaks down when considering almost sure winning strategies instead of just sure winning strategies: Games under delayed control for which Player O wins the corresponding delay game, are no longer either undetermined or won by environment, but may well be won by controller almost surely.

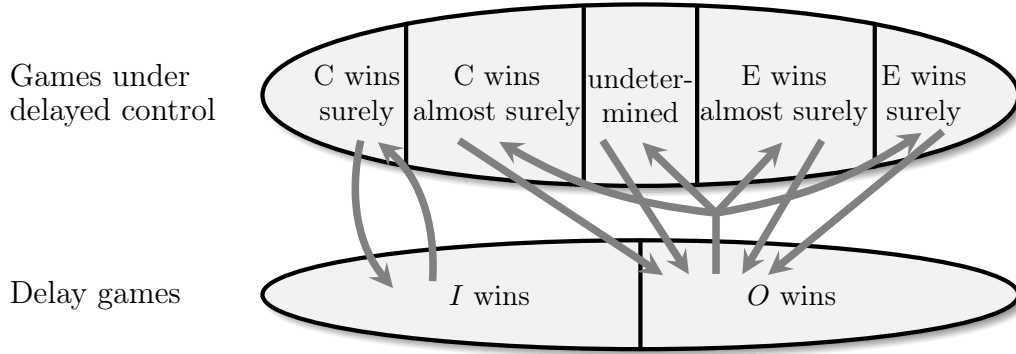


Figure 7: The relation between safety games under delayed control and delay games with Borel winning conditions. The upper ellipsis contains pairs (\mathcal{G}, δ) consisting of a game \mathcal{G} under delayed control and a fixed delay δ ; the lower one contains delay games $\Gamma_k(L)$ for some fixed k . The arrows represent the two transformations described in Section 3.

This consequently refines the correspondence between games under delayed control and delay games shown in Figure 5 as follows.

Theorem 5.3. *Given a game \mathcal{G} with Borel winning condition and an even $\delta \geq 2$, the following correspondences between \mathcal{G} and the corresponding delay game $\Gamma_k(\overline{L(\mathcal{G})})$ for $k = \frac{\delta}{2}$ hold:*

- (1) *Controller surely wins \mathcal{G} under delay δ if and only if Player I surely wins $\Gamma_k(\overline{L(\mathcal{G})})$.*
- (2) *If controller almost surely wins \mathcal{G} under delay δ but cannot surely win \mathcal{G} under delay δ then Player O surely wins $\Gamma_k(\overline{L(\mathcal{G})})$.*
- (3) *If environment surely or almost surely wins \mathcal{G} under delay δ then Player O wins $\Gamma_k(\overline{L(\mathcal{G})})$.*
- (4) *If \mathcal{G} is undetermined under delay δ with respect to almost sure winning strategies then Player O wins $\Gamma_k(\overline{L(\mathcal{G})})$.*
- (5) *All the aforementioned classes are non-empty, i.e., there exist games under delayed control where controller wins, where controller wins almost surely (but not surely), where environment wins surely, where environment wins almost surely (but not surely), and games which are undetermined with respect to almost-sure winning strategies.*

The above correspondences are depicted in Figure 7.

Proof. Equivalence (1) has been shown in Section 4. Implications (2)–(4) follow immediately from (1) and determinacy of the delay game $\Gamma_k(\overline{L(\mathcal{G})})$: whenever controller fails to have a sure winning strategy then Player I does not have one either, and Player O consequently wins surely due to determinacy of $\Gamma_k(\overline{L(\mathcal{G})})$.

For (5), one observes that games where either controller or environment surely win exist trivially, as universal as well as empty winning conditions are permitted. Existence of a game under delayed control that can be won by controller almost surely, but not surely, is shown by the example in Figure 6. Complementation of its winning condition, i.e., declaring environment the winner if and only if the black vertex is visited, converts this game into one that environment wins almost surely, but not surely. Existence of a game that is undetermined with respect to almost sure winning strategies is witnessed by the game depicted in Figure 4, wherein under delay ≥ 2 , neither controller nor environment can secure a win with probability




Figure 8: A reachability game that, under any positive delay, is won by controller with probability $\frac{7}{8}$ via the simple randomized strategy of balanced coin tossing (thus randomly generating head and tail events h and t with probability $\frac{1}{2}$ each), but won by player O surely if interpreted as a delay game due to the lookahead on Player I 's actions granted to Player O . The objective for controller (Player I , resp.) in this game is to reach some black state.

strictly larger than $\frac{1}{2}$. This game consequently remains undetermined with respect to almost sure winning strategies. \square

Item (2) of Theorem 5.3 is of particular interest, as it expresses a delay-related strengthening of controller relative to Player I , letting controller win almost surely where Player I loses for sure. The correspondence between controller and Player I observed in the deterministic setting thus breaks down when almost sure winning is considered and mixed strategies are permitted.

6. FURTHER REFINING THE CORRESPONDENCE: WINNING WITH PROBABILITY

The differences between delay games and games under delayed control become even more pronounced when we ask for strategies that win with at least a given probability $\theta \in]0, 1[$ instead of asking for an almost-sure win.

Figure 8 depicts a reachability game under delayed control that controller can win with probability $\frac{1}{2} + \left(\frac{1}{2}\right)^2 + \left(\frac{1}{2}\right)^3 = \frac{7}{8}$ when the delay is non-zero, while Player I in the corresponding delay game loses surely.

Such a difference is not particular to reachability games: Note that controller would also win the Büchi game defined by the black states being controller's winning set with probability $\frac{7}{8}$, and likewise controller would win the safety game where it has to avoid the black states forever with probability $\frac{1}{8}$, while all these games will surely be lost by Player I . Note that we can furthermore also generate a safety game that is won by controller with probability $\frac{7}{8}$ (and likewise a reachability game or a Büchi game which are won by controller with probability $\frac{1}{8}$) by simply swapping the acceptance status of the four bottom connected components in the game graph of Figure 8. These games are still lost by Player I for sure.

Our next result shows that a game with these properties exists for every rational number in $[0, 1]$.

Lemma 6.1. *Let $\theta \in [0, 1]$ be rational. Then there exists a game under delayed control with ω -regular winning condition that controller wins with probability θ .*

Proof. The case of $\theta = 1$ has been covered in Section 3 and Section 5, as all games won by Player I as well as the game in Figure 6 are won with probability 1 by controller. Hence, w.l.o.g., $\theta = \frac{n}{m}$ with n, m being integer and $0 \leq n < m$. Consider the game where first controller picks a number $a \in \{0, \dots, m-1\}$ and then environment picks a number $b \in \{0, \dots, m-1\}$. The game then is over (i.e., the states now reached are all sinks) and controller wins if and only if $(a+b) \bmod m < n$.

When not subject to delay, this game is won by the environment by observing a and picking $b = m-1-a$, which gives $(a+b) \bmod m = (a+m-1-a) \bmod m = m-1 \geq n$, violating controller's winning condition.

But under any positive delay, the best strategy played by controller obviously is to draw a random choice uniformly from $\{0, \dots, m-1\}$, as any biased choice (no matter whether deterministic or random) would render its otherwise unobservable (due to the delay) choice predictable to the environment and thus give the environment an advantage. But with a being chosen uniformly from $\{0, \dots, m-1\}$, the random variable $(a+b) \bmod m$ also is uniformly distributed over $\{0, \dots, m-1\}$, no matter what strategy the environment plays, as the environmental strategy has to be independent of the particular choice of a due to the delay. The chance for controller to win, given its winning condition $(a+b) \bmod m < n$ and the uniform distribution of $(a+b) \bmod m$ over $\{0, \dots, m-1\}$, consequently is $\frac{n}{m}$. Hence, this game under delayed control features a win probability of exactly $\frac{n}{m} = \theta$ for controller. \square

Remark 6.2. By the symmetry between controller and environment inherent to games under delayed control with respect to a fixed delay [CFL⁺21, Section 2], which subjects both to the same informedness constraints, environment wins with probability $1-\theta$ when controller wins with probability θ . Hence, for each $\theta \in [0, 1]$ being rational there also exists a game under delayed control with ω -regular winning condition that environment wins with probability θ .

As delay games with Borel winning conditions are determined, there obviously is no counterpart to the property expressed in Lemma 6.1 for delay games: in such delay games, the winning probabilities are, under optimal strategies, always either 0 or 1; in fact, these losses or wins are even sure and not only almost sure.

Lemma 6.3. *If a game under delayed control with Borel winning condition is won by controller with probability $\theta < 1$ then the corresponding delay game is won by Player O surely.*

Proof. As $\theta < 1$, the controller cannot win the game surely. By the results from Section 3, in particular Lemma 3.7, this implies that Player I cannot win the corresponding delay game. As delay games with Borel winning conditions are determined, Player O then wins (for sure). \square

Please note that in case of a win probability of $\theta = 1$ for controller, either a win of Player I or of Player O may apply, as has been shown in Section 3 and Section 5. The transformation from Section 3 shows that games won by Player I are surely won by controller, while Figure 6 provides an example of a game won surely by Player O while won with probability 1 by controller.

As Lemma 6.1 demonstrates that the win probabilities of games under delayed control are dense in $[0, 1]$, the following refined statement about correspondence between the different types of games under delay makes sense.




Figure 9: A zoom on the area labeled “undetermined” in Figure 7. For every rational number $\theta \in]0, 1[$ there is a game under delayed control \mathcal{G} and a delay δ such that controller wins \mathcal{G} under delay δ with value θ .

Theorem 6.4. *Let \mathcal{G} be a game with Borel winning condition, $\delta = 2k$ with $k \in \mathbb{N} \setminus \{0\}$ a delay, and $\theta \in [0, 1]$ the probability of winning \mathcal{G} under delay δ by controller with an optimal mixed strategy.*

- (1) *If $\theta < 1$ then Player O wins the corresponding delay game $\Gamma_k(\overline{L(\mathcal{G})})$ with delay k , and this win is for sure.*
- (2) *If $\theta = 1$, yet controller does not win \mathcal{G} surely, then Player O wins in the corresponding delay game $\Gamma_k(\overline{L(\mathcal{G})})$ with delay k for sure.*
- (3) *If controller wins \mathcal{G} surely under delay δ then Player I wins the corresponding delay game $\Gamma_k(\overline{L(\mathcal{G})})$ for sure under delay k .*
- (4) *If environment wins \mathcal{G} surely under delay δ then Player O wins the corresponding delay game $\Gamma_k(\overline{L(\mathcal{G})})$ for sure under delay k .*
- (5) *All the aforementioned classes are non-empty, i.e., there exist games under delayed control where controller wins, where controller wins almost surely (but not surely), where controller wins with some probability $\theta \in]0, 1[$, where environment wins almost surely (but not surely), and where environment wins surely.*

Figure 9 refines the are labeled “undetermined” in Figure 7.

Proof. Lemma 6.3 states property (1). Properties (2) to (4) are reformulations of the corresponding properties from Theorem 5.3. The last property (5) follows from the corresponding property in Theorem 5.3 together with Lemma 6.1. \square

7. CONCLUSION

We have compared delay games [KZ16a] and games under delayed control [CFL⁺21], two types of infinite games aiming to model asynchronicity in reactive synthesis, and have exhibited the differences in definitions and charted the relation between them with respect to both deterministic and randomized strategies: When considering sure winning and deterministic strategies, one can efficiently transform a game under delayed control into a delay game such that controller wins the game under delayed control with delay δ by a deterministic strategy if and only if Player I wins the resulting delay game with lookahead of size $\frac{\delta}{2}$. Dually, one can efficiently transform a delay game into a game under delayed control such that Player I wins the delay game with lookahead of size k if and only if controller wins the resulting game under delayed control with delay $2k$ by a deterministic strategy. These results allow us to transfer known complexity results and bounds on the amount of delay from delay games

to games under delayed control, for which no such results were known (when considering deterministic strategies).

Analogous results fail in the setting of randomized strategies and almost sure or even quantitative winning conditions. Here, it is completely open who can win the game under delayed control and with which probability, unless Player I can win the delay game surely. For the case that Player I does not have a sure winning strategy, i.e., whenever Player O has a sure winning strategy (due to determinacy of delay games with Borel winning conditions), we can construct games that controller, normally corresponding to Player I , nevertheless wins almost surely, wins with any rational probability in $]0, 1]$, loses almost surely, or loses surely. These findings refine our original result from [CFL⁺21] that games under delayed control are not determined. They also expose a profound difference between the problem of control via delaying channels, which games under delayed control capture, and delay games.

Acknowledgements: Martin Fränzle has been supported by Deutsche Forschungsgemeinschaft under grant no. DFG FR 2715/5-1 “Konfliktresolution und kausale Inferenz mittels integrierter sozio-technischer Modellbildung”. Paul Kröger received funding from Germany’s Federal Ministry of Education and Research strategic “MANNHEIM” project “AutoDevSafeOps—Integrated development and operation of safe automotive systems”. The majority of the research was carried out while Sarah Winter has been affiliated with Université libre de Bruxelles, Belgium and has been supported by the Fonds National de la Recherche Scientifique – F.R.S.-FNRS. Martin Zimmermann has been supported by DIREC—Digital Research Centre Denmark.

REFERENCES

- [BL69] J. Richard Büchi and Lawrence H. Landweber. Solving sequential conditions by finite-state strategies. *Trans. Amer. Math. Soc.*, 138:pp. 295–311, 1969. doi:10.2307/1994916.
- [CFL⁺21] Mingshuai Chen, Martin Fränzle, Yangjia Li, Peter Nazier Mosaad, and Naijun Zhan. Indecision and delays are the parents of failure - taming them algorithmically by synthesizing delay-resilient control. *Acta Informatica*, 58(5):497–528, 2021. doi:10.1007/s00236-020-00374-7.
- [FWZ23] Martin Fränzle, Sarah Winter, and Martin Zimmermann. Strategies resilient to delay: Games under delayed control vs. delay games. In Antonis Achilleos and Dario Della Monica, editors, *GandALF 2023*, volume 390 of *EPTCS*, pages 220–235. Open Publishing Association, 2023. doi:10.4204/EPTCS.390.14.
- [GS53] David Gale and Frank M. Stewart. Infinite games with perfect information. *Annals of Mathematics*, 28:245–266, 1953. doi:10.1515/9781400881970-014.
- [GTW02] Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *LNCS*. Springer, 2002. doi:10.1007/3-540-36387-4.
- [HKT12] Michael Holtmann, Lukasz Kaiser, and Wolfgang Thomas. Degrees of lookahead in regular infinite games. *Log. Methods Comput. Sci.*, 8(3), 2012. doi:10.2168/LMCS-8(3:24)2012.
- [HL72] Frederick A. Hosch and Lawrence H. Landweber. Finite delay solutions for sequential conditions. In Maurice Nivat, editor, *ICALP 1972*, pages 45–60. North-Holland, Amsterdam, 1972.
- [KZ15] Felix Klein and Martin Zimmermann. What are strategies in delay games? Borel determinacy for games with lookahead. In Stephan Kreutzer, editor, *CSL 2015*, volume 41 of *LIPICs*, pages 519–533. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPICs.CSL.2015.519.
- [KZ16a] Felix Klein and Martin Zimmermann. How much lookahead is needed to win infinite games? *Log. Methods Comput. Sci.*, 12(3), 2016. doi:10.2168/LMCS-12(3:4)2016.
- [KZ16b] Felix Klein and Martin Zimmermann. Prompt delay. In Akash Lal, S. Akshay, Saket Saurabh, and Sandeep Sen, editors, *FSTTCS 2016*, volume 65 of *LIPICs*, pages 43:1–43:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICs.FSTTCS.2016.43.

- [McN00] Robert McNaughton. Playing infinite games in finite time. In Arto Salomaa, Derick Wood, and Sheng Yu, editors, *A Half-Century of Automata Theory: Celebration and Inspiration*, pages 73–91. World Scientific, 2000.
- [Pnu77] Amir Pnueli. The temporal logic of programs. In *FOCS 1977*, pages 46–57. IEEE Computer Society, 1977. doi:10.1109/SFCS.1977.32.
- [Zim22] Martin Zimmermann. Approximating the minimal lookahead needed to win infinite games. *Information Processing Letters*, 177:106264, 2022. doi:10.1016/j.ip1.2022.106264.