arXiv:1509.07495v2 [cs.GT] 26 Jan 2016

Unbounded Lookahead in WMSO+U Games*

Martin Zimmermann

Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany
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. We consider delay games with winning conditions expressed in weak monadic second order logic with the unbounding quantifier (WMSO+U), which is able to express (un)boundedness properties. It is decidable whether the delaying player is able to win such a game with bounded lookahead, i.e., if she only skips a finite number of moves. However, bounded lookahead is not always sufficient: we present a game that can be won with unbounded lookahead, but not with bounded lookahead. Then, we consider WMSO+U delay games with unbounded lookahead and show that the exact evolution of the lookahead is irrelevant: the winner is always the same, as long as the initial lookahead is large enough and the lookahead tends to infinity.

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 Büchi and Landweber [10] 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.

Ever since, this result was extended along different dimensions, e.g., the number of players, the type of arena, the type of winning condition, the type of interaction between the players (alternation or concurrency), zero-sum or non-zero-sum, and complete or incomplete information. In this work, we consider two of these dimensions, namely more expressive winning conditions and the possibility for one player to delay her moves.

WMSO+U Recall that the ω-regular languages are exactly those that are definable in monadic second order logic (MSO) [9]. Recently, Bojańczyk has started a program [1,2,3,4,5,6,7,8,19] investigating the logic MSO+U, MSO extended with

* Supported by the project "TriCS" (ZI 1516/1-1) of the German Research Foundation (DFG).

the unbounding quantifier U. A formula $\mathrm{U}X\varphi(X)$ is satisfied, if there are arbitrarily large *finite* sets X such that $\varphi(X)$ holds. MSO+U is able to express all ω-regular languages as well as non-regular ones such as $L = \{a^{n_0}ba^{n_1}ba^{n_2}b\cdots \mid \limsup_i n_i = \infty\}$. Decidability of MSO+U remained an open problem until recently: satisfiability of MSO+U on infinite words is undecidable [6].

Even before this undecidability result was shown, much attention was being paid to fragments of the logic obtained by restricting the power of the second-order quantifiers. In particular, considering weak[1] MSO with the unbounding quantifier (denoted by prepending a W) turned out to be promising: WMSO+U on infinite words [2] and on infinite trees [8] and WMSO+U with the path quantifier (WMSO+UP) on infinite trees [3] have equivalent automata models with decidable emptiness. Hence, these logics are decidable.

For WMSO+U on infinite words, these automata are called max-automata, deterministic automata with counters whose acceptance conditions are a boolean combination of conditions "counter c is bounded during the run". While processing the input, a counter may be incremented, reset to zero, or the maximum of two counters may be assigned to it (hence the name max-automata). In this work, we continue the investigation of delay games with winning conditions given by max-automata, so-called max-regular conditions.

Delay Games In such a delay game, one of the players can postpone her moves for some time, thereby obtaining a lookahead on her opponent's moves. This allows her to win some games which she loses 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 when transmission of data in networks or components with buffers are modeled.

From a more theoretical point of view, uniformization of relations by continuous functions [23,24,25] can be expressed and analyzed using delay games. We consider games in which two players pick letters from alphabets Σ_I and Σ_O, respectively, thereby producing $\alpha \in \Sigma_I^\omega$ and $\beta \in \Sigma_O^\omega$. Thus, a strategy for the second player induces a mapping $\tau\colon \Sigma_I^\omega \to \Sigma_O^\omega$. It is winning for the second player if $(\alpha, \tau(\alpha))$ is contained in the winning condition $L \subseteq \Sigma_I^\omega \times \Sigma_O^\omega$ for every α. Then, we say that τ uniformizes L.

In the classical setting of infinite games, in which the players pick letters in alternation, the n-th letter of $\tau(\alpha)$ depends only on the first n letters of α, i.e., τ satisfies a very strong notion of continuity. 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 arbitrary lookahead induces a continuous function (or equivalently, a uniformly continuous function, as Σ^ω is compact).

Hosch and Landweber proved that it is decidable whether a game with ω-regular winning condition can be won with bounded lookahead [18]. This result was improved by Holtmann, Kaiser, and Thomas who showed that if a

[1] Here, the second-order quantifiers are restricted to finite sets.

player wins a game with arbitrary lookahead, then she wins already with doubly-exponential bounded lookahead, and gave a streamlined decidability proof yielding an algorithm with doubly-exponential running time [17]. Again, these results were improved by giving an exponential upper bound on the necessary lookahead and showing ExpTime-completeness of the solution problem [21]. Going beyond ω-regular winning conditions by considering context-free conditions leads to undecidability and non-elementary lower bounds on the necessary lookahead, even for very weak fragments [15]. In contrast, studying delay games with WMSO+U turned out to be more fruitful [26]: the winner of such a game w.r.t. bounded lookahead is decidable, i.e., the Hosch-Landweber Theorem holds for max-regular conditions, too.

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.

In another line of work, Carayol and Löding considered the case of finite words [13], and Löding and Winter [22] considered the case of finite trees, which are both decidable. However, the nonexistence of MSO-definable choice functions on the infinite binary tree [12,16] implies that uniformization fails for such trees.

Another application of delay games concerns the existence of Wadge reductions between max-regular languages [11], which can be expressed as a max-regular delay game.

Our Contribution Here, we continue the investigation of delay games with max-regular winning conditions, which was started by proving the analogue of the Hosch-Landweber Theorem [26]: the winner of a game w.r.t. bounded lookahead is decidable. In particular, we are interested in the analogue of the Holtmann-Kaiser-Thomas Theorem (is bounded lookahead sufficient?). Not surprisingly, our first result (which was already announced, but not proved, in [26]) shows that this does not hold: unbounded lookahead is more powerful when it comes to unboundedness conditions.

We complement this by showing that the ability of Player O to win a max-regular delay game does not depend on the growth rate, only on the fact that it grows without bound and a sufficiently large initial lookahead. This is, to the best of our knowledge, the first such result and should be contrasted with the case of ω-context-free winning conditions, for which a non-elementary growth rate might be necessary for Player O to win [15].

As the analogue of the Holtmann-Kaiser-Thomas Theorem fails, determining the winner of max-regular delay games with respect to arbitrary delay functions does not coincide with determining the winner with respect to bounded delay functions. Hence, we investigate the former problem: we give lower bounds on the complexity and discuss some obstacles one encounters when trying the extend the decidability proof for the bounded case and the undecidability proof for MSO+U satisfiability.

2 Preliminaries

The set of non-negative integers is denoted by $\mathbb{N}$. An alphabet Σ is a non-empty finite set of letters, and Σ^* (Σ^n, Σ^ω) denotes the set of finite words (words of length n, infinite words) over Σ. The empty word is denoted by ε, 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. Given two infinite words $\alpha \in \Sigma_I^\omega$ and $\beta \in \Sigma_O^\omega$ we write $\binom{\alpha}{\beta}$ for the word $\binom{\alpha(0)}{\beta(0)}\binom{\alpha(1)}{\beta(1)}\binom{\alpha(2)}{\beta(2)}\cdots \in (\Sigma_I \times \Sigma_O)^\omega$. Analogously, we write $\binom{x}{y}$ for finite words x and y, provided they are of equal length. Finally, the index of an equivalence relation $\equiv$, i.e., the number of its equivalence classes, is denoted by $\mathrm{idx}(\equiv)$.

Max-Automata Given a finite set C of counters storing non-negative integers,

$$\mathrm{Ops}(C) = \{c := c + 1, c := 0, c := \max(c_0, c_1) \mid c, c_0, c_1 \in C\}$$

is the set of counter operations over C. A counter valuation over C is a mapping $\nu\colon C \to \mathbb{N}$. By $\nu\pi$ we denote the counter valuation that is obtained by applying a finite sequence $\pi \in \mathrm{Ops}(C)^*$ of counter operations to ν, which is defined as implied by the operations' names.

A max-automaton $\mathcal{A} = (Q, C, \Sigma, q_I, \delta, \ell, \varphi)$ consists of a finite set Q of states, a finite set C of counters, an input alphabet Σ, an initial state q_I, a (deterministic and complete) transition function $\delta\colon Q \times \Sigma \to Q$, a transition labeling[2] $\ell\colon \delta \to \mathrm{Ops}(C)^*$ which labels each transition by a (possibly empty) sequence of counter operations, and an acceptance condition φ, which is a boolean formula over C.

A run of $\mathcal{A}$ on $\alpha \in \Sigma^\omega$ is an infinite sequence

$$\rho = (q_0, \alpha(0), q_1)\,(q_1, \alpha(1), q_2)\,(q_2, \alpha(2), q_3)\cdots \in \delta^\omega \tag{1}$$

with $q_0 = q_I$. Runs on finite words are defined analogously, i.e.,

$$(q_0, \alpha(0), q_1)\cdots(q_{n-1}, \alpha(n-1), q_n)$$

is the run of $\mathcal{A}$ on $\alpha(0)\cdots\alpha(n-1)$ starting in q_0. We say that this run ends with q_n. As δ is deterministic, $\mathcal{A}$ has a unique run on every finite or infinite word.

Let ρ be as in (1) and define $\pi_n = \ell(q_n, \alpha(n), q_{n+1})$, i.e., π_n is the label of the n-th transition of ρ. Given an initial counter valuation ν and a counter $c \in C$, we define the sequence $\rho_c = \nu(c)\,,\ \nu\pi_0(c)\,,\ \nu\pi_0\pi_1(c)\,,\ \nu\pi_0\pi_1\pi_2(c)\,,\ldots$ of counter values of c reached on the run after applying *all* operations of a transition label. The run ρ of $\mathcal{A}$ on α is accepting, if the acceptance condition φ is satisfied by the variable valuation that maps a counter c to true if and only if $\limsup \rho_c$ is finite. Thus, φ can intuitively be understood as a boolean combination of conditions "$\limsup \rho_c < \infty$". Note that the limit superior of ρ_c is independent of the initial valuation used to define ρ_c, which is the reason it is not part of the description of $\mathcal{A}$. We denote the language accepted by $\mathcal{A}$ by $L(\mathcal{A})$ and say that it is max-regular.

[2] Here, and later whenever convenient, we treat δ as relation $\delta \subseteq Q \times \Sigma \times Q$.

A parity condition (say min-parity) can be expressed in this framework using a counter for each color that is incremented every time this color is visited and employing the acceptance condition to check that the smallest color whose associated counter is unbounded, is even. Hence, the class of ω-regular languages is contained in the class of max-regular languages.

Delay Games A delay function is a mapping $f\colon \mathbb{N} \to \mathbb{N} \setminus \{0\}$, which is said to be bounded, if $f(i) = 1$ for almost all i. Otherwise, f is unbounded. A special case of the bounded delay functions are the constant ones: delay functions f with $f(i) = 1$ for every $i > 0$.

Fix an input alphabet Σ_I and an output alphabet Σ_O. Given a delay function f and an ω-language $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$, the game $\Gamma_f(L)$ is played by two players (Player I and Player O) in rounds $i = 0, 1, 2, \ldots$ 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 $(u_0, v_0), (u_1, v_1), (u_2, v_2), \ldots$ as a play of $\Gamma_f(L)$. Player O wins the play if and only if the outcome $\binom{u_0 u_1 u_2 \cdots}{v_0 v_1 v_2 \cdots}$ is in L, otherwise Player I wins.

Given a delay function f, a strategy for Player I is a mapping $\tau_I\colon \Sigma_O^* \to \Sigma_I^*$ such that $|\tau_I(w)| = f(|w|)$, and a strategy for Player O is a mapping $\tau_O\colon \Sigma_I^* \to \Sigma_O$. Consider a play $(u_0, v_0), (u_1, v_1), (u_2, v_2), \ldots$ of $\Gamma_f(L)$. Such a play is consistent with τ_I, if $u_i = \tau_I(v_0 \cdots v_{i-1})$ for every i; it is consistent with τ_O, if $v_i = \tau_O(u_0 \cdots u_i)$ for every i. A strategy τ for Player p is winning for her, if every play that is consistent with τ is won by Player p. In this case, we say Player p wins $\Gamma_f(L)$.

Given a max-automaton $\mathcal{A}$, we want to determine whether Player O has a winning strategy for $\Gamma_f(L(\mathcal{A}))$ for some f, and, if yes, what kind of f is sufficient to win. Note that due to monotonicity, Player O wins a delay game for an arbitrary winning condition w.r.t. a bounded delay function if and only if she wins the game w.r.t. a constant delay function.

3 Bounded Lookahead is not Always Sufficient

The winner of a max-regular delay game w.r.t. constant delay functions can be determined effectively by a reduction to delay-free games with max-regular winning conditions, i.e., the following problem is decidable [26]: given a max-automaton $\mathcal{A}$, does Player O win $\Gamma_f(L(\mathcal{A}))$ for some constant (equivalently, bounded) delay function f? However, in this section, we show that bounded and thus constant lookahead does not suffice to win every delay game that Player O can win with arbitrary lookahead.

Theorem 1. *There is a max-regular language L such that Player O wins $\Gamma_f(L)$ for every unbounded f, but not for any bounded f.*

Proof. Let $\Sigma_I = \{0, 1, \#\}$ and $\Sigma_O = \{0, 1, *\}$. An input block is a word $\#w$ with $w \in \{0, 1\}^+$. An output block is a word $\binom{\#}{\alpha(n)}\binom{\alpha(1)}{*}\binom{\alpha(2)}{*}\cdots\binom{\alpha(n-1)}{*}\binom{\alpha(n)}{\alpha(n)} \in$

$(\Sigma_I \times \Sigma_O)^+$ with $\alpha(j) \in \{0,1\}$ for all j in the range $1 \leq j \leq n$. Note that the first and last letter in an output block are the only ones whose second component is not an $*$, and that these letters have to be equal to the first component of the block's last letter. Also, note that neither an input block nor an output block has to be maximal in the sense that it has to end with a # (in the first component). Every input block of length n can be extended to an output block of length n and projecting an output block to its first components yields an input block.

Let $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ be the language of words $\binom{\alpha}{\beta}$ satisfying the following property: if α contains infinitely many # and arbitrarily long input blocks, then $\binom{\alpha}{\beta}$ contains arbitrarily long output blocks. It is easy to come up with a WMSO+U formula defining L by formalizing the definitions of input and output blocks in first-order logic.

Now, consider L as winning condition for a delay game. Intuitively, Player O has to specify arbitrarily long output blocks, provided Player I produces arbitrarily long input blocks. The challenge for Player O is that she has to specify at the beginning of every output block whether she ends the block in a position where Player I has picked a 0 or a 1.

First, consider $\Gamma_f(L)$ for an unbounded delay function f. The following strategy is winning for Player O: whenever she has to pick $\beta(i)$ at a position where Player I picked $\alpha(i) = \#$, she picks the last letter of the longest input block in the lookahead that starts with the current #. Then, she completes the output block by picking $*$ until the end of the input block, where she copies $\beta(i)$, which completes the output block. At every other position, she picks an arbitrary letter. Now, consider a play consistent with this strategy: if Player I picks infinitely many # and arbitrarily large input blocks, then Player O will see arbitrarily large input blocks in her lookahead, i.e., her strategy picks arbitrarily large output blocks. Thus, the strategy is indeed winning.

It remains to show that Player I wins $\Gamma_f(L)$ for every bounded delay function f. Fix such a function and define $\ell = \sum_{i:\, f(i)>0} f(i) - 1$, i.e., ℓ is the maximal lookahead size Player O will achieve. Player I produces longer and longer input blocks of the following form: he starts picking # followed by 0's until Player O has picked an answer at the position of the last #. If she picked a 0, then Player I finishes the input block by picking 1's; if she picked a 1 (or an $*$), then he finishes the input block by picking 0's. Thus, the length of every output block is at most ℓ, since Player O has to determine the answer to every # after seeing at most the next ℓ letters picked by Player I. Thus, Player I picks infinitely many # and arbitrarily long input blocks, while the length of the output blocks is bounded. Hence, the strategy is winning for Player I.

4 Any Unbounded Lookahead is Sufficient

In this section, we complement the result of the previous section, showing that bounded lookahead is not always sufficient for max-regular delay games, by showing that any unbounded lookahead is sufficient for Player O, provided some lookahead allows her to win at all. If Player O wins a game with respect to some

delay function f, then she also wins with respect to every f' that grants her at every round larger lookahead[3]. The hard part of the proof is to show that she also wins for *smaller* functions f' that grant her less lookahead.

To this end, in Subsection 4.1, we introduce equivalence relations that capture the behavior of a max-automaton up to a certain precision. Then, in Subsection 4.2, we define an infinite-state game $\mathcal{G}$ based on these equivalence relations. Intuitively, the players' pick equivalence classes and Player I is in charge of increasing the precision of the approximation of the automaton's behavior, i.e., there is no explicit delay function in the definition of $\mathcal{G}$. This game allows to prove that smaller, but unbounded, lookahead is also sufficient.

4.1 Equivalence Relations for Max-Automata

Fix $\mathcal{A} = (Q, C, \Sigma, q_I, \delta, \ell, \varphi)$. We generalize notions introduced in [2] and [26] to define equivalences over sequences of counter operations and over words over Σ to capture the behavior of $\mathcal{A}$ up to a given precision. To this end, we need to introduce some notation to deal with runs of $\mathcal{A}$. Given a state q and $w \in \Sigma^* \cup \Sigma^\omega$, let $\rho(q, w)$ be the run of $\mathcal{A}$ on w starting in q. If w is finite, then $\delta^*(q, w)$ denotes the state $\rho(q, w)$ ends with. The transition profile of $w \in \Sigma^*$ is the mapping $q \mapsto \delta^*(q, w)$.

Now, we define inductively what it means for a sequence $\pi \in \mathrm{Ops}(C)^*$ to transfer a counter c to a counter d. The empty sequence and the operation $c := c + 1$ transfer every counter to itself. The operation $c := 0$ transfers every counter $c' \neq c$ to itself and the operation $c := \max(c_0, c_1)$ transfers every counter but c to itself and transfers c_0 and c_1 to c. Finally, if π_0 transfers c to e and π_1 transfers e to d, then $\pi_0\pi_1$ transfers c to d. If π transfers c to d, then we have $\nu\pi(d) \geq \nu(c)$ for every counter valuation ν, i.e., the value of d after executing π is larger or equal to the value of c before executing π, independently of the initial counter values.

Furthermore, a sequence of counter operations π transfers c to d with $m \geq 0$ increments, if there are counters $e_1, \ldots, e_m$ and a decomposition

$$\pi = \pi_0 \, (e_1 := e_1 + 1) \, \pi_1 \, (e_2 := e_2 + 1) \, \pi_2 \cdots \pi_{m-1} \, (e_m := e_m + 1) \, \pi_m$$

of π such that π_0 transfers c to e_1, π_j transfers e_j to e_{j+1} for every j in the range $1 \leq j < m$, and π_m transfers e_m to d. If π transfers c to d with m increments, then we have $\nu\pi(d) \geq \nu(c) + m$ for every counter valuation ν. Also note that if π transfers c to d with $m > 0$ increments, then it also transfers c to d with m' increments for every $m' \leq m$. Finally, we say that π is a c-trace of length m, if there is a counter c' such that π transfers c' to c with m increments. Thus, if π is a c-trace of length m, then $\nu\pi(c) \geq m$ for every valuation ν.

As only counter values reached after executing all counter operations of a transition label are considered in the semantics of max-automata, we treat $\Lambda = \{\ell(q, a, q') \mid (q, a, q') \in \delta\}$ as an alphabet. Every word $\lambda \in \Lambda^*$ can be flattened

[3] This holds for every winning condition, not only max-regular ones.

to a word in $\mathrm{Ops}(C)^*$, which is denoted by $\mathrm{flat}(\lambda)$. However, infixes, prefixes, or suffixes of λ are defined with respect to the alphabet Λ. We define $\ell(q, w) \in \Lambda^*$ to be the sequence of elements in Λ labeling the run $\rho(q, w)$.

Let ρ be a finite run of $\mathcal{A}$ and let $\pi \in \mathrm{Ops}(C)^*$. We say that ρ ends with π, if π is a suffix of $\mathrm{flat}(\ell(\rho))$. A finite or infinite run contains π, if it has a prefix that ends with π.

Lemma 1 ([2]). *Let ρ be a run of $\mathcal{A}$ and c a counter. Then, $\limsup \rho_c = \infty$ if and only if ρ contains arbitrarily long c-traces.*

We use the notions of transfer (with increment) to define the equivalence relations that capture $\mathcal{A}$'s behavior. Fix some $m \geq 0$. We say that $\lambda, \lambda' \in \Lambda^*$ are m-equivalent, denoted by $\lambda \equiv^m_{\mathrm{ops}} \lambda'$, if for all counters c and d and for all m' in the range $0 \leq m' \leq m$:

1. λ has an infix whose flattening has a suffix that is a c-trace of length m' if and only if λ' has an infix whose flattening has a suffix that is a c-trace of length m',
2. the flattening of λ has a suffix that is a c-trace of length m' if and only if the flattening of λ' has a suffix that is a c-trace of length m',
3. the flattening of λ transfers c to d with m' increments if and only if the flattening of λ' transfers c to d with m' increments, and
4. λ has a prefix whose flattening transfers c to d with m' increments if and only if λ' has a prefix whose flattening transfers c to d with m' increments.

Using this, we define two words $x, x' \in \Sigma^*$ to be m-equivalent, denoted by $x \equiv^m_{\mathcal{A}} x'$, if they have the same transition profile and if $\ell(q, x) \equiv^m_{\mathrm{ops}} \ell(q, x')$ for all states q.

Recall that a congruence is an equivalence relation $\equiv$ over Σ^* such that $x \equiv y$ implies $xz \equiv yz$ for every $z \in \Sigma^*$.

Lemma 2. *Let $\mathcal{A}$ be a max-automaton with n states and k counters and let $m \in \mathbb{N}$.*

1. *$\lambda \equiv^m_{\mathrm{ops}} \lambda'$ implies $\lambda \equiv^{m'}_{\mathrm{ops}} \lambda'$ for every $m' \leq m$.*
2. *$x \equiv^m_{\mathcal{A}} x'$ implies $x \equiv^{m'}_{\mathcal{A}} x'$ for every $m' \leq m$.*
3. *$\equiv^m_{\mathrm{ops}}$ is a congruence.*
4. *$\equiv^m_{\mathcal{A}}$ is a congruence.*
5. *The index of $\equiv^m_{\mathrm{ops}}$ is at most $2^{2(k^2+k)\log(m+2)}$.*
6. *The index of $\equiv^m_{\mathcal{A}}$ is at most $2^{n(\log(n)+2(k^2+k)\log(m+2))}$.*

Proof. The first two items follow trivially from the definition of $\equiv^m_{\mathrm{ops}}$. Thus, we only consider the latter four items.

3. Let $\lambda \equiv^m_{\mathrm{ops}} \lambda'$ and let $\pi \in \Lambda$ (note that we treat π as a letter from λ, although it is also a sequence of counter operations). We show $\lambda\pi \equiv^m_{\mathrm{ops}} \lambda'\pi$. An inductive application proves that $\equiv^m_{\mathrm{ops}}$ is an equivalence.

First, assume $\lambda\pi$ has an infix λ_0 whose flattening has a suffix π_0 that is a c-trace of length m' for some $m' \leq m$. If λ_0 is an infix of λ, then $\lambda \equiv^m_{\mathrm{ops}} \lambda'$ implies

that λ' has an infix with the same property. The other trivial case is when λ_0 is equal to π. Thus, it remains to consider the case where λ_0 is a suffix of $\lambda\pi$ of length at least two (recall that we treat π as one letter, i.e., λ_0 contains at least one letter from λ). Thus, λ_0 can be decomposed into two parts, one that is a c'-trace of length m_0 and is a suffix of the flattening of λ, and another one that is equal to π (treated as a sequence of counter operations now), which transfers c' to c with m_1 increments. Furthermore, we have $m_0 + m_1 = m' \leq m$.

Due to $\lambda \equiv^m_{\text{ops}} \lambda'$, we conclude that the flattening of λ' has a suffix that is a c'-trace of length m_0. Combining this suffix with π, we obtain a suffix of the flattening of $\lambda'\pi$ that is a c-trace of length m'. This is also an infix of $\lambda'\pi$ whose flattening has a suffix that is a c-trace of length m'.

The argument where $\lambda'\pi$ has such an infix is symmetric and the reasoning for the other three properties in the definition of $\equiv^m_{\text{ops}}$ is analogous.

4. Having the same transition profile is a congruence, since $\delta^*(q, xz) = \delta^*(\delta^*(q, x), z)$. This, and $\equiv^m_{\text{ops}}$ being a congruence imply that $\equiv^m_{\mathcal{A}}$ is a congruence as well.

5. An equivalence class of $\equiv^m_{\text{ops}}$ is uniquely characterized by the following properties:

- for every counter c, whether its elements have an infix whose flattening has a suffix that is a c-trace, and if yes by the largest $m' \leq m$ such that the length of such a c-trace is m'.
- For every counter c, whether its elements have a suffix that is a c-trace, and if yes by the largest $m' \leq m$ such that the length of such a trace is m'.
- For every pair (c, d) of counters, whether the flattenings of its elements transfer c to d, and if yes by the largest $m' \leq m$ such that the transfer has m' increments.
- For every pair (c, d) of counters, whether its elements have a prefix whose flattening transfers c to d, and if yes by the largest $m' \leq m$ such that the transfer has m' increments.

Thus, an equivalence class is induced by two mappings from C to $\{\bot, 0, 1, \ldots, m\}$ and two mappings from C^2 to $\{\bot, 0, 1, \ldots, m\}$, where $\bot$ encodes that no such trace or transfer exists. The number of quadruples of such mappings is bounded by $(m+2)^{2(k^2+k)} = 2^{2(k^2+k)\log(m+2)}$.

6. An equivalence class of $\equiv^m_{\mathcal{A}}$ is uniquely characterized by a transition profile and, for every state q, by the $\equiv^m_{\text{ops}}$ equivalence class of the sequence of counter operations encountered along the run starting in q. Thus, the class is characterized by a mapping from Q to pairs of a state and an $\equiv^m_{\text{ops}}$ class. Thus, the index of $\equiv^m_{\mathcal{A}}$ is bounded by the number of such mappings, i.e., by

$$(n \cdot \text{idx}(\equiv^m_{\text{ops}}))^n = 2^{\log\left((n2^{2(k^2+k)\log(m+2)})^n\right)} = 2^{n\left(\log(n)+2(k^2+k)\log(m+2)\right)}. \qquad \square$$

Next, we show that we take any infinite word $x_0x_1x_2\cdots$ with $x_i \in \Sigma^*$ and replace each x_i by an *equivalent* x'_i without changing membership in $L(\mathcal{A})$. To capture the evolution of the counters properly with the imprecise equivalence

relations $\equiv^m_{\mathcal{A}}$, we require that the x_i and the x'_i are $\equiv^m_{\mathcal{A}}$-equivalent for m tending to infinity. Formally, a sequence $(r_i)_{i\in\mathbb{N}}$ of natural numbers is a (convergence) rate, if it is weakly increasing and unbounded, i.e., $r_i \leq r_{i+1}$ for every i and $\sup_i r_i = \infty$.

Lemma 3. *Let $(x_i)_{i\in\mathbb{N}}$ and $(x'_i)_{i\in\mathbb{N}}$ be two sequences of words over Σ^* and let $(r_i)_{i\in\mathbb{N}}$ be a rate such that $x_i \equiv^{r_i}_{\mathcal{A}} x'_i$ for all i. Then, $x = x_0x_1x_2\cdots \in L(\mathcal{A})$ if and only if $x' = x'_0x'_1x'_2\cdots \in L(\mathcal{A})$.*

Proof. Let $\rho = (q_0, \alpha(0), q_1)(q_1, \alpha(1), q_2)\cdots$ be the run of $\mathcal{A}$ on x and let $\rho' = (q'_0, \alpha'(0), q'_1)(q'_1, \alpha'(1), q'_2)\cdots$ be the run of $\mathcal{A}$ on x', i.e., $x = \alpha(0)\alpha(1)\alpha(2)\cdots$ and $x' = \alpha'(0)\alpha'(1)\alpha'(2)\cdots$. Furthermore, let $n_i = |x_0\cdots x_{i-1}|$ and $n'_i = |x'_0\cdots x'_{i-1}|$. By definition of $\equiv^m_{\mathcal{A}}$, we obtain $q_{n_i} = q'_{n'_i}$ for every $i \geq 0$. Finally, we have $\ell(q_{n_i}, x_i) \equiv^{r_i}_{\text{ops}} \ell(q'_{n'_i}, x'_i)$ for every i, due to $x_i \equiv^{r_i}_{\mathcal{A}} x'_i$ and $q_{n_i} = q'_{n'_i}$.

We show that ρ contains arbitrarily long c-traces if and only if ρ' contains arbitrarily long c-traces. Due to Lemma 1, this suffices to show that ρ is accepting if and only if ρ' is accepting. Furthermore, due to symmetry, it suffices to show one direction of the equivalence. Thus, assume ρ contains arbitrarily long c-traces and pick $m \in \mathbb{N}$ arbitrarily. We show the existence of a c-trace of length m contained in ρ'.

To this end, fix a c-trace of length m in ρ. We can assume w.l.o.g. that the trace is contained in a run infix $\rho(q_{n_i}, x_i\cdots x_{i'})$ (which ends with $q_{n_{i'+1}}$) for some $i \leq i'$ with $m \leq r_i \leq r_{i'}$. Furthermore, we assume w.l.o.g. that i (i') is maximal (minimal) with this property for the fixed trace.

If $i = i'$, then the complete c-trace is contained in $\rho(q_{n_i}, x_i)$, i.e., $\ell(q_{n_i}, x_i)$ has an infix whose flattening has a suffix that is a c-trace of length $m \leq r_i$. Thus, the first requirement in the definition of $\equiv^{r_i}_{\text{ops}}$ yields an infix of $\ell(q'_{n'_i}, x'_i)$ whose flattening has a suffix that is a c-trace of length m. Thus, ρ' contains a c-trace of length m.

If $i < i'$, then the maximality of i and the minimality of i' imply that there are counters d_0, d_1 and non-negative numbers $m_0 + m_1 + m_2 = m$ such that

- the flattening of $\ell(q_{n_i}, x_i)$ has a suffix that is a d_0-trace of length m_0,
- the flattening of $\ell(q_{n_i+1}, x_{i+1}\cdots x_{i'-1})$ transfers d_0 to d_1 with m_1 increments, and
- $\ell(q_{n'_i}, x_{i'})$ has a prefix whose flattening transfers d_1 to c with m_2 increments.

The latter three requirements in the definition of $\equiv^{r_i}_{\text{ops}}$ imply the existence of the same transfers and traces in $\ell(q'_{n'_i}, x'_i)$, $\ell(q'_{n'_i+1}, x'_{i+1}\cdots x'_{i'-1})$, and $\ell(q'_{n'_{i'}}, x'_{i'})$, respectively. Hence, ρ' contains a c-trace of length m.

The $\equiv^m_{\mathcal{A}}$ classes are regular and *trackable* on-the-fly by a finite automaton $\mathcal{T}_m$ due to $\equiv^m_{\mathcal{A}}$ being a congruence.

Lemma 4. *There is a deterministic finite automaton $\mathcal{T}_m$ with set of states $\Sigma/\equiv^m_{\mathcal{A}}$ such that the run of $\mathcal{T}_m$ on $w \in \Sigma^*$ ends with state $[w]_{\equiv^m_{\mathcal{A}}}$.*

Proof. Define $\mathcal{T}_m = (\Sigma/\equiv^m_{\mathcal{A}}, \Sigma, [\varepsilon]_{\equiv^m_{\mathcal{A}}}, \delta_{\mathcal{T}_m}, \emptyset)$ where $\delta_{\mathcal{T}_m}([x]_{\equiv^m_{\mathcal{A}}}, a) = [xa]_{\equiv^m_{\mathcal{A}}}$, which is independent of the representative x and based on the fact that $\equiv^m_{\mathcal{A}}$ is a congruence. A straightforward induction over $|w|$ shows that $\mathcal{T}_m$ has the desired properties.

In particular, every $\equiv^m_{\mathcal{A}}$ equivalence class is regular and recognized by the DFA obtained from $\mathcal{T}_m$ by making the class to be recognized the only final state.

For the remainder of this section, we assume $\Sigma = \Sigma_I \times \Sigma_O$. We denote the projection of $\Sigma_I \times \Sigma_O$ to Σ_I by $\pi_I(\cdot)$, an operation we lift to words and languages over $\Sigma_I \times \Sigma_O$ in the usual way. Now, for each equivalence relation $\equiv^m_{\mathcal{A}}$ over $(\Sigma_I \times \Sigma_O)^*$ we define its projection[4] $=^m_{\mathcal{A}}$ over Σ_I^* via $x =^m_{\mathcal{A}} x'$ if and only if for all $\equiv^m_{\mathcal{A}}$ classes S: $x \in \pi_I(S)$ if and only if $x' \in \pi_I(S)$.

Remark 2 $\mathrm{idx}(=^m_{\mathcal{A}}) \leq 2^{\mathrm{idx}(\equiv^m_{\mathcal{A}})}$.

Furthermore, every $=^m_{\mathcal{A}}$ equivalence class is regular: we have

$$[x]_{=^m_{\mathcal{A}}} = \bigcap_{S \in (\Sigma_I \times \Sigma_O)^*/\equiv^m_{\mathcal{A}} \,:\, x \in \pi_I(S)} \pi_I(S) \cap \bigcap_{S \in (\Sigma_I \times \Sigma_O)^*/\equiv^m_{\mathcal{A}} \,:\, x \notin \pi_I(S)} \Sigma_I^* \setminus \pi_I(S),$$

where each projection $\pi_I(S)$ and each complemented projection $\Sigma_I^* \setminus \pi_I(S)$ is recognized by a DFA of size $2^{\mathrm{idx}(\equiv^m_{\mathcal{A}})}$. Thus, $[x]_{=^m_{\mathcal{A}}}$ is recognized by a DFA of size $2^{\mathrm{idx}(\equiv^m_{\mathcal{A}})^2}$. In particular, we have the following bound that will be applied in the next subsection.

Remark 3 *Let x be in a finite equivalence class of $=^0_{\mathcal{A}}$. Then, we have* $|x| < 2^{2^{2n(\log(n)+2(k^2+k))}}$.

4.2 A Game on Equivalence Classes

In this section, we show that the winner of a delay game with max-regular winning condition does not depend on the exact delay function under consideration, as long as it is unbounded and $f(0)$ is *large enough*. Note that this is true for the game analyzed in the previous section: Player O wins for every unbounded delay function.

Theorem 4. *Let $\mathcal{A}$ be a max-automaton with n states and k counters and let $d = 2^{2^{2n(\log(n)+2(k^2+k))}}$. The following are equivalent:*

1. *Player O wins $\Gamma_f(L(\mathcal{A}))$ for some f.*
2. *Player O wins $\Gamma_f(L(\mathcal{A}))$ for every unbounded f with $f(0) \geq 2d$.*

This result is proven by defining a delay-free game $\mathcal{G}(\mathcal{A})$ where Player I picks equivalence classes of $=^m_{\mathcal{A}}$ for increasing m and Player O constructs a run of $\mathcal{A}$ on a word over $\Sigma_I \times \Sigma_O$ that is compatible with the choices of Player I. Furthermore, Player I is always one move ahead to account for the delay.

[4] The notation $=^m_{\mathcal{A}}$ should not be understood as denoting equality, but merely as having projected away one bar from $\equiv^m_{\mathcal{A}}$.

Fix $\mathcal{A} = (Q, C, \Sigma_I \times \Sigma_O, q_I, \delta, \ell, \varphi)$ with $|Q| = n$ and $|C| = k$. We define the game $\mathcal{G}(\mathcal{A})$ between Player I and Player O played in rounds $i = 0, 1, 2, \ldots$ as follows: In round 0, Player I picks natural numbers r_0, r_1 and picks infinite equivalence classes $[x_0]_{=_{\mathcal{A}}^{r_0}}$ and $[x_1]_{=_{\mathcal{A}}^{r_1}}$. Then, Player O picks an equivalence class $[\binom{x_0}{y_0}]_{\equiv_{\mathcal{A}}^{r_0}}$. Note that this choice is independent of the representative x_0. Now, consider round $i > 0$: Player I picks $r_{i+1} \in \mathbb{N}$ and an infinite equivalence class $[x_{i+1}]_{=_{\mathcal{A}}^{r_{i+1}}}$. Afterwards, Player O picks an equivalence class $[\binom{x_i}{y_i}]_{\equiv_{\mathcal{A}}^{r_i}}$, whose choice is again independent of the representative x_i.

Thus, the players produce a play

$$[x_0]_{=_{\mathcal{A}}^{r_0}} [\binom{x_0}{y_0}]_{\equiv_{\mathcal{A}}^{r_0}} [x_1]_{=_{\mathcal{A}}^{r_1}} [\binom{x_1}{y_1}]_{\equiv_{\mathcal{A}}^{r_1}} [x_2]_{=_{\mathcal{A}}^{r_2}} [\binom{x_2}{y_2}]_{\equiv_{\mathcal{A}}^{r_2}} \cdots .$$

Player O wins, if $(r_i)_{i\in\mathbb{N}}$ is not a rate or if $\binom{x_0}{y_0}\binom{x_1}{y_1}\binom{x_2}{y_2}\cdots \in L(\mathcal{A})$. Otherwise, i.e., if $(r_i)_{i\in\mathbb{N}}$ is a rate and $\binom{x_0}{y_0}\binom{x_1}{y_1}\binom{x_2}{y_2}\cdots \notin L(\mathcal{A})$, Player I wins. By Lemma 3, winning does not depend on the choice of representatives x_i and y_i. Strategies and winning strategies for $\mathcal{G}(\mathcal{A})$ are defined as expected, taking into account that Player I is always one equivalence class ahead.

The following lemma about the relation between $\Gamma_f(L(\mathcal{A}))$ and $\mathcal{G}(\mathcal{A})$ implies Theorem 4.

Lemma 5. *The following are equivalent:*

1. *Player O wins $\Gamma_f(L(\mathcal{A}))$ for some f.*
2. *Player O wins $\Gamma_f(L(\mathcal{A}))$ for every unbounded f with $f(0) \geq 2d$.*
3. *Player O wins $\mathcal{G}(\mathcal{A})$.*

Proof. It suffices to show that 1. implies 3. and that 3. implies 2., as 2. implies 1. is trivially true. For the sake of readability, we will write Γ instead of $\Gamma_f(L(\mathcal{A}))$, as long as f is clear from context. Similarly, we will write $\mathcal{G}$ instead of $\mathcal{G}(\mathcal{A})$.

Let Player O win $\Gamma_f(L(\mathcal{A}))$ for some f, say with winning strategy τ_O. We construct a winning strategy τ'_O for her in $\mathcal{G}$ by simulating a play in Γ that is consistent with τ_O.

In round 0 of $\mathcal{G}$, Player I picks r_0,r_1, $[x_0]_{=_{\mathcal{A}}^{r_0}}$, and $[x_1]_{=_{\mathcal{A}}^{r_1}}$. As both equivalence classes are infinite, we can assume without loss of generality $|x_0| \geq f(0)$ and $|x_1| \geq \sum_{j=1}^{|x_0|-1} f(j)$. Now, assume Player I picks in Γ the prefix of x_0x_1 of length $\sum_{j=0}^{|x_0|-1} f(j)$ during the first $|x_0|$ rounds. Let y_0 of length $|x_0|$ be the answer of Player O to these choices determined by the winning strategy τ_O. We define τ'_O such that it picks $[\binom{x_0}{y_0}]_{\equiv_{\mathcal{A}}^{r_0}}$ as answer to Player I picking r_0,r_1, $[x_0]_{=_{\mathcal{A}}^{r_0}}$, and $[x_1]_{=_{\mathcal{A}}^{r_1}}$ in round 0.

Now, we are in the following situation for $i = 1$: in $\mathcal{G}$, Player I has picked natural numbers $r_0, \ldots, r_i$ and $[x_0]_{=_{\mathcal{A}}^{r_0}}, \ldots, [x_i]_{=_{\mathcal{A}}^{r_i}}$ with $|x_0| \geq f(0)$, $|x_1| \geq \sum_{j=1}^{|x_0|-1} f(j)$, and $|x_{i'}| \geq \sum_{j=0}^{|x_{i'-1}|-1} f(|x_0 \cdots x_{i'-1}| + j)$ for every i' with $1 < i' \leq i$ (which is empty for $i = 1$). Player O has picked $[\binom{x_0}{y_0}]_{\equiv_{\mathcal{A}}^{r_0}}, \ldots, [\binom{x_{i-1}}{y_{i-1}}]_{\equiv_{\mathcal{A}}^{r_{i-1}}}$. Further, in Γ, Player I has picked the prefix of $x_0 \cdots x_i$ of length $\sum_{j=0}^{|x_0\cdots x_{i-1}|-1} f(j)$

during the first $|x_0 \cdots x_{i-1}|$ rounds, which was answered by Player O according to τ_O by picking $y_0 \cdots y_{i-1}$.

In this situation, it is Player I's turn in $\mathcal{G}$, i.e., he picks r_{i+1} and $[x_{i+1}]_{=_{\mathcal{A}}^{r_{i+1}}}$. Again, as the class is infinite, we can assume $|x_{i+1}| \geq \sum_{j=0}^{|x_i|-1} f(|x_0 \cdots x_i| + j)$. Thus, we continue the play in Γ by letting Player I pick letters such that he has picked the prefix of $x_0 \cdots x_{i+1}$ of length $\sum_{j=0}^{|x_0 \cdots x_i|-1} f(j)$ during the first $|x_0 \cdots x_i|$ rounds. Again, this is answered by Player I by picking $y_0 \cdots y_i$ such that $|y_{i'}| = |x_{i'}|$ according to τ_O. Now, we define τ'_O such that it picks $[\binom{x_i}{y_i}]_{\equiv_{\mathcal{A}}^{r_i}}$ as next move. Thus, we are in the situation described above for $i+1$.

Let $w' = [x_0]_{=_{\mathcal{A}}^{r_0}} [\binom{x_0}{y_0}]_{\equiv_{\mathcal{A}}^{r_0}} [x_1]_{=_{\mathcal{A}}^{r_1}} [\binom{x_1}{y_1}]_{\equiv_{\mathcal{A}}^{r_1}} [x_2]_{=_{\mathcal{A}}^{r_2}} [\binom{x_2}{y_2}]_{\equiv_{\mathcal{A}}^{r_2}} \cdots$ be a play in $\mathcal{G}$ that is consistent with τ'_O. Consider the outcome $w = \binom{x_0}{y_0}\binom{x_1}{y_1}\binom{x_2}{y_2} \cdots$ of the play in Γ constructed during the simulation. It is consistent with τ_O, hence $w \in L(\mathcal{A})$. Accordingly, Player O wins the play w'. Thus, τ'_O is indeed a winning strategy for Player O in $\mathcal{G}$.

Now, consider the second implication to be proven: assume Player O has a winning strategy τ'_O for $\mathcal{G}$ and let f be an arbitrary unbounded delay function with $f(0) \geq 2d$. We construct a winning strategy τ_O for Player O in Γ by simulating a play of $\mathcal{G}$.

To this end, we define a strictly increasing auxiliary rate $(d_i)_{i \in \mathbb{N}}$ recursively as follows: let d_0 be minimal with the property that every word of length at least d_0 is in some infinite equivalence class of $=_{\mathcal{A}}^0$, i.e., $d_0 \leq d = 2^{2^{2n(\log n + 2(k^2+k))}}$ due to Remark 3. Now, we define d_{i+1} to be the minimal integer strictly greater than d_i such that every word of length at least d_{i+1} is in some infinite equivalence class of $=_{\mathcal{A}}^{i+1}$.

Let Player I pick $x_0 x_1$ of length $f(0) \geq 2 \cdot d_0$ in round 0 of Γ (the exact decomposition into x_0 and x_1 is irrelevant, we just use it to keep the notation consistent). Now, decompose $x_0 x_1 = x'_0 x'_1 \beta_1$ such that $|x'_0| = |x'_1| = d_0$. We simulate these moves by letting Player I pick $r_0 = r_1 = 0$, $[x'_0]_{=_{\mathcal{A}}^{r_0}}$, and $[x'_1]_{=_{\mathcal{A}}^{r_1}}$ in round 0 of $\mathcal{G}$, which are legal moves by the choice of d_0.

Thus, we are in the following situation for $i = 1$: in Γ, Player I has picked $x_0 \cdots x_i$ and Player O has picked $y_0 \cdots y_{i-2}$. Furthermore, in $\mathcal{G}$, Player I has picked $[x'_0]_{=_{\mathcal{A}}^{r_0}}, \ldots, [x'_i]_{=_{\mathcal{A}}^{r_i}}$ and there is a buffer $\beta_i \in \Sigma_I^*$ such that $x_0 \cdots x_i = x'_0 \cdots x'_i \beta_i$. Finally, Player O has picked $[\binom{x'_0}{y_0}]_{\equiv_{\mathcal{A}}^{r_0}} \cdots [\binom{x'_{i-2}}{y_{i-2}}]_{\equiv_{\mathcal{A}}^{r_{i-2}}}$.

In this situation, it is Player O's turn and τ'_O returns a class $[\binom{x'_{i-1}}{y_{i-1}}]_{\equiv_{\mathcal{A}}^{r_{i-1}}}$. Thus, we define τ_O such that it picks y_{i-1} during the next rounds, in which Player I picks letters forming x_{i+1} satisfying $|x_{i+1}| \geq |y_{i-1}|$. We consider two cases to simulate these in $\mathcal{G}$:

1. If $|\beta_i x_{i+1}| \geq 2d_{r_i+1} - d_{r_i}$, then Player I picks $r_{i+1} = r_i + 1$ and $[x'_{i+1}]_{=_{\mathcal{A}}^{r_{i+1}}}$, where x'_{i+1} is the prefix of $\beta_i x_{i+1}$ of length $d_{r_{i+1}}$. This is an infinite equivalence class by the choice of $d_{r_{i+1}}$. The remaining suffix of $\beta_i x_{i+1}$ is stored in the buffer β_{i+1}, i.e., we have $\beta_i x_{i+1} = x'_{i+1} \beta_{i+1}$.
2. Now, consider the case $|\beta_i x_{i+1}| < 2d_{r_i+1} - d_{r_i}$: we show $|\beta_i x_{i+1}| \geq d_{r_i}$. Then, Player I picks $r_{i+1} = r_i$ and $[x'_{i+1}]_{=_{\mathcal{A}}^{r_{i+1}}}$, where x'_{i+1} is the prefix of $\beta_i x_{i+1}$ of

length $d_{r_{i+1}} = d_{r_i}$, which is again an infinite equivalence class by the choice of $d_{r_{i+1}}$. The remaining suffix of $\beta_i x_{i+1}$ is stored in the buffer β_{i+1}, i.e., we have $\beta_i x_{i+1} = x'_{i+1}\beta_{i+1}$.
To show $|\beta_i x_{i+1}| \geq d_{r_i}$, we again consider two cases: if $r_{i-1} = r_i$, then we have

$$|\beta_i x_{i+1}| \geq |x_{i+1}| \geq |y_{i-1}| = |x'_{i-1}| = d_{r_{i-1}} = d_{r_i}.$$

On the other hand, if $r_{i-1} < r_i$, which implies $r_{i-1} + 1 = r_i$, as we are in the second case above, then we have $|\beta_{i-1}x_i| \geq 2d_{r_{i-1}+1} - d_{r_{i-1}}$ and x'_i is the prefix of length $d_{r_i} = d_{r_{i-1}+1}$ of $\beta_{i-1}x_i$, which implies $|\beta_i| \geq d_{r_{i-1}+1} - d_{r_{i-1}}$, as it is the remaining suffix of $\beta_{i-1}x_i$. Finally, we have $|x_{i+1}| \geq |y_{i-1}| = |x'_{i-1}| = d_{r_{i-1}}$. Altogether, we obtain

$$|\beta_i x_{i+1}| \geq (d_{r_{i-1}+1} - d_{r_{i-1}}) + d_{r_{i-1}} = d_{r_{i-1}+1} = d_{r_i}.$$

In both cases, we are back in the situation described above for $i + 1$.

Let $w = \binom{x_0x_1x_2\cdots}{y_0y_1y_2\cdots}$ be the outcome of a play in Γ that is consistent with τ_O. The play $[x'_0]_{=^{r_0}_{\mathcal{A}}} [\binom{x'_0}{y_0}]_{\equiv^{r_0}_{\mathcal{A}}} [x'_1]_{=^{r_1}_{\mathcal{A}}} [\binom{x'_1}{y_1}]_{\equiv^{r_1}_{\mathcal{A}}} [x'_2]_{=^{r_2}_{\mathcal{A}}} [\binom{x'_2}{y_2}]_{\equiv^{r_2}_{\mathcal{A}}} \cdots$ in $\mathcal{G}$ constructed during the simulation is consistent with τ'_O. As f is unbounded, $(r_i)_{i\in\mathbb{N}}$ is unbounded as well and thus a rate. Hence, we conclude $\binom{x'_0}{y_0}\binom{x'_1}{y_1}\binom{x'_2}{y_2} \in L(\mathcal{A})$, as τ'_O is a winning strategy. Also, a straightforward induction shows $x_0x_1x_2\cdots = x'_0x'_1x'_2\cdots$. Thus, $w \in L(\mathcal{A})$, i.e., τ_O is a winning strategy for Player O in Γ.

5 Solving Max-regular Delay games with Unbounded Lookahead

Unlike for ω-regular delay games, bounded lookahead is not always sufficient for Player O to win a max-regular delay game. Hence, determining the winner with respect to arbitrary delay functions is not equivalent to determining the winner with respect to bounded delay functions, which is known to be decidable [26]. We refer to the former problem as "solving max-regular delay games". In this concluding section, we discuss some obstacles one has to overcome in order to extend the decidability result for bounded lookahead to unbounded one. Furthermore, we give straightforward lower bounds on the complexity.

Proving upper bounds, e.g., decidability of determining the winner of max-regular delay games with respect to arbitrary delay functions, is complicated by the need for unbounded lookahead. All known decidability results are for the case of bounded lookahead [17,21,26]. In particular, decidability of max-regular delay games with respect to bounded lookahead [26] is based on a game similar to the game $\mathcal{G}(\mathcal{A})$ presented in Section 4, but where only equivalence classes of $=^1_{\mathcal{A}}$ and $\equiv^1_{\mathcal{A}}$ are picked by the players. This results in a finite delay-free game with max-regular winning condition, which is effectively solvable [3]. Correctness follows from the fact that the error introduced by the imprecise equivalence relation $\equiv^1_{\mathcal{A}}$ is bounded, if the lookahead is bounded. As we are only interested in (un)boundedness, this error is negligible.

However, for unbounded lookahead, the error is unbounded as well. In particular, the example presented in Section 3 shows that bounded counters might grow arbitrarily large during different plays: the winning condition L described in the proof of Theorem 1 is recognized by a max-automaton with four counters: c_i counts the length of input blocks and is reset at every #, c'_o is incremented during prefixes of possible output blocks and reset at the end of such a block. Furthermore, the value of c'_o is copied to c_o every time the requirement on the first and last letter of an output block is met. Finally, a counter $c_\#$ counts the number of #'s in the word. The acceptance condition of the automaton recognizing L is given by the formula

$$\text{“}\limsup \rho_{c_\#} < \infty\text{”} \ \vee \ \text{“}\limsup \rho_{c_i} < \infty\text{”} \ \vee \ \text{“}\limsup \rho_{c_o} = \infty\text{”}.$$

As already argued, Player O has a winning strategy for $\Gamma_f(L)$, provided f is unbounded. However, she does not have a strategy that bounds the counters $c_\#$ and c_i to some fixed value among all consistent plays that are won due to $c_\#$ or c_i being bounded: for example, Player I can pick any finite number of #'s and then stop doing so. This implies that $c_\#$ is bounded, but with an arbitrarily large value among different plays. The lack of such a uniform bound in itself is not surprising, but entails that one has to deal with arbitrarily large counter values when trying to extend the approach described above for the setting with bounded lookahead. In particular, it is not enough to replace $=^1_{\mathcal{A}}$ and $\equiv^1_{\mathcal{A}}$ by $=^m_{\mathcal{A}}$ and $\equiv^m_{\mathcal{A}}$ for some fixed m that only depends on the winning condition.

Two other possible approaches follow from the results mentioned in this paper: first, one could show that $\mathcal{G}(\mathcal{A})$ can be solved effectively. However, the game is of infinite size and not in one of the classes of effectively solvable infinite games, e.g., pushdown games. Second, one can pick any unbounded delay function with $f(0)$ large enough and solve $\Gamma_f(L(\mathcal{A}))$, as winning with respect to one such function is equivalent to winning with respect to all of them. However, $\Gamma_f(L(\mathcal{A}))$ is again infinite and not in in one of the classes of effectively solvable infinite games.

Finally, there is a class of winning conditions for which solving delay games is indeed known to be undecidable, namely (very restricted fragments of) ω-context-free conditions [15]. However, this result is based on the language $\{a^n b^n \mid n \in \mathbb{N}\}$ being context-free, which suffices to encode two-counter machines. As max-automata have no mechanism to compare arbitrarily large numbers exactly, this simple encoding of two-counter machines cannot be captured in a delay game with max-regular winning condition.

This can be overcome by allowing quantification over arbitrary sets: recently, and after being an open problem for more than a decade, satisfiability of MSO+U over infinite words was shown to be undecidable [6] by capturing termination of two-counter machines by MSO+U formulas based on a specially tailored encoding. However, the resulting formulas have six alternations between existential and universal set quantifiers and then a block of (negated) unbounding quantifiers. To adapt this proof to show undecidability of max-regular delay games with respect to arbitrary delay functions, one has to replace the set quantifiers by the interaction between the players, which seems unlikely to achieve.

On the other hand, one can prove some straightforward lower bounds. As usual, solving delay games with max-regular winning conditions (given by max-automata) is at least as hard as solving the universality problem for max-automata: given such an automaton $\mathcal{A}$ over some alphabet Σ, we change the alphabet to $\Sigma \times \Sigma$ by replacing each letter a on a transition by the letter $\binom{a}{a}$. Call the resulting automaton $\mathcal{A}'$. The game $\Gamma_f(\mathcal{A}')$ is won by Player O if and only if $L(\mathcal{A})$ is universal, independently of f: if $L(\mathcal{A})$ is not universal, when Player I can produce some $\alpha \notin L(\mathcal{A})$ and thereby win; if it is indeed universal, then Player O can mimic the choices of Player I and thereby win.

Proposition 1. *Solving max-regular delay games is at least as hard as solving the universality problem for max-automata.*

The best known lower bound on the universality problem for max-automata is PSPACE-hardness, which stems from max-automata being closed under complementation and the emptiness problem being PSPACE-hard [7]. The exact complexity of the emptiness problem for max-automata is, to the best of our knowledge, an open problem.

Another lower bound is obtained by considering delay games with weaker winning conditions: solving delay games with winning conditions recognized by deterministic safety automata is EXPTIME-complete [21]. Such automata can be transformed into max-automata without increasing the number of states: turn the non-safe states into sinks and increment a designated counter c on every transition not leading into a non-safe state. Then, the max-automaton with acceptance condition "$\limsup \rho_c = \infty$" recognizes the same language as the original safety automaton. Hence, we obtain the following lower bound.

Theorem 5. *Solving max-regular delay games is* EXPTIME*-hard.*

This lower bound is oblivious to the intricate acceptance condition of max-automata and relies solely on the transition structure. This is in line with results for ω-regular games: solving delay games with winning conditions given by deterministic parity automata is in EXPTIME, i.e., it matches the lower bound for the special case of safety. It is open whether moving to more concise acceptance conditions for deterministic ω-automata, e.g., Rabin, Streett, and Muller, increases the complexity. These results would directly transfer to max-automata as well. Another aspect that is not exploited by this reduction is the unbounded lookahead: the safety delay game is always winnable with bounded lookahead. We are currently investigating whether these two aspects can be exploited to improve the bounds.

References

1. Mikołaj Bojańczyk. A bounding quantifier. In Jerzy Marcinkowski and Andrzej Tarlecki, editors, *CSL 2004*, volume 3210 of *LNCS*, pages 41–55. Springer, 2004.
2. Mikołaj Bojańczyk. Weak MSO with the unbounding quantifier. *Theory Comput. Syst.*, 48(3):554–576, 2011.

3. Mikołaj Bojańczyk. Weak MSO+U with path quantifiers over infinite trees. In Esparza et al. [14], pages 38–49.
4. Mikołaj Bojańczyk and Thomas Colcombet. Bounds in w-regularity. In *LICS 2006*, pages 285–296. IEEE Computer Society, 2006.
5. Mikołaj Bojańczyk, Tomasz Gogacz, Henryk Michalewski, and Michał Skrzypczak. On the decidability of MSO+U on infinite trees. In Esparza et al. [14], pages 50–61.
6. Mikołaj Bojańczyk, Paweł Parys, and Szymon Toruńczyk. The MSO+U theory of $(\mathbb{N}, <)$ is undecidable, 2015. ArXiv 1502.04578. To appear at STACS 2016.
7. Mikołaj Bojańczyk and Szymon Toruńczyk. Deterministic automata and extensions of weak MSO. In Kannan and Kumar [20], pages 73–84.
8. Mikołaj Bojańczyk and Szymon Toruńczyk. Weak MSO+U over infinite trees. In Christoph Dürr and Thomas Wilke, editors, *STACS 2012*, volume 14 of *LIPIcs*, pages 648–660. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2012.
9. J. Richard Büchi. On a Decision Method in Restricted Second-Order Arithmetic. In *International Congress on Logic, Methodology, and Philosophy of Science*, pages 1–11. Stanford University Press, 1962.
10. J. Richard Büchi and Lawrence H. Landweber. Solving sequential conditions by finite-state strategies. *Trans. Amer. Math. Soc.*, 138:pp. 295–311, 1969.
11. Jérémie Cabessa, Jacques Duparc, Alessandro Facchini, and Filip Murlak. The wadge hierarchy of max-regular languages. In Kannan and Kumar [20], pages 121–132.
12. 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.
13. Arnaud Carayol and Christof Löding. Uniformization in automata theory. In Peter Schroeder-Heister, Gerhard Heinzmann, Wilfrid Hodges, and Pierre Edouard Bour, editors, *International Congress of Logic, Methodology and Philosophy of Science*. College Publications, London, 2012. To appear.
14. Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors. *ICALP 2014, Part II*, volume 8573 of *LNCS*. Springer, 2014.
15. 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 für Informatik, 2011.
16. Yuri Gurevich and Saharon Shelah. Rabin's uniformization problem. *The Journal of Symbolic Logic*, 48:1105–1119, 1983.
17. Michael Holtmann, Łukasz Kaiser, and Wolfgang Thomas. Degrees of lookahead in regular infinite games. *LMCS*, 8(3), 2012.
18. Frederick A. Hosch and Lawrence H. Landweber. Finite delay solutions for sequential conditions. In *ICALP 1972*, pages 45–60, 1972.
19. Szczepan Hummel and Michał Skrzypczak. The topological complexity of MSO+U and related automata models. *Fundam. Inform.*, 119(1):87–111, 2012.
20. Ravi Kannan and K. Narayan Kumar, editors. *FSTTCS 2009*, volume 4 of *LIPIcs*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2009.
21. 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.
22. 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.

23. Wolfgang Thomas. Infinite games and uniformization. In Mohua Banerjee and Anil Seth, editors, *ICLA*, volume 6521 of *LNCS*, pages 19–21. Springer, 2011.
24. 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.
25. 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.
26. 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.