HALF-POSITIONAL OBJECTIVES
RECOGNIZED BY DETERMINISTIC BÜCHI AUTOMATA [*]

PATRICIA BOUYER [a], ANTONIO CASARES [b], MICKAEL RANDOUR [c],
AND PIERRE VANDENHOVE [a,c]

[a] Université Paris-Saclay, CNRS, ENS Paris-Saclay, Laboratoire Méthodes Formelles, 91190, Gif-sur-Yvette, France

[b] LaBRI, Université de Bordeaux, Bordeaux, France

[c] F.R.S.-FNRS & UMONS – Université de Mons, Mons, Belgium

ABSTRACT. In two-player games on graphs, the simplest possible strategies are those that can be implemented without any memory. These are called *positional strategies*. In this paper, we characterize objectives recognizable by deterministic Büchi automata (a subclass of *ω-regular objectives*) that are *half-positional*, that is, for which the protagonist can always play optimally using positional strategies (both over finite and infinite graphs). Our characterization consists of three natural conditions linked to the language-theoretic notion of *right congruence*. Furthermore, this characterization yields a polynomial-time algorithm to decide half-positionality of an objective recognized by a given deterministic Büchi automaton.

1. INTRODUCTION

Graph games and reactive synthesis. We study *zero-sum turn-based games on graphs* in which two players (a protagonist and its opponent) confront each other. They interact by moving a pebble in turns through the edges of a graph for an infinite amount of time. Each vertex belongs to a player, and the player controlling the current vertex decides on the next state of the game. Edges of the graph are labeled with *colors*, and the interaction of the two players therefore produces an infinite sequence of them. The objective of the game is specified by a subset of infinite sequences of colors, and the protagonist wins if the produced sequence belongs to this set. We are interested in finding a *winning strategy* for the protagonist, that is, a function indicating how the protagonist should move in any situation, guaranteeing the achievement of the objective.

This game-theoretic model is particularly fitted to study the *reactive synthesis problem* [BCJ18]: a system (the protagonist) wants to satisfy a specification (the objective) while interacting continuously with its environment (the opponent). The goal is to build a controller for the system satisfying the specification, whenever possible. This comes down to finding a winning strategy for the protagonist in the derived game.

Half-positionality. In order to obtain a controller for the system that is simple to implement, we are interested in finding the simplest possible winning strategy. Here, we focus on the amount of information that winning strategies have to remember. The simplest strategies are then arguably *positional* (also called *memoryless*) strategies, which do not remember anything about the past and base their decisions solely on the current state of the game. We intend to understand for which objectives positional strategies suffice for the protagonist to play optimally (i.e., to win whenever it is possible) — we call these objectives *half-positional*. We distinguish half-positionality from *bipositionality* (sometimes called *memoryless-determinacy* or just *positionality*), which refers to objectives for which positional strategies suffice to play optimally for *both* players.

Many natural objectives have been shown to be bipositional over games on finite and sometimes infinite graphs: e.g., discounted sum [Sha53], mean-payoff [EM79], parity [EJ91], total payoff [GZ04], energy [BFL+08], or average-energy games [BMR+18]. Bipositionality can be established using general criteria and characterizations, over games on both finite [GZ04, GZ05, AR17] and infinite [CN06] graphs. Yet, there exist many objectives and combinations thereof for which one player, but not both, has positional optimal strategies (Rabin conditions [KK91, Kla94], mean-payoff parity [CHJ05], energy parity [CD12], some window objectives [CDRR15, BHR16], energy mean-payoff [BHRR19]...), and to which these results do not apply.

Various attempts have been made to understand common underlying properties of half-positional objectives and provide sufficient conditions [Kop06, Kop07, Kop08, BFMM11], but little more was known until the recent work of Ohlmann [Ohl23] (discussed below). These sufficient conditions are not general enough to prove half-positionality of some very simple objectives, even in the well-studied class of *ω-regular objectives* [BFMM11, Lemma 13]. Furthermore, multiple questions concerning half-positionality remain open. For instance, in [Kop08], Kopczyński conjectured that *prefix-independent* half-positional objectives are closed under finite union (this conjecture was recently refuted for games on finite graphs [Koz22a], but is still unsolved for games on infinite graphs). Also, Kopczyński showed that given a deterministic parity automaton recognizing a prefix-independent objective W, we can decide if W is half-positional over finite arenas [Kop07]. However, the time complexity of his algorithm is $\mathcal{O}(n^{\mathcal{O}(n^2)})$, where n is the number of states of the automaton. It is unknown whether this can be done in polynomial time, and no algorithm is known for deciding half-positionality over infinite arenas or in the non-prefix-independent case.

ω-regular objectives and deterministic Büchi automata. A central class of objectives, whose half-positionality is not yet completely understood, is the class of *ω-regular objectives*. There are multiple equivalent definitions for them: they are the objectives defined, e.g., by ω-regular expressions, by non-deterministic Büchi automata [McN66], and by deterministic parity automata [Mos84]. These objectives coincide with the class of objectives defined by monadic second-order formulas [Büc62], and they encompass linear-time temporal logic (LTL) specifications [Pnu77]. Part of their interest is due to the landmark result that finite-state

machines are sufficient to implement optimal strategies in ω-regular games [BL69, GH82], implying the decidability of the monadic second-order theory of natural numbers with the successor relation [Büc62] and the decidability of the synthesis problem under LTL specifications [PR89].

In this paper, we focus on the subclass of ω-regular objectives recognized by *deterministic Büchi automata* (DBA), that we call *DBA-recognizable*. DBA-recognizable objectives correspond to the ω-regular objectives that can be written as a countable intersection of open objectives (for the product topology, that is, that are G_δ-sets of the Borel hierarchy); or equivalently, that are the limit of a regular language of finite words [Lan69, PP04]. Deciding the winner of a game with a DBA-recognizable objective is doable in polynomial time in the size of the arena and the DBA (by solving a Büchi game on the product of the arena and the DBA [BCJ18]).

We now discuss two technical tools at the core of our approach: *universal graphs* and *right congruences*.

Universal graphs. One recent breakthrough in the study of half-positionality is the introduction of *well-monotonic universal graphs*, combinatorial structures that can be used to provide a witness of winning strategies in games with a half-positional objective. Recently, Ohlmann [Ohl23] has shown that the existence of a *well-monotonic universal graph* for an objective W exactly characterizes half-positionality (under minor technical assumptions on W). Moreover, under these assumptions, a wide class of algorithms, called *value iteration algorithms*, can be applied to solve any game with a half-positional objective [CFGO22, Ohl23].

Although it brings insight into the structure of half-positional objectives, showing half-positionality through the use of universal graphs is not always straightforward, and has not yet been applied in a systematic way to ω-regular objectives.

Right congruence. Given an objective W, the *right congruence $\sim_W$ of W* is an equivalence relation on finite words: two finite words w_1 and w_2 are equivalent for $\sim_W$ if for all infinite continuations w, $w_1w \in W$ if and only if $w_2w \in W$. There is a natural automaton classifying the equivalence classes of the right congruence, which we refer to as the *prefix-classifier* [Sta83, MS97].

In the case of languages of *finite* words, a straightforward adaptation of the right congruence recovers the known Myhill-Nerode congruence. This equivalence relation characterizes the regular languages (a language is regular if and only if its congruence has finitely many equivalence classes), and the prefix-classifier is exactly the smallest deterministic finite automaton recognizing a language — this is the celebrated Myhill-Nerode theorem [Ner58].

Objectives are languages of *infinite* words, for which the situation is not so clear-cut. In particular, some ω-regular objectives cannot be recognized by their prefix-classifier along with a natural acceptance condition (Büchi, coBüchi, parity, Muller...) [MS97, AF21].

Contributions. Our main contribution is a *characterization* of half-positionality for DBA-recognizable objectives through a conjunction of three easy-to-check conditions (Theorem 3.11).

(1) The equivalence classes of the right congruence are *totally* ordered w.r.t. inclusion of their winning continuations.

(2) Whenever the set of winning continuations of a finite word w_1 is a proper subset of the set of winning continuations of a concatenation w_1w_2, the word $w_1(w_2)^\omega$ produced by repeating infinitely often w_2 is winning.

(3) The objective has to be recognizable by a DBA using the structure of its prefix-classifier.

A few examples of simple DBA-recognizable objectives that were not encompassed by previous half-positionality criteria [Kop06, BFMM11] are, e.g., reaching a color twice [BFMM11, Lemma 13] and weak parity [Tho08]. We also refer to Example 3.7, which is half-positional but not bipositional, and whose half-positionality is straightforward using our characterization.

Various corollaries with practical and theoretical interest follow from our characterization.

• We obtain a painless path to show (by checking each of the three conditions) that given a deterministic Büchi automaton, the half-positionality (over both finite and infinite arenas) of the objective it recognizes is decidable in time $\mathcal{O}(k \cdot n^4)$, where k is the number of colors and n is the number of states of the DBA (Section 3.3).

• Prefix-independent DBA-recognizable half-positional objectives are exactly the very simple *Büchi conditions*, which consist of all the infinite words containing infinitely many occurrences of colors from some fixed subset (Proposition 3.12). In particular, Kopczyński's conjecture trivializes for DBA-recognizable objectives (the finite union of Büchi conditions is a Büchi condition).

• We obtain a *finite-to-infinite* and *one-to-two-player* lift result (Proposition 3.15): in order to check that a DBA-recognizable objective is half-positional over arbitrary — possibly two-player and infinite — graphs, it suffices to check the existence of positional optimal strategies over *finite* graphs where all the vertices are controlled by the protagonist.

Technical overview. Condition (1) turns out to be equivalent to earlier properties used to study bipositionality and half-positionality [GZ05, BFMM11] (details in Appendix A). Condition (2), to the best of our knowledge, is a novel condition. A property similar to Condition (3) has been studied multiple times in the language-theoretic literature, both for itself and for minimization and learning algorithms [Sta83, Le 90, MS97, AF21]. As an example, all objectives defined by deterministic *weak* automata (a restriction on DBA) satisfy Condition (3) [Sta83, AF21]. This property is sometimes called having an *informative right congruence* (for some given acceptance condition). However, its links with positionality had never been explored.

Conditions (1) and (2) are necessary for respectively bipositionality and half-positionality of general objectives. Condition (3) is necessary for half-positionality of DBA-recognizable objectives, but not for all (even ω-regular) objectives in general (see Example 3.10). The proof of its necessity is more involved than for the first two conditions, and will build on automata-theoretic ideas introduced for good-for-games coBüchi automata [AK22]. Together, the three conditions are sufficient for half-positionality of DBA-recognizable objectives: the proof of sufficiency uses the theory of universal graphs, and consists of building a family of well-monotonic universal graphs [Ohl23] for objectives satisfying the three properties.

Other related works. We have discussed relevant literature on half-positionality [Kop06, Kop07, BFMM11, Ohl23] and bipositionality [GZ04, GZ05, CN06, AR17]. A more general quest is to understand *memory requirements* when positional strategies are not powerful enough: e.g., [LPR18, BLO$^+$22, BORV21, BRV23].

Memory requirements have been precisely characterized for some classes of ω-regular objectives (not encompassing the class of DBA-recognizable objectives), such as Muller conditions [DJW97, Zie98, Cas22, CCL22] and safety specifications, i.e., objectives that are closed for the product topology [CFH14]. The latter also uses the order of the equivalence classes of the right congruence as part of its characterization.

Recently, a link between the prefix-classifier, the memory requirements, and the recognizability of ω-regular objectives was established [BRV23]. However, this result does not provide optimal bounds on the strategy complexity, and is thereby insufficient to study half-positionality.

Our article is an extended version (with complete proofs and additional examples and remarks) of a preceding conference version [BCRV22].

Structure of the paper. Notations and definitions are introduced in Section 2. Our main contributions are presented in Section 3: we introduce and discuss the three conditions used in our results, then we state our main characterization (Theorem 3.11) and some corollaries, and we end with an explanation on how to use the characterization to decide half-positionality of DBA-recognizable objectives in polynomial time. Section 4 and Section 5 contain the proof of Theorem 3.11: the former shows the necessity of the three conditions for half-positionality of DBA-recognizable objectives, and the latter shows their sufficiency through the use of universal graphs.

Conference version. Our article extends a work already published as a conference version [BCRV22] with many details. It contains in particular the full proofs of the statements and additional explanations, remarks, and examples.

2. Preliminaries

In the whole article, letter C refers to a (finite or infinite) non-empty set of *colors*. Given a set A, we write respectively A^*, A^+, and A^ω for the set of finite, non-empty finite, and infinite sequences of elements of A. We denote by ε the empty word.

2.1. **Games and positionality.**

Graphs. An *(edge-colored) graph* $\mathcal{G} = (V, E)$ is given by a non-empty set of *vertices* V (of any cardinality) and a set of *edges* $E \subseteq V \times C \times V$. We write $v \xrightarrow{c} v'$ if $(v, c, v') \in E$. We assume graphs to be *non-blocking*: for all $v \in V$, there exists $(v, c, v') \in E$. We allow graphs with infinite branching. For $v \in V$, an *infinite path of $\mathcal{G}$ from v* is an infinite sequence of edges $\pi = (v, c_1, v_1)(v_1, c_2, v_2)(v_2, c_3, v_3) \ldots \in E^\omega$. A *finite path of $\mathcal{G}$ from v* is a finite prefix in E^* of an infinite path of $\mathcal{G}$ from v. For convenience, we assume that there is a distinct *empty path starting from v* for every $v \in V$. If $\gamma = (v_0, c_1, v_1) \ldots (v_{n-1}, c_n, v_n)$ is a non-empty finite path of $\mathcal{G}$, we define $\mathsf{last}(\gamma) = v_n$. For an empty path γ_v starting from v, we define $\mathsf{last}(\gamma_v) = v$. An infinite (resp. finite) path $(v_0, c_1, v_1)(v_1, c_2, v_2) \ldots$ (resp. $(v_0, c_1, v_1) \ldots (v_{n-1}, c_n, v_n)$) is sometimes denoted $v_0 \xrightarrow{c_1} v_1 \xrightarrow{c_2} \ldots$ (resp. $v_0 \xrightarrow{c_1} \ldots \xrightarrow{c_n} v_n$). A graph $\mathcal{G} = (V, E)$ is *finite* if both V and E are finite. A graph is *strongly connected* if for every pair of vertices $(v, v') \in V \times V$ there is a path from v to v'. A *strongly connected component* of $\mathcal{G}$ is a maximal strongly connected subgraph.

Arenas and strategies. We consider two players $\mathcal{P}_1$ and $\mathcal{P}_2$. An *arena* is a tuple $\mathcal{A} = (V, V_1, V_2, E)$ such that (V, E) is a graph and V is the disjoint union of V_1 and V_2. Intuitively, vertices in V_1 are controlled by $\mathcal{P}_1$ and vertices in V_2 are controlled by $\mathcal{P}_2$. An arena $\mathcal{A} = (V, V_1, V_2, E)$ is a *one-player arena of $\mathcal{P}_1$* (resp. *of $\mathcal{P}_2$*) if $V_2 = \emptyset$ (resp. $V_1 = \emptyset$). Finite paths of (V, E) are called *histories of $\mathcal{A}$*. For $i \in \{1, 2\}$, we denote by $\mathsf{Hists}_i(\mathcal{A})$ the set of histories γ of $\mathcal{A}$ such that $\mathsf{last}(\gamma) \in V_i$.

Let $i \in \{1, 2\}$. A *strategy of $\mathcal{P}_i$ on $\mathcal{A}$* is a function $\sigma_i \colon \mathsf{Hists}_i(\mathcal{A}) \to E$ such that for all $\gamma \in \mathsf{Hists}_i(\mathcal{A})$, the first component of $\sigma_i(\gamma)$ coincides with $\mathsf{last}(\gamma)$. Given a strategy σ_i of $\mathcal{P}_i$, we say that an infinite path $\pi = e_1 e_2 \ldots$ is *consistent with σ_i* if for all finite prefixes $\gamma = e_1 \ldots e_j$ of π such that $\mathsf{last}(\gamma) \in V_i$, $\sigma_i(\gamma) = e_{j+1}$. A strategy σ_i is *positional* (also called *memoryless* in the literature) if its outputs only depend on the current vertex and not on the whole history, i.e., if there exists a function $f \colon V_i \to E$ such that for $\gamma \in \mathsf{Hists}_i(\mathcal{A})$, $\sigma_i(\gamma) = f(\mathsf{last}(\gamma))$.

Objectives. An *objective* is a set $W \subseteq C^\omega$ (subsets of C^ω are sometimes also called *languages of infinite words*, *ω-languages*, or *winning conditions* in the literature). When an objective W is clear in the context, we say that an infinite word $w \in C^\omega$ is *winning* if $w \in W$, and *losing* if $w \notin W$. We write $\overline{W}$ for the complement $C^\omega \setminus W$ of an objective W. An objective W is *prefix-independent* if for all $w \in C^*$ and $w' \in C^\omega$, $w' \in W$ if and only if $ww' \in W$. An objective that we will often consider is the *Büchi condition*: given a subset $F \subseteq C$, we denote by $\mathsf{Büchi}(F)$ the set of infinite words containing infinitely many occurrences of colors in F. Such an objective is prefix-independent. A *game* is a tuple $(\mathcal{A}, W)$ of an arena $\mathcal{A}$ and an objective W.

Optimality and half-positionality. Let $\mathcal{A} = (V, V_1, V_2, E)$ be an arena, $(\mathcal{A}, W)$ be a game, and $v \in V$. We say that *a strategy σ_1 of $\mathcal{P}_1$ is winning from v* if for all infinite paths $v_0 \xrightarrow{c_1} v_1 \xrightarrow{c_2} \ldots$ from v consistent with σ_1, $c_1 c_2 \ldots \in W$.

A strategy of $\mathcal{P}_1$ is *optimal for $\mathcal{P}_1$ in $(\mathcal{A}, W)$* if it is winning from all the vertices from which $\mathcal{P}_1$ has a winning strategy. We often write *optimal for $\mathcal{P}_1$ in $\mathcal{A}$* if the objective W is clear from the context. We stress that this notion of optimality requires a *single* strategy to be winning from *all* the winning vertices (a property sometimes called *uniformity*). However, we do not require an optimal strategy to win when starting from some losing vertex, even if winning becomes possible eventually due to opponent's mistakes.

An objective W is *half-positional* if for all arenas $\mathcal{A}$, there exists a positional strategy of $\mathcal{P}_1$ on $\mathcal{A}$ that is optimal for $\mathcal{P}_1$ in $\mathcal{A}$. We sometimes only consider half-positionality on a restricted set of arenas (typically, finite and/or one-player arenas). For a class of arenas $\mathcal{X}$, an objective W is *half-positional over $\mathcal{X}$* if for all arenas $\mathcal{A} \in \mathcal{X}$, there exists a positional strategy of $\mathcal{P}_1$ on $\mathcal{A}$ that is optimal for $\mathcal{P}_1$ in $\mathcal{A}$.

Remark 2.1 (ε-edges)**.** Sometimes, arenas are considered to be colored over the alphabet $C \cup \{\varepsilon\}$, adding the restriction that no cycle is entirely labeled with ε [DJW97, Zie98, Kop07, Hor09, Cas22]. In that case, an infinite word in $(C \cup \{\varepsilon\})^\omega$ labeling a path is winning if the word obtained by removing the occurrences of ε belongs to W. In this paper, we consider arenas without ε-edges, but all our results apply to this other setting (cf. Remark 5.3). In general, allowing for arenas with ε-edges has an effect on strategy complexity [Cas22]. ◁

2.2. Büchi automata.

Automaton structures and Büchi automata. A *non-deterministic automaton structure (on C)* is a tuple $\mathcal{S} = (Q, C, Q_{\mathsf{init}}, \Delta)$ such that Q is a finite set of *states*, $Q_{\mathsf{init}} \subseteq Q$ is a non-empty set of *initial states* and $\Delta \subseteq Q \times C \times Q$ is a set of *transitions*. We assume that all states of automaton structures are reachable from an initial state in Q_{init} by taking transitions in Δ.

A *(transition-based) non-deterministic Büchi automaton* (NBA) is an automaton structure $\mathcal{S}$ together with a set of transitions $\widehat{\alpha} \subseteq \Delta$. The transitions in $\widehat{\alpha}$ are called *Büchi transitions*. Given an NBA $\mathcal{B} = (Q, C, Q_{\mathsf{init}}, \Delta, \widehat{\alpha})$, a *(finite or infinite) run of $\mathcal{B}$ on a (finite or infinite) word* $w = c_1 c_2 \ldots \in C^* \cup C^\omega$ is a sequence $(q_0, c_1, q_1)(q_1, c_2, q_2) \ldots \in \Delta^* \cup \Delta^\omega$ such that $q_0 \in Q_{\mathsf{init}}$. An infinite run $(q_0, c_1, q_1)(q_1, c_2, q_2) \ldots \in \Delta^\omega$ of $\mathcal{B}$ is *accepting* if for infinitely many $i \geq 0$, $(q_i, c_{i+1}, q_{i+1}) \in \widehat{\alpha}$. A word $w \in C^\omega$ is *accepted* by $\mathcal{B}$ if there exists an accepting run of $\mathcal{B}$ on w — if not, it is *rejected*. We denote the set of infinite words accepted by $\mathcal{B}$ by $\mathcal{L}(\mathcal{B})$, and we then say that $\mathcal{L}(\mathcal{B})$ is the objective *recognized by $\mathcal{B}$*.

Here, we take the definition of an *ω-regular objective* as an objective that can be recognized by an NBA (the classical definition uses *ω-regular expressions*, but our definition is well-known to be equivalent [McN66]). Given an automaton structure $\mathcal{S} = (Q, C, Q_{\mathsf{init}}, \Delta)$, we say that an NBA $\mathcal{B}$ is *built on top of $\mathcal{S}$* if there exists $\widehat{\alpha} \subseteq \Delta$ such that $\mathcal{B} = (Q, C, Q_{\mathsf{init}}, \Delta, \widehat{\alpha})$.

Remark 2.2. Notice that for generality, we allow the set of colors C to be infinite (as opposed to the finite set of states Q). This is uncommon for ω-regular objectives and automata, but has no real impact: when an objective is specified by an NBA, there will be at most finitely many equivalence classes of colors for the equivalence relation "inducing exactly the same transitions in the NBA". In Section 4.3.2, it will be helpful to consider infinite sets of colors.   ◁

Deterministic automata. An automaton structure $\mathcal{S} = (Q, C, Q_{\mathsf{init}}, \Delta)$ is *deterministic* if $|Q_{\mathsf{init}}| = 1$ and, for each $q \in Q$ and $c \in C$, there is exactly one $q' \in Q$ such that $(q, c, q') \in \Delta$ (we remark that, without loss of generality, we define deterministic automaton structures so that for each state and each color there is one outgoing transition — such automata are sometimes called *complete* or *total*). For deterministic automaton structures, we use the simplified notation $\mathcal{S} = (Q, C, q_{\mathsf{init}}, \delta)$, which corresponds to structure $(Q, C, Q_{\mathsf{init}}, \Delta)$ where q_{init} is the unique element of Q_{init} and $\delta \colon Q \times C \to Q$ is the *update function* that associates to $(q, c) \in Q \times C$ the unique $q' \in Q$ such that $(q, c, q') \in \Delta$. We denote by δ^* the natural extension of δ to finite words — by induction, the function $\delta^* \colon Q \times C^* \to Q$ is such that $\delta^*(q, \varepsilon) = q$, and for $w \in C^*$, $c \in C$, $\delta^*(q, wc) = \delta(\delta^*(q, w), c)$.

A *deterministic Büchi automaton* (DBA) is an NBA whose underlying automaton structure is deterministic. For DBA, we use the simplified notation $\mathcal{B} = (Q, C, q_{\mathsf{init}}, \delta, \alpha)$, which corresponds to NBA $(Q, C, Q_{\mathsf{init}}, \Delta, \widehat{\alpha})$ where q_{init} and δ are defined as for deterministic structures and $\alpha \subseteq Q \times C$ is the projection of $\widehat{\alpha}$ to the first two components (which does not lose information thanks to determinism).

For a DBA $\mathcal{B}$, a state $q \in Q$, and a word $w = c_1 c_2 \ldots \in C^* \cup C^\omega$, we denote by $\mathcal{B}(q, w) = (q, c_1, q_1)(q_1, c_2, q_2) \ldots \in \Delta^* \cup \Delta^\omega$ the only run on w starting from q.

An objective W is *DBA-recognizable* if there exists a DBA $\mathcal{B}$ such that $W = \mathcal{L}(\mathcal{B})$. For $F \subseteq C$, notice that $\mathsf{Büchi}(F)$ is DBA-recognizable: it is recognized by the DBA $(\{q_{\mathsf{init}}\}, C, q_{\mathsf{init}}, \delta, \alpha)$ with a *single* state such that $(q_{\mathsf{init}}, c) \in \alpha$ if and only if $c \in F$.

Remark 2.3. The fact that an automaton with a single state suffices to recognize $\mathsf{Büchi}(F)$ relies on the assumption that our DBA are *transition-based* and not *state-based* (α is a set of transitions, not of states). Indeed, apart from the trivial cases $F = \emptyset$ and $F = C$, a state-based DBA recognizing $\mathsf{Büchi}(F)$ requires two states.

For standard acceptance conditions, including the Büchi one, there is no difference of expressivity between state-based and transition-based automata: every state-based automaton can be converted to an equivalent transition-based automaton (with as many states in general) and every transition-based automaton can be converted to an equivalent state-based automaton (with more states in general). Considering transition-based acceptance is especially useful in our paper as it enables to state Condition (3) of the upcoming characterization (Theorem 3.11); an equivalent formulation for state-based automata would arguably be more technical and less intuitive. For instance, Condition (3) does not apply as stated to the simple $\mathsf{Büchi}(F)$ if we only consider state-based DBA. Moreover, the use of transition-based automata allows the use of techniques such as the saturation of automata with Büchi transitions, which is used repeatedly via the application of Lemma 2.9. ◁

In this paper, all automata will be deterministic, and the term "automaton" will stand for "deterministic automaton" by default.

Remark 2.4. Deterministic Büchi automata recognize a proper subset of the ω-regular objectives. That is, not every non-deterministic Büchi automaton can be converted into a deterministic one recognizing the same objective [Wag79]. ◁

Right congruence and prefix preorder. Let $W \subseteq C^\omega$ be an objective. For a finite word $w \in C^*$, we write $w^{-1}W = \{w' \in C^\omega \mid ww' \in W\}$ for the set of *winning continuations of* w. We define the *right congruence* $\sim_W \subseteq C^* \times C^*$ *of* W as $w_1 \sim_W w_2$ if $w_1^{-1}W = w_2^{-1}W$ (meaning that w_1 and w_2 have the same winning continuations). Relation $\sim_W$ is an equivalence relation. When W is clear from the context, we write $\sim$ for $\sim_W$ and, for $w \in C^*$, we denote by $[w] \subseteq C^*$ its equivalence class of $\sim$.

When $\sim$ has finitely many equivalence classes, we can associate a natural deterministic automaton structure $\mathcal{S}_\sim = (Q_\sim, C, \widetilde{q}_{\mathsf{init}}, \delta_\sim)$ to $\sim$ such that $Q_\sim$ is the set of equivalence classes of $\sim$, $\widetilde{q}_{\mathsf{init}} = [\varepsilon]$, and $\delta_\sim([w], c) = [wc]$ [Sta83, MS97]. The transition function $\delta_\sim$ is well-defined since it follows from the definition of $\sim$ that if $w_1 \sim w_2$, then for all $c \in C$, $w_1 c \sim w_2 c$. Hence, the choice of representatives for the equivalence classes does not have an impact on this definition. We call the automaton structure $\mathcal{S}_\sim$ the *prefix-classifier of* W.

Remark 2.5. Equivalence relation $\sim_W$ has only one equivalence class if and only if W is prefix-independent. In particular, an objective has a prefix-classifier with a single state if and only if it is prefix-independent. ◁

An important property of $\sim$ is the following.

Lemma 2.6. *Let* $w_1, w_2 \in C^*$. *If* $w_1 \sim w_2$, *then for all* $w \in C^*$, $w_1 w \sim w_2 w$.

Proof. If w_1 and w_2 have the same winning continuations, they have in particular the same winning continuations starting with w. □

We define the *prefix preorder* $\preceq_W$ *of* W: for $w_1, w_2 \in C^*$, we write $w_1 \preceq_W w_2$ if $w_1^{-1}W \subseteq w_2^{-1}W$ (meaning that any continuation that is winning after w_1 is also winning after w_2). Intuitively, $w_1 \preceq_W w_2$ means that a game starting with w_2 is always preferable

to a game starting with w_1 for $\mathcal{P}_1$, as there are more ways to win after w_2. When W is clear from the context, we write $\preceq$ for $\preceq_W$. Relation $\preceq \subseteq C^* \times C^*$ is a preorder. Notice that $\sim$ is equal to $\preceq \cap \succeq$. We also define the strict preorder $\prec = \preceq \setminus \sim$.

Given a DBA $\mathcal{B} = (Q, C, q_{\mathsf{init}}, \delta, \alpha)$ recognizing the objective W, observe that for $w, w' \in C^*$ such that $\delta^*(q_{\mathsf{init}}, w) = \delta^*(q_{\mathsf{init}}, w')$, we have $w \sim w'$. In this case, equivalence relation $\sim$ has at most $|Q|$ equivalence classes. For $q \in Q$, we write $\mathcal{B}^q$ for the automaton $(Q, C, q, \delta, \alpha)$, with q as the initial state and everything else as in $\mathcal{B}$. Observe that objective $\mathcal{L}(\mathcal{B}^q)$ equals $w^{-1}W$ for any word $w \in C^*$ such that $\delta^*(q_{\mathsf{init}}, w) = q$. This means that as long as q is reachable from q_{init}, objective $\mathcal{L}(\mathcal{B}^q)$ is a set of winning continuations of W. We extend the equivalence relation $\sim$ and preorder $\preceq$ to elements of Q (we sometimes write $\sim_{\mathcal{B}}$ and $\preceq_{\mathcal{B}}$ to avoid any ambiguity): $q \sim q'$ if $\mathcal{L}(\mathcal{B}^q) = \mathcal{L}(\mathcal{B}^{q'})$, and $q \preceq q'$ if $\mathcal{L}(\mathcal{B}^q) \subseteq \mathcal{L}(\mathcal{B}^{q'})$. Transitions of $\mathcal{B}$ following a given word are non-decreasing w.r.t. states.

Lemma 2.7. *Let* $\mathcal{B} = (Q, C, q_{\mathsf{init}}, \delta, \alpha)$ *be a DBA. For every* $w \in C^*$, *function* $\delta^*(\cdot, w)$ *is non-decreasing for* $\preceq_{\mathcal{B}}$ *(i.e., for* $q_1, q_2 \in Q$, *if* $q_1 \preceq_{\mathcal{B}} q_2$, *then* $\delta^*(q_1, w) \preceq_{\mathcal{B}} \delta^*(q_2, w)$).

Proof. Let W be the objective that $\mathcal{B}$ recognizes. Let $w \in C^*$ and $q_1, q_2 \in Q$ be such that $q_1 \preceq_{\mathcal{B}} q_2$. Let $q_1' = \delta^*(q_1, w)$ and $q_2' = \delta^*(q_2, w)$. We show that $q_1' \preceq_{\mathcal{B}} q_2'$, i.e., that $\mathcal{L}(\mathcal{B}^{q_1'}) \subseteq \mathcal{L}(\mathcal{B}^{q_2'})$. Let $w' \in \mathcal{L}(\mathcal{B}^{q_1'})$. This implies that $ww' \in \mathcal{L}(\mathcal{B}^{q_1})$. As $q_1 \preceq_{\mathcal{B}} q_2$, we also have that $ww' \in \mathcal{L}(\mathcal{B}^{q_2})$. This implies that $w' \in \mathcal{L}(\mathcal{B}^{q_2'})$. $\qquad\square$

α-free words. Let $\mathcal{B} = (Q, C, q_{\mathsf{init}}, \delta, \alpha)$ be a DBA. We say that a finite run ϱ of $\mathcal{B}$ is *α-free* if it does not contain any transition from α. For $q \in Q$, we define

$$\alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q) = \{w \in C^* \mid \mathcal{B}(q, w) \text{ is } \alpha\text{-free}\},$$
$$\alpha\text{-}\mathsf{FreeCycles}_{\mathcal{B}}(q) = \{w \in C^* \mid w \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q) \text{ and } \delta^*(q, w) = q\}.$$

We call the words in the first set the *α-free words from q*, and the words in the second set the *α-free cycles from q*. We state an important property of α-free words.

Lemma 2.8. *Let* $q_1, q_2 \in Q$ *be such that* $\alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_1) = \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_2)$. *Then for all* $w \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_1)$, $\alpha\text{-}\mathsf{Free}_{\mathcal{B}}(\delta^*(q_1, w)) = \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(\delta^*(q_2, w))$.

Proof. Let $w \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_1)$. Let $q_1' = \delta^*(q_1, w)$, $q_2' = \delta^*(q_2, w)$, and $w' \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_1')$. Since both runs $\mathcal{B}(q_1, w)$ and $\mathcal{B}(q_1', w')$ are α-free, we have $ww' \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_1) = \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_2)$. Therefore, the finite run $\mathcal{B}(q_2, ww')$ is α-free, so the run $\mathcal{B}(q_2', w')$ is α-free as well and $w' \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_2')$. $\qquad\square$

Saturation of DBA with Büchi transitions. In what follows, we will make extensive use of a "normal form" of Büchi automata satisfying that any α-free path can be extended to an α-free cycle. Such a normal form can be produced by saturating a given DBA $\mathcal{B}$ with Büchi transitions [KS15, AK22]. To do so, we add to α all transitions that do not appear in an α-free cycle of $\mathcal{B}$. These transitions can be easily identified: we remove all Büchi transitions from $\mathcal{B}$, and decompose the resulting automaton in strongly connected components; transitions that we can add to α are exactly those that are not included in any of those SCCs.

We say that $\mathcal{B} = (Q, C, q_{\mathsf{init}}, \delta, \alpha)$ is *saturated* if for every $\alpha' \supsetneq \alpha$, the automaton obtained by replacing α with α' does not recognize $\mathcal{L}(\mathcal{B})$.

An α-*free component* of $\mathcal{B}$ is a strongly connected component of the graph obtained by removing the Büchi transitions from $\mathcal{B}$. That is, it is a strongly connected component of the graph $(Q, \Delta_{\alpha\text{-free}})$, with $\Delta_{\alpha\text{-free}} = \{(q, c, \delta(q, c)) \in Q \times C \times Q \mid (q, c) \notin \alpha\}$.

Lemma 2.9. *Let $\mathcal{B} = (Q, C, q_{\text{init}}, \delta, \alpha)$ be a DBA. There is a unique set $\alpha_{\text{sat}} \subseteq Q \times C$ such that the automaton $\mathcal{B}_{\text{sat}} = (Q, C, q_{\text{init}}, \delta, \alpha_{\text{sat}})$ satisfies that:*

(1) $\mathcal{L}(\mathcal{B}_{\text{sat}}) = \mathcal{L}(\mathcal{B})$.

(2) $\mathcal{B}_{\text{sat}}$ *is saturated.*

Moreover, α_{sat} is the set of transitions not appearing in any α-free component of $\mathcal{B}$, and it can be computed in $\mathcal{O}(|C| \cdot |Q|)$ time when C is finite.

Proof. We first prove the existence of such an α_{sat}. Let $(Q_1, \Delta_1), \ldots, (Q_k, \Delta_k)$ be the α-free components of $\mathcal{B}$. We consider the automaton $\mathcal{B}_{\text{sat}}$ whose Büchi transitions are those that do not appear in any α-free component (Q_i, Δ_i), that is, we let

$$\alpha_{\text{sat}} = (Q \times C) \setminus \{(q, c) \mid (q, c, \delta(q, c)) \in \bigcup_{i=1}^{k} \Delta_i\} \text{ and } \mathcal{B}_{\text{sat}} = (Q, C, q_{\text{init}}, \delta, \alpha_{\text{sat}}).$$

We show that $\mathcal{L}(\mathcal{B}_{\text{sat}}) = \mathcal{L}(\mathcal{B})$. Since $\alpha \subseteq \alpha_{\text{sat}}$, it holds that $\mathcal{L}(\mathcal{B}) \subseteq \mathcal{L}(\mathcal{B}_{\text{sat}})$. For the other inclusion, let $w \notin \mathcal{L}(\mathcal{B})$. There are $w_0 \in C^*, w' \in C^\omega$ such that $w = w_0 w'$ and the finite run $\mathcal{B}(q_0, w')$ produced by reading w' from $q_0 = \delta^*(q_{\text{init}}, w_0)$ does not visit any Büchi transition. In particular, $\mathcal{B}(q_0, w')$ is an infinite path in the finite graph $(Q, \Delta_{\alpha\text{-free}})$. This implies that eventually, run $\mathcal{B}(q_0, w')$ reaches and stays in the same strongly connected component of graph $(Q, \Delta_{\alpha\text{-free}})$. Formally, there are $w_1 \in C^*, w_2 \in C^\omega$ such that $w' = w_1 w_2$ and the infinite run $\mathcal{B}(q_1, w_2)$ produced by reading w_2 from $q_1 = \delta^*(q_0, w_1)$ lies entirely in some α-free component (Q_i, Δ_i). Thus, all transitions in $\mathcal{B}(q_1, w_2)$ lie in Δ_i. Therefore, $\mathcal{B}_{\text{sat}}(q_{\text{init}}, w)$ is also a rejecting run of $\mathcal{B}_{\text{sat}}$ and $w \notin \mathcal{L}(\mathcal{B}_{\text{sat}})$.

We prove that $\mathcal{B}_{\text{sat}}$ is saturated and the uniqueness of α_{sat} at the same time. Let α' be another acceptance set such that $\alpha' \nsubseteq \alpha_{\text{sat}}$ and let $\mathcal{B}'$ be the automaton obtained by replacing α_{sat} with α' in $\mathcal{B}_{\text{sat}}$. Let $(q, c) \in \alpha' \setminus \alpha_{\text{sat}}$. Since $(q, c) \notin \alpha_{\text{sat}}$, there is an α-free component (Q_i, Δ_i) such that $(q, c, \delta(q, c)) \in \Delta_i$. We can therefore consider a word $w \in \alpha\text{-}\mathsf{FreeCycles}_{\mathcal{B}_{\text{sat}}}(q)$ labeling an α-free cycle including the transition $(q, c, \delta(q, c))$. Let $w_0 \in C^*$ such that $\delta^*(q_{\text{init}}, w_0) = q$. Then, $w_0 w^\omega \notin \mathcal{L}(\mathcal{B}_{\text{sat}})$, whereas $w_0 w^\omega \in \mathcal{L}(\mathcal{B}')$, so $\mathcal{B}'$ does not recognize the same objective as $\mathcal{B}_{\text{sat}}$.

When C is finite, the set of transitions $\mathcal{B}_{\text{sat}}$ can be computed in time $\mathcal{O}(|C| \cdot |Q|)$, as it consists of decomposing a graph with at most $|C| \cdot |Q|$ transitions into strongly connected components [Tar72]. □

Figure 1: A DBA (left) and its unique saturation (right). Transitions labeled with a • symbol are the Büchi transitions. In figures, automaton states are depicted with diamonds.

In Figure 1, we show an example of the saturation process presented in the proof of Lemma 2.9.

The following simple lemma follows, which holds true in saturated DBA: every word that is α-free from a state can be completed into an α-free cycle from the same state. This will be a key technical lemma in the upcoming proofs.

Lemma 2.10. *Let $\mathcal{B} = (Q, C, q_{\mathsf{init}}, \delta, \alpha)$ be a saturated DBA. Let $q \in Q$ and $w \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q)$. There exists $w' \in C^*$ such that $ww' \in \alpha\text{-}\mathsf{FreeCycles}_{\mathcal{B}}(q)$.*

Proof. Let $q' = \delta^*(q, w)$. Thanks to the saturation property and by Lemma 2.9, α contains all transitions that do not belong to an α-free component. This implies that any two states connected by an α-free run are in the same α-free component. In particular, as q' is reachable from q through an α-free run, q' must belong to the same α-free component as q. Therefore, there exists an α-free run from q' to q. Taking the word w' labeling this run, we obtain the desired result. □

3. Half-positionality characterization for DBA-recognizable objectives

In this section, we present our main contribution in Theorem 3.11, by giving three conditions that exactly characterize half-positional DBA-recognizable objectives. These conditions are presented in Section 3.1. Theorem 3.11 and several consequences of it are stated in Section 3.2 (the proof of Theorem 3.11 is postponed to Sections 4 and 5). In Section 3.3, we use this characterization to show that we can decide the half-positionality of a DBA in polynomial time.

3.1. **Three conditions for half-positionality.** We define the three conditions of objectives at the core of our characterization.

Condition 3.1 (Total prefix preorder)**.** We say that an objective $W \subseteq C^\omega$ has a *total prefix preorder* if for all $w_1, w_2 \in C^*$, $w_1 \preceq_W w_2$ or $w_2 \preceq_W w_1$.

Example 3.2 (Not total prefix preorder)**.** Let $C = \{a, b\}$. We consider the objective W recognized by the DBA $\mathcal{B}$ depicted in Figure 2 (left). It consists of the infinite words starting with aa or bb. This objective does not have a total prefix preorder: words a and b are incomparable for $\preceq_W$. Indeed, a^ω is winning after a but not after b, and b^ω is winning after b but not after a. In terms of automaton states, we have that q_a and q_b are incomparable for $\preceq_{\mathcal{B}}$. This objective is not half-positional, as witnessed by the arena on the right of Figure 2. In this arena, $\mathcal{P}_1$ is able to win when the game starts in v_1 by playing a in v_3, and when the game starts in v_2 by playing b. However, no positional strategy wins from both v_1 and v_2. ◁

Remark 3.3. The prefix preorder of an objective W is total if and only if the prefix preorder of its complement $\overline{W}$ is total. ◁

Remark 3.4. Having a total prefix preorder is equivalent to the *strong monotony* notion [BFMM11] in general, and equivalent to *monotony* [GZ05] for ω-regular objectives. We discuss in more depth the relation between the conditions appearing in the characterization and other properties from the literature studying half-positionality in Appendix A. ◁

A straightforward result for an objective W recognized by a DBA $\mathcal{B}$ is that it has a total prefix preorder if and only if the (reachable) states of $\mathcal{B}$ are totally ordered for $\preceq_{\mathcal{B}}$.

Figure 2: DBA $\mathcal{B}$ recognizing objective $W = (aa + bb)C^\omega$ (left), and an arena in which positional strategies do not suffice for $\mathcal{P}_1$ to play optimally for this objective (right). In figures, circles represent arena vertices controlled by $\mathcal{P}_1$.

Condition 3.5 (Progress-consistency). We say that an objective W is *progress-consistent* if for all $w_1 \in C^*$ and $w_2 \in C^+$ such that $w_1 \prec w_1w_2$, we have $w_1(w_2)^\omega \in W$.

Intuitively, this means that whenever a word w_2 can be used to make progress after starting with a word w_1 (in the sense of getting to a position in which more continuations are winning), then repeating this word has to be winning.

Example 3.6 (Not progress-consistent). Let $C = \{a, b\}$. We consider the objective $W = C^*aaC^\omega$ recognized by the DBA with three states in Figure 3 (left). This objective consists of the words containing, at some point, twice the color a in a row. Notice that the prefix preorder of this objective is total ($q_{\mathsf{init}} \prec q_a \prec q_{aa}$). This objective is not progress-consistent: we have $\varepsilon \prec ba$, but $(ba)^\omega \notin W$. This objective is not half-positional: if $\mathcal{P}_1$ plays in an arena with a choice among two cycles ba and ab depicted in Figure 3 (right), it is possible to win by playing ba and then ab, but a positional strategy can only achieve words $(ba)^\omega$ or $(ab)^\omega$, which are both losing.   ◁


Figure 3: A DBA recognizing the set of words containing aa at some point (left), an arena in which positional strategies do not suffice for $\mathcal{P}_1$ to play optimally for this objective (right). Squiggly arrows indicate a sequence of edges or transitions (here, a sequence of two edges).

Example 3.7 (Progress-consistent objective). We consider a slight modification of the previous example by adding two Büchi transitions: see the DBA in Figure 4. The objective recognized by this DBA is $W = \mathsf{Büchi}(\{a\}) \cup C^*aaC^\omega$: W is the set of words containing infinitely many occurrences of a, or containing a twice in a row at some point. The sets of winning continuations of W are $\mathcal{L}(\mathcal{B}^{q_{\mathsf{init}}}) = W$, $\mathcal{L}(\mathcal{B}^{q_a}) = aC^\omega \cup W$ and $\mathcal{L}(\mathcal{B}^{q_{aa}}) = C^\omega$. This objective is progress-consistent: any word reaching q_{aa} is straightforwardly accepted when repeated infinitely often, and any word w such that $\delta^*(q_{\mathsf{init}}, w) = q_a$ necessarily contains at least one a, and thus is accepted when repeated infinitely often. Objective W is half-positional, which will be readily shown with our upcoming characterization (Theorem 3.11).

Here, notice that the complement $\overline{W}$ of W is not progress-consistent. Indeed, $a \prec_{\overline{W}} a(bab)$, but $a(bab)^\omega \notin \overline{W}$. Unlike having a total prefix preorder, progress-consistency can hold for an objective but not its complement.

Note that half-positionality of W cannot be shown using existing half-positionality criteria [Kop06, BFMM11] (it is neither prefix-independent nor *concave*) or bipositionality criteria, as it is simply not bipositional.   ◁


Figure 4: A DBA recognizing the set of words containing infinitely many occurrences of a, or containing aa at some point.

Condition 3.8 (Recognizability by the prefix-classifier). Being recognized by a Büchi automaton built on top of the prefix-classifier is our third condition. In other words, for a DBA-recognizable objective $W \subseteq C^\omega$ and its prefix-classifier $\mathcal{S}_\sim = (Q_\sim, C, \widetilde{q}_{\mathsf{init}}, \delta_\sim)$, this condition requires that there exists $\alpha_\sim \subseteq Q_\sim \times C$ such that W is recognized by DBA $(Q_\sim, C, \widetilde{q}_{\mathsf{init}}, \delta_\sim, \alpha_\sim)$.

We show an example of a DBA-recognizable objective that satisfies the first two conditions (having a total prefix preorder and progress-consistency), but not this third condition, and which is not half-positional.

Example 3.9 (Not recognizable by the prefix-classifier). Let $C = \{a, b\}$. We consider the objective $W = \mathsf{Büchi}(\{a\}) \cap \mathsf{Büchi}(\{b\})$ recognized by the DBA in Figure 5. This objective is prefix-independent: as such (Remark 2.5), there is only one equivalence class of $\sim$. This implies that the prefix preorder is total, and that W is progress-consistent (the premise of the progress-consistency property can never be true). This objective is not half-positional, as witnessed by the arena in Figure 5 (right): $\mathcal{P}_1$ has a winning strategy from v, but it needs to take infinitely often both a and b.

Any DBA recognizing this objective has at least two states, but all their (reachable) states are equivalent for $\sim$ — no matter the state we choose as an initial state, the recognized objective is the same (by prefix-independence). As it is prefix-independent, its prefix-classifier $\mathcal{S}_\sim$ has only one state.   ◁

Figure 5: DBA recognizing the objective $\mathsf{Büchi}(\{a\}) \cap \mathsf{Büchi}(\{b\})$ (left), and an arena in which positional strategies do not suffice for $\mathcal{P}_1$ to play optimally for this objective (right).

As will be shown formally, being recognized by a DBA built on top of the prefix-classifier is necessary for half-positionality of *DBA-recognizable* objectives over finite one-player arenas. Unlike the two other conditions, it is in general not necessary for half-positionality of general objectives, including objectives recognized by other standard classes of ω-automata.

Example 3.10. We consider the complement $\overline{W}$ of the objective $W = \mathsf{Büchi}(\{a\}) \cap \mathsf{Büchi}(\{b\})$ of Example 3.9, which consists of the words ending with a^ω or b^ω. Objective $\overline{W}$ is not DBA-recognizable (a close proof can be found in [BK08, Theorem 4.50]). Still, it is recognizable by a *deterministic coBüchi automaton* similar to the automaton in Figure 5, but which accepts infinite words that visit transitions labeled with $\bullet$ only finitely often. This objective is half-positional, which can be shown using [DJW97, Theorem 6]. However, its prefix-classifier has just one state, and there is no way to recognize $\overline{W}$ by building a coBüchi (or even parity) automaton on top of it. ◁

3.2. Characterization and corollaries.
We have now defined the three conditions required for our characterization.

Theorem 3.11. *Let $W \subseteq C^\omega$ be a DBA-recognizable objective. Objective W is half-positional (over all arenas) if and only if*

- *its prefix preorder $\preceq$ is total,*
- *it is progress-consistent, and*
- *it can be recognized by a Büchi automaton built on top of its prefix-classifier $\mathcal{S}_\sim$.*

Proof. The proof of the necessity of the three conditions can be found in Section 4, respectively in Propositions 4.2, 4.3, and 4.4. The proof of the sufficiency of the conjunction of the three conditions can be found in Section 5, Proposition 5.4. □

This characterization is valuable to prove (and disprove) half-positionality of DBA-recognizable objectives. Examples 3.2, 3.6, and 3.9 are all not half-positional, and each of them falsifies exactly one of the three conditions from the statement. On the other hand, Example 3.7 is half-positional. We have already discussed its progress-consistency, but it is also straightforward to verify that its prefix preorder is total and that it is recognizable by its prefix-classifier: the right congruence has three totally ordered equivalence classes corresponding to the states of the DBA in Figure 4.

We state two notable consequences of Theorem 3.11 and of its proof technique. The first one is the specialization of Theorem 3.11 to prefix-independent objectives. It states that all prefix-independent, DBA-recognizable objectives that are half-positional are of the kind $\mathsf{Büchi}(F)$ for some $F \subseteq C$. Prefix-independence of objectives is a frequent assumption in the literature [Kop06, CN06, GK14a, CFGO22] — we show that under this assumption, half-positionality of DBA-recognizable objectives is very easy to understand and characterize.

Proposition 3.12. *Let $W \subseteq C^\omega$ be a prefix-independent, DBA-recognizable objective. Objective W is half-positional if and only if there exists $F \subseteq C$ such that $W = \mathsf{Büchi}(F)$.*

Proof. The right-to-left implication follows from the known half-positionality of objectives of the kind $\mathsf{Büchi}(F)$ (this is a special case of a parity game [EJ91]). For the left-to-right implication, we assume that W is a prefix-independent, DBA-recognizable, half-positional objective. By Theorem 3.11, it is recognized by a DBA $\mathcal{B}$ built on top of $\mathcal{S}_\sim$. As W is prefix-independent (Remark 2.5), its prefix-classifier has just one state, and there is a single transition from and to this single state for each color. Hence, $W = \mathsf{Büchi}(F)$, where F is the set of colors whose only transition is a Büchi transition of $\mathcal{B}$. □

Remark 3.13. A corollary of this result is that when W is prefix-independent, DBA-recognizable, and half-positional, we also have that $\overline{W}$ is half-positional. Indeed, the complement of objective $W = \mathsf{B\ddot{u}chi}(F)$ is a so-called *coBüchi objective*, which is also known to be half-positional [EJ91]. This statement does not hold in general when W is not prefix-independent, as was shown in Example 3.7. Moreover, the reciprocal of the statement also does not hold, as was shown in Example 3.10.   ◁

Remark 3.14. A second corollary is that prefix-independent DBA-recognizable half-positional objectives are closed under finite union (since a finite union of Büchi conditions is a Büchi condition). This settles Kopczyński's conjecture for DBA-recognizable objectives. ◁

A second consequence of Theorem 3.11 and its proof technique shows that half-positionality of DBA-recognizable objectives can be reduced to half-positionality over the restricted class of *finite, one-player arenas*. Results reducing strategy complexity in two-player arenas to the easier question of strategy complexity in one-player arenas are sometimes called *one-to-two-player lifts* and appear in multiple places in the literature [GZ05, BLO+22, Koz22b, BRV23].

Proposition 3.15 (One-to-two-player and finite-to-infinite lift)**.** *Let $W \subseteq C^\omega$ be a DBA-recognizable objective. If objective W is half-positional over finite one-player arenas, then it is half-positional over all arenas (of any cardinality).*

Proof. When showing the necessity of the three conditions for half-positionality of DBA-recognizable objectives in Section 4 (Propositions 4.2, 4.3, and 4.4), we actually show their necessity for half-positionality over *finite one-player* arenas. Hence, assuming half-positionality over finite one-player arenas, we have the three conditions from the characterization of Theorem 3.11, so we have half-positionality over all arenas.   □

One-to-two-player lifts from the literature all require an assumption on the strategy complexity of *both* players, and are either stated solely over finite arenas, or solely over infinite arenas. Proposition 3.15, albeit set in the more restricted context of DBA-recognizable objectives, displays stronger properties than the known one-to-two-player lifts.

- It is *asymmetric* in the sense that we simply need a hypothesis on *one* player: *half*-positionality of DBA-recognizable objectives over one-player arenas implies their *half*-positionality over two-player arenas.
- It shows that half-positionality of DBA-recognizable objectives over *finite* arenas implies half-positionality over *infinite* arenas.

Both these properties do not hold for general objectives.

- Some objectives are half-positional over one-player but not over two-player arenas [GK14b, Section 7] — we have that this is not possible for DBA-recognizable objectives.
- Some objectives are half-positional over finite but not over infinite arenas (see, e.g., the mean-payoff objective [EM79, Put94]) — we have that this is not possible for DBA-recognizable objectives.

3.3. **Deciding half-positionality in polynomial time.** In this section, we assume that C is finite. We show that the problem of deciding, given a DBA $\mathcal{B} = (Q, C, q_{\mathsf{init}}, \delta, \alpha)$ as an input, whether $\mathcal{L}(\mathcal{B})$ is half-positional can be solved in polynomial time, and more precisely in time $\mathcal{O}(|C| \cdot |Q|^4)$. We will use the following classical result about the complexity of checking the inclusion of languages recognized by DBA.

Lemma 3.16. *Let $\mathcal{B} = (Q, C, q_{\mathsf{init}}, \delta, \alpha)$ and $\mathcal{B}' = (Q', C, q'_{\mathsf{init}}, \delta', \alpha')$ be two DBA on the same finite set of colors C. Deciding whether $\mathcal{L}(\mathcal{B}) \subseteq \mathcal{L}(\mathcal{B}')$ can be done in time $\mathcal{O}(|C| \cdot |Q| \cdot |Q'|)$.*

Proof. Observe first that $\mathcal{L}(\mathcal{B}) \subseteq \mathcal{L}(\mathcal{B}')$ if and only if $\mathcal{L}(\mathcal{B}) \cap \overline{\mathcal{L}(\mathcal{B}')} = \emptyset$. Therefore, we first want to complement $\mathcal{B}'$ (as an NBA $\overline{\mathcal{B}'}$), and then check the emptiness of the intersection of the languages of $\mathcal{B}$ and $\overline{\mathcal{B}'}$.

We use the fact that DBA can be easily complemented (as NBA): there is an NBA $\overline{\mathcal{B}'}$ with at most $2 \cdot |Q'|$ states and at most two non-deterministic transitions for each element of $C \times Q'$ such that $\mathcal{L}(\overline{\mathcal{B}'}) = \overline{\mathcal{L}(\mathcal{B}')}$ [Kur87].[1]

Then, we use the fact that we can construct an NBA $\mathcal{B}_\cap$ with at most $2 \cdot |Q| \cdot (2 \cdot |Q'|) = \mathcal{O}(|Q| \cdot |Q'|)$ states and $\mathcal{O}(|C| \cdot |Q| \cdot |Q'|)$ transitions such that $\mathcal{L}(\mathcal{B}_\cap) = \mathcal{L}(\mathcal{B}) \cap \mathcal{L}(\overline{\mathcal{B}'})$ [Bok18, Section 3.1].

The emptiness of the language of $\mathcal{B}_\cap$ can finally be checked in time $\mathcal{O}(|Q| \cdot |Q'| + |C| \cdot |Q| \cdot |Q'|) = \mathcal{O}(|C| \cdot |Q| \cdot |Q'|)$ by decomposing the graph of the NBA into strongly connected components [Tar72]. $\qquad\square$

We investigate how to verify each property used in the characterization of Theorem 3.11. Let $\mathcal{B} = (Q, C, q_{\mathsf{init}}, \delta, \alpha)$ be a DBA (we assume w.l.o.g. that all states in Q are reachable from q_{init}) and $W = \mathcal{L}(\mathcal{B})$ be the objective it recognizes. Our algorithm first verifies that the prefix preorder is total and recognizability by $\mathcal{S}_\sim$, and then, under these first two assumptions, progress-consistency. For each condition, we sketch an algorithm to decide it, and we discuss the time complexity of this algorithm.

Total prefix preorder. To check that W has a total prefix preorder, it suffices to check that the states of $\mathcal{B}$ are totally preordered by $\preceq_\mathcal{B}$. We start by computing, for each pair of states $q, q' \in Q$, whether $q \preceq_\mathcal{B} q'$, $q' \preceq_\mathcal{B} q$, or none of these. This can be rephrased as an *inclusion problem* for two DBA-recognizable objectives: if $\mathcal{B}^q = (Q, C, q, \delta, \alpha)$ and $\mathcal{B}^{q'} = (Q, C, q', \delta, \alpha)$, we have that $q \preceq_\mathcal{B} q'$ if and only if $\mathcal{L}(\mathcal{B}^q) \subseteq \mathcal{L}(\mathcal{B}^{q'})$. By Lemma 3.16, this can therefore be decided in time $\mathcal{O}(|C| \cdot |Q|^2)$. We can therefore know for all $|Q|^2$ pairs $q, q' \in Q$ whether $q \preceq_\mathcal{B} q'$, $q' \preceq_\mathcal{B} q$, $q' \sim_\mathcal{B} q$ (as $\sim_\mathcal{B} = \preceq_\mathcal{B} \cap \succeq_\mathcal{B}$), or none of these in time $\mathcal{O}(|Q|^2 \cdot (|C| \cdot |Q|^2)) = \mathcal{O}(|C| \cdot |Q|^4)$. In particular, the prefix preorder is total if and only if for all $q, q' \in Q$, we have $q \preceq_\mathcal{B} q'$ or $q' \preceq_\mathcal{B} q$.

Recognizability by the prefix-classifier. After all the relations $\preceq_\mathcal{B}$ and $\sim_\mathcal{B}$ between pairs of states are computed in the previous step, we can compute the states and transitions of the prefix-classifier $\mathcal{S}_\sim = (Q_\sim, C, \widetilde{q}_{\mathsf{init}}, \delta_\sim)$ by merging all the equivalence classes of $\sim_\mathcal{B}$. We assume for simplicity that $Q_\sim = Q/_{\sim_\mathcal{B}}$.

We now wonder whether it is possible to recognize W with a Büchi automaton built on top of $\mathcal{S}_\sim$, i.e., if it possible to find a set $\alpha_\sim \subseteq Q_\sim \times C$ of transitions of $\mathcal{S}_\sim$ such that the language of DBA $(Q_\sim, C, \widetilde{q}_{\mathsf{init}}, \delta_\sim, \alpha_\sim)$ is W. We simplify the search for such a set with the following result, which shows that when $\mathcal{B}$ is saturated, it suffices to try with one specific set $\alpha_\sim$. We can then simply check whether $W = \mathcal{L}((Q_\sim, C, \widetilde{q}_{\mathsf{init}}, \delta_\sim, \alpha_\sim))$, an equivalence

---

[1]The construction to complement a DBA as an NBA is simple: keep the original DBA, add a copy of it in which the transitions in α' are removed, add non-deterministic transitions from the original to the copy mimicking the original transitions, and define as Büchi transitions only the transitions left in the copy. In this way, accepted words in the new NBA are the ones that, from some point on (guessed non-deterministically), move to the copy and only use transitions that were not Büchi transitions of the original DBA.

query which can be performed in time $\mathcal{O}(|C| \cdot |Q|^2)$ by checking two inclusion queries as in Lemma 3.16.

Lemma 3.17. *We assume that $\mathcal{B}$ is saturated and that W is recognized by a DBA built on top of the prefix-classifier $\mathcal{S}_\sim = (Q_\sim, C, \widetilde{q}_{\mathsf{init}}, \delta_\sim)$. We define*

$$\alpha_\sim = \{([q], c) \in Q_\sim \times C \mid \forall q' \in [q], (q', c) \in \alpha\}.$$

Then, W is recognized by $(Q_\sim, C, \widetilde{q}_{\mathsf{init}}, \delta_\sim, \alpha_\sim)$.

Proof. We assume that W is recognized by a DBA built on top of $\mathcal{S}_\sim$. We start by saturating this DBA, which yields a set of Büchi transitions α' such that W is also recognized by the saturated DBA $\mathcal{B}' = (Q_\sim, C, \widetilde{q}_{\mathsf{init}}, \delta_\sim, \alpha')$ (Lemma 2.9). To prove the claim, we show that $\alpha' = \alpha_\sim$.

We first show that $\alpha' \subseteq \alpha_\sim$. Let $([q], c) \notin \alpha_\sim$ — we show that $([q], c) \notin \alpha'$. As $([q], c) \notin \alpha_\sim$, by definition of $\alpha_\sim$, there is $q' \in Q$ such that $(q', c) \notin \alpha$. As $\mathcal{B}$ is saturated, by Lemma 2.10, there exists $w' \in C^*$ such that $cw' \in \alpha\text{-}\mathsf{FreeCycles}_\mathcal{B}(q')$. By construction of the prefix-classifier, $\delta_\sim^*([q], cw') = [q]$. Also, as $W = \mathcal{L}(\mathcal{B}')$, word $(cw')^\omega$ must be rejected from $[q]$ in $\mathcal{B}'$. Therefore, $([q], c)$ cannot be a Büchi transition of $\mathcal{B}'$ and is not in α'.

We now show that $\alpha_\sim \subseteq \alpha'$. Let $([q], c) \notin \alpha'$ — we show that $([q], c) \notin \alpha_\sim$. As $\mathcal{B}'$ is saturated, by Lemma 2.10, there exists $w' \in C^*$ such that $cw' \in \alpha\text{-}\mathsf{FreeCycles}_{\mathcal{B}'}([q])$. As $W = \mathcal{L}(\mathcal{B}')$, word $(cw')^\omega$ is rejected from any state in $[q]$ in $\mathcal{B}$. If for all $q' \in [q]$, (q', c) was in α, $(cw')^\omega$ would be accepted from all states in $[q]$ in $\mathcal{B}$. Hence, there exists $q' \in [q]$ such that $(q', c) \notin \alpha$. We conclude that $([q], c) \notin \alpha_\sim$. $\qquad\square$

Progress-consistency. We assume that we have already checked that W is recognizable by a Büchi automaton built on top of $\mathcal{S}_\sim$, and that we know the (total) ordering of the states. We show that checking progress-consistency, under these two hypotheses, can be done in polynomial time. We prove a lemma reducing the search for words witnessing that W is not progress-consistent to a problem computationally easier to investigate.

Lemma 3.18. *We assume that $\mathcal{B}$ is built on top of the prefix-classifier $\mathcal{S}_\sim$ and that the prefix preorder of W is total. Then, W is progress-consistent if and only if for all $q, q' \in Q$ with $q \prec_\mathcal{B} q'$,*

$$\{w \in C^+ \mid \delta^*(q, w) = q'\} \cap \alpha\text{-}\mathsf{FreeCycles}_\mathcal{B}(q') = \emptyset.$$

Proof. We prove the left-to-right implication by showing the contrapositive. We assume that there exist $q, q' \in Q$ with $q \prec_\mathcal{B} q'$ and $w \in C^+$ such that $\delta^*(q, w) = q'$ and $w \in \alpha\text{-}\mathsf{FreeCycles}_\mathcal{B}(q')$. Let $w_q \in C^*$ be a word such that $\delta^*(q_{\mathsf{init}}, w_q) = q$. We have that $w_q \prec w_q w$, but $w_q w^\omega$ is not accepted by $\mathcal{B}$ as w is a cycle on q' that does not contain any Büchi transition. Hence, W is not progress-consistent.

For the right-to-left implication, we again prove the contrapositive. We assume that W is not progress-consistent. Thus, there exist $w' \in C^*$ and $w \in C^+$ such that $w' \prec w'w$ and $w'w^\omega \notin W$. Let $q_1 = \delta^*(q_{\mathsf{init}}, w')$ and $q_2 = \delta^*(q_1, w)$ — we have $q_1 \prec q_2$. As $q_1 \prec q_2$, by Lemma 2.7, we have $\delta^*(q_1, w) = q_2 \preceq \delta^*(q_2, w)$. We distinguish two cases, using the fact that there is exactly one state per equivalence class of $\sim_\mathcal{B}$. We represent what happens in Figure 6.

- If $q_2 = \delta^*(q_2, w)$, we then have that $w \in \alpha\text{-}\mathsf{FreeCycles}_\mathcal{B}(q_2)$, and we have what we want with $q = q_1$ and $q' = q_2$.

- If not, we have that $q_2 \prec \delta^*(q_2, w)$. Let $q_3 = \delta^*(q_2, w)$. We can repeat the argument on q_2 and q_3: either $w \in \alpha\text{-FreeCycles}_{\mathcal{B}}(q_3)$ and we are done, or $q_3 \prec \delta^*(q_3, w)$. As there are finitely many states, this process necessarily ends with two states $q = q_n$ and $q' = q_{n+1}$ such that $\delta^*(q, w) = q'$ and $w \in \alpha\text{-FreeCycles}_{\mathcal{B}}(q')$. $\square$


Figure 6: Situation in the proof of Lemma 3.18.

Notice that for each pair of states $q, q' \in Q$, the sets $\{w \in C^+ \mid \delta(q, w) = q'\}$ and $\alpha\text{-FreeCycles}_{\mathcal{B}}(q')$ are both regular languages recognized by deterministic finite automata with at most $|Q|$ states. The emptiness of their intersection can be decided in time $\mathcal{O}(|C| \cdot |Q|^2)$ (by solving a reachability problem in the product of the two automata, using an argument similar to but easier than Lemma 3.16) [RS59]. Thanks to Lemma 3.18, we can therefore decide whether $\mathcal{B}$ is progress-consistent in time $\mathcal{O}(|Q|^2 \cdot (|C| \cdot |Q|^2)) = \mathcal{O}(|C| \cdot |Q|^4)$: for all $|Q|^2$ pairs of states $q, q' \in Q$, if $q \prec q'$, we test the emptiness of the intersection of these two regular languages.

Complexity wrap-up. By checking the three conditions as explained and in this order, the time complexities are respectively $\mathcal{O}(|C| \cdot |Q|^4)$, $\mathcal{O}(|C| \cdot |Q|^2)$, and $\mathcal{O}(|C| \cdot |Q|^4)$. This yields a time complexity of $\mathcal{O}(|C| \cdot |Q|^4)$ for the whole algorithm.

4. NECESSITY OF THE CONDITIONS

We prove that each of the three conditions introduced in Section 3.1 is necessary for half-positionality over finite one-player arenas of DBA-recognizable objectives. Each condition gets its own devoted subsection. For the first two conditions (having a total prefix preorder and progress-consistency), we also show for completeness that they are necessary for half-positionality of general objectives over countably infinite one-player arenas, and necessary for half-positionality of ω-regular objectives over finite one-player arenas. This distinction is worthwhile, as there exist half-positional objectives that are not ω-regular (see, e.g., *finitary Büchi* objectives [CF13]).

We start with a well-known lemma about ω-regular objectives: if two ω-regular objectives are not equal, then they are distinguished by an ultimately periodic word. Ultimately periodic words can easily be finitely represented, and this lemma will be used throughout this section to force some behaviors to appear in *finite* arenas.

Lemma 4.1. *Let W_1, W_2 be two ω-regular objectives. If $W_1 \neq W_2$, then there exist $w_1 \in C^*$ and $w_2 \in C^+$ such that either $w_1(w_2)^\omega \in W_1 \setminus W_2$, or $w_1(w_2)^\omega \in W_2 \setminus W_1$. In particular, if $W_1 \not\subseteq W_2$, then there exist $w_1 \in C^*$ and $w_2 \in C^+$ such that $w_1(w_2)^\omega \in W_1 \setminus W_2$.*

Proof. The first statement is standard and follows from McNaughton's theorem [McN66]. For the second statement, if $W_1 \not\subseteq W_2$, then $W_1 \setminus W_2 \neq \emptyset$. Objective $W_1 \setminus W_2 \neq \emptyset$ is ω-regular (as ω-regular objectives are closed by complement and intersection) and so is $\emptyset$. By the first statement, we take $w_1 \in C^*$ and $w_2 \in C^+$ such that either $w_1(w_2)^\omega \in (W_1 \setminus W_2) \setminus \emptyset$, or $w_1(w_2)^\omega \in \emptyset \setminus (W_1 \setminus W_2)$. As $\emptyset \setminus (W_1 \setminus W_2) = \emptyset$, we have $w_1(w_2)^\omega \in (W_1 \setminus W_2) \setminus \emptyset = W_1 \setminus W_2$. $\square$

4.1. **Total prefix preorder.** Having a total prefix preorder is necessary in general for half-positionality over countably infinite arenas, and even over finite arenas for ω-regular objectives.

Proposition 4.2. *Let $W \subseteq C^\omega$ be an objective. If W is half-positional over countably infinite one-player arenas, then its prefix preorder is total. If W is ω-regular and half-positional over finite one-player arenas, then its prefix preorder is total.*

Proof. By contrapositive, we assume that the prefix preorder of W is not total. Then, there exist two finite words $w_1, w_2 \in C^*$ such that $w_1 \not\preceq w_2$ and $w_2 \not\preceq w_1$. We can find two infinite continuations $w_1', w_2' \in C^\omega$ such that $w_1 w_1' \in W$, $w_2 w_1' \notin W$, $w_2 w_2' \in W$, and $w_1 w_2' \notin W$. Using these four words, we build a countably infinite one-player arena depicted in Figure 7 (left) for which $\mathcal{P}_1$ has no positional optimal strategy. Indeed, if the game started with w_1, $\mathcal{P}_1$ needs to reply with w_1' in v_3 to win, but if the game started with w_2, $\mathcal{P}_1$ needs to reply with w_2' in v_3 to win.

Moreover, if W is ω-regular, so are $w_1^{-1}W$ and $w_2^{-1}W$. By Lemma 4.1, we can therefore assume w.l.o.g. that $w_1' = x_1(y_1)^\omega$ and $w_2' = x_2(y_2)^\omega$ are ultimately periodic, and we can carry out similar arguments with the finite arena depicted in Figure 7 (right). $\square$


Figure 7: Arenas in which $\mathcal{P}_1$ cannot play optimally with a positional strategy used in the proof of Proposition 4.2. The arena on the right is used for the ω-regular case.

4.2. **Progress-consistency.** Progress-consistency is necessary for half-positionality over countably infinite arenas in general, and even over finite arenas for ω-regular objectives.

Proposition 4.3. *Let $W \subseteq C^\omega$ be an objective. If W is half-positional over countably infinite one-player arenas, then it is progress-consistent. If W is ω-regular and half-positional even over finite one-player arenas, then it is progress-consistent.*

Proof. By contrapositive, we assume that W is not progress-consistent. Then there exist $w_1 \in C^*$ and $w_2 \in C^+$ such that $w_1 \prec w_1 w_2$, but $w_1(w_2)^\omega \notin W$. As $w_1 \prec w_1 w_2$, there exists an infinite continuation $w' \in C^\omega$ such that $w_1 w' \notin W$ and $w_1 w_2 w' \in W$. Using these three words, we build a countably infinite one-player arena depicted in Figure 8 (left). In this arena, from vertex v_1, a positional strategy can only achieve words $w_1(w_2)^\omega$ or $w_1 w'$, which are both losing. However, there is a (non-positional) winning strategy achieving word $w_1 w_2 w'$.

If W is additionally ω-regular, using Lemma 4.1, we can assume w.l.o.g. that $w' = xy^\omega$ is ultimately periodic, and we can carry out similar arguments with the finite arena depicted in Figure 8 (right). $\square$

Figure 8: Arenas in which $\mathcal{P}_1$ cannot play optimally with a positional strategy used in the proof of Proposition 4.3. The arena on the right is used in the ω-regular case.

4.3. **Recognizability by the prefix-classifier.** We now prove that for a DBA-recognizable objective, being recognized by a Büchi automaton built on top of its prefix-classifier $\mathcal{S}_\sim$ is necessary for half-positionality.

Proposition 4.4. *Let $W \subseteq C^\omega$ be a DBA-recognizable objective that is half-positional over finite one-player arenas. Then, W is recognized by a Büchi automaton built on top of $\mathcal{S}_\sim$.*

The rest of Section 4.3 is devoted to the proof of this result, which is more involved than the proofs in Sections 4.1 and 4.2. We fix an objective $W \subseteq C^\omega$ recognized by a DBA $\mathcal{B} = (Q, C, q_{\mathsf{init}}, \delta, \alpha)$. We make the assumption that W is half-positional over finite one-player arenas. Our goal is to show that W can be defined by a Büchi automaton built on top of $\mathcal{S}_\sim$. We assume w.l.o.g. that $\mathcal{B}$ is saturated. Many upcoming arguments heavily rely on this assumption through the use of Lemma 2.10 (any α-free word can be completed into an α-free cycle).

Our proof first assumes in Section 4.3.1 that $\mathcal{B}$ recognizes a prefix-independent objective. We will then build on this first case to conclude for the general case in Section 4.3.2. We provide a proof sketch at the start of each subsection.

4.3.1. *Prefix-independent case.* We assume that the objective W recognized by $\mathcal{B}$ is prefix-independent, so all the states of $\mathcal{B}$ are equivalent for $\sim$. We want to show that W can be recognized by a Büchi automaton built on top of $\mathcal{S}_\sim$, and in this case, the automaton structure $\mathcal{S}_\sim$ has just one state. Therefore, we want to find $F \subseteq C$ such that $W = \mathsf{Büchi}(F)$. We start with a high level description of the proof technique.

Proof sketch. The goal is to find a suitable definition for F. To do so, we exhibit a state q_{max} of $\mathcal{B}$ that is "the most rejecting state of the automaton": it has the property that the set of α-free words from q_{max} contains the α-free words from all the other states (q_{max} is then called an α-*free-maximum*) and that the set of α-free cycles on q_{max} contains the α-free cycles on all the other states (it is also an α-*free-cycle-maximum*). We define F as the set of colors c such that $(q_{\mathsf{max}}, c) \in \alpha$.

We first show that if an α-free-maximum exists, we can assume w.l.o.g. that it is unique (Lemma 4.5). In Lemmas 4.7, 4.8 and 4.9, we show the existence of an α-free-cycle-maximum. This part of the proof relies on the half-positionality over finite one-player arenas of W. Finally, defining F using q_{max} as explained above, we prove that $W = \mathsf{Büchi}(F)$ (Lemma 4.10). □

We call a state $q_{\mathsf{max}} \in Q$ of $\mathcal{B}$ an α-*free-maximum* (resp. an α-*free-cycle-maximum*) if for all $q \in Q$, we have $\alpha\text{-}\mathsf{Free}_\mathcal{B}(q) \subseteq \alpha\text{-}\mathsf{Free}_\mathcal{B}(q_{\mathsf{max}})$ (resp. if for all $q \in Q$, we have $\alpha\text{-}\mathsf{FreeCycles}_\mathcal{B}(q) \subseteq \alpha\text{-}\mathsf{FreeCycles}_\mathcal{B}(q_{\mathsf{max}})$). We remark that if $\mathcal{B}$ is saturated, an α-free-cycle-maximum is also an α-free-maximum (this can be shown using Lemma 2.10).

We first show that we can remove states from $\mathcal{B}$, while still recognizing the same objective, until it has at most one α-free-maximum.

Lemma 4.5. *There exists a DBA $\mathcal{B}'$ recognizing W with at most one α-free-maximum.*

Proof. Assume that $q_{max}^1, q_{max}^2 \in Q$ are distinct α-free-maxima. In particular, we have that $\alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{max}^1) = \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{max}^2)$. We show that in such a situation, the objective recognized by $\mathcal{B}$ can be recognized by an automaton with one less state, in which we discard one of the two α-free-maxima. To simplify the upcoming arguments, we assume that $q_{init} = q_{max}^1$ (which is without loss of generality as all states of $\mathcal{B}$ are equivalent for $\sim$).

We define a new automaton in which we remove q_{max}^2 and redirect all its incoming transitions to q_{max}^1. Formally, let $\mathcal{B}' = (Q', C, q'_{init}, \delta', \alpha')$ with update function δ' be such that

- $Q' = Q \setminus \{q_{max}^2\}$, $q'_{init} = q_{max}^1$,
- for $q \in Q'$ and $c \in C$, if $\delta(q, c) = q_{max}^2$, then $\delta'(q, c) = q_{max}^1$; otherwise, $\delta'(q, c) = \delta(q, c)$,
- for $q \in Q'$ and $c \in C$, $(q, c) \in \alpha'$ if and only if $(q, c) \in \alpha$.

We also assume that states that are not reachable from q'_{init} in $\mathcal{B}'$ are removed from Q'.

We show that this automaton with (at least) one less state recognizes the same objective as $\mathcal{B}$. Let $w = c_1 c_2 \ldots \in C^\omega$ be an infinite word. We show that w is accepted by $\mathcal{B}$ if and only if it is accepted by $\mathcal{B}'$.

Let $\Delta_{\not\rightarrow q_{max}^2} = \{(q, c) \in Q' \times C \mid \delta(q, c) = q_{max}^2\}$ be the transitions of $\mathcal{B}'$ that were directed to q_{max}^2 in $\mathcal{B}$ but are now redirected to q_{max}^1 in $\mathcal{B}'$. Let ϱ be the infinite run of $\mathcal{B}$ on w, and $\varrho' = (q'_0, c_1)(q'_1, c_2) \ldots$ be the infinite run of $\mathcal{B}'$ on w. The two runs start by taking corresponding transitions, but differ once a transition in $\Delta_{\not\rightarrow q_{max}^2}$ is taken.

We first assume that ϱ' uses transitions in $\Delta_{\not\rightarrow q_{max}^2}$ only finitely many times. Then, there exists $k \in \mathbb{N}$ such that $q'_k = q_{max}^1$ and for all $l \geq k$, $(q'_l, c_{l+1}) \notin \Delta_{\not\rightarrow q_{max}^2}$. Let $w_{>k} = c_{k+1} c_{k+2} \ldots$ be the infinite word consisting of the colors taken after the last occurrence of a transition in $\Delta_{\not\rightarrow q_{max}^2}$. We have that

$$\mathcal{B} \text{ accepts } w \Longleftrightarrow \mathcal{B} \text{ accepts } w_{>k} \quad \text{as } \mathcal{B} \text{ recognizes a prefix-independent objective}$$

$$\Longleftrightarrow \mathcal{B}' \text{ accepts } w_{>k} \quad \text{as } w_{>k} \text{ visits exactly the same transitions as in } \mathcal{B}$$

$$\Longleftrightarrow \mathcal{B}' \text{ accepts } w \quad \text{as } c_1 \ldots c_k \text{ is a cycle on the initial state } q_{max}^1 \text{ of } \mathcal{B}'.$$

We now assume that ϱ' uses transitions in $\Delta_{\not\rightarrow q_{max}^2}$ infinitely many times. We decompose ϱ' into infinitely many finite runs $\varrho'_1, \varrho'_2, \ldots$ such that $\varrho' = \varrho'_1 \varrho'_2 \ldots$ and every run ϱ'_i contains exactly one transition in $\Delta_{\not\rightarrow q_{max}^2}$ as its last transition. This implies that all these finite runs start in state q_{max}^1. We represent run ϱ' in Figure 9. We define words $w_1, w_2, \ldots$ as the respective projection of runs $\varrho'_1, \varrho'_2, \ldots$ to their colors (we have $w = w_1 w_2 \ldots$). Notice that

$$\forall i \geq 1, w_i \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{max}^1) \Longleftrightarrow w_i \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}'}(q_{max}^1), \tag{4.1}$$

as the transitions used by w_i from q_{max}^1 in $\mathcal{B}'$ correspond to the transitions used by w_i from q_{max}^1 in $\mathcal{B}$ (this property is not true for all words, but this is true for these words that read only one transition in $\Delta_{\not\rightarrow q_{max}^2}$ as their last transition). We also have by construction that

$$\forall i \geq 1, \delta^*(q_{max}^1, w_i) = q_{max}^2. \tag{4.2}$$

We distinguish whether w is accepted or rejected by $\mathcal{B}'$.

Assume w is accepted by $\mathcal{B}'$. Then, we know that for infinitely many $i \in \mathbb{N}$, $w_i \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}'}(q_{max}^1)$. This implies that for these indices i, $w_i \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{max}^1)$ by (4.1). As q_{max}^1 is

an α-free-maximum, for all $q \in Q$, $w_i \notin \alpha\text{-Free}_\mathcal{B}(q)$ (this is simply the contrapositive of the definition of α-free-maximum). Hence, for infinitely many $i \in \mathbb{N}$, when w_i is read in $\mathcal{B}$ (no matter from where), a Büchi transition is used, so w is accepted by $\mathcal{B}$.

Assume w is rejected by $\mathcal{B}'$. Then there exists $k \in \mathbb{N}$ such that for all $l \geq k$, $w_l \in \alpha\text{-Free}_{\mathcal{B}'}(q^1_{\max})$. As $\mathcal{B}$ is prefix-independent, up to removing the start of w, we assume w.l.o.g. that $k = 1$. We show by induction that

$$\forall i \geq 1, \alpha\text{-Free}_\mathcal{B}(q^1_{\max}) = \alpha\text{-Free}_\mathcal{B}(\delta^*(q^1_{\max}, w_1 \ldots w_i)).$$

This is true for $i = 1$, as $\delta^*(q^1_{\max}, w_1) = q^2_{\max}$ by (4.2) and the fact that $q^1_{\max}$ and $q^2_{\max}$ are both α-free-maxima. Assume $\alpha\text{-Free}_\mathcal{B}(q^1_{\max}) = \alpha\text{-Free}_\mathcal{B}(\delta^*(q^1_{\max}, w_1 \ldots w_{i-1}))$ for some $i \geq 2$. Then, by Lemma 2.8, as $w_i \in \alpha\text{-Free}_\mathcal{B}(q^1_{\max})$, we have $\alpha\text{-Free}_\mathcal{B}(\delta^*(q^1_{\max}, w_i)) = \alpha\text{-Free}_\mathcal{B}(\delta^*(q^1_{\max}, w_1 \ldots w_{i-1}w_i))$. Property (4.2) gives $\alpha\text{-Free}_\mathcal{B}(\delta^*(q^1_{\max}, w_i)) = \alpha\text{-Free}_\mathcal{B}(q^2_{\max})$, which is itself equal to $\alpha\text{-Free}_\mathcal{B}(q^1_{\max})$. We now know that for all $i \geq 1$, $w_i \in \alpha\text{-Free}_\mathcal{B}(q^1_{\max})$ by (4.1). Therefore, we conclude that for all $i \geq 1$, $w_i \in \alpha\text{-Free}_\mathcal{B}(\delta^*(q^1_{\max}, w_1 \ldots w_{i-1}))$. In particular, w uses no Büchi transition when read from $q^1_{\max}$ in $\mathcal{B}$ and is also rejected by $\mathcal{B}$.

We have shown that $\mathcal{B}'$ is a DBA with fewer states than $\mathcal{B}$ recognizing W. If $\mathcal{B}'$ still has two or more α-free-maxima, we repeat our construction until there is at most one left. □


Figure 9: Features of run ϱ' when it takes infinitely many transitions in $\Delta_{\not\to q^2_{\max}}$.

Thanks to Lemma 4.5, we now assume w.l.o.g. that $\mathcal{B}$ has at most one α-free-maximum. We intend to show that there exists an α-free-cycle-maximum. To do so, we exhibit an (infinite) arena in which $\mathcal{P}_1$ has no winning strategy, which we prove by using half-positionality of W over finite one-player arenas. We then prove that the non-existence of an α-free-cycle-maximum would imply that $\mathcal{P}_1$ has a winning strategy in this arena (Lemma 4.9).

Let $\mathcal{A}_\mathcal{B}$ be the infinite one-player arena of $\mathcal{P}_1$ depicted in Figure 10. This arena consists of one vertex v with a choice to make among all non-empty words that are α-free cycles from some state of $\mathcal{B}$. Vertex v is the only vertex with multiple outgoing edges. The goal of the next three short lemmas is to show that in this arena, $\mathcal{P}_1$ has no winning strategy.


Figure 10: Infinite one-player arena $\mathcal{A}_\mathcal{B}$ of $\mathcal{P}_1$, with choices from v among every non-empty word in $\bigcup_{q \in Q} \alpha\text{-FreeCycles}_\mathcal{B}(q)$.

Lemma 4.6. *If $\mathcal{P}_1$ has a winning strategy in $\mathcal{A}_\mathcal{B}$, then $\mathcal{P}_1$ has a positional winning strategy.*

Proof. Suppose that there is a winning strategy of $\mathcal{P}_1$ in $\mathcal{A}_\mathcal{B}$. Let $w = w_1 w_2 \ldots$ be an infinite winning word such that for $i \geq 1$, $w_i \in \bigcup_{q \in Q} \alpha\text{-FreeCycles}_\mathcal{B}(q) \setminus \{\varepsilon\}$. Let $q_0 = q_{\mathsf{init}}$, and for $i \geq 1$, let $q_i = \delta^*(q_{\mathsf{init}}, w_1 \ldots w_i)$ be the current automaton state after reading the first i finite words composing w. As there are only finitely many automaton states and w is winning, there are $k, l \geq 1$ with $k < l$ such that $q_k = q_l$ and $w_{k+1} \ldots w_l \notin \alpha\text{-Free}_\mathcal{B}(q_k)$. Word $w_1 \ldots w_k (w_{k+1} \ldots w_l)^\omega$ is also a winning word and uses only finitely many different words in $\bigcup_{q \in Q} \alpha\text{-FreeCycles}_\mathcal{B}(q) \setminus \{\varepsilon\}$.

Hence, there is a finite restriction ("subarena") $\mathcal{A}'_\mathcal{B}$ of the arena $\mathcal{A}_\mathcal{B}$ with at most l choices in v in which $\mathcal{P}_1$ has a winning strategy. Arena $\mathcal{A}'_\mathcal{B}$ being finite and one-player, half-positionality of W over finite one-player arenas implies that $\mathcal{P}_1$ has a positional winning strategy in $\mathcal{A}'_\mathcal{B}$. This positional winning strategy can also be played in $\mathcal{A}_\mathcal{B}$ (as every choice available in $\mathcal{A}'_\mathcal{B}$ is also available in $\mathcal{A}_\mathcal{B}$). $\qquad\square$

Lemma 4.7. *No positional strategy of $\mathcal{P}_1$ is winning in $\mathcal{A}_\mathcal{B}$.*

Proof. Any positional strategy of $\mathcal{P}_1$ generates a word w^ω, where $w \in \alpha\text{-FreeCycles}_\mathcal{B}(q) \setminus \{\varepsilon\}$ for some $q \in Q$. In particular, word w^ω is rejected when it is read from state q. As all the states in Q are equivalent for $\sim$ (as we assume that W is prefix-independent), we have $q \sim q_{\mathsf{init}}$, so w^ω is also rejected when read from the initial state q_{init} of the automaton. $\qquad\square$

Using Lemmas 4.6 and 4.7, we deduce the desired result.

Lemma 4.8. *No strategy of $\mathcal{P}_1$ is winning in $\mathcal{A}_\mathcal{B}$.*

We use the statement of Lemma 4.8 to show the existence of an α-free-cycle-maximum.

Lemma 4.9. *There exists a unique α-free-cycle-maximum in $\mathcal{B}$.*

Proof. We first show that there exists an α-free-cycle-maximum. Let us assume by contradiction that there is no α-free-cycle-maximum. We show how to build a winning strategy of $\mathcal{P}_1$ in $\mathcal{A}_\mathcal{B}$, contradicting Lemma 4.8. To do so, we build an infinite word accepted by $\mathcal{B}$ by combining finite words that are α-free cycles from some state.

We claim that for every $q \in Q$ which is not an α-free-maximum, there exists

$$w_q \in \bigcup_{q' \in Q} \alpha\text{-FreeCycles}_\mathcal{B}(q') \setminus \{\varepsilon\} \text{ such that } w_q \notin \alpha\text{-Free}_\mathcal{B}(q).$$

Let $q' \in Q$ be such that $\alpha\text{-Free}_\mathcal{B}(q') \not\subseteq \alpha\text{-Free}_\mathcal{B}(q)$, which exists as q is not an α-free-maximum. Let $w_1 \in \alpha\text{-Free}_\mathcal{B}(q') \setminus \alpha\text{-Free}_\mathcal{B}(q)$. By Lemma 2.10, there exists $w_2 \in C^*$ such that $w_1 w_2 \in \alpha\text{-FreeCycles}_\mathcal{B}(q')$ (this holds as we have assumed w.l.o.g. that $\mathcal{B}$ is saturated). As $w_1 \notin \alpha\text{-Free}_\mathcal{B}(q)$, we also have $w_1 w_2 \notin \alpha\text{-Free}_\mathcal{B}(q)$. Taking $w_q = w_1 w_2$ proves the claim. For $q \in Q$ not an α-free-maximum, we fix $w_q \in C^+$ such that $w_q \in \bigcup_{q' \in Q} \alpha\text{-FreeCycles}_\mathcal{B}(q') \setminus \{\varepsilon\}$ and $w_q \notin \alpha\text{-Free}_\mathcal{B}(q)$.

Thanks to Lemma 4.5, we can assume that there is at most one α-free-maximum in $\mathcal{B}$. If it exists, let $q_{\mathsf{max}} \in Q$ be a unique α-free-maximum. Using our initial assumption, q_{max} is not an α-free-cycle-maximum: there is $q \in Q$ such that $\alpha\text{-FreeCycles}_\mathcal{B}(q) \not\subseteq \alpha\text{-FreeCycles}_\mathcal{B}(q_{\mathsf{max}})$. Let $w_{\mathsf{max}} \in \alpha\text{-FreeCycles}_\mathcal{B}(q) \setminus \alpha\text{-FreeCycles}_\mathcal{B}(q_{\mathsf{max}})$. Notice that as $w_{\mathsf{max}} \in \alpha\text{-Free}_\mathcal{B}(q)$ and q_{max} is an α-free-maximum, $w_{\mathsf{max}} \in \alpha\text{-Free}_\mathcal{B}(q_{\mathsf{max}})$. Therefore, w_{max} cannot be a cycle on q_{max}, i.e., $\delta^*(q_{\mathsf{max}}, w_{\mathsf{max}}) \neq q_{\mathsf{max}}$.

We build iteratively an infinite winning word that can be played by $\mathcal{P}_1$ in $\mathcal{A}_\mathcal{B}$ (our construction works whether there is an α-free-maximum or not). As $\mathcal{P}_1$ plays, we keep track

in parallel of the current automaton state. The game starts in v, with current automaton state $q_0 = q_{\mathsf{init}}$. Let $n \geq 0$. We distinguish two cases.

- If q_n is not an α-free-maximum, then $\mathcal{P}_1$ plays word w_{q_n}. As $w_{q_n} \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_n)$, a Büchi transition is used along the way. The current automaton state becomes $q_{n+1} = \delta^*(q_n, w_{q_n})$.
- If $q_n = q_{\mathsf{max}}$ is the α-free-maximum, then $\mathcal{P}_1$ plays w_{max}. The current automaton state becomes $q_{n+1} = \delta(q_{\mathsf{max}}, w_{\mathsf{max}})$, which is not equal to q_{max}.

For infinitely many $i \geq 0$, the automaton state q_i is not an α-free-maximum (as there is at most one α-free-maximum in $\mathcal{B}$, and it cannot appear twice in a row). Therefore, we have described a winning strategy for $\mathcal{P}_1$, since the corresponding run over $\mathcal{B}$ visits infinitely often a Büchi transition.

We have obtained the existence of an α-free-cycle-maximum by contradiction. Using Lemma 2.10, we can show that an α-free-cycle-maximum is also an α-free-maximum. As we have assumed there is at most one α-free-maximum (through Lemma 4.5), we also obtain the uniqueness of the α-free-cycle-maximum. $\qquad\square$

We now know that there exists a unique α-free-cycle-maximum $q_{\mathsf{max}} \in Q$. We show how to use the outgoing transitions of q_{max} in order to realize W as an objective of the kind $\mathsf{Büchi}(F)$ for some $F \subseteq C$, which is the goal of the current subsection.

Lemma 4.10. *There exists $F \subseteq C$ such that $W = \mathsf{Büchi}(F)$.*

Proof. Let $F = \{c \in C \mid (q_{\mathsf{max}}, c) \in \alpha\}$ — equivalently, if we consider colors as words with one letter, F is the set of colors c such that $c \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{\mathsf{max}})$.

We first show that $\mathsf{Büchi}(F) \subseteq W$. Let $c \in F$. Then, $c \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{\mathsf{max}})$. As q_{max} is an α-free-maximum, for all $q \in Q$, $c \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q)$. Therefore, any word containing infinitely many occurrences of colors in F uses infinitely many Büchi transitions and is accepted by $\mathcal{B}$.

We now show that $W \subseteq \mathsf{Büchi}(F)$. By contrapositive, let $w = c_1 c_2 \ldots \notin \mathsf{Büchi}(F)$ be an infinite word with only finitely many colors in F. We show that $w \notin W$. As W is prefix-independent, we may assume w.l.o.g. that w has no color in F, i.e., that for all $i \geq 1$, $c_i \in C \setminus F$. We claim that when read from q_{max}, word w uses no Büchi transition and is thus rejected. This implies that $w \notin W$ as $q_{\mathsf{max}} \sim q_{\mathsf{init}}$.

Assume by contradiction that there is some Büchi transition when reading w from q_{max}, i.e., there exists $k \geq 0$ such that for $w_{\leq k} = c_1 \ldots c_k$, $w_{\leq k} \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{\mathsf{max}})$, but $w_{\leq k} c_{k+1} \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{\mathsf{max}})$. We will deduce that $(q_{\mathsf{max}}, c_{k+1})$ is a Büchi transition, contradicting that $c_{k+1} \in C \setminus F$.

We depict the situation in Figure 11. Let $q_1 = \delta^*(q_{\mathsf{max}}, w_{\leq k})$ (whether q_1 equals q_{max} or not does not matter). By Lemma 2.10, there exists $w' \in C^*$ such that $w_{\leq k} w' \in \alpha\text{-}\mathsf{FreeCycles}_{\mathcal{B}}(q_{\mathsf{max}})$. By construction, we have $w' w_{\leq k} \in \alpha\text{-}\mathsf{FreeCycles}_{\mathcal{B}}(q_1)$. As q_{max} is an α-free-cycle-maximum, we have $\alpha\text{-}\mathsf{FreeCycles}_{\mathcal{B}}(q_1) \subseteq \alpha\text{-}\mathsf{FreeCycles}_{\mathcal{B}}(q_{\mathsf{max}})$, so we also have $w' w_{\leq k} \in \alpha\text{-}\mathsf{FreeCycles}_{\mathcal{B}}(q_{\mathsf{max}})$. Let $q_2 = \delta^*(q_{\mathsf{max}}, w')$. Notice that $q_{\mathsf{max}} = \delta^*(q_2, w_{\leq k})$ and $w_{\leq k} \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_2)$. As q_{max} is an α-free-maximum and $w_{\leq k} c_{k+1} \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{\mathsf{max}})$, we also have that $w_{\leq k} c_{k+1} \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_2)$. Therefore, transition $(q_{\mathsf{max}}, c_{k+1})$ must be a Büchi transition, which contradicts that $c_{k+1} \in C \setminus F$. $\qquad\square$

4.3.2. *General case.* We now remove the prefix-independence assumption on W. We still assume that W is half-positional over finite one-player arenas, and show that W can be recognized by a Büchi automaton built on top of $\mathcal{S}_\sim$. If $\mathcal{B}$ has exactly one state per

Figure 11: Situation in the proof of Lemma 4.10, with $w_{\leq k} \in \alpha\text{-Free}_{\mathcal{B}}(q_{\text{max}})$ but $w_{\leq k}c_{k+1} \notin \alpha\text{-Free}_{\mathcal{B}}(q_{\text{max}})$.

equivalence class of $\sim$, it means that it is built on top of $\mathcal{S}_{\sim}$, and we are done. If not, let $q_{\sim} \in Q$ be a state such that $|[q_{\sim}]| \geq 2$.

We briefly sketch the proof of this section.

Proof sketch. Our proof will show how to modify $\mathcal{B}$ by "merging" all states in equivalence class $[q_{\sim}]$ into a single state, while still recognizing the same objective W. The main technical argument is to build a variant $W_{[q_{\sim}]}$ of objective W on a new set of colors $C_{[q_{\sim}]}$, that turns out to also be half-positional over finite one-player arenas and DBA-recognizable, but which is *prefix-independent*. We can therefore use Lemma 4.10 from Section 4.3.1 and find $F_{[q_{\sim}]} \subseteq C_{[q_{\sim}]}$ such that $W_{[q_{\sim}]} = \text{Büchi}(F_{[q_{\sim}]})$. Then, we exhibit a state $q_{\text{max}} \in [q_{\sim}]$ whose α-free words are tightly linked to the elements of $F_{[q_{\sim}]}$ (Lemma 4.11 and Corollary 4.12). Finally, akin to the way we removed an α-free-maximum in Lemma 4.5, we show that it is still possible to recognize W while keeping only state q_{max} in $[q_{\sim}]$ (Lemma 4.13).

Once we know how to merge the equivalence class $[q_{\sim}]$ into a single state, we can simply repeat the operation for each equivalence class with multiple states, until we obtain a DBA built on top of $\mathcal{S}_{\sim}$. □

We define a new set of colors $C_{[q_{\sim}]}$ using finite words in C^+ such that

$$C_{[q_{\sim}]} = \{w = c_1 \ldots c_i \in C^+ \mid \delta^*(q_{\sim}, w) \sim q_{\sim} \wedge \forall j \text{ s.t. } 1 \leq j < i, \delta^*(q_{\sim}, c_1 \ldots c_j) \not\sim q_{\sim}\}.$$

This (possibly infinite) set contains all the finite words that, read from $q_{\sim}$, come back to a state in $[q_{\sim}]$ for the first time after being read. By Lemma 2.6, for all $q \in [q_{\sim}]$, for all $w \in C_{[q_{\sim}]}$, we also have that $\delta^*(q, w) \sim \delta^*(q_{\sim}, w) \sim q_{\sim}$. The set $C_{[q_{\sim}]}$ therefore corresponds to the set of words with the seemingly stronger property that, when read from any state in $[q_{\sim}]$, come back to a state in $[q_{\sim}]$ for the first time. We define an objective $W_{[q_{\sim}]}$ of infinite words on this new set of colors such that

$$W_{[q_{\sim}]} = \{w_1 w_2 \ldots \in C_{[q_{\sim}]}^{\omega} \mid w_1 w_2 \ldots \in \mathcal{L}(\mathcal{B}^{q_{\sim}})\}.$$

We show that $W_{[q_{\sim}]}$ has the three conditions allowing us to apply Lemma 4.10 to it.

- Objective $W_{[q_{\sim}]}$ is DBA-recognizable: we consider the DBA $\mathcal{B}_{[q_{\sim}]} = ([q_{\sim}], C_{[q_{\sim}]}, q_{\sim}, \delta', \alpha')$, whose update function δ' is the restriction of δ^* to $[q_{\sim}] \times C_{[q_{\sim}]}$, and $\alpha' = \{(q, w) \in [q_{\sim}] \times C_{[q_{\sim}]} \mid w \notin \alpha\text{-Free}_{\mathcal{B}}(q)\}$.
- Objective $W_{[q_{\sim}]}$ is prefix-independent, as adding or removing a finite number of cycles on $q_{\sim}$ does not affect the accepted status of a word in $\mathcal{L}(\mathcal{B}^{q_{\sim}})$.
- Half-positionality of $W_{[q_{\sim}]}$ over finite one-player arenas is implied by half-positionality of W over finite one-player arenas. Indeed, every (finite one-player) arena $\mathcal{A}_{[q_{\sim}]}$ using colors in $C_{[q_{\sim}]}$ can be transformed into a (finite one-player) arena $\mathcal{A}$ with similar properties using colors in C. Two transformations are applied: *(i)* we replace every $C_{[q_{\sim}]}$-colored edge in

$\mathcal{A}_{[q_\sim]}$ by a corresponding finite chain of C-colored edges, and (ii) for every vertex v of $\mathcal{A}_{[q_\sim]}$, we prefix it with a finite chain of C-colored edges starting from a vertex v' reading a word $w_{q_\sim}$ such that $\delta^*(q_{\mathsf{init}}, w_{q_\sim}) = q_\sim$. We then have that $\mathcal{P}_1$ has a winning strategy from a vertex v in $\mathcal{A}_{[q_\sim]}$ if and only if $\mathcal{P}_1$ has a winning strategy from v' in $\mathcal{A}$, and a positional winning strategy from v' in $\mathcal{A}$ can be transformed into a positional winning strategy from v in $\mathcal{A}_{[q_\sim]}$.

By Lemma 4.10, there exists a set $F_{[q_\sim]} \subseteq C_{[q_\sim]}$ such that $W_{[q_\sim]} = \mathsf{Büchi}(F_{[q_\sim]})$.

We now show links between α-free words from states of $[q_\sim]$ and the words in $F_{[q_\sim]}$. The arguments once again rely on the saturation of $\mathcal{B}$.

Lemma 4.11.

- *Let $w \in C_{[q_\sim]}$. If $w \in F_{[q_\sim]}$, then for all $q \in [q_\sim]$, $w \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q)$.*
- *There exists $q \in [q_\sim]$ such that, for all $w \in C_{[q_\sim]} \setminus F_{[q_\sim]}$, $w \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q)$.*

Proof. For the first statement, we prove the contrapositive. We assume that there exists a state $q \in [q_\sim]$ such that $w \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q)$. By Lemma 2.10, there is $w' \in C^*$ such that $ww' \in \alpha\text{-}\mathsf{FreeCycles}_{\mathcal{B}}(q)$. In particular, $(ww')^\omega \notin \mathcal{L}(\mathcal{B}^q) = \mathcal{L}(\mathcal{B}^{q_\sim})$. We can assume w.l.o.g. that $w' \in C^+$ (if $w' = \varepsilon$, then we can simply replace it with $w' = w$). As w' is read from a state in $[q_\sim]$ to a state in $[q_\sim]$, it can be written uniquely as $w' = w'_1 \ldots w'_n$ for some words $w'_i \in C_{[q_\sim]}$. Therefore, $(ww'_1 \ldots w'_n)^\omega$ is also an infinite word on $C_{[q_\sim]}$, and we have $(ww'_1 \ldots w'_n)^\omega \notin W_{[q_\sim]}$ since $(ww'_1 \ldots w'_n)^\omega \notin \mathcal{L}(\mathcal{B}^{q_\sim})$. As $W_{[q_\sim]} = \mathsf{Büchi}(F_{[q_\sim]})$, clearly $w \notin F_{[q_\sim]}$, which ends the proof of the first statement.

For the second statement, assume by contradiction that for all $q \in [q_\sim]$, there exists $w_q \in C_{[q_\sim]} \setminus F_{[q_\sim]}$ such that $w_q \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q)$. As there are only finitely many states in $[q_\sim]$, it is then possible to build a word $w = w_{q_1} \ldots w_{q_n}$ such that for all $1 \le i < n$, $\delta^*(q_i, w_{q_i}) = q_{i+1}$, $\delta^*(q_n, w_{q_n}) = q_1$, and for all $1 \le i \le n$, $w_{q_i} \in C_{[q_\sim]} \setminus F_{[q_\sim]}$ and $w_{q_i} \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_i)$. Word w^ω is accepted from q_1 as it uses infinitely many Büchi transitions, so it is in $\mathcal{L}(\mathcal{B}^{q_1}) = \mathcal{L}(\mathcal{B}^{q_\sim})$. However, if we consider $(w_{q_1} \ldots w_{q_n})^\omega$ as an infinite word on $C_{[q_\sim]}$, then it is not in $W_{[q_\sim]} = \mathsf{Büchi}(F_{[q_\sim]})$ as every letter of the word is in $C_{[q_\sim]} \setminus F_{[q_\sim]}$. This yields a contradiction. $\square$

We use the previous result in a straightforward way to exhibit a state q_{max} whose non-α-free words in $C_{[q_\sim]}$ are exactly the words in $F_{[q_\sim]}$. The reader may notice that, echoing the proof of the prefix-independent case (Section 4.3.1), the state q_{max} given by Corollary 4.12 is actually an α-free-maximum among states in $[q_\sim]$.

Corollary 4.12. *There exists $q_{\mathsf{max}} \in [q_\sim]$ such that for all $w \in C_{[q_\sim]}$, $w \in F_{[q_\sim]}$ if and only if $w \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{\mathsf{max}})$.*

Proof. By the second item of Lemma 4.11, we take $q_{\mathsf{max}} \in [q_\sim]$ such that, for all $w \in C_{[q_\sim]} \setminus F_{[q_\sim]}$, $w \in \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{\mathsf{max}})$. Let $w \in C_{[q_\sim]}$. The property we already have on q_{max} gives us by contrapositive that for $w \in C_{[q_\sim]}$, $w \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{\mathsf{max}})$ implies that $w \in F_{[q_\sim]}$. Reciprocally, the first item of Lemma 4.11 gives us that if $w \in F_{[q_\sim]}$, then $w \notin \alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{\mathsf{max}})$. $\square$

From now on, we assume that $q_{\mathsf{max}} \in [q_\sim]$ is a state having the property of Corollary 4.12. We show that W can be recognized by a smaller DBA consisting of DBA $\mathcal{B}$ in which all the states in $[q_\sim]$ have been merged into the single state q_{max}, by redirecting all incoming transitions of $[q_\sim]$ to q_{max}. We assume w.l.o.g. that if $q_{\mathsf{init}} \in [q_\sim]$, then $q_{\mathsf{init}} = q_{\mathsf{max}}$ (this does not change the objective recognized by $\mathcal{B}$, and will be convenient in the upcoming construction). We consider DBA $\mathcal{B}' = (Q', C, q'_{\mathsf{init}}, \delta', \alpha')$ with

- $Q' = (Q \setminus [q_\sim]) \cup \{q_{\mathsf{max}}\}$, $q'_{\mathsf{init}} = q_{\mathsf{init}}$,
- for $q \in Q'$ and $c \in C$, if $\delta(q,c) \in [q_\sim]$, then $\delta'(q,c) = q_{\mathsf{max}}$; otherwise, $\delta'(q,c) = \delta(q,c)$,
- for $q \in Q'$ and $c \in C$, $(q,c) \in \alpha'$ if and only if $(q,c) \in \alpha$.

We also assume that states that are not reachable from q'_{init} in $\mathcal{B}'$ are removed from Q'.

Lemma 4.13. *The objective recognized by $\mathcal{B}'$ is also W.*

Proof. Let $w = c_1 c_2 \ldots \in C^\omega$, $\varrho = (q_0, c_1, q_1)(q_1, c_2, q_2) \ldots$ be the infinite run of $\mathcal{B}$ on w, and $\varrho' = (q'_0, c_1, q'_1)(q'_1, c_2, q'_2) \ldots$ be the infinite run of $\mathcal{B}'$ on w. We have that $q_0 = q'_0$, but states in both runs may not coincide after a state in $[q_\sim]$ has been reached. Yet, we show inductively that

$$\forall i \geq 0, q_i \sim_{\mathcal{B}} q'_i. \tag{4.3}$$

It is true for $i = 0$ (we even have equality in this case), and if $q_n \sim_{\mathcal{B}} q'_n$, then by construction of the transitions of $\mathcal{B}'$ and Lemma 2.6, we still have $q_{n+1} \sim_{\mathcal{B}} q'_{n+1}$.

We want to show that w is accepted by $\mathcal{B}$ if and only if it is accepted by $\mathcal{B}'$. This is clear if w never goes through a state in $[q_\sim]$ (as the same transitions are then taken in $\mathcal{B}$ and $\mathcal{B}'$).

We first assume that run ϱ visits $[q_\sim]$ finitely many times, and that the last visit to $[q_\sim]$ happens in q_n for some $n \geq 0$. Notice that run ϱ' also visits $[q_\sim]$ for the last time in q'_n by (4.3). Therefore, word $c_{n+1} c_{n+2} \ldots$ is accepted from q'_n in $\mathcal{B}'$ if and only if it is accepted from q'_n in $\mathcal{B}$: all the subsequent transitions coincide. As $q_n \sim_{\mathcal{B}} q'_n$, we have moreover that $c_{n+1} c_{n+2} \ldots$ is accepted from q'_n in $\mathcal{B}'$ if and only if it is accepted from q_n in $\mathcal{B}$. This implies that w is accepted by $\mathcal{B}$ if and only if it is accepted by $\mathcal{B}'$.

We now assume that run ϱ visits $[q_\sim]$ infinitely many times. We decompose w inductively into a prefix $w_{q_\sim}$ reaching $[q_\sim]$ followed by cycles $w_1, w_2, \ldots$ on $[q_\sim]$. Formally, let $w_{q_\sim}$ be any finite prefix of w such that $\delta^*(q_{\mathsf{init}}, w_{q_\sim}) \in [q_\sim]$. We then proceed by induction: we assume that $w_{q_\sim} w_1 \ldots w_n$ is a prefix of w such that $\delta^*(q_{\mathsf{init}}, w_{q_\sim} w_1 \ldots w_n) \in [q_\sim]$. We define $w_{n+1} \in C^+$ as the shortest non-empty word such that $w_{q_\sim} w_1 \ldots w_n w_{n+1}$ is a prefix of w and $\delta^*(q_{\mathsf{init}}, w_{q_\sim} w_1 \ldots w_n w_{n+1}) \in [q_\sim]$. In particular, $w_{n+1} \in C_{[q_\sim]}$. We have $w = w_{q_\sim} w_1 w_2 \ldots$ by construction. By (4.3), we also have that for all $n \geq 0$, $(\delta')^*(q'_{\mathsf{init}}, w_q w_1 \ldots w_n) \sim_{\mathcal{B}} q_\sim$, so $(\delta')^*(q'_{\mathsf{init}}, w_q w_1 \ldots w_n) = q_{\mathsf{max}}$ as this is the only state left in that class in $\mathcal{B}'$.

If w is accepted by $\mathcal{B}$, then for infinitely many $i \geq 1$, word w_i is in $F_{[q_\sim]}$. By Corollary 4.12, all these infinitely many words are not in $\alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{\mathsf{max}})$ and therefore use a Büchi transition when read from q_{max}, so w is also accepted by $\mathcal{B}'$.

If w is rejected by $\mathcal{B}$, then there exists $n \geq 1$ such that for all $n' \geq n$, $w_{n'} \in C_{[q_\sim]} \setminus F_{[q_\sim]}$. By Corollary 4.12, for all $n' \geq n$, $w_{n'}$ is in $\alpha\text{-}\mathsf{Free}_{\mathcal{B}}(q_{\mathsf{max}})$ and therefore does not use a Büchi transition when read from q_{max}, so w is also rejected by $\mathcal{B}'$. $\qquad\square$

We have all the arguments to show our goal for the section (Proposition 4.4), that is, to show that W can be recognized by a Büchi automaton built on top of $\mathcal{S}_\sim$.

Proof of Proposition 4.4. We have shown in Lemma 4.13 how to merge an equivalence class of $\mathcal{B}$ into a single state, while still recognizing the same objective. Repeating this construction for each equivalence class with two or more states, we end up with a DBA with exactly one state per equivalence class of $\sim$ still recognizing W. By definition of $\mathcal{S}_\sim$, this DBA is necessarily built on top of (some automaton isomorphic to) $\mathcal{S}_\sim$. $\qquad\square$

5. Sufficiency of the conditions

We show that a DBA $\mathcal{B}$ with the three conditions from Section 3.1 (recognizing a progress-consistent objective having a total prefix preorder and being recognizable by a Büchi automaton built on top of $\mathcal{S}_\sim$) recognizes a half-positional objective. As these three conditions have been shown to be necessary for the half-positionality of objectives recognized by a DBA, this will imply a characterization of half-positionality.

Our main technical tool is to construct, thanks to these three conditions, a family of *completely well-monotonic universal graphs*. The existence of such objects implies, thanks to recent results [Ohl23], that $\mathcal{P}_1$ has positional optimal strategies, even in two-player arenas of arbitrary cardinality.

5.1. Completely well-monotonic universal graphs. We fix extra terminology about graphs only used in Section 5, and recall the relevant results from [Ohl23].

Extra preliminaries on graphs. Let $\mathcal{G} = (V, E)$ be a graph and $W \subseteq C^\omega$ be an objective. A vertex v of $\mathcal{G}$ *satisfies* W if for all infinite paths $v_0 \xrightarrow{c_1} v_1 \xrightarrow{c_2} \ldots$ from v, we have $c_1 c_2 \ldots \in W$.

Given two graphs $\mathcal{G} = (V, E)$ and $\mathcal{G}' = (V', E')$, a *(graph) morphism from $\mathcal{G}$ to $\mathcal{G}'$* is a function $\phi \colon V \to V'$ such that $(v_1, c, v_2) \in E$ implies $(\phi(v_1), c, \phi(v_2)) \in E'$.

A morphism ϕ from $\mathcal{G}$ to $\mathcal{G}'$ is *W-preserving* if for all $v \in V$, v satisfies W implies that $\phi(v)$ satisfies W. Notice that if $\phi(v)$ satisfies W, then v satisfies W, as any path $v \xrightarrow{c_1} v_1 \xrightarrow{c_2} \ldots$ of $\mathcal{G}$ implies the existence of a path $\phi(v) \xrightarrow{c_1} \phi(v_1) \xrightarrow{c_2} \ldots$ of $\mathcal{G}'$ — there are "more paths" in $\mathcal{G}'$. For κ a cardinal, a graph $\mathcal{U}$ is *(κ, W)-universal* if for all graphs $\mathcal{G}$ of cardinality $< \kappa$, there is a W-preserving morphism from $\mathcal{G}$ to $\mathcal{U}$.

We consider a graph $\mathcal{G} = (V, E)$ along with a total order $\leq$ on its vertex set V. We say that $\mathcal{G}$ is *monotonic* if for all $v, v', v'' \in V$, for all $c \in C$,

- $(v \xrightarrow{c} v'$ and $v' \geq v'') \implies v \xrightarrow{c} v''$, and
- $(v \geq v'$ and $v' \xrightarrow{c} v'') \implies v \xrightarrow{c} v''$.

This means that (i) whenever there is an edge $v \xrightarrow{c} v'$, there is also an edge with color c from v to all vertices smaller than v' for $\leq$, and (ii) whenever $v \geq v'$, then v has at least the same outgoing edges as v'. Graph $\mathcal{G}$ is *well-monotonic* if it is monotonic and the total order $\leq$ is a well-order (i.e., any set of vertices has a minimum). Graph $\mathcal{G}$ is *completely well-monotonic* if it is well-monotonic and there exists a vertex $\top \in V$ maximum for $\leq$ such that for all $v \in V$, $c \in C$, $\top \xrightarrow{c} v$.

Example 5.1. We provide an example illustrating these notions with $C = \{a, b\}$ and $W = \mathsf{Büchi}(\{a\})$. This is a special case of our upcoming construction, and it is already discussed in more depth in [Ohl23, Section 4]. For θ an ordinal, we define a graph $\mathcal{U}_\theta$ with vertices $U_\theta = \theta \cup \{\top\}$, and such that for all ordinals $\lambda, \lambda' < \theta$,

- $\lambda \xrightarrow{a} \lambda'$, and
- $\lambda \xrightarrow{b} \lambda'$ if and only if $\lambda' < \lambda$.

Moreover, for all $v \in U_\theta$ (including the case $v = \top$), we define edges $\top \xrightarrow{a} v$ and $\top \xrightarrow{b} v$. We order vertices using the natural order on the ordinals, and with $\lambda < \top$ for all $\lambda < \theta$.

Vertex $\top$ does not satisfy W, as there is an infinite path $\top \xrightarrow{b} \top \xrightarrow{b} \ldots$, and $b^\omega \notin W$. All other vertices satisfy W by construction: they cannot reach $\top$, and as reading b decreases

the current ordinal, a path cannot have an infinite suffix using only color b (there is no infinite decreasing sequence of ordinals).

This graph is (κ, W)-universal for $\kappa = |\theta|$. Intuitively, for any graph $\mathcal{G}$ with less than κ vertices, a W-preserving morphism from $\mathcal{G}$ to $\mathcal{U}_\theta$ can be defined by mapping vertices not satisfying W to $\top$, and vertices satisfying W to an ordinal λ that depends on how long it may take to guarantee reading an a from them.

Graph $\mathcal{U}_\theta$ is monotonic, which can be quickly checked with the definition. As $\leq$ is a well-order and there is a vertex $\top$ with the right properties, $\mathcal{U}_\theta$ is even completely well-monotonic. The properties of $\mathcal{U}_\theta$ imply half-positionality of $\mathsf{Büchi}(\{a\})$, thanks to the following theorem. ◁

We state an important result linking half-positionality and completely well-monotonic universal graphs from [Ohl23].

Theorem 5.2 (One direction of [Ohl23, Theorem 3.1]). *Let $W \subseteq C^\omega$ be an objective. If for all cardinals κ, there exists a completely well-monotonic (κ, W)-universal graph, then W is half-positional (over all arenas).*

A more precise result from Ohlmann's thesis [Ohl21, Theorem 1.1] can actually be instantiated on more precise classes of arenas. However, we use it to prove here half-positionality of a family of objectives over *all* arenas, so the above result turns out to be sufficient.

Remark 5.3. Our approach also implies the half-positionality of W over arenas with ε-edges (cf. Remark 2.1). Indeed, we can add a fresh color e to C and define an objective $W_e \subseteq (C \cup \{e\})^\omega$ such that for $w \in (C \cup \{e\})^\omega$,

$$w \in W_e \iff \begin{cases} \text{the word obtained by removing } e \text{ is infinite and belongs to } W, \text{ or} \\ w = w'e^\omega \text{ for some } w' \in (C \cup \{e\})^*. \end{cases}$$

Then, the existence of a family of completely well-monotonic (κ, W)-universal graphs for W implies that W_e is half-positional [Ohl23] (in the terminology of [Ohl23], e is a *strongly neutral letter*). Therefore, we can label the ε-edges of any arena with e and obtain an equivalent game with objective W_e. ◁

5.2. **Universal graphs for Büchi automata.** We show that for a DBA-recognizable objective, the three conditions that were shown to be necessary for half-positionality in Section 4 are actually sufficient.

Proposition 5.4. *Let $W \subseteq C^\omega$ be an objective that has a total prefix preorder, is progress-consistent, and is recognizable by a Büchi automaton built on top of $\mathcal{S}_\sim$. Then, W is half-positional.*

The rest of the section is devoted to the proof of this result, using Theorem 5.2. Let $W \subseteq C^\omega$ be an objective with a total prefix preorder, that is progress-consistent, and that is recognized by a DBA $\mathcal{B} = (Q, C, q_{\mathsf{init}}, \delta, \alpha)$ built on top of $\mathcal{S}_\sim$ for the rest of this section. We assume as in previous sections that $\mathcal{B}$ is saturated (in particular, by Lemma 2.10, any α-free path can be extended to an α-free cycle). As $\mathcal{B}$ is built on top of $\mathcal{S}_\sim$, for $q, q' \in Q$, $q \sim q'$ if and only if $q = q'$.

For θ an ordinal, we build a graph $\mathcal{U}_{\mathcal{B},\theta}$ in the following way.

• We set the vertices as $U_{\mathcal{B},\theta} = \{(q, \lambda) \mid q \in Q, \lambda < \theta\} \cup \{\top\}$.

- For every transition $\delta(q, c) = q'$ of $\mathcal{B}$,
  - if $(q, c) \in \alpha$, then for all ordinals λ, λ', we define an edge $(q, \lambda) \xrightarrow{c} (q', \lambda')$;
  - if $(q, c) \notin \alpha$, then for all ordinals λ, λ' s.t. $\lambda' < \lambda$, we define an edge $(q, \lambda) \xrightarrow{c} (q', \lambda')$.
  - for $q'' \prec q'$, then for all ordinals λ, λ'', we define an edge $(q, \lambda) \xrightarrow{c} (q'', \lambda'')$.
- For all $c \in C$, $v \in U_{\mathcal{B},\theta}$, we define an edge $\top \xrightarrow{c} v$.

We order the vertices lexicographically: $(q, \lambda) \le (q', \lambda')$ if $q \prec q'$ or ($q = q'$ and $\lambda \le \lambda'$), and we define $\top$ as the maximum for $\le$ (($q, \lambda) < \top$ for all $q \in Q$, $\lambda < \theta$).

Graph $\mathcal{U}_{\mathcal{B},\theta}$ is built such that on the one hand, it is sufficiently large and has sufficiently many edges so that there is a morphism from any graph $\mathcal{G}$ (of cardinality smaller than some function of $|\theta|$) to $\mathcal{U}_{\mathcal{B},\theta}$. On the other hand, for the morphism to be W-preserving, at least some vertices of $\mathcal{U}_{\mathcal{B},\theta}$ need to satisfy W, which imposes a restriction on the infinite paths from vertices. Graph $\mathcal{U}_{\mathcal{B},\theta}$ is actually built so that for any automaton state $q \in Q$ and ordinal $\lambda < \theta$, the vertex (q, λ) satisfies $\mathcal{L}(\mathcal{B}^q)$ (see Lemma 5.9). The intuitive idea is that for a non-Büchi transition $(q, c) \notin \alpha$ of the automaton such that $\delta(q, c) = q'$, a c-colored edge from a vertex (q, λ) in the graph either (i) reaches a vertex with first component q', in which case the ordinal must decrease on the second component, or (ii) reaches a vertex with first component $q'' \prec q'$, with no restriction on the second component, but therefore with fewer winning continuations. Using progress-consistency and the fact that there is no infinitely decreasing sequence of ordinals, we can show that this implies that no infinite path in $\mathcal{U}_{\mathcal{B},\theta}$ corresponds to an infinite run in the automaton visiting only non-Büchi transitions.

We state two properties that directly follow from the definition of $\mathcal{U}_{\mathcal{B},\theta}$:

$$\text{if } (q, \lambda) \xrightarrow{c} (q', \lambda'), \text{ then } q' \preceq \delta(q, c); \tag{5.1}$$

$$\text{if } (q, \lambda) \xrightarrow{c} (q', \lambda') \text{ and } \lambda'' \le \lambda', \text{ then } (q, \lambda) \xrightarrow{c} (q', \lambda''). \tag{5.2}$$

Example 5.5. We consider again the DBA $\mathcal{B}$ from Example 3.7, recognizing the words containing infinitely many occurrences of a, or a twice in a row at some point. We represent the graph $\mathcal{U}_{\mathcal{B},\theta}$, with $\theta = \omega$ in Figure 12. ◁


Figure 12: The graph $\mathcal{U}_{\mathcal{B},\omega}$, where $\mathcal{B}$ is the automaton from Example 3.7 ($\mathcal{L}(\mathcal{B}) = \mathsf{B\ddot{u}chi}(\{a\}) \cup C^* a a C^\omega$). The dashed edge with color b indicates that $(q_{\mathsf{init}}, \lambda) \xrightarrow{b} (q_{\mathsf{init}}, \lambda')$ if and only if $\lambda' < \lambda$ (it corresponds to the only non-Büchi transition of $\mathcal{B}$). Elsewhere, an edge labeled with color c between two rectangles q, q' means that for all natural numbers (i.e., ordinals less than ω) λ, λ', $(q, \lambda) \xrightarrow{c} (q', \lambda')$. Thick edges correspond to the original transitions of $\mathcal{B}$. There are edges from $\top$ to all vertices of the graph with colors a and b. Vertices are totally ordered from left to right.

In order to use Theorem 5.2, we show that the graph $\mathcal{U}_{\mathcal{B},\theta}$ is completely well-monotonic (Lemma 5.6) and, for all cardinals κ, is (κ, W)-universal for sufficiently large θ (Proposition 5.10).

Lemma 5.6. *Graph $\mathcal{U}_{\mathcal{B},\theta}$ is completely well-monotonic.*

Proof. The order $\leq$ on the vertices is a well-order, and there exists a vertex $\top \in U_{\mathcal{B},\theta}$ maximum for $\leq$ such that for all $v \in U_{\mathcal{B},\theta}$, $c \in C$, $\top \xrightarrow{c} v$. To show that $\mathcal{U}_{\mathcal{B},\theta}$ is completely well-monotonic, it now suffices to show that $\mathcal{U}_{\mathcal{B},\theta}$ is monotonic.

The first item of the monotonicity definition follows from the construction of the graph. We assume that $(q, \lambda) \xrightarrow{c} (q', \lambda')$ and $(q', \lambda') \geq (q'', \lambda'')$, and we show that $(q, \lambda) \xrightarrow{c} (q'', \lambda'')$. By (5.1), we have $q' \preceq \delta(q, c)$. The inequality $(q', \lambda') \geq (q'', \lambda'')$ means by definition that $q'' \prec q'$ or ($q'' = q'$ and $\lambda'' \leq \lambda'$). If $q'' \prec q'$, we obtain that $q'' \prec \delta(q, c)$, so by construction of the graph (third case for transitions) we also have $(q, \lambda) \xrightarrow{c} (q'', \lambda'')$. If $q'' = q'$ and $\lambda'' \leq \lambda'$, then we also have $(q, \lambda) \xrightarrow{c} (q', \lambda'')$ by (5.2), which is what we want as $q' = q''$.

The second item of the monotonicity definition is slightly more involved and follows from progress-consistency, the fact that the prefix preorder is total, and the saturation of $\mathcal{B}$. We assume that $(q, \lambda) \geq (q', \lambda')$ and $(q', \lambda') \xrightarrow{c} (q'', \lambda'')$, and we show that $(q, \lambda) \xrightarrow{c} (q'', \lambda'')$. The assumption $(q, \lambda) \geq (q', \lambda')$ implies that $q \succeq q'$, so $\delta(q, c) \succeq \delta(q', c)$ by Lemma 2.7. Moreover, $(q', \lambda') \xrightarrow{c} (q'', \lambda'')$ implies that $\delta(q', c) \succeq q''$ by (5.1). Hence, $\delta(q, c) \succeq q''$. If $\delta(q, c) \succ q''$, then $(q, \lambda) \xrightarrow{c} (q'', \lambda'')$ by definition of the graph. The same holds if $\delta(q, c) = q''$ and $(q, c) \in \alpha$.

It is left to discuss the case $\delta(q, c) = q''$ and $(q, c) \notin \alpha$. By the above inequalities, this implies that we also have $\delta(q', c) = q''$.

- If $q = q'$, then $\lambda \geq \lambda'$. Moreover, as $(q', c) = (q, c) \notin \alpha$, the existence of edge $(q', \lambda') \xrightarrow{c} (q'', \lambda'')$ implies that $\lambda' > \lambda''$. So $\lambda > \lambda''$ and we also have $(q, \lambda) \xrightarrow{c} (q'', \lambda'')$.
- We show that $q' \prec q$ is not possible — we assume it holds and draw a contradiction. As $(q, c) \notin \alpha$, we have $c \in \alpha\text{-Free}_{\mathcal{B}}(q)$. By Lemma 2.10, there is $w \in C^*$ such that $cw \in \alpha\text{-FreeCycles}_{\mathcal{B}}(q)$. As $\delta(q, c) = \delta(q', c)$ and $\delta^*(q, cw) = q$, we have $\delta^*(q', cw) = q$. As $q' \prec q$, by progress-consistency, the word $(cw)^\omega$ must be accepted by $\mathcal{B}$ when read from q'. It must therefore also be accepted when read from q (as $q' \prec q$), which contradicts that $cw \in \alpha\text{-FreeCycles}_{\mathcal{B}}(q)$. $\qquad\square$

We now intend to show, for all cardinals κ, (κ, W)-universality of some $\mathcal{U}_{\mathcal{B},\theta}$ with a sufficiently large θ. Lemmas 5.7 and 5.8 give insight into properties of the paths of $\mathcal{U}_{\mathcal{B},\theta}$, to then establish which vertices of $\mathcal{U}_{\mathcal{B},\theta}$ satisfy W (Lemma 5.9). Understanding which vertices satisfy W is useful to later define a W-preserving morphism into $\mathcal{U}_{\mathcal{B},\theta}$. We first show that paths in this graph "underapproximate" corresponding runs in the automaton: a finite path $\gamma = (q_0, \lambda_0) \xrightarrow{c_1} \ldots \xrightarrow{c_n} (q_n, \lambda_n)$ in $\mathcal{U}_{\mathcal{B},\theta}$ visits vertices labeled with automaton states at most as large (for $\preceq$) as the corresponding states visited by the finite run from q_0 on $c_1 \ldots c_n$.

Lemma 5.7. *Let $\gamma = (q_0, \lambda_0) \xrightarrow{c_1} \ldots \xrightarrow{c_n} (q_n, \lambda_n)$ be a finite path of $\mathcal{U}_{\mathcal{B},\theta}$ and $w = c_1 \ldots c_n$. Then, $q_n \preceq \delta^*(q_0, w)$.*

Proof. We proceed by induction on the length n of γ. If $n = 0$, then $w = \varepsilon$, so $q_n = q_0 = \delta^*(q_0, w)$, and the result holds. If $n \geq 1$, we assume by induction hypothesis that $q_{n-1} \preceq \delta^*(q_0, c_1 \ldots c_{n-1})$. By Lemma 2.7, we have $\delta(q_{n-1}, c_n) \preceq \delta(\delta^*(q_0, c_1 \ldots c_{n-1}), c_n) = \delta^*(q_0, w)$. By (5.1), if $(q_{n-1}, \lambda_{n-1}) \xrightarrow{c_n} (q_n, \lambda_n)$, then $q_n \preceq \delta(q_{n-1}, c_n)$. By transitivity, $q_n \preceq \delta^*(q_0, w)$. $\qquad\square$

We now show that in $\mathcal{U}_{\mathcal{B},\theta}$, a finite path that goes back to its initial value w.r.t. the first component without decreasing the ordinal necessarily induces an accepted word when repeated.

Lemma 5.8. *Let* $\gamma = (q_0, \lambda_0) \xrightarrow{c_1} \ldots \xrightarrow{c_n} (q_n, \lambda_n)$ *be a finite path of* $\mathcal{U}_{\mathcal{B},\theta}$ *with* $n \geq 1$, $q_0 = q_n$, *and* $\lambda_0 \leq \lambda_n$. *Let* $w = c_1 \ldots c_n$. *Then,* $w^\omega \in \mathcal{L}(\mathcal{B}^{q_0})$.

Proof. Let $\varrho = \mathcal{B}(q_0, w) = (q'_0, c_1, q'_1) \ldots (q'_{n-1}, c_n, q'_n)$ be the finite run of $\mathcal{B}$ obtained by reading w from q_0. States $q_0, \ldots, q_n$ correspond to the first component of the vertices visited by γ in $\mathcal{U}_{\mathcal{B},\theta}$, whereas $q'_0, \ldots, q'_n$ are the states visited by the finite word w in $\mathcal{B}$. We have that $q_0 = q'_0$, but the subsequent states may or may not correspond. By Lemma 5.7, we still know that for all i, $0 \leq i \leq n$, $q_i \preceq q'_i$. In particular, $q_0 = q_n \preceq \delta^*(q_0, w) = q'_n$. We distinguish three cases, depending on whether $q_0 \prec \delta^*(q_0, w)$ or $q_0 \sim \delta^*(q_0, w)$ (which implies $q_0 = \delta^*(q_0, w)$ as $\mathcal{B}$ is built on top of its prefix-classifier), and depending on whether $q_i = q'_i$ for all $0 \leq i \leq n$ or not.

- If $q_0 \prec \delta^*(q_0, w)$, by progress-consistency, $w^\omega \in \mathcal{L}(\mathcal{B}^{q_0})$.
- If $q_0 = \delta^*(q_0, w)$ and for all i, $0 \leq i \leq n$, $q_i = q'_i$ (i.e., γ only uses edges that directly correspond to transitions of the automaton $\mathcal{B}$), then for the ordinal to be greater than or equal to its starting value, some Büchi transition has to be taken since non-Büchi transitions strictly decrease the ordinal on the second component. Hence, w is not an α-free cycle from q_0, so $w^\omega \in \mathcal{L}(\mathcal{B}^{q_0})$.
- If $q_0 = \delta^*(q_0, w)$ and for some index i, $1 \leq i < n$, we have $q_i \prec q'_i$ (i.e., γ takes at least one edge that does not correspond to a transition of the automaton). We represent the situation in Figure 13. We have that $\delta^*(q'_i, c_{i+1} \ldots c_n) = \delta^*(q_0, w) = q_0$. By Lemma 2.7, as $q_i \prec q'_i$, this implies that $\delta^*(q_i, c_{i+1} \ldots c_n) \preceq q_0$. Also, in the graph, there is a path from q_i to q_0 with colors $c_{i+1} \ldots c_n$. Therefore, by Lemma 5.7, $q_0 \preceq \delta^*(q_i, c_{i+1} \ldots c_n)$. Thus, $q_0 \sim \delta^*(q_i, c_{i+1} \ldots c_n)$, and as $\mathcal{B}$ is built on top of $\mathcal{S}_\sim$, $q_0 = \delta^*(q_i, c_{i+1} \ldots c_n)$. Therefore, $q_i \prec \delta^*(q_i, c_{i+1} \ldots c_n c_1 \ldots c_i) = q'_i$. By progress-consistency, $(c_{i+1} \ldots c_n c_1 \ldots c_i)^\omega \in \mathcal{L}(\mathcal{B}^{q_i})$. As $\delta^*(q_i, c_{i+1} \ldots c_n) = q_0$, this implies that $(c_1 \ldots c_n)^\omega = w^\omega \in \mathcal{L}(\mathcal{B}^{q_0})$. $\square$


Figure 13: Situation in the last case of the proof of Lemma 5.8.

We show that every vertex (q_0, λ_0) such that $q_0 \preceq q_{\mathsf{init}}$ satisfies W.

Lemma 5.9. *Let* $q_0 \in Q$. *For all ordinals* $\lambda_0 < \theta$, *the vertex* (q_0, λ_0) *of* $\mathcal{U}_{\mathcal{B},\theta}$ *satisfies* $\mathcal{L}(\mathcal{B}^{q_0})$. *In particular, if* $q_0 \preceq q_{\mathsf{init}}$, (q_0, λ_0) *satisfies* W.

Proof. Let $\pi = (q_0, \lambda_0) \xrightarrow{c_1} (q_1, \lambda_1) \xrightarrow{c_2} \ldots$ be an infinite path of $\mathcal{U}_{\mathcal{B},\theta}$ from (q_0, λ_0) and $w = c_1 c_2 \ldots$ be the sequence of colors along its edges. We show that $w \in \mathcal{L}(\mathcal{B}^{q_0})$.

Let $q'_i = \delta^*(q_0, c_1 \ldots c_i)$ be the state of $\mathcal{B}$ reached after reading the first i colors of w. We claim that there are two states $q, q' \in Q$ occurring infinitely often in the sequences $(q_i)_{i \geq 0}$, $(q'_i)_{i \geq 0}$, respectively, and an increasing sequence of indices $(i_k)_{k \geq 1}$ satisfying that for all $k \geq 1$,

- $q_{i_k} = q$, $q'_{i_k} = q'$, and
- $\lambda_{i_k} \leq \lambda_{i_{k+1}}$.

Indeed, we can first choose a state q appearing infinitely often in $(q_i)_{i \geq 0}$ and pick a sequence $(i_j)_{j \geq 1}$ such that $q_{i_j} = q$ and $(\lambda_{i_j})_{j \geq 1}$ is not decreasing (this is possible since there is no infinite decreasing sequence of ordinals). Then, we can just pick q' appearing infinitely often in $(q'_{i_j})_{j \geq 1}$ and extract the subsequence corresponding to its occurrences.

Let $q, q' \in Q$, $(i_k)_{k \geq 1}$ satisfying the above properties. Let $w_0 = c_1 \ldots c_{i_1}$ and for $k \geq 1$, let $w_k = c_{i_k+1} \ldots c_{i_{k+1}} \in C^+$ be the colors over the edges in π from (q_{i_k}, λ_{i_k}) to $(q_{i_{k+1}}, \lambda_{i_{k+1}})$.

By Lemma 5.7, it holds that $q \preceq q' = \delta^*(q_0, w_0)$. Moreover, since $\lambda_{i_k} \leq \lambda_{i_{k+1}}$, by Lemma 5.8 we have that $w_k^\omega \in \mathcal{L}(\mathcal{B}^q) \subseteq \mathcal{L}(\mathcal{B}^{q'})$ for all $k \geq 1$. We conclude that for all $k \geq 1$, the word w_k labels a cycle over q' in $\mathcal{B}$ visiting some Büchi transition, and therefore $w = w_0 w_1 w_2 \ldots \in \mathcal{L}(\mathcal{B}^{q_0})$.   □

We now have all the tools to show, for all cardinals κ, (κ, W)-universality of $\mathcal{U}_{\mathcal{B}, \theta}$ for sufficiently large θ.

Proposition 5.10. *Let κ be a cardinal, and θ' be an ordinal such that $\kappa \leq |\theta'|$. Let $\theta = |Q| \cdot \theta'$. Graph $\mathcal{U}_{\mathcal{B}, \theta}$ is (κ, W)-universal.*

Proof. Let $\mathcal{G} = (V, E)$ be a graph such that $|V| < \kappa$ (in particular, $|V| < |\theta'|$). For $v \in V$, let $q_v \in Q \cup \{\top\}$ be the smallest automaton state (for $\preceq$) such that v satisfies $\mathcal{L}(\mathcal{B}^{q_v})$ (i.e., such that all infinite paths from v are in objective $\mathcal{L}(\mathcal{B}^{q_v})$), or $\top$ if it satisfies none of them. We remark that $q_v \preceq q$ if and only if v satisfies $\mathcal{L}(\mathcal{B}^q)$. To show that there is a W-preserving morphism from $\mathcal{G}$ to $\mathcal{U}_{\mathcal{B}, \theta}$, we follow the six steps outlined in [Ohl23, Lemma 4.3].

(i) In this first step, we classify and order vertices of $\mathcal{G}$ in an inductive way, which will later be used to map them to vertices of $\mathcal{U}_{\mathcal{B}, \theta}$. For $q \in Q$ and λ an ordinal, we define by transfinite induction

$$V_\lambda^q = \left\{ v \in V \mid q_v \preceq q, \text{ and } \forall c \in C \left(v \xrightarrow{c} v' \implies [(q, c) \in \alpha \text{ or } \exists \eta < \lambda, v' \in V_\eta^{\delta(q,c)}] \right) \right\}.$$

Intuitively, for v to be in V_λ^q, it has to satisfy $\mathcal{L}(\mathcal{B}^q)$ and to guarantee that a Büchi transition is visited "quickly" when colors of paths from v are read from q in $\mathcal{B}$ (how quickly depends on the value of λ). We remark that for each state $q \in Q$, the sequence $(V_\lambda^q)_\lambda$ is non-decreasing: for $\lambda \leq \lambda'$, $V_\lambda^q \subseteq V_{\lambda'}^q$. We illustrate this induction on a concrete case in Example 5.11. We also discuss another example in Remark 5.12 to illustrate why we need to consider the product $\theta = |Q| \cdot \theta'$ rather than just taking $\theta = \theta'$ for our construction. The subsequent steps mostly follow from this definition.

(ii) Let $V^q = \bigcup_\lambda V_\lambda^q$. We show that if v satisfies $\mathcal{L}(\mathcal{B}^q)$, then it is in V^q. Assume that $v \notin V^q$. If $q \prec q_v$, then we immediately have that v does not satisfy $\mathcal{L}(\mathcal{B}^q)$. If $q_v \preceq q$, then v has an outgoing edge $v \xrightarrow{c} v'$ such that $(q, c) \notin \alpha$ and $v' \notin \bigcup_\lambda V_\lambda^{\delta(q,c)}$. By induction, we build an infinite path from v whose projection in $\mathcal{B}$ only uses non-Büchi transitions, so v does not satisfy $\mathcal{L}(\mathcal{B}^q)$.

(iii) In this step and the next one, we show that there is no use in considering ordinals beyond θ in our construction. We first show that if for all $q \in Q$, $V_\lambda^q = V_{\lambda+1}^q$, then for all $q \in Q$ and all $\lambda' \geq \lambda$, $V_\lambda^q = V_{\lambda'}^q$. For $\lambda \leq \lambda'$, we always have $V_\lambda^q \subseteq V_{\lambda'}^q$. For the other inclusion, we assume by transfinite induction that $V_\lambda^{q'} = V_\eta^{q'}$ for all $q' \in Q$ and for all η such that $\lambda \leq \eta < \lambda'$. Let $v \in V_{\lambda'}^q$. Every edge $v \xrightarrow{c} v'$ either satisfies

$(q, c) \in \alpha$, or there exists $\eta < \lambda'$ such that $v' \in V_\eta^{\delta(q,c)}$. Since $V_\eta^{\delta(q,c)} \subseteq V_\lambda^{\delta(q,c)}$ by induction hypothesis, we have $v' \in V_\lambda^{\delta(q,c)}$. Hence, $v \in V_{\lambda+1}^q = V_\lambda^q$.

(iv) We prove that there exists $\lambda < \theta$ such that for all $q \in Q$, $V_\lambda^q = V_{\lambda+1}^q$. If not, using the axiom of choice, we can build a map $\psi \colon \theta \to Q \times V$ such that for $\lambda < \theta$, $\psi(\lambda) = (q, v)$ for some $q \in Q$ and $v \in V_{\lambda+1}^q \setminus V_\lambda^q$. This map is injective, as any pair (q, v) can be chosen at most once (as $(V_\lambda^q)_\lambda$ is non-decreasing). This implies that $|\theta| = |Q| \cdot |\theta'| \leq |Q| \cdot |V|$, a contradiction since $|V| < |\theta'|$.

Using additionally Item (iii), we deduce that there exists $\lambda < \theta$ such that for all $\lambda' \geq \lambda$, $V_\lambda^q = V_{\lambda'}^q$.

(v) Let $\phi \colon V \to U_{\mathcal{B},\theta}$ be such that

$$\phi(v) = \begin{cases} (q_v, \min\{\lambda \mid v \in V_\lambda^{q_v}\}) & \text{if } q_v \prec \top, \\ \top & \text{if } q_v = \top. \end{cases}$$

By Item (ii), for all $v \in V$, there exists λ such that $v \in V_\lambda^{q_v}$, so $\{\lambda \mid v \in V_\lambda^{q_v}\}$ is non-empty. By Item (iv), we have that if $\phi(v) = (q, \lambda)$, then $\lambda < \theta$, so the image of ϕ is indeed in $U_{\mathcal{B},\theta}$. We show that ϕ is W-preserving: if v satisfies W, then $q_v \preceq q_{\text{init}}$, so by Lemma 5.9, $\phi(v)$ also satisfies W.

(vi) We show that ϕ is a graph morphism. Let $v \xrightarrow{c} v'$ be an edge of $\mathcal{G}$ — we need to show that $\phi(v) \xrightarrow{c} \phi(v')$ is an edge of $U_{\mathcal{B},\theta}$. If $\phi(v) = \top$, this is clear as there are all possible outgoing edges from $\top$. If not, we denote $\phi(v) = (q, \lambda)$ and $\phi(v') = (q', \lambda')$. We have that v satisfies $\mathcal{L}(\mathcal{B}^q)$. Thus, v' must satisfy $\mathcal{L}(\mathcal{B}^{\delta(q,c)})$. This implies that $q' \preceq \delta(q, c)$. We distinguish two cases.

- If $(q, c) \in \alpha$, then by construction of $U_{\mathcal{B},\theta}$, there are c-colored edges from (q, λ) to $(\delta(q, c), \lambda'')$ for all ordinals λ''. By monotonicity, as $q' \preceq \delta(q, c)$, there is also a c-colored edge from (q, λ) to (q', λ').
- We assume that $(q, c) \notin \alpha$. If $q' = \delta(q, c)$, this means that there exists $\eta < \lambda$ such that $v' \in V_\eta^{q'}$. Therefore, $\lambda' \leq \eta < \lambda$. So the edge $(q, \lambda) \xrightarrow{c} (q', \lambda')$ exists by construction of $U_{\mathcal{B},\theta}$. If $q' \prec \delta(q, c)$, then by construction of $U_{\mathcal{B},\theta}$, there are c-colored edges from (q, λ) to (q', λ'') for all ordinals λ''.

We have shown in Item (v) and Item (vi) that ϕ is a W-preserving morphism from $\mathcal{G}$ to $U_{\mathcal{B},\theta}$. $\qquad\square$

Example 5.11. We consider the objective $W = \mathsf{Büchi}(\{a\}) \cup C^* aa C^\omega$ recognized by the DBA $\mathcal{B}$ from Example 3.7, for which graph $U_{\mathcal{B},\omega}$ was shown in Example 5.5. We discuss how our construction maps the vertices of graph $\mathcal{G}$ in Figure 14. Notice that v_1, v_2 and v_3 satisfy W, but not v_4 and v_5.

We build explicitly the sets V_λ^q from Item (i) of the proof of Proposition 5.10. All five vertices are in $V_\lambda^{q_{aa}}$ for every ordinal λ, as all vertices satisfy $\mathcal{L}(\mathcal{B}^{q_{aa}}) = C^\omega$, and any transition from q_{aa} is a Büchi transition. For the same reasons, v_1, v_2, v_3, and v_4 are in $V_\lambda^{q_a}$ for all λ. We determine the sets $V_\lambda^{q_{\text{init}}}$ using the inductive definition. First, v_3 is in $V_\lambda^{q_{\text{init}}}$ for all λ, since $(q_{\text{init}}, a) \in \alpha$. Therefore, $v_2 \in V_\lambda^{q_{\text{init}}}$ for $\lambda \geq 1$ (the b-colored edge from v_2 leads to v_3, and $\delta(q_{\text{init}}, b) = q_{\text{init}}$). Finally, $v_1 \in V_\lambda^{q_{\text{init}}}$ for $\lambda \geq 2$ for the same reason. The morphism ϕ from Item (v) assigns $\phi(v_1) = (q_{\text{init}}, 2)$, $\phi(v_2) = (q_{\text{init}}, 1)$, $\phi(v_3) = (q_{\text{init}}, 0)$, $\phi(v_4) = (q_a, 0)$, $\phi(v_5) = (q_{aa}, 0)$. $\qquad\triangleleft$

Figure 14: Graph $\mathcal{G}$ used in Example 5.11.

Remark 5.12. We illustrate why, in the statement of Proposition 5.10, we use $|Q| \cdot \theta'$ vertices for each automaton state and not just θ'.

Let $C = \{a, b, c\}$, W be the objective recognized by the DBA in Figure 15 (left), and $\mathcal{G}$ be the graph with two vertices in Figure 15 (right). This objective has a total prefix preorder, is progress-consistent, is recognized by a DBA built on top of its prefix-classifier, and the DBA in Figure 15 is saturated. Both vertices of $\mathcal{G}$ satisfy $W = \mathcal{L}(\mathcal{B}^{q_1})$ but do not satisfy $\mathcal{L}(\mathcal{B}^{q_0})$, so $q_{v_1} = q_{v_2} = q_1$. We have $v_1 \in V_0^{q_3}$ (the base case of the induction), and this inductively implies that $v_2 \in V_1^{q_3}$, $v_1 \in V_2^{q_2}$, $v_2 \in V_3^{q_2}$, $v_1 \in V_4^{q_1}$, $v_2 \in V_5^{q_1}$. We represent these steps in the figure, below the graph. Thus, $\phi(v_1) = (q_1, 4)$ and $\phi(v_2) = (q_1, 5)$. We see here that two copies of each automaton state in our universal graph construction would not have sufficed; we actually used six different indices to fully understand the situation, which is due to the interlacing between the graph and the structure of the automaton. $\triangleleft$


Figure 15: DBA (left) and graph $\mathcal{G}$ (right) used in Remark 15 to illustrate the need for $|Q| \cdot \theta'$ vertices for each automaton state in the universal graph used in the construction of Proposition 5.10.

We conclude this section by proving Proposition 5.4, showing that W is half-positional under the three conditions from Theorem 3.11.

Proof of Proposition 5.4. Using Lemma 5.6 and Proposition 5.10, we have for all cardinals κ that there exists a completely well-monotonic (κ, W)-universal graph. By Theorem 5.2, this implies that W is half-positional.   □

6. Conclusion and future work

We have provided a characterization of half-positionality for DBA-recognizable objectives. This novel result is necessary to advance toward a full understanding of half-positionality of ω-regular objectives.

Besides ω-regular objectives, one could study the half-positionality of classes of objectives defined by topological properties; for instance, those corresponding to different levels of

the Borel hierarchy. In particular, a natural step in this direction would be to study the half-positionality of Π_2^0 objectives, which correspond to those that can be recognized by deterministic, but infinite, Büchi automata.

Another interesting extension is to characterize the memory requirements of DBA-recognizable objectives. An intermediate and already seemingly difficult step would be to characterize the memory requirements of objectives recognized by deterministic *weak* automata [Wag79, Sta83, Löd01], generalizing the characterization for safety specifications [CFH14].

Acknowledgments

References

[AF21]   Dana Angluin and Dana Fisman. Regular ω-languages with an informative right congruence. *Information and Computation*, 278:104598, 2021. `doi:10.1016/j.ic.2020.104598`.

[AK22]   Bader Abu Radi and Orna Kupferman. Minimization and canonization of GFG transition-based automata. *Logical Methods in Computer Science*, 18(3), 2022. `doi:10.46298/lmcs-18(3:16)2022`.

[AR17]   Benjamin Aminof and Sasha Rubin. First-cycle games. *Information and Computation*, 254:195–216, 2017. `doi:10.1016/j.ic.2016.10.008`.

[BCJ18]   Roderick Bloem, Krishnendu Chatterjee, and Barbara Jobstmann. Graph games and reactive synthesis. In Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors, *Handbook of Model Checking*, pages 921–962. Springer, 2018. `doi:10.1007/978-3-319-10575-8_27`.

[BCRV22] Patricia Bouyer, Antonio Casares, Mickael Randour, and Pierre Vandenhove. Half-positional objectives recognized by deterministic Büchi automata. In Bartek Klin, Sławomir Lasota, and Anca Muscholl, editors, *Proceedings of the 33rd International Conference on Concurrency Theory, CONCUR 2022, Warsaw, Poland, September 12–16, 2022*, volume 243 of *LIPIcs*, pages 20:1–20:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. `doi:10.4230/LIPIcs.CONCUR.2022.20`.

[BFL$^+$08] Patricia Bouyer, Ulrich Fahrenberg, Kim G. Larsen, Nicolas Markey, and Jirí Srba. Infinite runs in weighted timed automata with energy constraints. In Franck Cassez and Claude Jard, editors, *Proceedings of the 6th International Conference on Formal Modeling and Analysis of Timed Systems, FORMATS 2008, Saint Malo, France, September 15–17, 2008*, volume 5215 of *Lecture Notes in Computer Science*, pages 33–47. Springer, 2008. `doi:10.1007/978-3-540-85778-5_4`.

[BFMM11] Alessandro Bianco, Marco Faella, Fabio Mogavero, and Aniello Murano. Exploring the boundary of half-positionality. *Annals of Mathematics and Artificial Intelligence*, 62(1-2):55–77, 2011. `doi:10.1007/s10472-011-9250-1`.

[BHR16]   Véronique Bruyère, Quentin Hautem, and Mickael Randour. Window parity games: an alternative approach toward parity games with time bounds. In Domenico Cantone and Giorgio Delzanno, editors, *Proceedings of the 7th International Symposium on Games, Automata, Logics, and Formal Verification, GandALF 2016, Catania, Italy, September 14–16, 2016*, volume 226 of *EPTCS*, pages 135–148, 2016. `doi:10.4204/EPTCS.226.10`.

[BHRR19] Véronique Bruyère, Quentin Hautem, Mickael Randour, and Jean-François Raskin. Energy mean-payoff games. In Wan J. Fokkink and Rob van Glabbeek, editors, *Proceedings of the 30th International Conference on Concurrency Theory, CONCUR 2019, Amsterdam, the Netherlands, August 27–30, 2019*, volume 140 of *LIPIcs*, pages 21:1–21:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. `doi:10.4230/LIPIcs.CONCUR.2019.21`.

[BK08]   Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.

[BL69]   J. Richard Büchi and Lawrence H. Landweber. Definability in the monadic second-order theory of successor. *Journal of Symbolic Logic*, 34(2):166–170, 1969. `doi:10.2307/2271090`.

[BLO+22]　Patricia Bouyer, Stéphane Le Roux, Youssouf Oualhadj, Mickael Randour, and Pierre Vandenhove. Games where you can play optimally with arena-independent finite memory. *Logical Methods in Computer Science*, 18(1), 2022. `doi:10.46298/lmcs-18(1:11)2022`.

[BMR+18]　Patricia Bouyer, Nicolas Markey, Mickael Randour, Kim G. Larsen, and Simon Laursen. Average-energy games. *Acta Informatica*, 55(2):91–127, 2018. `doi:10.1007/s00236-016-0274-1`.

[Bok18]　Udi Boker. Why these automata types? In Gilles Barthe, Geoff Sutcliffe, and Margus Veanes, editors, *Proceedings of the 22nd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, LPAR 2022, Awassa, Ethiopia, 16–21 November 2018*, volume 57 of *EPiC Series in Computing*, pages 143–163. EasyChair, 2018. `doi:10.29007/c3bj`.

[BORV21]　Patricia Bouyer, Youssouf Oualhadj, Mickael Randour, and Pierre Vandenhove. Arena-independent finite-memory determinacy in stochastic games. In Serge Haddad and Daniele Varacca, editors, *Proceedings of the 32nd International Conference on Concurrency Theory, CONCUR 2021, Virtual Conference, August 24–27, 2021*, volume 203 of *LIPIcs*, pages 26:1–26:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. `doi:10.4230/LIPIcs.CONCUR.2021.26`.

[BRV23]　Patricia Bouyer, Mickael Randour, and Pierre Vandenhove. Characterizing omega-regularity through finite-memory determinacy of games on infinite graphs. *TheoretiCS*, 2:1–48, 2023. `doi:10.46298/theoretics.23.1`.

[Büc62]　J. Richard Büchi. On a decision method in restricted second order arithmetic. *Proceedings of the International Congress on Logic, Methodology and Philosophy of Science*, pages 1–11, 1962. `doi:10.1007/978-1-4613-8928-6_23`.

[Cas22]　Antonio Casares. On the minimisation of transition-based Rabin automata and the chromatic memory requirements of Muller conditions. In Florin Manea and Alex Simpson, editors, *Proceedings of the 30th EACSL Annual Conference on Computer Science Logic, CSL 2022, Göttingen, Germany, February 14–19, 2022*, volume 216 of *LIPIcs*, pages 12:1–12:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. `doi:10.4230/LIPIcs.CSL.2022.12`.

[CCL22]　Antonio Casares, Thomas Colcombet, and Karoliina Lehtinen. On the size of good-for-games Rabin automata and its link with the memory in Muller games. In Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff, editors, *Proceedings of the 49th International Colloquium on Automata, Languages, and Programming, ICALP 2022, Paris, France, July 4–8, 2022*, volume 229 of *LIPIcs*, pages 117:1–117:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. `doi:10.4230/LIPIcs.ICALP.2022.117`.

[CD12]　Krishnendu Chatterjee and Laurent Doyen. Energy parity games. *Theoretical Computer Science*, 458:49–60, 2012. `doi:10.1016/j.tcs.2012.07.038`.

[CDRR15]　Krishnendu Chatterjee, Laurent Doyen, Mickael Randour, and Jean-François Raskin. Looking at mean-payoff and total-payoff through windows. *Information and Computation*, 242:25–52, 2015. `doi:10.1016/j.ic.2015.03.010`.

[CF13]　Krishnendu Chatterjee and Nathanaël Fijalkow. Infinite-state games with finitary conditions. In Simona Ronchi Della Rocca, editor, *Proceedings of the 27th International Workshop on Computer Science Logic, CSL 2013, 22nd Annual Conference of the EACSL, Torino, Italy, September 2–5, 2013*, volume 23 of *LIPIcs*, pages 181–196. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2013. `doi:10.4230/LIPIcs.CSL.2013.181`.

[CFGO22]　Thomas Colcombet, Nathanaël Fijalkow, Pawel Gawrychowski, and Pierre Ohlmann. The theory of universal graphs for infinite duration games. *Logical Methods in Computer Science*, 18(3), 2022. `doi:10.46298/lmcs-18(3:29)2022`.

[CFH14]　Thomas Colcombet, Nathanaël Fijalkow, and Florian Horn. Playing safe. In Venkatesh Raman and S. P. Suresh, editors, *Proceedings of the 34th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2014, New Delhi, India, December 15–17, 2014*, volume 29 of *LIPIcs*, pages 379–390. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2014. `doi:10.4230/LIPIcs.FSTTCS.2014.379`.

[CHJ05]　Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdziński. Mean-payoff parity games. In *Proceedings of the 20th Annual IEEE Symposium on Logic in Computer Science, LICS 2005, Chicago, IL, USA, June 26–29, 2005*, pages 178–187. IEEE Computer Society, 2005. `doi:10.1109/LICS.2005.26`.

[CN06]　Thomas Colcombet and Damian Niwiński. On the positional determinacy of edge-labeled games. *Theoretical Computer Science*, 352(1-3):190–196, 2006. `doi:10.1016/j.tcs.2005.10.046`.

[DJW97] Stefan Dziembowski, Marcin Jurdziński, and Igor Walukiewicz. How much memory is needed to win infinite games? In *Proceedings of the 12th Annual IEEE Symposium on Logic in Computer Science, LICS 1997, Warsaw, Poland, June 29 – July 2, 1997*, pages 99–110. IEEE Computer Society, 1997. `doi:10.1109/LICS.1997.614939`.

[EJ91]   E. Allen Emerson and Charanjit S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *Proceedings of the 32nd Annual Symposium on Foundations of Computer Science, FOCS 1991, San Juan, Puerto Rico, October, 1991*, pages 368–377. IEEE Computer Society, 1991. `doi:10.1109/SFCS.1991.185392`.

[EM79]   Andrzej Ehrenfeucht and Jan Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, 1979. `doi:10.1007/BF01768705`.

[GH82]   Yuri Gurevich and Leo Harrington. Trees, automata, and games. In Harry R. Lewis, Barbara B. Simons, Walter A. Burkhard, and Lawrence H. Landweber, editors, *Proceedings of the 14th Annual ACM Symposium on Theory of Computing, STOC 1982, San Francisco, CA, USA, May 5–7, 1982*, pages 60–65. ACM, 1982. `doi:10.1145/800070.802177`.

[GK14a] Hugo Gimbert and Edon Kelmendi. Submixing and shift-invariant stochastic games. *CoRR*, abs/1401.6575, 2014. `doi:10.48550/arXiv.1401.6575`.

[GK14b] Hugo Gimbert and Edon Kelmendi. Two-player perfect-information shift-invariant submixing stochastic games are half-positional. *CoRR*, abs/1401.6575v2, 2014. `arXiv:1401.6575v2`.

[GZ04]   Hugo Gimbert and Wiesław Zielonka. When can you play positionally? In Jiří Fiala, Václav Koubek, and Jan Kratochvíl, editors, *Proceedings of the 29th International Symposium on Mathematical Foundations of Computer Science, MFCS 2004, Prague, Czech Republic, August 22–27, 2004*, volume 3153 of *Lecture Notes in Computer Science*, pages 686–697. Springer, 2004. `doi:10.1007/978-3-540-28629-5_53`.

[GZ05]   Hugo Gimbert and Wiesław Zielonka. Games where you can play optimally without any memory. In Martín Abadi and Luca de Alfaro, editors, *Proceedings of the 16th International Conference on Concurrency Theory, CONCUR 2005, San Francisco, CA, USA, August 23–26, 2005*, volume 3653 of *Lecture Notes in Computer Science*, pages 428–442. Springer, 2005. `doi:10.1007/11539452_33`.

[Hor09] Florian Horn. Random fruits on the Zielonka tree. In Susanne Albers and Jean-Yves Marion, editors, *Proceedings of the 26th International Symposium on Theoretical Aspects of Computer Science, STACS 2009, Freiburg, Germany, February 26–28, 2009*, volume 3 of *LIPIcs*, pages 541–552. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 2009. `doi:10.4230/LIPIcs.STACS.2009.1848`.

[KK91]   Nils Klarlund and Dexter Kozen. Rabin measures and their applications to fairness and automata theory. In *Proceedings of the 6th Annual IEEE Symposium on Logic in Computer Science, LICS 1991, Amsterdam, The Netherlands, July 15–18, 1991*, pages 256–265. IEEE Computer Society, 1991. `doi:10.1109/LICS.1991.151650`.

[Kla94] Nils Klarlund. Progress measures, immediate determinacy, and a subset construction for tree automata. *Annals of Pure and Applied Logic*, 69(2-3):243–268, 1994. `doi:10.1016/0168-0072(94)90086-8`.

[Kop06] Eryk Kopczyński. Half-positional determinacy of infinite games. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *Proceedings (Part II) of the 33rd International Colloquium on Automata, Languages and Programming, ICALP 2006, Venice, Italy, July 10–14, 2006*, volume 4052 of *Lecture Notes in Computer Science*, pages 336–347. Springer, 2006. `doi:10.1007/11787006_29`.

[Kop07] Eryk Kopczyński. Omega-regular half-positional winning conditions. In Jacques Duparc and Thomas A. Henzinger, editors, *Proceedings of the 21st International Workshop on Computer Science Logic, CSL 2007, 16th Annual Conference of the EACSL, Lausanne, Switzerland, September 11–15, 2007*, volume 4646 of *Lecture Notes in Computer Science*, pages 41–53. Springer, 2007. `doi:10.1007/978-3-540-74915-8_7`.

[Kop08] Eryk Kopczyński. *Half-positional Determinacy of Infinite Games*. PhD thesis, Warsaw University, 2008.

[Koz22a] Alexander Kozachinskiy. Energy games over totally ordered groups. *CoRR*, abs/2205.04508, 2022. `doi:10.48550/arXiv.2205.04508`.

[Koz22b] Alexander Kozachinskiy. One-to-two-player lifting for mildly growing memory. In Petra Berenbrink and Benjamin Monmege, editors, *Proceedings of the 39th International Symposium on Theoretical*

Aspects of Computer Science, STACS 2022, Marseille, France, March 15–18, 2022, volume 219 of *LIPIcs*, pages 43:1–43:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. `doi:10.4230/LIPIcs.STACS.2022.43`.

[KS15] Denis Kuperberg and Michal Skrzypczak. On determinisation of good-for-games automata. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *Proceedings (Part II) of the 42nd International Colloquium on Automata, Languages, and Programming, ICALP 2015, Kyoto, Japan, July 6–10, 2015*, volume 9135 of *Lecture Notes in Computer Science*, pages 299–310. Springer, 2015. `doi:10.1007/978-3-662-47666-6_24`.

[Kur87] Robert P. Kurshan. Complementing deterministic Büchi automata in polynomial time. *Journal of Computer and System Sciences*, 35(1):59–71, 1987. `doi:10.1016/0022-0000(87)90036-5`.

[Lan69] Lawrence H. Landweber. Decision problems for omega-automata. *Mathematical Systems Theory*, 3(4):376–384, 1969. `doi:10.1007/BF01691063`.

[Le 90] Bertrand Le Saëc. Saturating right congruences. *RAIRO – Theoretical Informatics and Applications*, 24:545–559, 1990. `doi:10.1051/ita/1990240605451`.

[Löd01] Christof Löding. Efficient minimization of deterministic weak omega-automata. *Information Processing Letters*, 79(3):105–109, 2001. `doi:10.1016/S0020-0190(00)00183-6`.

[LPR18] Stéphane Le Roux, Arno Pauly, and Mickael Randour. Extending finite-memory determinacy by Boolean combination of winning conditions. In Sumit Ganguly and Paritosh K. Pandya, editors, *Proceedings of the 38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2018, Ahmedabad, India, December 11–13, 2018*, volume 122 of *LIPIcs*, pages 38:1–38:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. `doi:10.4230/LIPIcs.FSTTCS.2018.38`.

[McN66] Robert McNaughton. Testing and generating infinite sequences by a finite automaton. *Information and Control*, 9(5):521–530, 1966. `doi:10.1016/S0019-9958(66)80013-X`.

[Mos84] Andrzej W. Mostowski. Regular expressions for infinite trees and a standard form of automata. In Andrzej Skowron, editor, *Proceedings of the 5th Symposium on Computation Theory, SCT 1984, Zaborów, Poland, December 3–8, 1984*, volume 208 of *Lecture Notes in Computer Science*, pages 157–168. Springer, 1984. `doi:10.1007/3-540-16066-3_15`.

[MS97] Oded Maler and Ludwig Staiger. On syntactic congruences for omega-languages. *Theoretical Computer Science*, 183(1):93–112, 1997. `doi:10.1016/S0304-3975(96)00312-X`.

[Ner58] Anil Nerode. Linear automaton transformations. *Proceedings of the American Mathematical Society*, 9(4):541–544, 1958. `doi:10.2307/2033204`.

[Ohl21] Pierre Ohlmann. *Monotonic graphs for parity and mean-payoff games*. PhD thesis, IRIF – Research Institute on the Foundations of Computer Science, 2021.

[Ohl23] Pierre Ohlmann. Characterizing positionality in games of infinite duration over infinite graphs. *TheoretiCS*, 2:1–51, 2023. `doi:10.46298/theoretics.23.3`.

[Pnu77] Amir Pnueli. The temporal logic of programs. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science, FOCS 1977, Providence, RI, USA, October 31 – November 1, 1977*, pages 46–57. IEEE Computer Society, 1977. `doi:10.1109/SFCS.1977.32`.

[PP04] Dominique Perrin and Jean-Eric Pin. *Infinite words – automata, semigroups, logic and games*, volume 141 of *Pure and applied mathematics series*. Elsevier Morgan Kaufmann, 2004.

[PR89] Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *Proceedings of the 16th Annual ACM Symposium on Principles of Programming Languages, POPL 1989, Austin, TX, USA, January 11–13, 1989*, pages 179–190. ACM Press, 1989. `doi:10.1145/75277.75293`.

[Put94] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. Wiley, 1994. `doi:10.1002/9780470316887`.

[RS59] Michael O. Rabin and Dana S. Scott. Finite automata and their decision problems. *IBM Journal of Research and Development*, 3(2):114–125, 1959. `doi:10.1147/rd.32.0114`.

[Sha53] Lloyd S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10):1095–1100, 1953. `doi:10.1073/pnas.39.10.1095`.

[Sta83] Ludwig Staiger. Finite-state ω-languages. *Journal of Computer and System Sciences*, 27(3):434–448, 1983. `doi:10.1016/0022-0000(83)90051-X`.

[Tar72] Robert E. Tarjan. Depth-first search and linear graph algorithms. *SIAM Journal on Computing*, 1(2):146–160, 1972. `doi:10.1137/0201010`.

[Tho08]   Wolfgang Thomas. Church's problem and a tour through automata theory. In Arnon Avron, Nachum Dershowitz, and Alexander Rabinovich, editors, *Pillars of Computer Science, Essays Dedicated to Boris (Boaz) Trakhtenbrot on the Occasion of His 85th Birthday*, volume 4800 of *Lecture Notes in Computer Science*, pages 635–655. Springer, 2008. `doi:10.1007/978-3-540-78127-1_35`.

[Wag79]   Klaus W. Wagner. On ω-regular sets. *Information and Control*, 43(2):123–177, 1979. `doi:10.1016/S0019-9958(79)90653-3`.

[Zie98]   Wiesław Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200(1-2):135–183, 1998. `doi:10.1016/S0304-3975(98)00009-7`.

Appendix A. Relation with other properties from the literature

We show precise links between the first two conditions we have defined in Section 3.1 and the *(strong) monotony* and *(strong) selectivity* notions from [GZ05, BFMM11].

Let $W \subseteq C^\omega$ be an objective. For $M \subseteq C^*$ a set of *finite* words, we write $[M]$ for the set of infinite words whose every finite prefix is a prefix of a finite word in M. For $w \in C^*$, we write wW for set of infinite words ww' with $w' \in W$. For two sets of infinite words $M, N \subseteq C^\omega$, we write $M \sqsubseteq_W N$ if $M \cap W \neq \emptyset$ implies $N \cap W \neq \emptyset$ (in other words, either all words in M are losing, or there exists a winning word in N). We write $M \sqsubset_W N$ if $M \cap W = \emptyset$ and $N \cap W \neq \emptyset$ (in other words, all words in M are losing and there exists a winning word in N).

Properties of prefixes. Objective W is *strongly monotone* if for all languages of infinite words $M, N \subseteq C^\omega$, if there exists $w \in C^*$ such that $wM \sqsubset_W wN$, then for all $w' \in C^*$, $w'M \sqsubseteq_W w'N$. Objective W is *monotone* if for all *regular* languages (of finite words) $M, N \subseteq C^*$, if there exists $w \in C^*$ such that $w[M] \sqsubset_W w[N]$, then for all $w' \in C^*$, $w'[M] \sqsubseteq_W w'[N]$.

We show that having a total prefix preorder is in general equivalent to strong monotony as defined in [BFMM11], and is even equivalent to monotony as defined in [GZ05] for ω-regular objectives. In particular, strong monotony and monotony coincide for ω-regular objectives.

Lemma A.1. *Let $W \subseteq C^\omega$ be an objective. Objective W has a total prefix preorder if and only if W is strongly monotone. If W is ω-regular, W has a total prefix preorder if and only if W is monotone.*

Proof. We first prove the first equivalence, with no restriction on W.

We assume that W has a total prefix preorder and we show that it is strongly monotone. Let $w \in C^*$, $M, N \subseteq C^\omega$ such that $wM \sqsubset_W wN$. Let $w' \in C^*$. As the prefix-preorder is total, we have that $w \preceq w'$ or $w' \preceq w$. If $w \preceq w'$, as w has a winning continuation in N, this continuation is also winning for w' — so $w'M \sqsubseteq_W w'N$. If $w' \preceq w$, then all the infinite words in $w'M$ are also losing, thus $w'M \sqsubseteq_W w'N$.

We assume that W is strongly monotone and we show that the prefix preorder is total. Let $x, y \in C^*$. We show that x and y are comparable for $\preceq$: we assume w.l.o.g. that $x \not\preceq y$, and we show that $y \preceq x$, i.e., that $y^{-1}W \subseteq x^{-1}W$. Let $y' \in y^{-1}W$. As $x \not\preceq y$, there is $x' \in C^\omega$ such that $xx' \in W$ but $yx' \notin W$. By taking $M = \{x'\}$ and $N = \{y'\}$, we have $yM \sqsubset_W yN$. Hence, $xM \sqsubseteq_W xN$ by strong monotony. As xM contains a winning word, so does xN, so $y' \in x^{-1}W$.

We now assume that W is ω-regular. The left-to-right implication still holds, as monotony is a weaker notion than strong monotony. We reprove the right-to-left implication with this extra assumption. We assume that W is monotone and we show that the prefix preorder is total. Let $x, y \in C^*$. We show that x and y are comparable for $\preceq$: we assume w.l.o.g. that $x \not\preceq y$, and we show that $y \preceq x$, i.e., that $y^{-1}W \subseteq x^{-1}W$. As $x^{-1}W$ and $y^{-1}W$ are also ω-regular, by Lemma 4.1, it suffices to show that all ultimately periodic words in $y^{-1}W$ are in $x^{-1}W$. Let $u \in C^*, v \in C^+$ such that $uv^\omega \in y^{-1}W$. As $x \not\preceq y$, once again by Lemma 4.1, there is $u' \in C^*, v' \in C^+$ such that $xu'(v')^\omega \in W$ but $yu'(v')^\omega \notin W$. By taking $M = u'(v')^*$ and $N = uv^*$ (which are regular languages), we have $[M] = \{u'(v')^\omega\}$ and $[N] = \{uv^\omega\}$. Therefore, $y[M] \sqsubset_W y[N]$. Hence, $xM \sqsubseteq_W xN$ by monotony. As xM contains a winning word, so does xN, so $uv^\omega \in x^{-1}W$. $\square$

As having a total prefix preorder holds symmetrically for an objective and its complement, we deduce that so does strong monotony, and so does monotony for ω-regular objectives. The latter is especially interesting for the study of ω-regular objectives, as monotony is not symmetric in general (there are objectives that are monotone but whose complement is not).

Properties of cycles. Objective W is *strongly selective* if for all languages of finite words $M, N, K \subseteq C^*$, for all $w \in C^*$, $w[(M \cup N)^*K] \sqsubseteq_W w[M^*] \cup w[N^*] \cup w[K]$. Selectivity is the same definition with M, N, K being restricted to being *regular* languages.

As for having a total prefix preorder and (strong) monotony, we can link progress-consistency and (strong) selectivity. We have that strong selectivity implies progress-consistency, and that selectivity implies progress-consistency when W is ω-regular.

Lemma A.2. *Let $W \subseteq C^\omega$ be an objective. If W is strongly selective, then it is progress-consistent. If W is ω-regular and selective, then it is progress-consistent.*

Proof. We prove the first implication by showing the contrapositive. We assume that W is not progress-consistent, i.e., there exist $w_1 \in C^*$ and $w_2 \in C^+$ such that $w_1 \prec w_1w_2$ and $w_1(w_2)^\omega \notin W$. Let $w_3 \in C^\omega$ be such that $w_1w_3 \notin W$ and $w_1w_2w_3 \in W$, which exists as $w_1 \prec w_1w_2$. Let $w = w_1$, $M = \{w_2\}$, $N = \emptyset$, and $K = \{w_3' \in C^* \mid w_3'$ is a prefix of $w_3\}$. We have that $w[(M \cup N)^*K]$ contains the winning word $w_1w_2w_3$. However, we have $w[M^*] = \{w_1(w_2)^\omega\}$, $w[N^*] = \emptyset$, and $w[K] = \{w_1w_3\}$: the set $w[M^*] \cup w[N^*] \cup w[K]$ contains only losing words. We do not have $w[(M \cup N)^*K] \sqsubseteq_W w[M^*] \cup w[N^*] \cup w[K]$.

For the second implication, we assume that W is ω-regular. The proof proceeds in the exact same way, except that we have to make sure that M, N, and K are regular. This is the case for M and N. To make K regular, we use Lemma 4.1 and obtain that there is an ultimately periodic continuation uv^ω witnessing that $w_1 \prec w_1w_2$. We can set $K = uv^*$ (which is regular) and the proof ends in the same way. $\square$