

HOW MUCH LOOKAHEAD IS NEEDED TO WIN INFINITE GAMES? *

FELIX KLEIN ^a AND MARTIN ZIMMERMANN ^b

^{a,b} Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany
e-mail address: {klein,zimmermann}@react.uni-saarland.de

ABSTRACT. Delay games are two-player games of infinite duration in which one player may delay her moves to obtain a lookahead on her opponent’s moves. For ω -regular winning conditions it is known that such games can be solved in doubly-exponential time and that doubly-exponential lookahead is sufficient.

We improve upon both results by giving an exponential time algorithm and an exponential upper bound on the necessary lookahead. This is complemented by showing EXPTIME-hardness of the solution problem and tight exponential lower bounds on the lookahead. Both lower bounds already hold for safety conditions. Furthermore, solving delay games with reachability conditions is shown to be PSPACE-complete.

This is a corrected version of the paper <https://arxiv.org/abs/1412.3701v4> published originally on August 26, 2016.

1. INTRODUCTION

Many of today’s problems in computer science are no longer concerned with programs that transform data and then terminate, but with non-terminating reactive systems which have to interact with a possibly antagonistic environment for an unbounded amount of time. The framework of infinite two-player games is a powerful and flexible tool to verify and synthesize such systems. The seminal theorem of Bchi and Landweber [1] states that the winner of an infinite game on a finite arena with an ω -regular winning condition can be determined and a corresponding finite-state winning strategy can be constructed effectively.

2012 ACM CCS: [Software and its engineering]: Software organization and properties—Software functional properties—Correctness / Formal methods—Software verification.

Key words and phrases: Infinite Games, Delay, ω -regular Languages.

* A preliminary version of this work appeared in the proceedings of ICALP 2015 [12].

^a Supported by the Transregional Collaborative Research Center “AVACS” (SFB/TR 14) of the German Research Foundation (DFG) and by an IMPRS-CS PhD scholarship.

^b Supported by the DFG project “TriCS” (ZI 1516/1-1).

1.1. Delay Games. In this work, we consider an extension of the classical framework: in a delay game, one player can postpone her moves for some time to obtain a lookahead on her opponent's moves. This allows her to win some games which she would lose without lookahead, e.g., if her first move depends on the third move of her opponent. Nevertheless, there are winning conditions that cannot be won with any finite lookahead, e.g., if her first move depends on every move of her opponent. Delay arises naturally if transmission of data in networks or components equipped with buffers are modeled.

From a more theoretical point of view, uniformization of relations by continuous functions [25, 26] can be expressed and analyzed using delay games [9, 24]. We consider games in which two players pick letters from alphabets Σ_I and Σ_O , respectively, thereby producing two infinite sequences α and β . Thus, a strategy for the second player induces a mapping $\tau: \Sigma_I^\omega \rightarrow \Sigma_O^\omega$. It is winning for her if $(\alpha, \tau(\alpha))$ is contained in the winning condition $L \subseteq \Sigma_I^\omega \times \Sigma_O^\omega$ for every α . If this is the case, we say that τ uniformizes L . In the classical setting, in which the players pick letters in alternation, the n -th letter of $\tau(\alpha)$ depends only on the first n letters of α . A strategy with bounded lookahead, i.e., only finitely many moves are postponed, induces a Lipschitz-continuous function τ (in the Cantor topology on Σ^ω) and a strategy with unbounded lookahead induces a continuous function (equivalently, a uniformly continuous function, as Σ^ω is compact (see, e.g., [9])).

1.2. Related Work. Hosch and Landweber proved that it is decidable whether a delay game with an ω -regular winning condition can be won with bounded lookahead [10]. Later, Holtmann, Kaiser, and Thomas revisited the problem and showed that if the delaying player wins such a game with unbounded lookahead, then she already wins it with doubly-exponential bounded lookahead, and gave a streamlined decidability proof yielding an algorithm with doubly-exponential running time [9]. Thus, the delaying player does not gain additional power from having unbounded lookahead, bounded lookahead is sufficient.

Going beyond ω -regularity by considering context-free winning conditions leads to undecidability and non-elementary lower bounds on the necessary lookahead, even for very weak fragments [6]. Nevertheless, there is another extension of the ω -regular winning conditions where one can prove the analogue of the Hosch-Landweber Theorem: it is decidable whether the delaying player wins a delay game with bounded lookahead, if the winning condition is definable in weak monadic second order logic with the unbounding quantifier (WMSO+U) [27]. Furthermore, doubly-exponential lookahead is sufficient for such winning conditions, provided the delaying player wins with bounded lookahead at all. However, bounded lookahead is not always sufficient to win such games, i.e., the analogue of the Holtmann-Kaiser-Thomas Theorem does not hold for WMSO+U winning conditions. In contrast, if the delaying player wins such a game, then she wins no matter how slowly the lookahead grows [28].

Due to these negative results for WMSO+U winning conditions, delay games with winning conditions in Prompt-LTL [15] have been considered: this logic extends LTL [20] with temporal operators whose scope is bounded in time. Determining the winner of such a game is 3EXPTIME-complete and triply-exponential bounded lookahead is necessary in general and always sufficient [14]. Furthermore, all lower bounds already hold for the special case of LTL.

Finally, all delay games with Borel winning conditions are determined, since they were shown to be reducible to delay-free games with Borel winning conditions [13].

Stated in terms of uniformization, Hosch and Landweber proved decidability of the uniformization problem for ω -regular relations by Lipschitz-continuous functions and Holtmann et al. proved the equivalence of the existence of a continuous uniformization function and the existence of a Lipschitz-continuous uniformization function for ω -regular relations. Furthermore, uniformization of context-free relations is undecidable, even with respect to Lipschitz-continuous functions, but uniformization of WMSO+U relations by Lipschitz-continuous functions is decidable.

In another line of work, Carayol and Lding considered the case of finite words [3], and Lding and Winter [17] considered the case of finite trees, which are both decidable. However, the non-existence of MSO-definable choice functions on the infinite binary tree [2, 8] implies that uniformization fails for such trees.

Although several extensions of ω -regular winning conditions for delay games have been studied, several gaps remain open, even for ω -regular winning conditions. Holtmann et al. proved that such games can be solved in doubly-exponential time, but only trivial lower bounds are known. Similarly, their algorithm yields a doubly-exponential upper bound on the necessary lookahead, but only straightforward linear lower bounds are known. Also, only deterministic parity automata were used to specify winning conditions, and the necessary lookahead and the solution complexity are measured in their size. It is open whether weaker automata models like reachability or safety automata have smaller lookahead requirements and allow for faster algorithms.

1.3. Our Contribution. We close these gaps and improve upon both results of Holtmann et al. by determining the exact complexity of ω -regular delay games and by giving tight bounds on the necessary lookahead, both for winning conditions specified by deterministic parity automata and for weaker models.

First, we present an exponential time algorithm for solving delay games with ω -regular winning conditions, an exponential improvement over the original doubly-exponential time algorithm. Both algorithms share some similarities: given a deterministic parity automaton \mathcal{A} recognizing the winning condition of the game, a parity game is constructed that is won by the delaying player if and only if she wins the delay game with winning condition $L(\mathcal{A})$. Furthermore, both parity games are induced by equivalence relations that capture the behavior of \mathcal{A} . However, our parity game is of exponential size while the one of Holtmann et al. is doubly-exponential. Also, they need an intermediate game, the so-called block game, to prove the equivalence of the delay game and the parity game, while our equivalence proof is direct. Thus, our algorithm and its correctness proof are even simpler than the ones of Holtmann et al.

Second, we show that solving delay games is EXPTIME-complete by proving the first non-trivial lower bound on the complexity of ω -regular delay games. The lower bound is proved by a reduction from the acceptance problem for alternating polynomial space Turing machines [4], which results in delay games with safety conditions. Thus, solving delay games with safety conditions is already EXPTIME-hard. Our reduction is inspired by the EXPTIME-hardness proof for continuous simulation games [11], a simulation game on Bchi automata where Duplicator is able to postpone her moves to obtain a lookahead on Spoiler's moves. However, this reduction is from a two-player tiling problem while we directly reduce from alternating Turing machines.

Third, we determine the exact amount of lookahead necessary to win delay games with ω -regular winning conditions. From our algorithm we derive an exponential upper bound,

which is again an exponential improvement. This upper bound is complemented by the first non-trivial lower bound on the necessary lookahead: there are reachability and safety conditions that are winning for the delaying player, but only with exponential lookahead, i.e., our upper bound is tight.

Fourth, we present the first results for fragments of ω -regular winning conditions. As already mentioned above, our lower bounds on complexity and necessary lookahead already hold for safety conditions, i.e., safety is already as hard as parity. Thus, the complexity of the problems manifests itself in the transition structure of the automaton, not in the acceptance condition. For reachability conditions, the situation is different: we show that solving delay games with reachability conditions is equivalent to universality of non-deterministic reachability automata and therefore PSPACE-complete. Thus, there is a gap between the complexity of solving delay games with reachability conditions and delay games with safety conditions (unless PSPACE = EXPTIME). Furthermore, as both reachability and safety already require exponential lookahead and are of high computational complexity it is natural to search for more tractable fragments. One such fragment is the class of winning conditions that are both reachability and safety, the so-called clopen languages. Here, we again prove tight bounds on the complexity and on the necessary lookahead for delay games with ω -regular clopen winning conditions: linear lookahead is necessary and sufficient and determining the winner is CO-NP-complete.

Our results are summarized in Figure 1, where an asterisk * denotes cases where the lower bounds hold for deterministic automata while the upper bounds hold for non-deterministic automata. All bounds are tight.

acceptance	lookahead	complexity
clopen*	linear	CO-NP-complete
reachability*	exponential	PSPACE-complete
det. safety	exponential	EXPTIME-complete
det. parity	exponential	EXPTIME-complete

Figure 1: Table of results.

2. PRELIMINARIES

The set of non-negative integers is denoted by \mathbb{N} . An alphabet Σ is a non-empty finite set of letters, Σ^* the set of finite words over Σ , Σ^n the set of words of length n , and Σ^ω the set of infinite words. The empty word is denoted by ε and the length of a finite word w by $|w|$. For $w \in \Sigma^* \cup \Sigma^\omega$ we write $w(n)$ for the n -th letter of w .

2.1. Automata. An automaton $\mathcal{A} = (Q, \Sigma, q_I, \Delta, \varphi)$ consists of a finite set Q of states with initial state $q_I \in Q$, an alphabet Σ , a non-deterministic transition function $\Delta: Q \times \Sigma \rightarrow 2^Q \setminus \{\emptyset\}$, and an acceptance condition φ , which is either a set $F \subseteq Q$ of accepting states, depicted by doubly-lined states, or a coloring $\Omega: Q \rightarrow \mathbb{N}$. Note that we require automata to be complete, i.e., every state has at least one outgoing transition labeled by a for every $a \in \Sigma$. We discuss this restriction after introducing the different acceptance conditions. The size of \mathcal{A} , denoted by $|\mathcal{A}|$, is the cardinality of Q . A run of \mathcal{A} on an infinite word $\alpha(0)\alpha(1)\alpha(2)\cdots$

over Σ is a sequence $q_0q_1q_2\cdots$ such that $q_0 = q_I$ and $q_{n+1} \in \Delta(q_n, \alpha(n))$ for every n . Runs on finite words are defined analogously.

An automaton is deterministic, if $|\Delta(q, a)| = 1$ for every q and a . In this case, we denote Δ by a function $\delta: Q \times \Sigma \rightarrow Q$. A deterministic automaton has a unique run on every finite or infinite word. A state q of \mathcal{A} is a sink, if $\Delta(q, a) = \{q\}$ for every $a \in \Sigma$.

Given an automaton \mathcal{A} over Σ with some set F of accepting states or with some coloring Ω , we consider the following acceptance conditions and their induced languages:

- $L_*(\mathcal{A}) \subseteq \Sigma^*$ denotes the set of finite words over Σ accepted by \mathcal{A} , i.e., the set of words that have a run ending in F .
- $L_\exists(\mathcal{A}) \subseteq \Sigma^\omega$ denotes the set of infinite words over Σ that have a run visiting an accepting state at least once, called reachability acceptance. Due to completeness, we have $L_\exists(\mathcal{A}) = L_*(\mathcal{A}) \cdot \Sigma^\omega$.
- Dually, $L_\forall(\mathcal{A}) \subseteq \Sigma^\omega$ denotes the set of infinite words over Σ that have a run only visiting accepting states, called safety acceptance.
- $L_p(\mathcal{A}) \subseteq \Sigma^\omega$ denotes the set of infinite words that have a run such that the maximal color visited infinitely often during this run is even. This is the classical (max)-parity acceptance condition.

A reachability (winning) condition is a language that is accepted by an automaton with reachability acceptance, called a reachability automaton. Safety (winning) conditions and parity (winning) conditions and safety and parity automata are defined similarly. Every (deterministic) reachability automaton can be turned into a (deterministic) parity automaton of the same size that recognizes the same language by turning the accepting states into sinks and by then defining an appropriate coloring. Similarly, one can turn a (deterministic) safety automaton into a (deterministic) parity automaton of the same size that recognizes the same language by turning the rejecting states into sinks. Deterministic parity automata recognize exactly the ω -regular languages.

Recall that we require automata to be complete. For safety and parity acceptance this is no restriction, since we can always add a fresh rejecting sink, i.e., one that is not in F in the case of safety and one with odd color in the case of parity, and lead all missing transitions to this sink. However, incomplete automata with reachability acceptance are strictly stronger than complete ones, as incompleteness can be used to check safety properties. We impose this restriction since we are interested in pure reachability conditions.

Given a language $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ we denote by $\text{pr}_I(L)$ its projection to the first component, i.e., $\text{pr}_I(L) = \{\alpha(0)\alpha(1)\alpha(2)\cdots \mid \binom{\alpha(0)}{\beta(0)} \binom{\alpha(1)}{\beta(1)} \binom{\alpha(2)}{\beta(2)} \cdots \in L\}$. Similarly, given an automaton \mathcal{A} over $\Sigma_I \times \Sigma_O$, we denote by $\text{pr}_I(\mathcal{A})$ the automaton obtained by projecting each letter to its first component, i.e., we replace the alphabet $\Sigma_I \times \Sigma_O$ by Σ_I and Δ by Δ' , where $\Delta'(q, a) = \bigcup_{b \in \Sigma_O} \Delta(q, \binom{a}{b})$ for all $q \in Q$ and all $a \in \Sigma_I$.

Remark 2.1. Let $\text{acc} \in \{*, \exists, \forall, p\}$, then $\text{pr}_I(L_{\text{acc}}(\mathcal{A})) = L_{\text{acc}}(\text{pr}_I(\mathcal{A}))$.

2.2. Games with Delay. A delay function is a mapping $f: \mathbb{N} \rightarrow \mathbb{N} \setminus \{0\}$, which is said to be constant, if $f(i) = 1$ for every $i > 0$. Given an ω -language $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ and a delay function f , the game $\Gamma_f(L)$ is played by two players, the input player “Player I ” and the output player “Player O ” in rounds $i = 0, 1, 2, \dots$ as follows: in round i , Player I picks a word $u_i \in \Sigma_I^{f(i)}$, then Player O picks one letter $v_i \in \Sigma_O$. We refer to the sequence $(u_0, v_0), (u_1, v_1), (u_2, v_2), \dots$ as a play of $\Gamma_f(L)$, which yields two infinite words $\alpha = u_0u_1u_2\cdots$

and $\beta = v_0 v_1 v_2 \dots$. Player O wins the play if and only if the outcome $(\alpha^{(0)}_{\beta^{(0)}})(\alpha^{(1)}_{\beta^{(1)}})(\alpha^{(2)}_{\beta^{(2)}}) \dots$ is in L , otherwise Player I wins.

Given a delay function f , a strategy for Player I is a mapping $\tau_I: \Sigma_O^* \rightarrow \Sigma_I^*$ where $|\tau_I(w)| = f(|w|)$, and a strategy for Player O is a mapping $\tau_O: \Sigma_I^* \rightarrow \Sigma_O$. Consider a play $(u_0, v_0), (u_1, v_1), (u_2, v_2), \dots$ of $\Gamma_f(L)$. Such a play is consistent with τ_I , if $u_i = \tau_I(v_0 \dots v_{i-1})$ for every $i \in \mathbb{N}$. It is consistent with τ_O , if $v_i = \tau_O(u_0 \dots u_i)$ for every $i \in \mathbb{N}$. A strategy τ for Player $P \in \{I, O\}$ is winning, if every play that is consistent with τ is winning for Player P . We say that a player wins $\Gamma_f(L)$, if she has a winning strategy.

We continue with two examples of delay games with ω -regular winning conditions.

Examples 2.2. Note that both conditions can be accepted by safety automata.

- (1) Consider L_1 over $\{a, b, c\} \times \{b, c\}$ with $(\alpha^{(0)}_{\beta^{(0)}})(\alpha^{(1)}_{\beta^{(1)}})(\alpha^{(2)}_{\beta^{(2)}}) \dots \in L_1$, if $\alpha(n) = a$ for every $n \in \mathbb{N}$ or if $\beta(0) = \alpha(n)$, where n is the smallest position with $\alpha(n) \neq a$. Intuitively, Player O wins, if the letter she picks in the first round is equal to the first letter other than a that Player I picks. Also, Player O wins, if there is no such letter.

We claim that Player I wins $\Gamma_f(L_1)$ for every delay function f : Player I picks $a^{f(0)}$ in the first round and assume Player O picks b afterwards (the case where she picks c is dual). Then, Player I picks a word starting with c in the second round. The resulting play is winning for Player I no matter how it is continued. Thus, Player I has a winning strategy in $\Gamma_f(L_1)$.

- (2) Now, consider L_2 over $\{a, b, c\} \times \{a, b, c\}$ where $(\alpha^{(0)}_{\beta^{(0)}})(\alpha^{(1)}_{\beta^{(1)}})(\alpha^{(2)}_{\beta^{(2)}}) \dots \in L_2$, if $\beta(n) = \alpha(n+2)$ for every $n \in \mathbb{N}$, i.e., Player O wins if the input is shifted two positions to the left.

Player O has a winning strategy for $\Gamma_f(L_2)$ for every f with $f(0) \geq 3$. In this case, Player O has at least three letters lookahead in each round, which suffices to shift the input of Player I two positions to the left. On the other hand, if $f(0) < 3$, then Player I has a winning strategy, since Player O has to pick $\beta(0)$ before $\alpha(0+2)$ has been picked by Player I .

Note that if a language L is recognizable by a (deterministic) parity automaton, then $\Gamma_f(L)$ is determined, i.e., exactly one of the players has a winning strategy, as a delay game with parity condition can be expressed as an explicit parity game in a countable arena, which is determined [5, 19]. This result has been recently generalized to the class of Borel winning conditions [13].

Also, note that universality of $\text{pr}_I(L)$ is a necessary condition for Player O to win $\Gamma_f(L)$. Otherwise, Player I could pick a word from $\Sigma_I^\omega \setminus \text{pr}_I(L)$, which is winning for him, no matter how Player O responds.

Proposition 2.3. *If Player O wins $\Gamma_f(L)$, then $\text{pr}_I(L)$ is universal.*

We are interested in solving delay games: given an automaton \mathcal{A} recognizing a language $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$, determine whether Player O wins $\Gamma_f(L)$ for some delay function f . Furthermore, we are interested in upper and lower bounds on the lookahead induced by such an f . We measure the complexity and the bounds in the size of the automation \mathcal{A} .

3. LOWER BOUNDS ON THE LOOKAHEAD

In this section, we prove lower bounds on the necessary lookahead for Player O to win delay games. We first give an exponential lower bound for reachability conditions, then we extend

this idea to provide an exponential lower bound for safety conditions. Consequently, the same bounds hold for more expressive acceptance conditions like Bchi, co-Bchi, and parity. They are complemented by an exponential upper bound for parity conditions in the next section. Note that both lower bounds already hold for deterministic automata.

Theorem 3.1. *For every $n > 1$ there is a language L_n such that*

- $L_n = L_{\exists}(\mathcal{A}_n)$ for some deterministic automaton \mathcal{A}_n with $|\mathcal{A}_n| \in \mathcal{O}(n)$,
- Player O wins $\Gamma_f(L_n)$ for some constant delay function f , but
- Player I wins $\Gamma_f(L_n)$ for every delay function f with $f(0) \leq 2^n$.

Proof. Let $\Sigma_I = \Sigma_O = \{1, \dots, n\}$. We say that w in Σ_I^* contains a bad j -pair, for $j \in \Sigma_I$, if there are two occurrences of j in w such that no $j' > j$ occurs in between. The automaton \mathcal{B}_j , depicted in Figure 2(a), accepts exactly the words with a bad j -pair. Now, consider the language L over Σ_I defined by

$$L = \bigcap_{1 \leq j \leq n} \{w \in \Sigma_I^* \mid w \text{ contains no bad } j\text{-pair}\}.$$

First, we show that every $w \in L$ satisfies $|w| < 2^n$. To this end, we prove the stronger statement $|w| < 2^m$, where m is the maximal letter occurring in w , by induction over m . The induction base $m = 1$ is trivial, so let $m > 1$. There cannot be two occurrences of m in w , as they would constitute a bad m -pair. Accordingly, there is exactly one m in w , i.e., we can decompose w into $w = w_{\triangleleft} m w_{\triangleright}$ such that w_{\triangleleft} and w_{\triangleright} contain no occurrence of m . Thus, the induction hypothesis is applicable and shows $|w_{\triangleleft}|, |w_{\triangleright}| < 2^{m-1}$, which implies $|w| < 2^m$.

Dually, there is a word $w_n \in L$ with $|w_n| = 2^n - 1$ which is defined inductively via $w_1 = 1$ and $w_m = w_{m-1} m w_{m-1}$ for $m > 1$. A simple induction shows $w_n \in L$ and $|w_n| = 2^n - 1$.




Figure 2: (a) Automaton \mathcal{B}_j for $j \in \Sigma_I$. (b) Construction of \mathcal{A}_n .

The winning condition L_n is defined as follows: $\binom{\alpha(0)}{\beta(0)} \binom{\alpha(1)}{\beta(1)} \binom{\alpha(2)}{\beta(2)} \dots$ is in L_n if $\alpha(1)\alpha(2)\alpha(3)\dots$ contains a bad $\beta(0)$ -pair, i.e., with her first move, Player O has to pick a j such that Player I has produced a bad j -pair. For technical reasons, the first letter picked by Player I is ignored. The construction of an automaton \mathcal{A}_n recognizing L_n is sketched in Figure 2(b), where $*$ denotes an arbitrary letter and $\mathcal{B}_j[a \setminus \binom{a}{*}]$ denotes \mathcal{B}_j where for each $a \in \Sigma_I$ every transition labeled by a is replaced by transitions labeled by $\binom{a}{b}$ for every $b \in \Sigma_O$. Clearly, we have $\mathcal{A}_n \in \mathcal{O}(n)$.

Player O wins $\Gamma_f(L_n)$ for every delay function with $f(0) > 2^n$. In the first round, Player I has to pick a word u_0 such that u_0 without its first letter is not in L , as it is too long for being in L . This allows Player O to find a bad j -pair for some j , i.e., she wins the play no matter how it is continued.

However, for f with $f(0) \leq 2^n$, Player I has a winning strategy by picking the prefix of $1w_n$ of length $f(0)$ in the first round. Player O has to answer with some $j \in \Sigma_O$. In this situation, Player I continues by playing some $j' \neq j$ ad infinitum, which ensures that the resulting sequence does not contain a bad j -pair. Thus, the play is winning for Player I . \square

For safety conditions, we use the same idea as in the reachability case, but we need to introduce a new letter $\#$ to give Player I the possibility to reach a non-accepting state.

Theorem 3.2. *For every $n > 1$ there is a language L'_n such that*

- $L'_n = L_{\forall}(\mathcal{A}'_n)$ for some deterministic automaton \mathcal{A}'_n with $|\mathcal{A}'_n| \in \mathcal{O}(n)$,
- Player O wins $\Gamma_f(L'_n)$ for some constant delay function f , but
- Player I wins $\Gamma_f(L'_n)$ for every delay function f with $f(0) \leq 2^n$.

Proof. Let $\Sigma_I = \Sigma_O = \{1, 2, \dots, n, \#\}$ and let w_n be defined as above. We introduce a new automaton $\mathcal{B}'_{\#}$ and extend every automaton \mathcal{B}_j from the previous proof to \mathcal{B}'_j as depicted in Figure 3. The automaton \mathcal{A}'_n is constructed as in the previous proof using the automata \mathcal{B}'_j and $\mathcal{B}'_{\#}$ instead of the \mathcal{B}_j .




Figure 3: (a) Automaton \mathcal{B}'_j for $j \in \Sigma_I \setminus \{\#\}$. (b) Automaton $\mathcal{B}'_{\#}$.

For $f(0) > 2^n$, Player O wins the game: assume Player I picks u_0 in the first round and let u'_0 be u_0 without its first letter. If u'_0 contains a $\#$ preceded by at most one n , then Player O answers with $\#$ in the first round. If there is more than one n before the first $\#$ in u'_0 , then she answers with n . Finally, if there is no $\#$ in u'_0 , she can pick a j such that u'_0 contains a bad j -pair. All outcomes are winning for Player O .

Player I still wins the game for a constant delay function f with $f(0) \leq 2^n$ by picking the prefix of $1w_n$ of length $f(0)$ in the first round: if Player O picks some $j \in \Sigma_O \setminus \{\#\}$ in the first round, then Player I just has to answer with $\#$. Otherwise, if Player O picks $\#$ in the first round, then Player I continues with n^ω . He wins in both situations. \square

The aforementioned constructions also work for constant-size alphabets, if we encode every $j \in \{1, \dots, n\}$ in binary with the most significant bit in the first position. Then, the natural ordering on $\{1, \dots, n\}$ is exactly the lexicographical ordering on the corresponding bit-string representation. Accordingly, we can encode every \mathcal{B}_j , \mathcal{B}'_j , and $\mathcal{B}'_{\#}$ in logarithmic size in n , as deciding whether the input represents j , is larger than j , or smaller than j can be checked bit-wise. Together with a binary decision tree of size $\mathcal{O}(n)$ for the initial choice of Player O we obtain deterministic automata \mathcal{A}_n and \mathcal{A}'_n whose sizes are in $\mathcal{O}(n \log n)$. It is open whether linear-sized automata and a constant-sized alphabet can be achieved simultaneously.

4. COMPUTATIONAL COMPLEXITY OF DELAY GAMES

In this section, we determine the computational complexity of solving delay games. First, we consider the special case of reachability conditions and prove such games to be PSPACE-complete. Then, we show that games with safety conditions are EXPTIME-hard. The latter bound is complemented by an EXPTIME-algorithm for solving delay games with parity conditions. From this algorithm, we also deduce an exponential upper bound on the necessary lookahead for Player O , which matches the lower bounds given in the previous section.

4.1. Reachability Conditions. Recall that universality of the projection to the first component of the winning condition is a necessary condition for Player O for having a winning strategy in a delay game. Our first result in this section states that universality is also sufficient in the case of reachability winning conditions. Thus, solving delay games with reachability conditions is equivalent, via linear time reductions, to the universality problem for non-deterministic reachability automata, which is PSPACE-complete (see Appendix A). Therefore, solving delay games with reachability conditions is PSPACE-complete as well. Also, our proof yields an exponential upper bound on the necessary lookahead.

Theorem 4.1. *Let $L = L_{\exists}(\mathcal{A})$, where \mathcal{A} is a non-deterministic reachability automaton. The following are equivalent:*

- (1) *Player O wins $\Gamma_f(L)$ for some delay function f .*
- (2) *Player O wins $\Gamma_f(L)$ for some constant delay function f with $f(0) \leq 2^{|\mathcal{A}|}$.*
- (3) *$\text{pr}_I(L)$ is universal.*

Proof. The implication (2) \Rightarrow (1) is trivial and the implication (1) \Rightarrow (3) is given by Proposition 2.3. It remains to show (3) \Rightarrow (2).

Let $\text{pr}_I(L)$ be universal. First, recall that the projection automaton $\text{pr}_I(\mathcal{A})$ recognizes $\text{pr}_I(L)$ and that $L_{\exists}(\text{pr}_I(\mathcal{A})) = L_*(\text{pr}_I(\mathcal{A})) \cdot \Sigma_I^{\omega}$. Using a pumping argument, we show that on every input $\text{pr}_I(\mathcal{A})$ has a run reaching an accepting state after at most $2^{|\mathcal{A}|}$ steps. Thus, given the first move $\alpha(0) \cdots \alpha(f(0) - 1)$ of Player I with $f(0) \geq 2^{|\mathcal{A}|}$, Player O wins by picking a suitable completion that forces the run of \mathcal{A} into an accepting state.

To formalize this, we assume w.l.o.g. that the accepting states of \mathcal{A} are sinks, which implies that $L_*(\text{pr}_I(\mathcal{A}))$ is suffix-closed, i.e., $w \in L_*(\text{pr}_I(\mathcal{A}))$ implies $ww' \in L_*(\text{pr}_I(\mathcal{A}))$ for every $w' \in \Sigma_I^*$. Furthermore, let \mathcal{A}^c be an automaton recognizing the complement of $L_*(\text{pr}_I(\mathcal{A}))$, which is prefix-closed, as it is the complement of a suffix-closed language. We can choose \mathcal{A}^c such that $|\mathcal{A}^c| \leq 2^{|\mathcal{A}|}$.

We claim that $L_*(\mathcal{A}^c)$ is finite. Assume it is infinite. Then, by Knig's Lemma there is an infinite word α whose prefixes are all in $L_*(\mathcal{A}^c)$. Due to universality, we have $\alpha \in L_{\exists}(\text{pr}_I(\mathcal{A}))$, i.e., there is a prefix of α in $L_*(\text{pr}_I(\mathcal{A}))$. Thus, the prefix is in $L_*(\text{pr}_I(\mathcal{A}))$ and in the complement $L_*(\mathcal{A}^c)$ yielding the desired contradiction. An automaton with n states with a finite language accepts words of length at most $n - 1$. Thus, $w \in L_*(\text{pr}_I(\mathcal{A}))$ for every $w \in \Sigma_I^*$ with $|w| \geq 2^{|\mathcal{A}|}$.

Using this, we show that Player O wins $\Gamma_f(L)$ if $f(0) \geq 2^{|\mathcal{A}|}$. Player I has to pick $f(0)$ letters with his first move, say $u_0 = \alpha(0) \cdots \alpha(f(0) - 1)$. As $f(0)$ is large enough, we have $u_0 \in L_*(\text{pr}_I(\mathcal{A}))$. Hence, there is a word $\beta(0) \cdots \beta(f(0) - 1) \in \Sigma_O^*$ such that $\binom{\alpha(0)}{\beta(0)} \cdots \binom{\alpha(f(0)-1)}{\beta(f(0)-1)} \in L_*(\mathcal{A})$. By picking $\beta(0), \dots, \beta(f(0) - 1)$ in the first $f(0)$ rounds, Player O wins the play, no matter how it is continued. Hence, she has a winning strategy. \square

The exponential upper bound on the necessary lookahead to win delay games with reachability conditions matches the lower bound presented in the previous section.

Theorem 4.1 shows that solving delay games with reachability conditions is equivalent to universality of non-deterministic reachability automata, which is PSPACE-complete (see Appendix A). Thus, we obtain the complexity of solving delay games with reachability conditions as a corollary of Theorem 4.1.

Corollary 4.2. *The following problem is PSPACE-complete: Given a non-deterministic reachability automaton \mathcal{A} , does Player O win $\Gamma_f(L_{\exists}(\mathcal{A}))$ for some f ?*

Another consequence of the proof of Theorem 4.1 concerns the strategy complexity of delay games with reachability conditions: if Player O wins for some delay function, then she has a winning strategy that receives exponentially many input letters and answers by also giving exponentially many output letters and thereby already guarantees a winning play, i.e., all later moves are irrelevant. Thus, the situation is similar to classical reachability games on graphs, in which positional attractor strategies allow a player to guarantee a win after a bounded number of moves. The strategy described above can be implemented by a lookup table that maps all minimal words in $L_*(\text{pr}_I(\mathcal{A}))$ to a word in Σ_O^* of the same length such that the combined word is accepted by \mathcal{A} .

Finally, note that both upper bounds, the one on the lookahead and the one on the complexity, hold for non-deterministic automata while the lower bounds already hold for deterministic automata.

4.2. Safety Conditions. Unsurprisingly, Example 2.2.1 shows that Theorem 4.1 does not hold for safety conditions: the projection $\text{pr}_I(L_1)$ is universal, but Player O has no winning strategy for any delay function. It turns out that safety conditions are even harder than reachability conditions (unless PSPACE equals EXPTIME): we show solving delay games with safety conditions to be EXPTIME-hard by a reduction from the acceptance problem for alternating polynomial space Turing machines [4].

Theorem 4.3. *The following problem is EXPTIME-hard: Given a deterministic safety automaton \mathcal{A} , does Player O win $\Gamma_f(L_{\forall}(\mathcal{A}))$ for some f ?*

Proof. Let $\mathcal{M} = (Q, Q_{\exists}, Q_{\forall}, \Sigma, q_I, \Delta, q_A, q_R)$ be an alternating polynomial space Turing machine, where $\Delta \subseteq Q \times \Sigma \times Q \times \Sigma \times \{-1, 0, 1\}$ is the transition relation, and let $x \in \Sigma^*$ be an input. For technical reasons, we assume the accepting state q_A and the rejecting state q_R to be equipped with a self-loop. Furthermore, let p be a polynomial that bounds \mathcal{M} 's space consumption. We construct a safety automaton \mathcal{A} of polynomial size in $|\Delta|$ and $p(|x|)$ such that \mathcal{M} rejects x if and only if Player O wins $\Gamma_f(L_{\forall}(\mathcal{A}))$ for some f . This is sufficient, since $\text{APSPACE} = \text{EXPTIME}$ [4] is closed under complement. Thus, we give Player I control over the existential states while Player O controls the universal ones. Additionally, Player I is in charge of producing all configurations with his moves. He can copy configurations in order to wait for Player O 's choices for the universal transitions, which are delayed due to the lookahead.

Formally, the input alphabet Σ_I contains $\Sigma \cup Q$ and two separators N and C while the output alphabet Σ_O contains Δ and two signals \times and \checkmark . Intuitively, Player I produces configurations of \mathcal{M} of length $p(|x|)$ preceded by either C or N to denote whether the configuration is a *copy* of the previous one or a *new* one. Copying configurations is necessary to bridge the lookahead while waiting for Player O to determine the transition that is applied

to a universal configuration. Player I could copy a configuration ad infinitum, but this will be losing for him, unless it is an accepting one. Player O chooses universal transitions at every separator¹ by picking a letter from Δ . At every other position, she has to pick a signal: \mathbf{X} allows her to claim an error in the configurations picked by Player I while \checkmark means that she does not claim an error at the current position.

The automaton \mathcal{A} is the product of safety automata checking the following properties of an input word $(\alpha_{\beta(0)}^{(0)} \alpha_{\beta(1)}^{(1)} \alpha_{\beta(2)}^{(2)} \dots) \in (\Sigma_I \times \Sigma_O)^\omega$:

- (1) $\alpha \in (\{N, C\} \cdot \text{Conf})^\omega$, where Conf is the set of encodings of configurations of length $p(|x|)$, i.e., words of length $p(|x|) + 1$ over $\Sigma \cup Q$ that contain exactly one letter from Q . If this is not the case, then the product automaton \mathcal{A} goes to an accepting sink, i.e., in order to win, Player I has to produce an α that satisfies the requirement.
- (2) $\beta \in (\Delta \cdot \{\checkmark, \mathbf{X}\}^{p(|x|)+1})^\omega$. If this is not the case, then the product automaton \mathcal{A} goes to a rejecting sink, i.e., in order to win, Player O has to produce a β that satisfies the requirement.
- (3) The first configuration picked by Player I is the initial one of \mathcal{M} on x . If this is not the case, the product automaton \mathcal{A} goes to an accepting sink.
- (4) If β contains a \mathbf{X} , then the automaton checks whether there is indeed an error by doing the following at the first occurrence of \mathbf{X} : it stores the previous, the current, and the next input letter, the transition picked by Player O at the last separator N , and whether the current configuration is existential or universal. Some of this information has to be stored continuously, since these letters appear before the first \mathbf{X} . This is possible using a set of states whose size is polynomial in $|\Sigma| + |Q| + |\Delta|$.

Then, the automaton processes $p(|x|) + 1$ letters (and remembers whether it traverses the separator N or C), and then checks whether the letter just reached is updated correctly or not:

- If the separator is C , then the current letter is updated correctly, if it is equal to the marked one.
- If the separator is N and the configuration in which the error was marked is existential, then the letter is updated correctly, if there is a transition of \mathcal{M} that is compatible with the current letter and the marked one.
- If the separator is N and the configuration in which the error was marked is universal, then the letter is updated correctly, if it is compatible with the transition picked by Player O at the last separator N before the \mathbf{X} , which is stored by the automaton. If she has picked a transition that is not applicable to the current configuration, the product automaton \mathcal{A} goes to a rejecting sink.

If the update is not correct, i.e., Player O has correctly claimed an error, then \mathcal{A} goes to an accepting sink. Otherwise, it goes to a rejecting sink, i.e., in order to win, Player O should only claim an error at an incorrect update of a configuration, but she wins if she correctly claims an error. All subsequent claims by Player O are ignored, i.e., after the first claim is evaluated, the play is either accepted or rejected, no matter how it is continued.

- (5) Finally, if α contains the accepting state of \mathcal{M} , then \mathcal{A} goes to a rejecting sink, unless Player O correctly claimed an error in a preceding configuration.

All these properties can be checked by deterministic safety automata whose sizes are polynomial in the size of \mathcal{M} and $p(|x|)$. All non-sink states of \mathcal{A} are accepting, i.e., as long

¹If the following configuration is existential or the separator is a C , then her choice is ignored.

as both players stick to their requirements on the format, Player I starts with the initial configuration, Player O does not incorrectly claim an error, and the accepting state of \mathcal{M} is not reached, then the input is accepted.

It remains to prove that \mathcal{M} rejects x if and only if Player O wins $\Gamma_f(L_{\forall}(\mathcal{A}))$ for some f .

“ \Rightarrow ”: Assume \mathcal{M} rejects x and let f be the constant delay function with $f(0) = p(|x|) + 3$. We show that Player O wins $\Gamma_f(L_{\forall}(\mathcal{A}))$. At every time, Player O has enough lookahead to correctly claim the first error introduced by Player I , if he introduces one. Furthermore, she has access to the whole encoding of each universal configuration whose successor she has to determine. This allows her to simulate the rejecting run of \mathcal{M} on x , which does not reach the accepting state q_A , no matter which transitions Player I picks. Thus, he has to introduce an error in order to win, which Player O can detect using the lookahead. If Player I does not introduce an error, the play proceeds ad infinitum by repeating a rejecting configuration forever. In every case, \mathcal{A} accepts the resulting play, i.e., Player O wins. Thus, Player O has a winning strategy for $\Gamma_f(L_{\forall}(\mathcal{A}))$.

“ \Leftarrow ”: We show the contrapositive. Assume that \mathcal{M} accepts x and let f be an arbitrary delay function. We show that Player I wins $\Gamma_f(L_{\forall}(\mathcal{A}))$. Player I starts with the initial configuration and picks the successor configuration of an existential one according to the accepting run, and copies universal configurations as often as necessary to obtain a play prefix in which Player O has to determine the transition she wants to apply in this configuration. Thus, he will eventually produce an accepting configuration of \mathcal{M} without ever introducing an error. Hence, either Player O incorrectly claims an error or the play reaches an accepting state. In either case, Player I wins the resulting play, i.e., he has a winning strategy for $\Gamma_f(L_{\forall}(\mathcal{A}))$. \square

It is noteworthy that the lower bound just proven does not require the full exponential lookahead that is in general necessary to win delay games with safety conditions: Player O wins the game constructed above with sublinear lookahead, as $p(|x|) + 3$ is smaller than the size of \mathcal{A} . Thus, determining the winner of a delay game with safety condition with respect to linearly bounded delay is already EXPTIME-hard.

Finally, the lower bound just proven already holds for deterministic safety automata.

4.3. Parity Conditions. In the previous two subsections, we showed solving delay games with reachability conditions to be PSPACE-complete and solving games with safety conditions to be EXPTIME-hard. To conclude this section, we complement the latter with an exponential time algorithm for solving delay games with parity conditions. Thus, delay games with safety or parity conditions are EXPTIME-complete. Also, we derive an exponential upper bound of the form $2^{(nk)^2}$ on the necessary lookahead from the algorithm, where n is the size and k the number of colors of the automaton. Finally, we lower the upper bound to 2^{nk} via a direct pumping argument, which matches the lower bound from the previous section. Note that all results only hold for deterministic automata.

Theorem 4.4. *The following problem is in EXPTIME: Given a deterministic parity automaton \mathcal{A} , does Player O win $\Gamma_f(L_p(\mathcal{A}))$ for some delay function f ?*

We proceed by constructing an exponentially-sized, delay-free parity game with the same number of colors as \mathcal{A} , which is won by Player O if and only if she wins $\Gamma_f(L_p(\mathcal{A}))$ for some delay function f . Intuitively, we assign to each potential lookahead $w \in \Sigma_I^*$ the behavior it induces on \mathcal{A} , which is given by a function $r: Q \rightarrow 2^{Q \times \Omega(Q)}$. If $(q', c) \in r(q)$, then \mathcal{A} has a

run from q to q' with maximal color c on a word over $\Sigma_I \times \Sigma_O$ whose projection to Σ_I is w . Having the same behavior gives rise to an equivalence relation over Σ_I^* of exponential index. In the parity game we construct, Player I picks equivalence classes of this relation and Player O constructs a run on representatives. Player O wins, if this run is accepting. To account for the delay in the original game, Player I is always one move ahead. This gives Player O a lookahead of one equivalence class, which can be stored in the state space of the parity game.

First, we adapt \mathcal{A} to keep track of the maximal color visited during a run. Let $\mathcal{A} = (Q, \Sigma_I \times \Sigma_O, q_I, \delta, \Omega)$ with $\Omega: Q \rightarrow \mathbb{N}$ and let $n = |Q \times \Omega(Q)|$. We define the color-tracking automaton $\mathcal{C} = (Q_{\mathcal{C}}, \Sigma_I \times \Sigma_O, q_I^{\mathcal{C}}, \delta_{\mathcal{C}}, \Omega_{\mathcal{C}})$ of size n where

- $Q_{\mathcal{C}} = Q \times \Omega(Q)$,
- $q_I^{\mathcal{C}} = (q_I, \Omega(q_I))$,
- $\delta_{\mathcal{C}}((q, c), a) = (\delta(q, a), \max\{c, \Omega(\delta(q, a))\})$, and
- $\Omega_{\mathcal{C}}(q, c) = c$.

Note that \mathcal{C} does not recognize $L_p(\mathcal{A})$. However, we are not interested in complete runs of \mathcal{C} , but only in runs on finite play infixes.

Remark 4.5. Let $w \in (\Sigma_I \times \Sigma_O)^*$ and let $(q_0, c_0)(q_1, c_1) \cdots (q_{|w|}, c_{|w|})$ be the run of \mathcal{C} on w from some state $(q_0, c_0) \in \{(q, \Omega(q)) \mid q \in Q\}$. Then, $q_0 q_1 \cdots q_{|w|}$ is the run of \mathcal{A} on w starting in q_0 and $c_{|w|} = \max\{\Omega(q_j) \mid 0 \leq j \leq |w|\}$.

In the following, we work with partial functions from $Q_{\mathcal{C}}$ to $2^{Q_{\mathcal{C}}}$, where we denote the domain of each such function r by $\text{dom}(r)$. Intuitively, we use r to capture the information encoded in the lookahead provided by Player I . Assume Player I has picked $\alpha(0) \cdots \alpha(j)$ and Player O has picked $\beta(0) \cdots \beta(i)$ for $i < j$ such that the lookahead is $w = \alpha(i+1) \cdots \alpha(j)$. Then, we can determine the state q that \mathcal{C} reaches after processing $\binom{\alpha(0)}{\beta(0)} \cdots \binom{\alpha(i)}{\beta(i)}$, but the automaton cannot process w , since Player O has not yet picked $\beta(i+1) \cdots \beta(j)$. However, we can determine the states Player O can enforce by picking an appropriate completion, which will be the ones contained in $r(q)$.

To formalize the functions capturing the lookahead picked by Player I , we define $\delta_{\mathcal{P}}: 2^{Q_{\mathcal{C}}} \times \Sigma_I \rightarrow 2^{Q_{\mathcal{C}}}$ via $\delta_{\mathcal{P}}(S, a) = \bigcup_{q \in S} \bigcup_{b \in \Sigma_O} \delta_{\mathcal{C}}(q, \binom{a}{b})$, i.e., $\delta_{\mathcal{P}}$ is the transition function of the powerset automaton of $\text{pr}_I(\mathcal{C})$. As usual, we extend $\delta_{\mathcal{P}}$ to $\delta_{\mathcal{P}}^*: 2^{Q_{\mathcal{C}}} \times \Sigma_I^* \rightarrow 2^{Q_{\mathcal{C}}}$ via $\delta_{\mathcal{P}}^*(S, \varepsilon) = S$ and $\delta_{\mathcal{P}}^*(S, wa) = \delta_{\mathcal{P}}(\delta_{\mathcal{P}}^*(S, w), a)$.

Let $D \subseteq Q_{\mathcal{C}}$ be non-empty and let $w \in \Sigma_I^*$. We define the function r_w^D with domain D as follows: for every $(q, c) \in D$, we have

$$r_w^D(q, c) = \delta_{\mathcal{P}}^*(\{(q, \Omega(q))\}, w).$$

Note that we apply $\delta_{\mathcal{P}}^*$ to $\{(q, \Omega(q))\}$, i.e., the second component is the color of q and not the color c from the argument to r_w^D . If $(q', c') \in r_w^D(q, c)$, then there is a word w' whose projection is w and such that the run of \mathcal{A} on w' leads from q to q' and has maximal color c' . Thus, if Player I has picked the lookahead w , then Player O could pick an answer such that the combined word leads \mathcal{A} from q to q' with maximal color c' .

We call w a witness for a partial function $r: Q_{\mathcal{C}} \rightarrow 2^{Q_{\mathcal{C}}}$, if we have $r = r_w^{\text{dom}(r)}$. Thus, we obtain a language $W_r \subseteq \Sigma_I^*$ of witnesses for each such function r . We define

$$\mathfrak{R} = \{r \mid \text{dom}(r) \neq \emptyset \text{ and } W_r \text{ is infinite}\}.$$

Lemma 4.6. *Let \mathfrak{R} be defined as above.*

- (1) Let $r \in \mathfrak{R}$. Then, $r(q) \neq \emptyset$ for every $q \in \text{dom}(r)$.
- (2) Let $r \neq r' \in \mathfrak{R}$ such that $\text{dom}(r) = \text{dom}(r')$. Then, $W_r \cap W_{r'} = \emptyset$.
- (3) Let r be a partial function from $Q_{\mathcal{C}}$ to $2^{Q_{\mathcal{C}}}$ with non-empty domain. Then, W_r is recognized by a deterministic finite automaton with at most 2^{n^2} states.
- (4) Let $D \subseteq Q_{\mathcal{C}}$ be non-empty and let $w \in \Sigma_I^*$ be such that $|w| \geq 2^{n^2}$. Then, there exists some $r \in \mathfrak{R}$ with $\text{dom}(r) = D$ and $w \in W_r$.

Proof. The first statement follows from completeness of the automata \mathcal{A} , \mathcal{C} , and \mathcal{P} while the second one follows from the definition of r_w^D , which is uniquely determined by w and D . Hence, a fixed w cannot witness two different functions r and r' with the same domain.

To prove the third statement, fix some partial function r from $Q_{\mathcal{C}}$ to $2^{Q_{\mathcal{C}}}$ with domain $D = \{(q_1, c_1), \dots, (q_{|D|}, c_{|D|})\}$. Then, the product of $|D|$ copies of the automaton \mathcal{P} with the initial state $\{(q_1, \Omega(q_1)), \dots, (q_{|D|}, \Omega(q_{|D|}))\}$ and the unique accepting state $(r(q_1, c_1), \dots, r(q_{|D|}, c_{|D|}))$ recognizes the witness language W_r . As $|D| \leq n$, the automaton has at most 2^{n^2} states.

For proving the last statement, we fix some non-empty D and some w of length at least 2^{n^2} . Define $r = r_w^D$, which implies $w \in W_r$ by definition. As just shown, there exists an automaton recognizing W_r with at most $2^{n^2} \leq |w|$ many states. Thus, the accepting run of the automaton on w contains a state-repetition. Hence, W_r is infinite, i.e., $r \in \mathfrak{R}$. \square

Now, we are able to define the equivalent delay-free parity game. As already alluded to, Player I picks elements from \mathfrak{R} while Player O produces a run on witnesses, which corresponds to picking suitable completions to witnesses of the functions picked by Player I . She wins, if the constructed run is accepting. Finally, to account for the lookahead, Player I is always one move ahead.

To keep the equivalence proof between the two games simple, we first give an abstract description of the delay-free game. Then, in the proof of our main theorem, we show how to model this game as a classical graph based parity game, which is solvable in exponential time.

Formally, the game $\mathcal{G}(\mathcal{A})$ is played between Player I and Player O in rounds $i = 0, 1, 2, \dots$ as follows: in each round, Player I picks a function from \mathfrak{R} and Player O answers by a state of \mathcal{C} subject to the following constraints. In the first round, Player I has to pick $r_0 \in \mathfrak{R}$ such that

$$\text{dom}(r_0) = \{q_I^{\mathcal{C}}\} \tag{C1}$$

and Player O has to answer by picking a state $q_0 \in \text{dom}(r_0)$, which implies $q_0 = q_I^{\mathcal{C}}$. Now, consider round $i > 0$: Player I has picked functions r_0, r_1, \dots, r_{i-1} and Player O has picked states q_0, q_1, \dots, q_{i-1} with $q_{i-1} \in \text{dom}(r_{i-1})$. Next, Player I has to pick a function $r_i \in \mathfrak{R}$ such that

$$\text{dom}(r_i) = r_{i-1}(q_{i-1}). \tag{C2}$$

Afterwards, Player O picks some state $q_i \in \text{dom}(r_i)$.

Both players can always move: Player I can move, as $r_{i-1}(q_{i-1})$ is always non-empty (Lemma 4.6.1) and thus the domain of some $r \in \mathfrak{R}$ (Lemma 4.6.4), Player O can move, as the domain of every $r \in \mathfrak{R}$ is non-empty by construction. The resulting play of $\mathcal{G}(\mathcal{A})$ is the sequence $r_0 q_0 r_1 q_1 r_2 q_2 \dots$, which is won by Player O if the maximal color occurring infinitely often in $\Omega_{\mathcal{C}}(q_0) \Omega_{\mathcal{C}}(q_1) \Omega_{\mathcal{C}}(q_2) \dots$ is even. Otherwise, Player I wins.

A strategy for Player I is a function τ_I' mapping the empty play prefix to a function r_0 satisfying (C1) and mapping a non-empty prefix $r_0 q_0 \dots r_{i-1} q_{i-1}$ to a function r_i satisfying

(C2). A strategy τ'_O for Player O maps a play prefix $r_0q_0 \cdots r_i$ to a state $q_i \in \text{dom}(r_i)$. A play $r_0q_0r_1q_1r_2q_2 \cdots$ is consistent with τ'_I , if $r_i = \tau'_I(r_0q_0 \cdots r_{i-1}q_{i-1})$ for every $i \in \mathbb{N}$ and it is consistent with τ'_O , if $q_i = \tau'_O(r_0q_0 \cdots r_i)$ for every $i \in \mathbb{N}$. A strategy τ' for Player $P \in \{I, O\}$ is winning, if every play that is consistent with τ' is winning for Player P . As usual, we say that a player wins $\mathcal{G}(\mathcal{A})$, if she has a winning strategy.

First, we show that the original delay game and the abstract game $\mathcal{G}(\mathcal{A})$ have the same winner. To this end, we fix a winning strategy for Player O in the delay game and simulate it in $\mathcal{G}(\mathcal{A})$, and vice versa, by translating real moves of Player I into abstract moves from \mathfrak{R} , and vice versa.

Lemma 4.7. *Player O wins $\Gamma_f(L_p(\mathcal{A}))$ for some delay function f if and only if Player O wins $\mathcal{G}(\mathcal{A})$.*

Proof. For the sake of readability, we denote $\mathcal{G}(\mathcal{A})$ by \mathcal{G} and $\Gamma_f(L_p(\mathcal{A}))$ by Γ , whenever f is clear from the context.

“ \Rightarrow ”: Let τ_O be a winning strategy for Player O in $\Gamma_f(L_p(\mathcal{A}))$ for some f , which we can assume² to be a constant [9]. We construct a winning strategy τ'_O for Player O in \mathcal{G} by simulating a play of \mathcal{G} by a play of Γ .

Let r_0 be the first move of Player I in \mathcal{G} , which has to be answered by Player O by picking $\tau'_O(r_0) = q_0^C$, and let r_1 be Player I 's response. As W_{r_0} and W_{r_1} are infinite by definition, we can choose witnesses $w_0 \in W_{r_0}$ and $w_1 \in W_{r_1}$ such that $f(0) \leq |w_0| \leq |w_1|$. We simulate this play prefix in Γ : Player I picks $w_0w_1 = \alpha(0) \cdots \alpha(\ell_1 - 1)$ in his first moves, which is enough to play at least $|w_0|$ rounds: w_0 is long enough to play the first round and w_1 is even long enough for at least the next $|w_0|$ rounds.

Now, let $\beta(0) \cdots \beta(\ell_1 - f(0))$ be the response of Player O according to τ_O . As we have played at least $|w_0|$ rounds, we obtain $|\beta(0) \cdots \beta(\ell_1 - f(0))| \geq |w_0|$. Thus, we are in the following situation for $i = 1$:

- in \mathcal{G} , we have constructed a play prefix $r_0q_0 \cdots r_{i-1}q_{i-1}r_i$,
- in Γ , Player I has picked $w_0 \cdots w_i = \alpha(0) \cdots \alpha(\ell_i - 1)$ and Player O has picked $\beta(0) \cdots \beta(\ell_i - f(0))$ according to τ_O . The moves of the players satisfy $|\beta(0) \cdots \beta(\ell_i - f(0))| \geq |w_0 \cdots w_{i-1}|$.
- Finally, w_j is a witness for r_j for every $j \leq i$.

Now, let $i \geq 1$ be arbitrary and let $q_{i-1} = (q'_{i-1}, c_{i-1})$. We define q_i to be the state of \mathcal{C} that is reached from $(q'_{i-1}, \Omega(q'_{i-1}))$ after processing w_{i-1} and the corresponding moves of Player O , i.e.,

$$\left(\begin{array}{c} \alpha(|w_0 \cdots w_{i-2}|) \\ \beta(|w_0 \cdots w_{i-2}|) \end{array} \right) \cdots \left(\begin{array}{c} \alpha(|w_0 \cdots w_{i-1}| - 1) \\ \beta(|w_0 \cdots w_{i-1}| - 1) \end{array} \right).$$

By definition of r_{i-1} , we have $q_i \in r_{i-1}(q_{i-1})$. Correspondingly, we can define

$$\tau'_O(r_0q_0 \cdots r_{i-1}q_{i-1}r_i) = q_i.$$

Now, let r_{i+1} be the next move of Player I in \mathcal{G} and let $w_{i+1} \in W_{r_{i+1}}$ be a witness with $|w_{i+1}| \geq |w_i|$. In Γ , let Player I pick $w_{i+1} = \alpha(\ell_i) \cdots \alpha(\ell_{i+1} - 1)$ as his next moves and let Player O respond by $\beta(\ell_i - f(0) + 1) \cdots \beta(\ell_{i+1} - f(0))$ according to τ_O . Thus, we are again in the aforementioned situation for $i + 1$, which concludes the definition of τ'_O .

It remains to show that τ'_O is winning. Consider a play $r_0q_0r_1q_1r_2q_2 \cdots$ that is consistent with τ'_O and let $w = \begin{pmatrix} \alpha(0) \\ \beta(0) \end{pmatrix} \begin{pmatrix} \alpha(1) \\ \beta(1) \end{pmatrix} \begin{pmatrix} \alpha(2) \\ \beta(2) \end{pmatrix} \cdots$ be the outcome in Γ constructed during the

²This assumption simplifies the proof, but can be avoided at the expense of introducing intricate notation.

simulation as defined above. Note that $\alpha(0)\alpha(1)\alpha(2)\cdots$ is equal to $w_0w_1w_2\cdots$, where each w_i is a witness of r_i . Furthermore, let $q_i = (q'_i, c_i)$.

A straightforward inductive application of Remark 4.5 shows that q'_{i+1} is the state that \mathcal{A} reaches after processing w_i and the corresponding moves of Player O starting in q'_i and that c_{i+1} is the maximal color seen on the run. Thus, the maximal color visited infinitely often by \mathcal{A} after processing w is the same as the maximal color of the sequence $c_0c_1c_2\cdots$, which is even, as w is consistent with a winning strategy and therefore accepted by \mathcal{A} . Hence, $r_0q_0r_1q_1r_2q_2\cdots$ is winning for Player O and τ'_O is a winning strategy.

“ \Leftarrow ”: Let τ'_O be a winning strategy for Player O in \mathcal{G} . We construct a winning strategy τ_O for her in $\Gamma_f(L_p(\mathcal{A}))$ for the constant delay function f with $f(0) = 2d$, where $d = 2^{n^2}$. The strategy τ_O is again constructed by simulating a play of Γ by a play of \mathcal{G} .

In the following, both players pick their moves in Γ in blocks of length d . We denote Player I 's blocks by \bar{a}_i and Player O 's blocks by \bar{b}_i , i.e., every \bar{a}_i is in Σ_I^d and every \bar{b}_i is in Σ_O^d .

Let $\bar{a}_0\bar{a}_1$ be the first move of Player I in Γ , define $q_0 = q_I^c$, and let $r_0 = r_{\bar{a}_0}^{\{q_0\}}$ and $r_1 = r_{\bar{a}_1}^{r_0(q_0)}$. Then, $r_0q_0r_1$ is a play prefix in \mathcal{G} that is consistent with τ'_O . Thus, we are in the following situation for $i = 1$:

- in Γ , Player I has picked blocks $\bar{a}_0 \cdots \bar{a}_i$ and Player O has picked $\bar{b}_0 \cdots \bar{b}_{i-2}$,
- in \mathcal{G} we have constructed a play prefix $r_0q_0 \cdots r_{i-1}q_{i-1}r_i$ that is consistent with τ'_O , and
- \bar{a}_j is a witness for r_j for every $j \leq i$.

Now, let $i \geq 1$ be arbitrary and $q_i = \tau'_O(r_0q_0 \cdots r_{i-1}q_{i-1}r_i)$. The rules of \mathcal{G} imply $q_i \in \text{dom}(r_i) = r_{i-1}(q_{i-1})$. Furthermore, as \bar{a}_{i-1} is a witness for r_{i-1} , there is some \bar{b}_{i-1} such that the automaton \mathcal{C} reaches q_i after processing $\left(\frac{\bar{a}_{i-1}}{\bar{b}_{i-1}}\right)$ from $(q'_{i-1}, \Omega(q'_{i-1}))$, where $q_{i-1} = (q'_{i-1}, c_{i-1})$. Player O 's strategy for Γ is to pick the letters of \bar{b}_{i-1} in the next d rounds. These are answered by Player I by d letters forming \bar{a}_{i+1} . This way, we obtain $r_{i+1} = r_{\bar{a}_{i+1}}^{r_i(q_i)}$ bringing us back to the aforementioned situation for $i + 1$, which concludes the definition of τ_O .

It remains to show that τ_O is winning for Player O . Let $w = \left(\frac{\bar{a}_0}{\bar{b}_0}\right)\left(\frac{\bar{a}_1}{\bar{b}_1}\right)\left(\frac{\bar{a}_2}{\bar{b}_2}\right)\cdots$ be the outcome of a play of Γ consistent with τ_O . Also, let $r_0q_0r_1q_1r_2q_2\cdots$ be the corresponding play of \mathcal{G} constructed as described in the simulation above, where each \bar{a}_i is a witness for r_i . Finally, let $q_i = (q'_i, c_i)$.

A straightforward inductive application of Remark 4.5 shows that q'_{i+1} is the state reached in \mathcal{A} after processing $\left(\frac{\bar{a}_i}{\bar{b}_i}\right)$ starting in q'_i and that c_{i+1} is the largest color seen on this run. As $r_0q_0r_1q_1r_2q_2\cdots$ is consistent with τ'_O , the sequence $\Omega_{\mathcal{C}}(q_0)\Omega_{\mathcal{C}}(q_1)\Omega_{\mathcal{C}}(q_2)\cdots = c_0c_1c_2\cdots$ satisfies the parity condition, i.e., the maximal color occurring infinitely often is even. Thus, the maximal color occurring infinitely often during the run of \mathcal{A} on w is even as well, i.e., w is winning for Player O . Thus, τ_O is a winning strategy for Player O . \square

Now, we can prove the main theorem of this section: determining whether Player O wins a delay game induced by a given parity automaton \mathcal{A} for some f is in EXPTIME. To this end, it suffices to model $\mathcal{G}(\mathcal{A})$ as a classical parity game and to show that it can be constructed and solved in exponential time. This is sufficient due to Lemma 4.7.

Proof of Theorem 4.4. First, we argue that \mathfrak{R} can be constructed in exponential time: to this end, one constructs for every partial function r from $Q_{\mathcal{C}}$ to 2^{Q_c} the automaton of

Lemma 4.6.3 recognizing W_r and tests it for recognizing an infinite language. There are exponentially many functions and each automaton is of exponential size, which yields the desired result.

Now, we can encode $\mathcal{G}(\mathcal{A})$ as a graph-based parity game³ $((V, V_I, V_O, E), \Omega')$ with $\Omega': V \rightarrow \mathbb{N}$ where

- $V = V_I \cup V_O$,
- $V_I = \{v_I\} \cup \mathfrak{X} \times Q_C$ for some fresh initial vertex v_I ,
- $V_O = \mathfrak{X}$,
- E is the union of the following sets of edges:
 - $\{(v_I, r) \mid \text{dom}(r) = \{q_I^c\}\}$: the initial moves of Player I .
 - $\{((r, q), r') \mid \text{dom}(r') = r(q)\}$: (regular) moves of Player I .
 - $\{(r, (r, q)) \mid q \in \text{dom}(r)\}$: moves of Player O , and
- $\Omega'(v) = \begin{cases} c & \text{if } v = (r, (q, c)) \in \mathfrak{X} \times Q_C, \\ 0 & \text{otherwise.} \end{cases}$

Then, Player O wins $\mathcal{G}(\mathcal{A})$ if and only if she has a winning strategy from v_I in the explicit parity game.

A parity game with n vertices, m edges, and k colors can be solved in time $\mathcal{O}(mn^{\frac{k}{3}})$ [22]. The parity game constructed above has at most $\mathcal{O}(2^{|\mathcal{A}|^4} \cdot |\mathcal{A}|^2)$ vertices and at most $|\mathcal{A}|$ colors. Hence, it can be solved in exponential time in the size of \mathcal{A} . \square

By applying both directions of the equivalence between $\Gamma_f(L_p(\mathcal{A}))$ and $\mathcal{G}(\mathcal{A})$, we obtain an exponential upper bound of the form $2^{(|\mathcal{A}|k)^2+1}$ on the lookahead necessary for Player O to win a delay game with a parity condition over k colors. However, this upper bound on the lookahead requires the introduction of the game $\mathcal{G}(\mathcal{A})$. To conclude this section, we present an alternative proof yielding a slightly better upper bound on the necessary lookahead for Player O in case she wins. The proof relies on a direct pumping argument for winning strategies for Player I : winning for the constant delay function f with $f(0) \leq 2^{|\mathcal{A}|^2k+1}$ allows him to cover arbitrarily larger lookahead⁴.

Theorem 4.8. *Let $L = L_p(\mathcal{A})$ where \mathcal{A} is a deterministic parity automaton with k colors. The following are equivalent:*

- (1) *Player O wins $\Gamma_f(L)$ for some constant delay function f with $f(0) \leq 2^{|\mathcal{A}|^2k+1}$.*
- (2) *Player O wins $\Gamma_{f'}(L)$ for some delay function f' .*

Proof. We only consider the non-trivial implication (2) \Rightarrow (1) and prove the contrapositive by turning a winning strategy for Player I in $\Gamma_f(L)$ into a winning strategy for games with arbitrarily larger lookahead. We use a pumping argument to show that every sufficiently large lookahead contains some repetitive behavior, which we pump to cover the larger lookahead. To this end, we construct a simulation of the original game to inductively define the new strategy for Player I . The simulation proceeds in big steps, i.e., the players pick sufficiently large blocks of input and output letters, respectively, which allows us to identify the repetitive behaviour in each such block. The structure of the simulation is similar to the game $\mathcal{G}(\mathcal{A})$ from the proof of Theorem 4.4, where Player I is always two functions ahead, after he moved, while Player O responds only to the first of these two functions. In this

³For a complete definition see, e.g., [7]

⁴The previous version of this paper claimed a bound of $2^{|\mathcal{A}|k+1}$, but the proof was incomplete. Here, we present a corrected proof for the bound $2^{|\mathcal{A}|^2k+1}$.

alternative proof, we do not need an abstract representation of the lookahead, which is why we represent it just by concrete sequences of input letters.

As in the previous proof, simulating the game in big steps requires us to extend the automaton recognizing L to accumulate the maximal color along a run. Let $\mathcal{A} = (Q, \Sigma_I \times \Sigma_O, q_I, \delta, \Omega)$ be given and let $\mathcal{C} = (Q_C, \Sigma_I \times \Sigma_O, q_I^C, \delta_C, \Omega_C)$ be the color-tracking extension of \mathcal{A} as defined before.

Now, let τ_I be a winning strategy for Player I for $\Gamma_f(L)$, let f' be an arbitrary delay function. Due to the results of Holtmann et al. [9], we can assume w.l.o.g. that f' is constant, which simplifies the proof, but is not essential. We define $d = 2^{|\mathcal{A}|^2 \cdot k}$, i.e., we have $f(0) = 2d$. We construct a winning strategy τ_I' for Player I in $\Gamma_{f'}(L)$ by simulating a play of $\Gamma_{f'}(L)$ via a play of $\Gamma_f(L)$. In the following, we will group the moves of the players in $\Gamma_{f'}(L)$ into blocks of length d . We denote Player I 's blocks by $\overline{a_i}$ and Player O 's blocks by $\overline{b_i}$. For the sake of readability, we denote $\Gamma_f(L)$ by Γ and $\Gamma_{f'}(L)$ by Γ' .

To begin, we have to give the first moves of Player I in Γ' . To this end, consider the first move $\tau_I(\varepsilon) = \overline{a_0 a_1}$ of Player I in Γ . By the choice of d and the fact that \mathcal{A} is complete, we can decompose $\overline{a_0}$ into $x_0 y_0 z_0$ with non-empty y_0 such that for every $q \in Q$

$$\delta_{\mathcal{P}}^*(\{(q, \Omega(q))\}, x_0) = \delta_{\mathcal{P}}^*(\{(q, \Omega(q))\}, x_0 y_0).$$

Similarly, we can decompose $\overline{a_1}$ into $x_1 y_1 z_1$ with non-empty y_1 such that for every $q \in Q$

$$\delta_{\mathcal{P}}^*(\{(q, \Omega(q))\}, x_1) = \delta_{\mathcal{P}}^*(\{(q, \Omega(q))\}, x_1 y_1).$$

Now, we pump y_0 sufficiently often to simulate the initial block in Γ' , i.e., let $\alpha(0) \cdots \alpha(\ell_0 - 1) = x_0 (y_0)^{h_0} z_0$ for the smallest $h_0 \geq 1$ such that $|x_0 (y_0)^{h_0} z_0| \geq f'(0)$. Similarly, we define $\alpha(\ell_0) \cdots \alpha(\ell_1 - 1) = x_1 (y_1)^{h_1} z_1$ for the smallest $h_1 \geq 1$ with $|x_1 (y_1)^{h_1} z_1| \geq |x_0 (y_0)^{h_0} z_0|$.

Now, we define τ_I' such that Player I picks $\alpha(0) \cdots \alpha(\ell_1 - 1)$ with his first moves in Γ' , which is answered by Player O by picking $\beta(0) \cdots \beta(\ell_1 - f'(0))$. By the choice of h_1 , we obtain $(\ell_1 - f'(0)) + 1 \geq \ell_0$.

Now, we are in the following situation for $i = 1$:

- In Γ' , Player I has picked $\alpha(0) \cdots \alpha(\ell_i - 1)$ such that for every $j \leq i$: $\alpha(\ell_{j-1}) \cdots \alpha(\ell_j - 1) = x_j (y_j)^{h_j} z_j$.
- In Γ' , Player O has picked $\beta(0) \cdots \beta(\ell_i - f'(0))$ such that $(\ell_i - f'(0)) + 1 \geq \ell_{i-1}$. Thus, Player O has provided an answer to $\alpha(\ell_{i-2}) \cdots \alpha(\ell_{i-1} - 1) = x_{i-1} (y_{i-1})^{h_{i-1}} z_{i-1}$.
- In Γ , Player I has picked blocks $\overline{a_0} \cdots \overline{a_i}$ and Player O has picked $\overline{b_0} \cdots \overline{b_{i-2}}$.

Now, let $i \geq 1$ be arbitrary and let q_{i-1} be the state reached by \mathcal{A} after processing $\left(\frac{\overline{a_0}}{b_0}\right) \cdots \left(\frac{\overline{a_{i-2}}}{b_{i-2}}\right)$. Further, let (q_{i-1}^*, c_{i-1}^*) be the state reached by \mathcal{C} after processing $x_{i-1} (y_{i-1})^{h_{i-1}-1}$ and the corresponding letters picked by Player O at these positions starting in $(q_{i-1}, \Omega(q_{i-1}))$. Due to the state repetition induced by the decomposition, which is independent of the starting state $(q, \Omega(q))$, we can find a word $x'_{i-1} \in \Sigma_O^{|\overline{a_{i-1}}|}$ such that \mathcal{C} reaches the same state (q_{i-1}^*, c_{i-1}^*) after processing $\left(\frac{x'_{i-1}}{x'_{i-1}}\right)$, when starting in $(q_{i-1}, \Omega(q_{i-1}))$.

We define $\overline{b_{i-1}} = x'_{i-1} y'_{i-1} z'_{i-1}$, where y'_{i-1} and z'_{i-1} are the letters picked by Player O at the positions of the last repetition of y_{i-1} and at the positions of z_{i-1} , respectively.

We continue the simulation in Γ by letting Player O pick the block $\overline{b_{i-1}}$ during the next d rounds, which is answered by Player I by picking the next block $\overline{a_{i+1}}$. Again, we can decompose $\overline{a_{i+1}}$ into $x_{i+1} y_{i+1} z_{i+1}$ with non-empty y_{i+1} such that for every $q \in Q$

$$\delta_{\mathcal{P}}^*(\{(q, \Omega(q))\}, x_1) = \delta_{\mathcal{P}}^*(\{(q, \Omega(q))\}, x_1 y_1).$$

As before, we define $\alpha(\ell_i) \cdots \alpha(\ell_{i+1} - 1) = x_{i+1}(y_{i+1})^{h_{i+1}} z_{i+1}$ for the smallest $h_{i+1} \geq 1$ such that $|x_{i+1}(y_{i+1})^{h_{i+1}} z_{i+1}| \geq |x_i(y_i)^{h_i} z_i|$. Now, we define τ'_I such that Player I picks $\alpha(\ell_i) \cdots \alpha(\ell_{i+1} - 1)$ with his next moves in Γ' , which is answered by Player O by picking $\beta((\ell_i - f'(0)) + 1) \cdots \beta(\ell_{i+1} - f'(0))$. By the choice of h_{i+1} , we obtain $(\ell_{i+1} - f'(0)) + 1 \geq \ell_i$. Thus, we are in the situation described above for $i + 1$, which completes the definition of τ'_I .

It remains to show that τ'_I is winning. Let $w' = \binom{\alpha(0)}{\beta(0)} \binom{\alpha(1)}{\beta(1)} \binom{\alpha(2)}{\beta(2)} \cdots$ be an outcome of a play that is consistent with τ'_I in Γ' and let $w = \binom{a_0}{b_0} \binom{a_1}{b_1} \binom{a_2}{b_2} \cdots$ be the simulating play of Γ as described above. Fix $(q_0, c_0) = (q_I, \Omega(q_I))$ and let (q_i, c_i) for $i \geq 1$ be the state reached by \mathcal{C} when processing

$$\binom{\alpha(\ell_{i-2})}{\beta(\ell_{i-2})} \cdots \binom{\alpha(\ell_{i-1} - 1)}{\beta(\ell_{i-1} - 1)},$$

starting in $(q_{i-1}, \Omega(q_{i-1}))$, using $\ell_{-1} = 0$. By construction, (q_i, c_i) is also the state reached by \mathcal{C} when processing $\binom{a_{i-1}}{b_{i-1}}$ starting in $(q_{i-1}, \Omega(q_{i-1}))$. Thus, applying Remark 4.5 inductively shows that \mathcal{A} accepts w if, and only if, it accepts w' . Since w is consistent with the winning strategy τ_I both are not accepted, i.e., τ'_I is indeed a winning strategy. \square

To conclude, let us briefly discuss the gap in the argumentation of the original proof. There, when decomposing the second block \bar{a}_1 , we only asked for a state repetition from the state of the powerset automaton reached after processing the first block \bar{a}_0 . However, this does not guarantee that the state (q_1^*, c_1^*) as defined above is reachable from $(q_1, \Omega(q_1))$, the state reached after processing $\binom{a_0}{b_0}$ from $(q_I, \Omega(q_I))$. One only obtains that there is a block \bar{b}_0 such that (q_1^*, c_1^*) is reachable from the state reached after processing $\binom{a_0}{b_0}$ from $(q_I, \Omega(q_I))$. Hence, if $\bar{b}_0^* \neq \bar{b}_0$, then the inductive argument breaks down. Having a state repetition for every possible such starting state resolves this issue.

5. WINNING CONDITIONS THAT ARE REACHABILITY AND SAFETY

Recall the exponential lower bounds on the necessary lookahead for reachability and safety conditions presented in Section 3. Both rely on the same construction and one can even turn the deterministic safety automata \mathcal{A}'_n exhibiting the lower bound into reachability automata exhibiting the same lower bound. To this end, one changes the set of accepting states, but leaves the transition structure of \mathcal{A}'_n unchanged. Nevertheless, these automata do not accept the same language. In the following, we prove that this is unavoidable: linear lookahead is necessary and sufficient for winning conditions that are both reachability and safety. Furthermore, solving such games is CO-NP-complete and thus simpler than solving delay games with general reachability or safety conditions (under standard complexity-theoretic assumptions).

Winning conditions that are both reachability and safety conditions are also known as ω -regular clopen conditions, as they are those ω -regular languages that are both closed and open in the Cantor topology. Note that the way we represent such conditions is a non-trivial issue that has influences on our results: we could specify an ω -regular clopen language either by a reachability or by a safety automaton or even give a pair of equivalent automata, one for each acceptance condition. However, we take another approach. A clopen language is fully characterized by a finite language of finite words. Thus, we use acyclic finite automata on finite words to represent ω -regular clopen winning conditions. For every deterministic

reachability and every deterministic safety automaton recognizing a clopen language there is an acyclic finite automaton of the same size that represents the same language, i.e., the change in representation does not incur a blowup, when starting with deterministic automata. The results mentioned above pertain to the representation of winning conditions via acyclic finite automata.

First, we formalize the characterization of ω -regular clopen languages by acyclic finite automata. In the following, we only consider (w.l.o.g.) automata whose states are all reachable from the initial state. Also, recall that we require all our automata to be complete. A state of an automaton is productive, if an accepting state is reachable from it, otherwise it is non-productive. In particular, every non-productive state is non-accepting. Finite languages of finite words are recognized by acyclic finite automata, i.e., automata that satisfy the following property: if q is on a cycle, then q is non-productive. The depth of an acyclic automaton \mathcal{A} is the length of the longest path from an initial to an accepting state, which is bounded by $|\mathcal{A}| - 1$.

Remark 5.1. Let \mathcal{A} be a non-deterministic automaton. Then, $L_*(\mathcal{A})$ is finite if and only if \mathcal{A} is acyclic.

By merging all non-productive states of a non-deterministic acyclic automaton into a single non-accepting sink, we obtain an equivalent acyclic automaton of the same depth where processing a word that is longer than the depth leads to the non-accepting sink.

Remark 5.2. Let \mathcal{A} be a non-deterministic acyclic automaton of depth k . Then, there is a non-deterministic acyclic automaton \mathcal{A}' of depth k with $|\mathcal{A}'| \leq |\mathcal{A}|$, $L_*(\mathcal{A}) = L_*(\mathcal{A}')$, and such that processing a word of length greater than k leads \mathcal{A}' to a (non-accepting) sink.

Now, we are able to state the characterization theorem for ω -regular clopen languages.

Lemma 5.3 ([16]). *Let $L \subseteq \Sigma^\omega$ be an ω -regular clopen language that is recognized by a deterministic automaton \mathcal{A} with reachability or safety acceptance. Then, $L = L_*(\mathcal{A}') \cdot \Sigma^\omega$ for some deterministic acyclic automaton \mathcal{A}' with $|\mathcal{A}'| \leq |\mathcal{A}|$.*

With this characterization at hand, we determine the complexity of solving delay games with ω -regular clopen winning conditions and then present tight upper and lower bounds on the necessary lookahead.

Theorem 5.4. *The following problem is CO-NP-complete: Given an acyclic automaton \mathcal{A} over $\Sigma_I \times \Sigma_O$, does Player O win $\Gamma_f(L_*(\mathcal{A}) \cdot (\Sigma_I \times \Sigma_O)^\omega)$ for some f ?*

The proof of the preceding theorem is split into the following two lemmata. We begin by showing CO-NP-hardness, which already holds for deterministic automata. Then, we show CO-NP-membership, which also holds for non-deterministic automata.

The following hardness proof is reminiscent of the proof of Theorem 4.3, but simpler since we only have to deal with universal polynomial time machines instead of alternating polynomial space machines.

Lemma 5.5. *The following problem is CO-NP-hard: Given an acyclic deterministic automaton \mathcal{A} over $\Sigma_I \times \Sigma_O$, does Player O win $\Gamma_f(L_*(\mathcal{A}) \cdot (\Sigma_I \times \Sigma_O)^\omega)$ for some f ?*

Proof. We show hardness by a reduction from the acceptance problem for universal polynomial time Turing machines. Let $\mathcal{M} = (Q, \Sigma, q_I, \Delta, q_A, q_R)$ be such a machine, where $\Delta \subseteq Q \times \Sigma \times Q \times \Sigma \times \{-1, 0, 1\}$ is the transition relation. We assume w.l.o.g. that the

accepting state q_A and the rejecting state q_R have self-loops. Let p be a polynomial that bounds the time-consumption of \mathcal{M} . Furthermore, let $x \in \Sigma^*$ be an input for \mathcal{M} .

We construct an acyclic deterministic automaton \mathcal{A} over $\Sigma_I \times \Sigma_O$ such that Player O wins $\Gamma_f(L_*(\mathcal{A}) \cdot (\Sigma_I \times \Sigma_O)^\omega)$ for some f if and only if \mathcal{M} accepts x , i.e., if and only if every run of \mathcal{M} on x is accepting. To this end, we construct a game where Player I produces sequences of configurations to simulate runs of \mathcal{M} on x and Player O can claim errors in this simulation. Player O wins, if she correctly claims an error or if the simulated run is accepting. Thus, we define $\Sigma_I = Q \cup \Sigma$ and $\Sigma_O = \{\checkmark, \boldsymbol{X}\}$. Player O uses \boldsymbol{X} to claim an error.

Let $t = p(|x|)$ and let Conf be the set of encodings of configurations of \mathcal{M} of length t , i.e., words over the alphabet $\Sigma \cup Q$ of length $t + 1$ containing exactly one letter from Q . Furthermore, let $c_0 \in \text{Conf}$ denote the encoding of the initial configuration of \mathcal{M} on x . In the following, we do not distinguish between a configuration and its encoding.

Now, consider the language L of words $\binom{u}{v} \in (\Sigma_I \times \Sigma_O)^*$ such that either

$$u \in \Sigma_I^{t \cdot (t+1)} \setminus c_0 \cdot (\text{Conf})^{t-1}$$

or, if u is in $c_0 \cdot (\text{Conf})^{t-1}$, then u either ends with an accepting configuration or the smallest position marked with \boldsymbol{X} in v indicates a cell which witnesses that the following configuration in u is not a successor configuration of the current one. It is straightforward to show that L is recognized by an automaton \mathcal{A} of polynomial size in $|\mathcal{M}| + t$, which is polynomial in the size of \mathcal{M} and the length of x . To this end, one relies on the fact that L only contains words of a fixed length, which is polynomial in t , and on the error checking routine described in the proof of Theorem 4.3. Finally, as L is finite, Remark 5.1 implies that \mathcal{A} is acyclic.

It remains to show that Player O wins $\Gamma_f(L \cdot (\Sigma_I \times \Sigma_O)^\omega)$ for some delay function f if and only if \mathcal{M} accepts x , i.e., every run of \mathcal{M} on x is accepting.

Thus, let x be accepted by \mathcal{M} , i.e., every run of \mathcal{M} on x reaches an accepting configuration after executing at most $t - 1$ transitions. Hence, in order to win a delay game with winning condition $L \cdot (\Sigma_I \times \Sigma_O)^\omega$, Player I has to introduce an error in the simulation of \mathcal{M} on x . With sufficiently large constant lookahead, Player O can catch him and correctly claim the error. Hence, she wins $\Gamma_f(L \cdot (\Sigma_I \times \Sigma_O)^\omega)$ for some f .

Conversely, assume x is not accepted by \mathcal{M} , i.e., there is a run of \mathcal{M} that does not reach an accepting configuration. Player I can simulate this run in $\Gamma_f(L \cdot (\Sigma_I \times \Sigma_O)^\omega)$ for every f . As he does not introduce an error, he wins the resulting play, i.e., Player O does not win $\Gamma_f(L \cdot (\Sigma_I \times \Sigma_O)^\omega)$ for any f . \square

Next, we show that the problem is in co-NP, which completes the proof of Theorem 5.4

Lemma 5.6. *The following problem is in co-NP: Given an acyclic non-deterministic automaton \mathcal{A} over $\Sigma_I \times \Sigma_O$, does Player O win $\Gamma_f(L_*(\mathcal{A}) \cdot (\Sigma_I \times \Sigma_O)^\omega)$ for some f ?*

Proof. Let $L = L_*(\mathcal{A}) \cdot (\Sigma_I \times \Sigma_O)^\omega$, which is in particular a reachability condition. Thus, Theorem 4.1 shows that Player O wins $\Gamma_f(L)$ if and only if $\text{pr}_I(L)$ is universal. Let k be the depth of \mathcal{A} . As every word of length greater than k leads to a non-accepting sink state of \mathcal{A} , membership in $L_*(\mathcal{A}) \cdot (\Sigma_I \times \Sigma_O)^\omega$ only depends on the prefix of length k . Hence, $\text{pr}_I(L)$ is universal if and only if $\text{pr}_I(\mathcal{A})$ accepts at least one prefix of every word of length $k \leq |\mathcal{A}|$. The latter problem can be solved in polynomial time by a universal Turing machine. \square

After resolving the complexity of delay games with ω -regular clopen winning conditions, we turn our attention to proving lower and upper bounds on the necessary lookahead in such games. We begin by showing that the depth of the automaton recognizing the winning

condition is an upper bound on the necessary lookahead, which follows from the reasoning presented in the proof of Lemma 5.6.

Theorem 5.7. *Let \mathcal{A} be an acyclic non-deterministic automaton over $\Sigma_I \times \Sigma_O$ and let k be the depth of \mathcal{A} . The following are equivalent:*

- (1) *Player O wins $\Gamma_f(L_*(\mathcal{A}) \cdot (\Sigma_I \times \Sigma_O)^\omega)$ for some delay function f .*
- (2) *Player O wins $\Gamma_f(L_*(\mathcal{A}) \cdot (\Sigma_I \times \Sigma_O)^\omega)$ for some constant delay function f with $f(0) \leq k$.*

Proof. We only consider the non-trivial implication (1) \Rightarrow (2). Let $L = L_*(\mathcal{A}) \cdot (\Sigma_I \times \Sigma_O)^\omega$.

Assume Player O wins $\Gamma_f(L)$ for some f . Then, by Theorem 4.1, $\text{pr}_I(L)$ is universal, which implies that $\text{pr}_I(L(\mathcal{A}))$ contains at least one prefix of every word of length k . Hence, for every possible first move $\alpha(0) \cdots \alpha(k-1) \in \Sigma_I^k$ of Player I in $\Gamma_f(L)$, there is a $\beta(0) \cdots \beta(k-1) \in \Sigma_O^k$ such that $\binom{\alpha(x)}{\beta(0)} \cdots \binom{\alpha(k-1)}{\beta(k-1)}$ has a prefix that is accepted by \mathcal{A} . Thus, every continuation of this prefix is in L . We conclude that Player O has a winning strategy for $\Gamma_f(L)$ whenever $f(0)$ is greater or equal than the depth of \mathcal{A} . \square

To conclude this section, we present a matching lower bound.

Theorem 5.8. *For every $n \geq 0$, there is a language L_n over $\Sigma_I = \Sigma_O = \{a, b\}$ such that*

- *$L_n = L_*(\mathcal{A}_n)$ for some deterministic acyclic automaton \mathcal{A}_n with $|\mathcal{A}_n| \in \mathcal{O}(n)$ and depth $n + 1$,*
- *Player O wins $\Gamma_f(L_n \cdot (\Sigma_I \times \Sigma_O)^\omega)$ for some constant delay function f , but*
- *Player I wins $\Gamma_f(L_n \cdot (\Sigma_I \times \Sigma_O)^\omega)$ for every delay function f with $f(0) \leq n$.*

Proof. Consider the language

$$L_n = \left\{ \binom{\alpha(0)}{\beta(0)} \cdots \binom{\alpha(n)}{\beta(n)} \mid \beta(0) = \alpha(n) \right\}.$$

It is straightforward to show that L_n is recognizable by a deterministic acyclic automaton of size $\mathcal{O}(n)$ and depth $n + 1$, and that Player O wins $\Gamma_f(L_n \cdot (\Sigma_I \times \Sigma_O)^\omega)$ if and only if $f(0) > n$. \square

Note that all upper bounds in this section hold for non-deterministic automata while the lower bounds hold for deterministic ones.

To conclude this section, let us come back to the lower bounds for reachability and safety conditions presented in Section 3. The deterministic safety automata \mathcal{A}'_n witnessing the exponential lower bound can be equipped with a reachability acceptance that witnesses an exponential lower bound for reachability conditions. However, these automata do not accept the same language. If they would, then Lemma 5.3 yields an acyclic automaton of depth $\mathcal{O}(n)$ that represents this language. Thus, Theorem 5.7 shows that lookahead $\mathcal{O}(n)$ is sufficient for Player O . However, this contradicts the lower bound proven in Theorem 3.2.

6. CONCLUSION

We gave the first algorithm that solves ω -regular delay games in exponential time, which is an exponential improvement over the previously known algorithms. We complemented this by showing the problem to be EXPTIME-complete, even for safety conditions. Also, we determined the exact amount of lookahead that is necessary to win ω -regular delay games by proving tight exponential bounds, which already hold for safety and reachability

conditions. The EXPTIME-completeness of solving delay games with safety conditions is contrasted by PSPACE-completeness of solving delay games with reachability conditions. Due to this gap in complexity and due to the exponential lower bounds on the lookahead, we also considered delay games with ω -regular clopen winning conditions. Here, linear lookahead suffices and is in general necessary. Furthermore, determining the winner in such games is CO-NP-complete. All our lower and upper bounds hold for deterministic automata while for reachability conditions (including clopen conditions) our results even hold for non-deterministic automata. To the best of our knowledge, these are the first non-trivial lower bounds on lookahead and complexity for delay games.

Thus, we completed the picture for deterministic automata and provided partial results for non-deterministic automata, e.g., for automata with reachability acceptance condition. One can trivially obtain upper bounds for the other types of non-deterministic (and universal) automata using determinization, but this incurs an exponential blowup. In current research, we show this to be unavoidable. Furthermore, for alternating automata, we have recently proven tight triply-exponential bounds on the necessary lookahead and shown 3EXPTIME-completeness of determining the winner of such a game [14]. Here, the lower bounds already hold for LTL specifications.

An open question concerns the influence on the necessary lookahead and the solution complexity when using different deterministic automata models that recognize the class of ω -regular conditions, e.g., Rabin, Streett, and Muller automata. Indeed, our construction used to prove Theorem 4.4 can be adapted to deal with these acceptance conditions, e.g., for conditions given by Muller automata, \mathcal{C} keeps track of the states visited on a run and $\mathcal{G}(\mathcal{A})$ is a Muller game. This yields upper bounds, but it is open whether these are optimal.

Another promising research direction is the study of restricted classes of strategies for delay games, e.g., finite-state strategies for games with ω -regular winning conditions. Our reduction presented in Section 4.3 implies the existence of such strategies, which in general need exponentially many memory states as well as access to the complete lookahead. This result has been proved and a more elegant construction based on the players picking blocks of letters (cf. the proof of Lemma 4.7), was presented by Salzmann [21].

Acknowledgments. We thank Bernd Finkbeiner for a fruitful discussion that lead to Theorem 3.1 and Theorem 4.1. Also, we thank the reviewers for their valuable feedback.

REFERENCES

- [1] J. Richard Büchi and Lawrence H. Landweber. Solving sequential conditions by finite-state strategies. *Trans. Amer. Math. Soc.*, 138:pp. 295–311, 1969.
- [2] Arnaud Carayol and Christof Löding. MSO on the infinite binary tree: Choice and order. In Jacques Duparc and Thomas A. Henzinger, editors, *CSL 2007*, volume 4646 of *LNCS*, pages 161–176. Springer, 2007.
- [3] Arnaud Carayol and Christof Lding. Uniformization in automata theory. In Peter Schroeder-Heister, Gerhard Heinzmann, Wilfrid Hodges, and Pierre Edouard Bour, editors, *CLMPS 2012*, pages 153–178. College Publications, London, 2015.
- [4] Ashok K. Chandra, Dexter Kozen, and Larry J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, 1981.
- [5] E. Allen Emerson and Charanjit S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *FOCS 1991*, pages 368–377. IEEE, 1991.
- [6] Wladimir Fridman, Christof Löding, and Martin Zimmermann. Degrees of lookahead in context-free infinite games. In Marc Bezem, editor, *CSL 2011*, volume 12 of *LIPICs*, pages 264–276. Schloss Dagstuhl - Leibniz-Zentrum fr Informatik, 2011.

- [7] Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *LNCS*. Springer, 2002.
- [8] Yuri Gurevich and Saharon Shelah. Rabin’s uniformization problem. *The Journal of Symbolic Logic*, 48:1105–1119, 1983.
- [9] Michael Holtmann, Lukasz Kaiser, and Wolfgang Thomas. Degrees of lookahead in regular infinite games. *LMCS*, 8(3), 2012.
- [10] Frederick A. Hosch and Lawrence H. Landweber. Finite delay solutions for sequential conditions. In *ICALP 1972*, pages 45–60, 1972.
- [11] Milka Hutagalung, Martin Lange, and Étienne Lozes. Buffered simulation games for Büchi automata. In Zoltán Ésik and Zoltán Fülöp, editors, *AFL 2014*, volume 151 of *EPTCS*, pages 286–300, 2014.
- [12] Felix Klein and Martin Zimmermann. How much lookahead is needed to win infinite games? In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *ICALP 2015 Part II*, volume 9135 of *LNCS*, pages 452–463. Springer, 2015.
- [13] 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.
- [14] Felix Klein and Martin Zimmermann. Prompt delay. *arXiv*, 1602.05045, 2016.
- [15] Orna Kupferman, Nir Piterman, and Moshe Y. Vardi. From liveness to promptness. *Form. Method. Syst. Des.*, 34(2):83–103, 2009.
- [16] Orna Kupferman and Moshe Y. Vardi. On bounded specifications. In Robert Nieuwenhuis and Andrei Voronkov, editors, *LPAR 2001*, volume 2250 of *LNCS*, pages 24–38. Springer, 2001.
- [17] Christof Löding and Sarah Winter. Synthesis of deterministic top-down tree transducers from automatic tree relations. In Adriano Peron and Carla Piazza, editors, *GandALF 2014*, volume 161 of *EPTCS*, pages 88–101, 2014.
- [18] Albert R. Meyer and Larry J. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. In *SWAT 1972*, pages 125–129. IEEE Computer Society, 1972.
- [19] Andrzej Mostowski. Games with forbidden positions. Technical Report 78, University of Gdańsk, 1991.
- [20] Amir Pnueli. The temporal logic of programs. In *FOCS 1977*, pages 46–57. IEEE, 1977.
- [21] Tobias Salzmann. How much memory is needed to win regular delay games? Master’s thesis, Saarland University, 2015.
- [22] Sven Schewe. Solving parity games in big steps. In Vikraman Arvind and Sanjiva Prasad, editors, *FSTTCS 2007*, volume 4855 of *LNCS*, pages 449–460. Springer, 2007.
- [23] A. Prasad Sistla, Moshe Y. Vardi, and Pierre Wolper. The complementation problem for Büchi automata with applications to temporal logic (extended abstract). In Wilfried Brauer, editor, *ICALP 1985*, volume 194 of *LNCS*, pages 465–474. Springer, 1985.
- [24] Wolfgang Thomas. Facets of synthesis: Revisiting church’s problem. In Luca de Alfaro, editor, *FOSSACS*, volume 5504 of *LNCS*, pages 1–14. Springer, 2009.
- [25] Wolfgang Thomas and Helmut Lescow. Logical specifications of infinite computations. In J. W. de Bakker, Willem P. de Roever, and Grzegorz Rozenberg, editors, *A Decade of Concurrency, Reflections and Perspectives, REX School/Symposium*, volume 803 of *LNCS*, pages 583–621. Springer, 1993.
- [26] B.A. Trakhtenbrot and I.M. Barzdin. *Finite Automata; Behavior and Synthesis*. Fundamental Studies in Computer Science, V. 1. North-Holland Publishing Company; New York: American Elsevier, 1973.
- [27] Martin Zimmermann. Delay games with WMSO+U winning conditions. In Lev D. Beklemishev and Daniil V. Musatov, editors, *CSR 2015*, volume 9139 of *LNCS*, pages 412–425. Springer, 2015.
- [28] Martin Zimmermann. Unbounded lookahead in WMSO+U games. *arXiv*, 1509.07495, 2015.

APPENDIX A. UNIVERSALITY OF NON-DETERMINISTIC REACHABILITY AUTOMATA

We show that universality of non-deterministic reachability automata is PSPACE-complete. The hardness proof is a small extension of the classical PSPACE-hardness proof for non-deterministic finite automata [18]. We give the proof here for the sake of completeness, as we could not find it in the literature (only some mentionings of the result).

Theorem A.1. *The universality problem for non-deterministic reachability automata is PSPACE-complete.*

Proof. Membership in PSPACE is straightforward: every reachability automaton can be turned into a Bchi automaton turning the accepting states into sinks such that the resulting automaton recognizes the same language⁵. Universality for Bchi automata is known to be in PSPACE [23].

We show hardness by a reduction from the acceptance problem for polynomial space Turing machines. Let $\mathcal{M} = (Q, \Sigma, q_I, \delta, q_A, q_R)$ be such a machine, where $\delta: Q \times \Sigma \rightarrow Q \times \Sigma \times \{-1, 0, 1\}$ is the transition function. We assume w.l.o.g. that the accepting state q_A and the rejecting state q_R have self-loops. Furthermore, let $x \in \Sigma^*$ be an input for \mathcal{M} .

Let p be a polynomial that bounds the space-consumption of \mathcal{M} . From p we can compute a polynomial p' such that $2^{p'}$ bounds the time-consumption of \mathcal{M} . Thus, $s = p(|x|)$ and $t = 2^{p'(|x|)}$ are upper bounds on the space- and time-consumption of \mathcal{M} on x , respectively.

We construct a reachability automaton \mathcal{A} such that x is accepted by \mathcal{M} if and only if \mathcal{A} is not universal. This suffices to prove our claim, since PSPACE is closed under complement.

A configuration of \mathcal{M} 's run on x is encoded by a word $c \in (Q \cup \Sigma)^{s+1}$ as usual. In the following, we do not distinguish between a configuration and its encoding. In the construction, we need to count the configurations of the run. To this end, let $\langle n \rangle_2 \in \{0, 1\}^*$ denote the binary encoding of n in the range $0 \leq n \leq t - 1$ using $\log_2(t)$ many bits. Now, consider the following three ω -languages over the alphabet $\Sigma' = Q \cup \Sigma \cup \{0, 1, \$, \#\}$, where we assume $(Q \cup \Sigma) \cap \{0, 1, \$, \#\} = \emptyset$.

- (1) L_1 is the set of ω -words that do *not* start with $\langle 0 \rangle_2 \$ c_0 \#$, where c_0 is the initial configuration of \mathcal{M} on x .
- (2) L_2 is the set of ω -words having an infix $\langle n \rangle_2 \$ c \#$ such that
 - $n < t - 1$,
 - the infix is *not* followed by $\langle n + 1 \rangle_2 \$ c' \#$, where c' is the successor configuration of c .
- (3) L_3 is the set of ω -words containing the infix $\langle t - 1 \rangle_2$ such that the first occurrence of the infix $\langle t - 1 \rangle_2$ is *not* followed by $\$ c \#$ for some accepting configuration c of length s .

We claim that all three languages can be accepted by non-deterministic reachability automata of polynomial size in $|\mathcal{M}| + |x|$. This is straightforward for L_1 and L_3 and for L_2 we use the fact that it suffices to find a single bit or tape-cell that is not updated correctly. Let \mathcal{A} be an automaton such that $L_{\exists}(\mathcal{A}) = L_1 \cup L_2 \cup L_3$, again of polynomial size in $|\mathcal{M}| + |x|$. Such an automaton exists, since the union of reachability languages can be recognized by the disjoint union of the corresponding automata with a fresh initial state. We claim that \mathcal{A} has the desired properties.

Assume \mathcal{M} accepts x and let c_0, c_1, \dots, c_k for $k \leq t - 1$ be the accepting run. Consider the word

$$\langle 0 \rangle_2 \$ c_0 \# \langle 1 \rangle_2 \$ c_1 \# \dots \langle k \rangle_2 \$ c_k \# \langle k + 1 \rangle_2 \$ c_k \# \dots \langle t - 1 \rangle_2 \$ c_k \# 0^\omega,$$

which is not in $L_{\exists}(\mathcal{A})$. Hence, \mathcal{A} is not universal.

Now, assume \mathcal{A} is not universal. Then, there is a word w that is not in $L_{\exists}(\mathcal{A}) = L_1 \cup L_2 \cup L_3$. Thus, w starts with $\langle 0 \rangle_2 \$ c_0 \#$, where c_0 is the initial configuration of \mathcal{M} on x . This is followed by $\langle 1 \rangle_2 \$ c_1 \#$, where c_1 is the successor configuration of c_0 . This is continued until we reach an infix $\langle t - 1 \rangle_2 \$ c_{t-1} \#$, which is the first occurrence of the infix $\langle t - 1 \rangle_2$.

⁵The correctness relies on the reachability automaton being complete, which we require of our automata.

Due to w not being in L_3 , we conclude that c_{t-1} is an accepting configuration. Therefore, \mathcal{M} accepts x , as we have constructed an accepting run. \square