The True Colors of Memory: A Tour of Chromatic-Memory Strategies in Zero-Sum Games on Graphs

Patricia Bouyer
Université Paris-Saclay, CNRS, ENS Paris-Saclay, Laboratoire Méthodes Formelles, 91190, Gif-sur-Yvette, France

Mickael Randour
F.R.S.-FNRS & UMONS – Université de Mons, Belgium

Pierre Vandenhove
F.R.S.-FNRS & UMONS – Université de Mons, Belgium
Université Paris-Saclay, CNRS, ENS Paris-Saclay, Laboratoire Méthodes Formelles, 91190, Gif-sur-Yvette, France

Abstract

Two-player turn-based zero-sum games on (finite or infinite) graphs are a central framework in theoretical computer science – notably as a tool for controller synthesis, but also due to their connection with logic and automata theory. A crucial challenge in the field is to understand *how complex* strategies need to be to play optimally, given a type of game and a winning objective. In this invited contribution, we give a tour of recent advances aiming to characterize games where finite-memory strategies suffice (i.e., using a limited amount of information about the past). We mostly focus on so-called chromatic memory, which is limited to using colors – the basic building blocks of objectives – seen along a play to update itself. Chromatic memory has the advantage of being usable in different game graphs, and the corresponding class of strategies turns out to be of great interest to both the practical and the theoretical sides.

2012 ACM Subject Classification Theory of computation → Formal languages and automata theory

Keywords and phrases two-player games on graphs, finite-memory strategies, chromatic memory, parity automata, ω-regularity

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2022.3

Category Invited Talk

Funding This work has been partially supported by the ANR Project MAVeriQ (ANR-20-CE25-0012), and the F.R.S.-FNRS Grant n° F.4520.18 (ManySynth). Mickael Randour is an F.R.S.-FNRS Research Associate and a member of the TRAIL Institute. Pierre Vandenhove is an F.R.S.-FNRS Research Fellow.

Acknowledgements Most of the results presented here [4, 6, 7, 3] are related to the F.R.S.-FNRS project FrontieRS, led by the authors. Some of these were obtained in collaboration with Antonio Casares, Stéphane Le Roux, and Youssouf Oualhadj. We express our utmost gratitude to our delightful co-authors.

1 Introduction

Two-player turn-based zero-sum games on graphs. We consider games between two players, $\mathcal{P}_1$ and $\mathcal{P}_2$, that are played on a (finite or infinite) graph, often called *arena*, whose set of vertices is partitioned into vertices controlled by $\mathcal{P}_1$ and vertices controlled by $\mathcal{P}_2$. The players interact by moving a pebble from vertex to vertex, ad infinitum, following edges of the graph. The game starts in a given vertex, and the owner of the current vertex decides where to send the pebble next. The infinite path thereby created is called a *play*.

We assume that the edges of the graph are labeled with *colors* from a finite or infinite set. These colors are used to define the *objective* of the game: an objective is simply a language of infinite color sequences [24]. This general view of objectives encompasses all classical notions from the literature, qualitative and quantitative objectives alike.

The goal of $\mathcal{P}_1$ is to create a play whose projection to colors belongs to the objective whereas $\mathcal{P}_2$ tries to prevent it; hence our games are *zero-sum*. They are also *turn-based* as the players take turns moving the pebble depending on the owner of the current vertex. These moves are chosen according to the *strategy* of the player, which, in general, might use memory (bounded or not) of the past moves to prescribe the next action.

This type of games has been studied for decades, for a plethora of objectives: see many examples in [4]. Interestingly, virtually all such games (i.e., for all reasonable objectives) are known to be *determined* since Martin's seminal result on Borel determinacy [36]. This means that for every vertex v, either $\mathcal{P}_1$ has a strategy that guarantees victory when the game starts in v, or $\mathcal{P}_2$ has one. Two natural questions follow: given a game, can we decide from which vertex each player can win, and what kind of strategy do they need to use?

Reactive synthesis. This survey focuses on the latter question, which is particularly relevant in the context of *controller synthesis* for reactive systems [25, 41, 8, 2]. This formal methods approach aims to automatically synthesize a provably-correct controller for a reactive system that operates within an uncontrollable environment. Through the game-theoretic metaphor, one can model the interaction between the system and its (possibly antagonistic) environment as a two-player zero-sum game on a graph: vertices of the graph model states of the system-environment pair whereas the specification of the system is encoded as a winning objective.

The goal of synthesis is to decide if $\mathcal{P}_1$ (the system) has a *winning strategy*, i.e., one that ensures the objective against all possible strategies of $\mathcal{P}_2$ (the environment), and to build such a strategy if it exists. Winning strategies are essentially formal blueprints for controllers to implement in practical applications: these will thus be correct by design.

A wide variety of objectives (and combinations thereof – e.g., [14, 33]) have been studied in the literature, notably to offer appropriate modeling power for applications in reactive synthesis. All can be expressed through the formalism of colors used in this paper.

Strategy complexity. Keeping in mind that strategies are used as blueprints for real-world controllers, one easily understands why their complexity is of utmost importance: the simpler the strategy, the easier and cheaper it will be to synthesize the corresponding controller and maintain it. On the theoretical level, comprehending which classes of strategies suffice to *play optimally* (i.e., win whenever winning is possible) for various classes of games is also a worthy venture as it may lead to more efficient solving algorithms and the identification of common grounds between these different classes of games.

Many classical objectives are known to be *memoryless-determined*: memoryless strategies, i.e., only using the current vertex as the basis for their decisions, suffice to play optimally. It is for example the case of mean-payoff (in finite graphs) [19] or parity (in finite and infinite graphs) [20, 45]. Yet, over the last decade, the need to model increasingly complex specifications has geared research toward games with more intricate objectives (e.g., [9, 5]) or objectives arising from the combination of simple ones (e.g., [33, 10]). When considering such rich objectives, memoryless strategies usually do not suffice, and one has to use strategies relying on a memory structure, which can be finite or infinite. A natural follow-up question is thus to quantify the memory that is needed for a given objective.

Two flavors of memory. Two models of strategies coexist in the recent literature. In both, a finite-memory strategy can be seen as the association of a memory structure – we will call this a *memory skeleton* – taking the form of a finite automaton and of a memoryless strategy defined, not on the arena, but on the product of the arena and this memory structure – this memoryless strategy is the *next-action* function in the classical Mealy machine model. The only but crucial difference is *how* the memory structure updates its state when an edge is taken in the arena.

In *chromatic memory* (e.g., [4, 7, 11]), the update function only considers the *color* of the edge, whereas in *chaotic memory* (e.g., [18, 31, 13]), the updates consider the actual edge of the arena. That is, in chaotic memory, two different edges bearing the same color may lead to different updates whereas this is not possible in chromatic memory. As such, chromatic memory can be seen as a restricted class of finite-memory strategies. Yet, interestingly, chromatic memory proves sufficient in most cases, and appears more robust with regard to general characterizations. An important feature of chromatic memory is that it allows to define a memory skeleton for an objective independently of an arena, which is impossible for chaotic memory as it explicitly uses the underlying arena. Indeed, chromatic memory coincides with the model used for *arena-independent* strategies in [4]. That being said, given a particular arena, chaotic memory may lead to smaller memory structures – state-wise, not necessarily true when counting transitions – as it can use additional knowledge of the graph; see examples in [11, 13, 32].

General characterizations vs. tight memory bounds. In this invited contribution, we survey recent advances in the study of chromatic-memory strategies, most stemming from a series of co-authored papers on the topic. One can identify two orthogonal yet complementary directions in our work.

Our first line of research aims to establish general characterizations for large classes of games. A typical example is our characterization of objectives for which arena-independent (chromatic) finite-memory strategies suffice [4] (in finite arenas), providing a finite-memory equivalent to Gimbert and Zielonka's seminal result [24]. The philosophy of this research direction is to identify the common grounds between various objectives and to pinpoint the underlying source of complexity, going beyond the use of ad-hoc proofs and techniques. In addition to its fundamental interest, this endeavor permits to establish amenable criteria that one can check to establish that finite-memory optimal strategies exist in a wide range of contexts. Of course, due to its generality, this approach might not lead to perfectly tight memory bounds in particular contexts.

Our second topic of interest can be seen as an answer to this limitation, as it aims to provide tight (lower and upper) memory bounds for specific objective classes. We focus on ω-regular objectives and rely on their representations through different classes of automata in our approach.

Outline. The structure of our paper follows these two lines of research. In Section 2, we introduce the main concepts and notations. Section 3 surveys our results on general characterizations, mainly [4] for finite arenas and [7] for infinite ones. Section 4 focuses on tight bounds for specific ω-regular objectives – we notably present the results of [3]. Finally, Section 5 discusses open questions and future work.

We highlight that this paper is meant as an introductory survey to the topic and, as such, does not give a fully-detailed presentation of all concepts and results. We settled on a high-level, hopefully intuitive, exposition of the field, trying to highlight the interest and limits of our current knowledge, along with connections between the different results. Interested readers may find all details in the corresponding full papers.

Additional related work. This paper only considers *deterministic* games and the corresponding results. In general, these do not carry over to stochastic (one-player or two-player) games, and specific techniques are needed, both to establish results on memory, but also to study the need for randomness – another aspect of the complexity of strategies arising in this context. We mention some references on the topic: [28, 27, 17, 6, 34].

2 Games on Graphs, Objectives, Chromatic Memory

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.

Arenas. 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 is a non-empty set of *vertices* and is the disjoint union of V_1 and V_2, and $E \subseteq V \times C \times V$ is a set of *(colored) edges*. Intuitively, vertices in V_1 are controlled by $\mathcal{P}_1$ and vertices in V_2 are controlled by $\mathcal{P}_2$. We assume arenas to be *non-blocking*: for all $v \in V$, there exists $(v, c, v') \in E$. For $v \in V$, a *play of $\mathcal{A}$ from v* is an infinite sequence of edges $\pi = (v_0, c_1, v_1)(v_1, c_2, v_2) \ldots \in E^\omega$ such that $v_0 = v$. A *history of $\mathcal{A}$ from v* is a finite prefix in E^* of a play of $\mathcal{A}$ from v. For convenience, we assume that there is a distinct *empty history* λ_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 history of $\mathcal{A}$, we define $\mathsf{last}(\gamma) = v_n$. For an empty history λ_v, we define $\mathsf{last}(\lambda_v) = v$. For $i \in \{1, 2\}$, we denote by $\mathsf{Hists}_i(\mathcal{A})$ the set of histories γ of $\mathcal{A}$ such that $\mathsf{last}(\gamma) \in V_i$. An arena is *finite* if V and E are finite. An arena $\mathcal{A} = (V, V_1, V_2, E)$ is a *one-player arena of $\mathcal{P}_1$* (resp. $\mathcal{P}_2$) if $V_2 = \emptyset$ (resp. $V_1 = \emptyset$).

Strategies. 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 a play $\pi = e_1 e_2 \ldots$ is *consistent with σ_i* if for all finite prefixes $\gamma = e_1 \ldots e_n$ of π such that $\mathsf{last}(\gamma) \in V_i$, $\sigma_i(\gamma) = e_{n+1}$. A strategy σ_i is *memoryless* (also called *positional* 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 all $\gamma \in \mathsf{Hists}_i(\mathcal{A})$, $\sigma_i(\gamma) = f(\mathsf{last}(\gamma))$.

Memory skeletons. A *(memory) skeleton* is a tuple $\mathcal{M} = (M, m_{\mathsf{init}}, \alpha_{\mathsf{upd}})$ such that M is a finite set of *states*, $m_{\mathsf{init}} \in M$ is an *initial state*, and $\alpha_{\mathsf{upd}} \colon M \times C \to M$ is an *update function*. Such skeletons are sometimes called *chromatic*, as their transitions are only based on the colors seen. We denote by α^*_{upd} the natural extension of α_{upd} to finite sequences of colors. We define the trivial skeleton $\mathcal{M}_{\mathsf{triv}}$ as the only skeleton with a single state.

Let $\mathcal{M}_1 = (M_1, m^1_{\mathsf{init}}, \alpha^1_{\mathsf{upd}})$ and $\mathcal{M}_2 = (M_2, m^2_{\mathsf{init}}, \alpha^2_{\mathsf{upd}})$ be two skeletons. Their *(direct) product* $\mathcal{M}_1 \otimes \mathcal{M}_2$ is the skeleton $(M, m_{\mathsf{init}}, \alpha_{\mathsf{upd}})$ where $M = M_1 \times M_2$, $m_{\mathsf{init}} = (m^1_{\mathsf{init}}, m^2_{\mathsf{init}})$, and for all $m_1 \in M_1$, $m_2 \in M_2$, $c \in C$, $\alpha_{\mathsf{upd}}((m_1, m_2), c) = (\alpha^1_{\mathsf{upd}}(m_1, c), \alpha^2_{\mathsf{upd}}(m_2, c))$.

For $\mathcal{M} = (M, m_{\mathsf{init}}, \alpha_{\mathsf{upd}})$ a skeleton, a strategy σ_i of $\mathcal{P}_i$ is *based on $\mathcal{M}$* if there exists a function $\alpha_{\mathsf{nxt}} \colon V \times M \to E$ such that for all vertices $v \in V_i$, $\sigma_i(\lambda_v) = \alpha_{\mathsf{nxt}}(v, m_{\mathsf{init}})$, and for all non-empty histories $\gamma = (v_0, c_1, v_1) \ldots (v_{n-1}, c_n, v_n) \in \mathsf{Hists}_i(\mathcal{A})$, $\sigma_i(\gamma) = \alpha_{\mathsf{nxt}}(\mathsf{last}(\gamma), \alpha^*_{\mathsf{upd}}(m_{\mathsf{init}}, c_1 \ldots c_n))$. Such a strategy is said to use *chromatic memory*. Notice that a strategy is memoryless if and only if it is based on $\mathcal{M}_{\mathsf{triv}}$.

Objectives. An *objective* is a set $W \subseteq C^\omega$ of infinite words. 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 $ww' \in W$. For a finite word $w \in C^*$, we write $w^{-1}W = \{w' \in C^\omega \mid ww' \in C^\omega\}$ for the *winning continuations* of w. We have in general that $\varepsilon^{-1}W = W$, and if W is prefix-independent, for all $w \in C^*$, we have $w^{-1}W = W$.

A *game* is a tuple $(\mathcal{A}, W)$, where $\mathcal{A}$ is an arena and W is an objective. In such a game, $\mathcal{P}_1$ wants to achieve an infinite word in W through the infinite interaction with $\mathcal{P}_2$ in $\mathcal{A}$, while $\mathcal{P}_2$ wants to achieve an infinite word in $\overline{W}$.

ω-regular objectives. A central class of objectives is the one of *ω-regular objectives.* They admit multiple equivalent definitions: they are the objectives that can be expressed using an *ω-regular expression*, a *non-deterministic Büchi automaton*, a *deterministic parity automaton*... In this contribution, we mostly use their representation as a deterministic parity automaton (DPA). A (transition-based) DPA is a tuple $\mathcal{D} = (Q, q_{\mathsf{init}}, \delta, p)$ where Q is a finite set of *states*, $q_{\mathsf{init}} \in Q$ is an *initial state*, $\delta\colon Q \times C \to Q$ is a *transition function*, and $p\colon Q \times C \to \mathbb{N}$ is a *priority function.* A DPA $\mathcal{D}$ recognizes an objective containing the words such that, when read in $\mathcal{D}$, the maximal priority that they see infinitely often is even. Notice that the first three components of a DPA are syntactically the same as a memory skeleton; for a memory skeleton $\mathcal{M} = (M, m_{\mathsf{init}}, \alpha_{\mathsf{upd}})$, if there is $p\colon M \times C \to \mathbb{N}$ such that $\mathcal{D} = (\mathcal{M}, p)$, we say that $\mathcal{D}$ is *built on top of $\mathcal{M}$*.

Optimality. 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 plays $(v_0, c_1, v_1)(v_1, c_2, v_2)\ldots$ from v consistent with σ_1, $c_1c_2\ldots \in W$. 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. This notion of optimality requires a *single* strategy to be winning from *all* the winning vertices (a property sometimes called *uniformity*).

Our games are *zero-sum*, hence the objective of $\mathcal{P}_2$ is formally $\overline{W}$. For the sake of readability, we still talk about winning and optimal strategies of $\mathcal{P}_2$ for W, taking into account the symmetric nature of their antagonistic role (i.e., an optimal strategy of $\mathcal{P}_2$ for W is an optimal strategy of $\mathcal{P}_1$ for $\overline{W}$ if the two players are swapped).

Let W be an objective and $i \in \{1, 2\}$. A memory skeleton $\mathcal{M}$ *suffices for $\mathcal{P}_i$ for W* (resp. in finite, one-player arenas) if $\mathcal{P}_i$ has an optimal strategy based on $\mathcal{M}$ in game $(\mathcal{A}, W)$ for all (resp. finite, one-player) arenas $\mathcal{A}$. We say that W is *$\mathcal{M}$-determined* (resp. in finite arenas) if $\mathcal{M}$ suffices for both $\mathcal{P}_1$ and $\mathcal{P}_2$ for W (resp. in finite arenas), and that W is *finite-memory-determined* if it is $\mathcal{M}$-determined for some $\mathcal{M}$. We call $\mathcal{M}_{\mathsf{triv}}$-determinacy *memoryless determinacy.* For consistency with the literature [30], we call the notion "$\mathcal{M}_{\mathsf{triv}}$ suffices to play optimally for $\mathcal{P}_1$ for W" *half-positionality of W*.

▶ Remark 1. We stress that the notion of finite-memory determinacy used throughout the paper is strong in several respects: it requires *chromatic memory* to be sufficient (the memory can only observe colors, and not actual edges that are taken during a play), and it requires the *same* memory skeleton to be sufficient in all arenas (that is, it is *arena-independent*). ⌟

3 Characterization of finite-memory-determined objectives

Many objectives are known to be memoryless-determined, that is, to require no memory except the knowledge of the current vertex to be won. For instance, Ehrenfeucht and Mycielski proved in 1979 that mean-payoff games are memoryless-determined in finite arenas [19], and Emerson and Jutla proved in 1991 that parity games are memoryless-determined [21].

It has therefore been a natural research direction to try to characterize the winning objectives that are memoryless-determined. Gimbert and Zielonka gave the first characterization of winning objectives that are memoryless-determined in finite arenas [24]; this characterization is presented in Section 3.1.1. Trying to extend that result to finite memory was very natural, but the extension could only be handled fifteen years later, and required to focus on chromatic arena-independent memory [4]; this is presented in Section 3.1.2.

Some objectives are memoryless-determined in finite arenas, but require infinite memory in infinite arenas; this is for instance the case of mean-payoff objectives. Hence the above characterizations do not carry over to infinite arenas. Inspired by a work by Colcombet and Niwiński [16] who focused on prefix-independent and memoryless-determined objectives, a full characterization of objectives that are finite-memory-determined in infinite arenas is given in [7], where it is proved that they actually coincide with ω-regular objectives; this is discussed in Section 3.2.

3.1 Finite graph games

3.1.1 Memoryless-determined objectives

The first complete characterization of objectives (or more generally *preference relations* – we do not define this notion here) admitting memoryless optimal strategies is established in [24]. By complete characterization, we mean sufficient and necessary conditions on the objectives. This result can be stated as follows.

▶ **Theorem 2** ([24, Theorem 2]). *An objective W is memoryless-determined if and only if both W and $\overline{W}$ are* monotone *and* selective*.*

Roughly, monotony for an objective says that if an infinite word uw_1^ω is winning and another one uw_2^ω is losing, then their "winning status" cannot be swapped by replacing prefix u, i.e., we cannot have $u'w_1^\omega$ losing and $u'w_2^\omega$ winning for any $u' \in C^*$. This property is obviously satisfied by prefix-independent objectives, but is more general.

Selectivity is defined with regard to cycle mixing: starting from two sequences of colors, it is impossible to create a third one by mixing the first two in such a way that the third one is winning while the first two are not. Similar-looking notions (*fairly mixing*, *concave* [23, 30]) had been defined in other attempts in the literature, but they slightly differ and are actually incomparable to selectivity.

A by-product of the proof of the above theorem is the following so-called *one-to-two-player lift*.

▶ **Corollary 3.** *Assume that for an objective W, memoryless strategies suffice for both $\mathcal{P}_1$ and $\mathcal{P}_2$ in their respective finite one-player arenas. Then W is memoryless-determined in finite arenas.*

Such a lifting corollary provides a neat and easy way to prove that an objective admits memoryless optimal strategies without proving monotony and selectivity at all: proving it in the two one-player subcases, which is generally much easier as it boils down to graph reasoning, and then lifting the result to the general two-player case through the corollary.

▶ **Example 4.** Let $C = \mathbb{Q}$. For $w = c_1c_2 \ldots \in C^\omega$, we define its *mean payoff*

$$\mathsf{MP}(w) = \liminf_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} c_i$$

as the limit (inferior) of the average of the colors of its finite prefixes. We define the objective "achieving a non-negative mean payoff" as $\mathsf{MP}_{\geq 0} = \{w \in C^\omega \mid \mathsf{MP}(w) \geq 0\}$. This objective is memoryless-determined in finite arenas, which was first proved in [19]. We argue informally that it is easy to recover this result using Corollary 3.

We take the point of view of $\mathcal{P}_1$. Let $\mathcal{A}$ be a finite *one-player* arena of $\mathcal{P}_1$. We consider the cycles of $\mathcal{A}$, and we say that a cycle is *non-negative* if the average of its colors is non-negative. Clearly, $\mathcal{P}_1$ can win for $\mathsf{MP}_{\geq 0}$ if a non-negative cycle is reachable: $\mathcal{P}_1$ can simply reach the cycle and loop around it. To win with a memoryless strategy, we simply observe that a non-negative cycle always contains a non-negative *simple* cycle (i.e., not going twice through the same vertex). A *simple* cycle can be reached and looped around with a *memoryless* strategy. This describes a way to win with a memoryless strategy if there is a non-negative cycle, and this is actually "complete": if there is no non-negative cycle, then $\mathcal{P}_1$ simply cannot win for $\mathsf{MP}_{\geq 0}$. We have shown that given a fixed initial vertex of $\mathcal{A}$, if $\mathcal{P}_1$ can win, then $\mathcal{P}_1$ can win with a memoryless strategy. Formally, we must still argue that we can build a single memoryless strategy winning "uniformly" from all the vertices from which $\mathcal{P}_1$ has a winning strategy, which we do not do here but is reasonably straightforward. These arguments show that memoryless strategies suffice for $\mathcal{P}_1$ in its one-player arenas, and the same reasoning works for $\mathcal{P}_2$ (using negative cycles instead). By Corollary 3, $\mathsf{MP}_{\geq 0}$ is memoryless-determined in finite arenas. ⌟

3.1.2 Finite-memory-determined objectives

Gimbert and Zielonka's result completely characterizes objectives which can be won with memoryless strategies in all games played on finite arenas. This paves the way to the quest for a similar characterization for finite-memory strategies. A first reasonable attempt for a generalization to finite-memory strategies could be: given an objective, if in all finite one-player arenas, players have finite-memory optimal strategies, does the same hold in finite two-player arenas? Unfortunately, this generalization does not actually hold: there are objectives for which both players have finite-memory optimal strategies in their respective finite one-player arenas, but for which there exists a finite two-player arena that requires infinite memory for a player to win – see [4, Figure 1].

However, a result similar to Theorem 2 can be proved for $\mathcal{M}$*-determinacy*. Note the subtlety here: $\mathcal{M}$-determinacy requires the memory skeleton $\mathcal{M}$ to be uniform with regard to the arena; the structure of the memory must be *arena-independent*, as opposed to the more general version sketched above where the memory may depend on the arena.

▶ **Theorem 5** ([4, Theorem 3.6]). *Let $\mathcal{M}$ be a memory skeleton. An objective W is $\mathcal{M}$-determined if and only if both W and its complement $\overline{W}$ are* $\mathcal{M}$-monotone *and* $\mathcal{M}$-selective.

The two concepts of $\mathcal{M}$-monotony and $\mathcal{M}$-selectivity are keys to the approach. Intuitively, they correspond to Gimbert and Zielonka's monotony and selectivity, modulo a memory skeleton. The more general concepts of $\mathcal{M}$-monotony and $\mathcal{M}$-selectivity serve the same purpose, but they only compare sequences of colors that are deemed equivalent by the memory skeleton. For the sake of illustration, take selectivity: it implies that one has no interest in mixing different cycles of the game arena. For its generalization, the memory skeleton is taken into account: $\mathcal{M}$-selectivity implies that one has no interest in mixing cycles of the game arena that are read as cycles on the same memory state in the skeleton $\mathcal{M}$. In particular, $\mathcal{M}_{\mathsf{triv}}$-monotony and $\mathcal{M}_{\mathsf{triv}}$-selectivity are respectively equivalent to the original notions of monotony and selectivity, as $\mathcal{M}_{\mathsf{triv}}$ never distinguishes sequences of colors.

The proof of this theorem mimics the one of [24] via the notion of *$\mathcal{M}$-covered arena*. An $\mathcal{M}$-covered arena is an arena which is compatible with $\mathcal{M}$, in the sense that there is a morphism from the arena to $\mathcal{M}$. Examples of $\mathcal{M}$-covered arenas are products of an arena with $\mathcal{M}$, but they are more general (notably, removing edges from these products preserves $\mathcal{M}$-coverability). Strategies based on $\mathcal{M}$ on arenas correspond to memoryless strategies on $\mathcal{M}$-covered arenas. The proof technique for memoryless strategies, relying on an induction on the size of arenas, can be made on $\mathcal{M}$-covered arenas to obtain the theorem for strategies based on $\mathcal{M}$.

As for memoryless strategies, a by-product of the proof of the above theorem is a one-to-two-player lift.

▶ **Corollary 6.** *Let $\mathcal{M}$ be a memory skeleton. Assume that for an objective W, strategies based on $\mathcal{M}$ suffice for both $\mathcal{P}_1$ and $\mathcal{P}_2$ in their respective finite one-player arenas. Then W is $\mathcal{M}$-determined in finite arenas.*

As in general $\mathcal{P}_1$ and $\mathcal{P}_2$ may need different memory structures, we may instantiate the result with two different memory structures $\mathcal{M}_1$ and $\mathcal{M}_2$: if $\mathcal{M}_1$ suffices for $\mathcal{P}_1$ and $\mathcal{M}_2$ suffices for $\mathcal{P}_2$ for W in finite one-player arenas, then $\mathcal{M}_1 \otimes \mathcal{M}_2$ suffices for both $\mathcal{P}_1$ and $\mathcal{P}_2$ in their finite one-player arenas (remembering more information cannot do harm), so W is $(\mathcal{M}_1 \otimes \mathcal{M}_2)$-determined in finite arenas. However, it is unknown whether $\mathcal{M}_1$ (resp. $\mathcal{M}_2$) alone could be sufficient for $\mathcal{P}_1$ (resp. $\mathcal{P}_2$) in finite (two-player) arenas, or if the product is required.

Figure 1 Memory skeletons $\mathcal{M}$ (left) and $\mathcal{M}'$ (right) for two-target reachability games. In figures, diamonds (resp. circles, squares) represent memory or automaton states (resp. arena vertices controlled by $\mathcal{P}_1$, arena vertices controlled by $\mathcal{P}_2$).

▶ **Example 7.** Consider the objective $W = (C^*aC^\omega) \cap (C^*bC^\omega)$ over alphabet $C = \{a, b, c\}$. The objective W requires that both colors a and b should be seen at least once. Consider the two skeletons $\mathcal{M}$ and $\mathcal{M}'$ in Figure 1.

Let us briefly explain why W is not $\mathcal{M}_{\mathsf{triv}}$-monotone (i.e., monotone) but is $\mathcal{M}$-monotone. On the one hand, $ab^\omega \in W$ is preferred to $aa^\omega \notin W$, but $ba^\omega \in W$ is preferred to $bb^\omega \notin W$; hence, W is not $\mathcal{M}_{\mathsf{triv}}$-monotone. On the other hand, skeleton $\mathcal{M}$ distinguishes a and b (in the sense that they reach two different states of skeleton $\mathcal{M}$), hence we do not need to compare their winning continuations $a^{-1}W$ and $b^{-1}W$. Also, we can prove that two pairs of continuations $u_1^{-1}W$ and $u_2^{-1}W$ such that u_1 and u_2 reach the same memory state of $\mathcal{M}$ are comparable (for the inclusion), hence W is $\mathcal{M}$-monotone.

Let us now briefly explain why W is not $\mathcal{M}_{\mathsf{triv}}$-selective (i.e., selective) but is $\mathcal{M}'$-selective. Notice that a and b are *losing cycles*, in the sense that if repeated infinitely often, they generate a losing word ($a^\omega, b^\omega \in \overline{W}$). But combining a and b, which are cycles on the state of $\mathcal{M}_{\mathsf{triv}}$ (as are all finite words), may result in a winning word. For instance, they can be used to make $(ab)^\omega \in W$. Hence, W is not $\mathcal{M}_{\mathsf{triv}}$-selective. On the other hand, W is $\mathcal{M}'$-selective, as $\mathcal{M}'$ distinguishes cycles in such a way that two losing cycles on the same memory state cannot be combined into a winning cycle. For m'_1: all cycles around m'_1 are losing (since they contain neither a nor b) and cannot be combined into a winning cycle. For m'_2: if one reaches m'_2, it means that one of a or b was seen, hence only one color remains to be seen. Any subsequent cycle on m'_2 is either useless (if it sees c or the color among a and b that was already seen) or immediately winning (if it sees the color among a and b that was not seen). There is therefore no advantage to combine multiple cycles once m'_2 is reached.

Overall, W is $\mathcal{M}$-monotone and $\mathcal{M}'$-selective, and we can show in a similar fashion that $\overline{W}$ is $\mathcal{M}$-monotone and $\mathcal{M}_{\mathsf{triv}}$-selective. Moreover, monotony and selectivity are "preserved by product" (intuitively, as having more information to our disposal is never harmful), so both W and $\overline{W}$ are $(\mathcal{M} \otimes \mathcal{M}')$-monotone and $(\mathcal{M} \otimes \mathcal{M}')$-selective. This allows to conclude by Theorem 5 that W is $(\mathcal{M} \otimes \mathcal{M}')$-determined in finite arenas. Skeleton $\mathcal{M} \otimes \mathcal{M}'$ has formally four states but only three reachable states, so both players can play optimally for W using three memory states, which is minimal [22]. ⌟

Applicability of these results goes beyond ω-regular objectives. For instance, the intersection of a reachability condition and a mean-payoff objective requires two memory states, and it suffices for both players to keep track of whether the reachability objective has already been satisfied, e.g., using skeleton $\mathcal{M}$ of Figure 1.

3.2 Infinite graph games

Unfortunately, memoryless determinacy in finite arenas does not carry over straightforwardly to infinite arenas. Indeed, there are objectives that are memoryless-determined in finite arenas but that require infinite memory in some infinite (even one-player) arenas: this is the case for the mean-payoff objective $\mathsf{MP}_{\geq 0}$, which we showed to be memoryless-determined in finite arenas in Example 4.

▶ **Example 8.** We consider objective $\mathsf{MP}_{\geq 0}$ and the infinite one-player arena of $\mathcal{P}_1$ in Figure 2 [40]. Despite the fact that all colors are negative, $\mathcal{P}_1$ has a winning strategy: the idea is to loop on state s_i sufficiently many times to bring the payoff close to $-\frac{1}{i}$, and then move to s_{i+1} and repeat. At the limit, the mean payoff is 0. However, this winning strategy requires memory (even *infinite* memory) to be implemented, and no memoryless strategy is winning. Hence, $\mathsf{MP}_{\geq 0}$ is not memoryless-determined in infinite arenas. ⌟

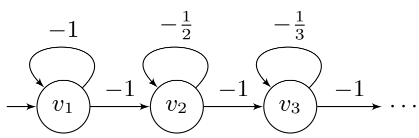

Figure 2 One-player arena requiring infinite memory for objective $\mathsf{MP}_{\geq 0}$.

An interesting result dealing with infinite arenas is the work by Colcombet and Niwiński [16], who proved that a prefix-independent objective W is memoryless-determined in infinite arenas if and only if W is a *parity condition*. A parity condition is an objective recognized by a deterministic parity automaton with a single state, in which each color is directly mapped to a priority in a finite set of integers. In [7], we generalized it to a characterization of objectives that are finite-memory-determined in infinite arenas. Formulating this generalization requires the language-theoretic notion of *right congruence*. Let $W \subseteq C^\omega$

be an objective. The *right congruence* $\sim_W \subseteq C^* \times C^*$ of W is defined 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). When $\sim_W$ has finitely many equivalence classes, we can associate to W a natural skeleton $\mathcal{M}_W$ whose set of states is the set of equivalence classes and with transitions defined in a natural way [43, 35]. We call skeleton $\mathcal{M}_W$ the *prefix classifier of* W, as two finite words reach the same state of $\mathcal{M}_W$ if and only if they are equivalent for $\sim_W$. We can now state the main result from [7].

▶ **Theorem 9** ([7]). *An objective is finite-memory-determined (in arenas of any cardinality) if and only if it is ω-regular. Furthermore, if W is $\mathcal{M}$-determined (in arenas of any cardinality), then W is recognized by a DPA built on top of $\mathcal{M} \otimes \mathcal{M}_W$.*

Before [7], ω-regular objectives were known to be finite-memory-determined [21, 45]: if it is possible to represent an objective with a DPA, then the structure of this automaton suffices to play optimally for both players. Indeed, keeping in memory the extra information from this DPA effectively reduces any game using this objective into a (larger) game using a (simpler) parity condition. This shows one implication of Theorem 9. The proof of the other implication goes through the following steps: if W is $\mathcal{M}$-determined for some skeleton $\mathcal{M}$, then

- W is $\mathcal{M}$*-cycle-consistent*: after any finite word, if we concatenate infinitely many winning (resp. losing) cycles on the skeleton state reached by that word, then it only produces winning (resp. losing) infinite words;
- $\mathcal{M}_W$ is finite, which implies that W is $\mathcal{M}_W$*-prefix-independent*: $\mathcal{M}_W$ classifies prefixes in such a way that two prefixes reaching the same memory state have the same winning continuations.

From this, we obtain in particular that W is $(\mathcal{M} \otimes \mathcal{M}_W)$-cycle-consistent and $(\mathcal{M} \otimes \mathcal{M}_W)$-prefix-independent. We can prove that under these properties, one can associate to transitions of $\mathcal{M} \otimes \mathcal{M}_W$ priorities such that W is recognized by a DPA built on top of $\mathcal{M} \otimes \mathcal{M}_W$. This part of the proof is rather technical, but relies on ordering the cycles according to "how good they are for winning"; order which can be used to assign priorities to transitions.

We recover the result of [16], since the prefix classifier has a single state in the case of a prefix-independent objective (that is, the prefix classifier is $\mathcal{M}_{\mathsf{triv}}$). Hence, under the assumption that W is prefix-independent and memoryless-determined (that is, $\mathcal{M}_{\mathsf{triv}}$-determined), we deduce that W is recognized by a DPA built on top of the skeleton $\mathcal{M}_{\mathsf{triv}} \otimes \mathcal{M}_{\mathsf{triv}}$ with a single state, that is, W is a parity condition.

As previously, we can extract a one-to-two-player lift from the proof of the above result.

▶ **Corollary 10.** *Let W be an objective. If $\mathcal{M}$ suffices for both $\mathcal{P}_1$ and $\mathcal{P}_2$ in their respective one-player arenas (of arbitrary cardinality), then W is $(\mathcal{M} \otimes \mathcal{M}_W)$-determined (in arenas of arbitrary cardinality). In particular, if W is prefix-independent, then under the previous hypotheses, W is $\mathcal{M}$-determined.*

We discuss two ω-regular objectives and illustrate that to represent them using DPAs, we may need some information given by their prefix classifier and some information given by a sufficient memory structure.

▶ **Example 11.** Consider the objective $W_1 = b^*ab^*aC^\omega$ over the alphabet $C = \{a, b\}$. This objective has a prefix classifier $\mathcal{M}_{W_1}$ with three states, corresponding to the three equivalence classes of finite words that have seen 0, 1, or 2 times the color a. This objective is also memoryless-determined that is, $\mathcal{M}_{\mathsf{triv}}$-determined; we do not prove it here (in finite arenas,

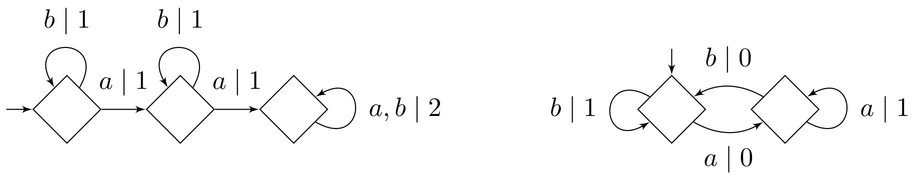

Figure 3 DPA recognizing $W_1 = b^* a b^* a C^\omega$ (left), which is built on top of its prefix classifier $\mathcal{M}_{W_1}$. DPA recognizing $W_2 = C^*(ab)^\omega$ (right), which is built on top of a minimal memory structure sufficient for $\mathcal{P}_1$ and $\mathcal{P}_2$. A transition from a state q to a state q' labeled with "$c \mid n$" means that $\delta(q, c) = q'$ and that $p(q, c) = n$.

it can be shown with Corollary 3). By Theorem 9, this means that W_1 can be recognized by a DPA built on top $\mathcal{M}_{\mathsf{triv}} \otimes \mathcal{M}_{W_1} = \mathcal{M}_{W_1}$. We represent such a DPA in Figure 3 (left): its structure is the same as $\mathcal{M}_{W_1}$, and we just added the right priorities to its transitions.

Consider the objective $W_2 = C^*(ab)^\omega$ over the alphabet $C = \{a, b\}$. It is prefix-independent, hence its prefix classifier is $\mathcal{M}_{\mathsf{triv}}$. It is furthermore $\mathcal{M}$-determined, where $\mathcal{M}$ is the skeleton with two states that remembers whether a or b was last seen; this skeleton has the same structure as the DPA in Figure 3 (right). By Theorem 9, W_2 can be recognized by a DPA built on top of $\mathcal{M} \otimes \mathcal{M}_{\mathsf{triv}} = \mathcal{M}$. We represent such a DPA in Figure 3 (right). ⌟

On the other hand, our results can also illustrate why an objective is not finite-memory-determined (or equivalently, ω-regular).

▶ **Example 12.** We go back to the mean-payoff objective $\mathsf{MP}_{\geq 0}$ defined in Example 4. It is a prefix-independent objective, hence $\mathcal{M}_{\mathsf{MP}_{\geq 0}} = \mathcal{M}_{\mathsf{triv}}$. However, it is not $\mathcal{M}_{\mathsf{triv}}$-cycle-consistent. Indeed, repeating ad infinitum the same color from the set $\{-\frac{1}{n} \mid n \in \mathbb{N}_{>0}\}$ results in a losing word; however, the sequence $(-1)(-\frac{1}{2})(-\frac{1}{3})\ldots$ combining multiple such losing colors is winning (its mean payoff is exactly 0). This argument can be extended to any *finite* memory skeleton $\mathcal{M}$ to show that $\mathsf{MP}_{\geq 0}$ is not $\mathcal{M}$-cycle-consistent. This means that $\mathsf{MP}_{\geq 0}$ is not finite-memory-determined over infinite arenas, and is not ω-regular. ⌟

In some classes of objectives, finite-memory determinacy/ω-regularity depends on the values of some parameters.

▶ **Example 13.** Let $C \subseteq \mathbb{Q}$ be a bounded set of colors and $\lambda \in (0, 1)$ be a *discount factor*. We consider the *discounted-sum objective* [42]

$$\mathsf{DS}_{\geq 0}^{C,\lambda} = \{c_1 c_2 \ldots \in C^\omega \mid \sum_{i=1}^{\infty} c_i \cdot \lambda^{i-1} \geq 0\}.$$

For all values of λ, it is possible to show that $\mathsf{DS}_{\geq 0}^{C,\lambda}$ is $\mathcal{M}_{\mathsf{triv}}$-cycle-consistent. This means by Theorem 9 that the prefix classifier of $\mathsf{DS}_{\geq 0}^{C,\lambda}$, *when finite*, can be used as a DPA to recognize objective $\mathsf{DS}_{\geq 0}^{C,\lambda}$. Proof details and a characterization of the values of C and λ such that $\mathsf{DS}_{\geq 0}^{C,\lambda}$ is ω-regular can be found in [7, Section 4.1]. We give a specific example of values that make $\mathsf{DS}_{\geq 0}^{C,\lambda}$ ω-regular: for $\lambda = \frac{1}{2}$ and $C = \{-2, -1, 0, 1, 2\}$, the prefix-classifier is finite and is depicted in Figure 4. ⌟

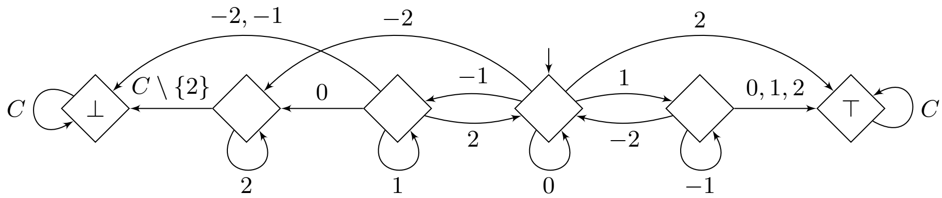

Figure 4 Prefix classifier of $\mathsf{DS}^{C,\lambda}_{\geq 0}$ for $\lambda = \frac{1}{2}$ and $C = \{-2,-1,0,1,2\}$. An infinite word if winning if and only if it does not reach state $\bot$.

4 Memory requirements of ω-regular objectives

Section 3.2 presented an equivalence between ω-regularity and a kind of finite-memory determinacy of two-player zero-sum games. This suggests that, in addition to their relevance in logic and in synthesis, understanding the *memory requirements of ω-regular objectives* is a natural stepping stone in order to study the strategy complexity of all two-player zero-sum games. We discuss known results about memory requirements of ω-regular objectives in this section.

As was stated in Theorem 9, one can relate the memory requirements of an ω-regular objective to its representation as a DPA. We may wonder how close this result brings us to characterizing precisely the memory requirements of ω-regular objectives. Albeit being quite general, there are still multiple questions about memory requirements that Theorem 9 is not able to answer precisely. We introduce these questions by highlighting two of its limitations.

The first limitation comes from the asymmetry of its two implications. In one direction, which is the novel contribution from Theorem 9, we start from a sufficient memory skeleton $\mathcal{M}$ for some objective and show that the objective can then be represented as a DPA built on the automatic structure $\mathcal{M} \otimes \mathcal{M}_W$. In the other direction, we simply use the known memoryless-determinacy of parity conditions, which we already mentioned. What does this tell us on the memory requirements of an ω-regular objective? We know two things: (i) a minimal memory skeleton has always at most as many states as any DPA representing the objective; (ii) a minimal memory skeleton and a minimal DPA differ at most by the factor $\mathcal{M}_W$. However, we do not know in general how to get minimal memory requirements from a representation as a DPA.

The second, perhaps more fundamental limitation is that it makes an assumption on the memory requirements of *both* players simultaneously, as it asks for a memory skeleton sufficient for both players. This assumption therefore conceals a possibly large gap between the individual memory requirements of the two players. This limitation also applies to Theorem 5 (which deals with games played on finite arenas), which cannot be used in general to give tight bounds about memory requirements of each individual player in two-player games.

We illustrate these two limitations on a small example: in this example, the representation of an objective as a DPA is an upper (but not a tight) bound on the memory requirements, and the memory requirements of each player differ.

▶ **Example 14.** Let $C = \{a, b\}$. We consider the objective $W = (b^*a)^\omega \cup (C^*aaC^\omega)$ of words that see a infinitely often or see a twice in a row at some point. This objective is recognized by the DPA with three states depicted in Figure 5 (left), and it is not possible to recognize it using a DPA with fewer states. We therefore know that both players can play optimally using three states of memory, using the memory skeleton underlying this automaton. The prefix classifier of this objective has three states corresponding to three classes of finite words, and its structure also corresponds to the underlying structure of the DPA in Figure 5. According to Theorem 9, this suggests that $\mathcal{P}_1$ and $\mathcal{P}_2$ may need between one and three states of memory to play optimally.

It turns out here that $\mathcal{P}_1$ can play optimally with just one state of memory in all arenas (i.e., W is half-positional), which can be proved using results from [3] (discussed below). Meanwhile, $\mathcal{P}_2$ cannot play optimally with just one state of memory, as is witnessed by the arena in Figure 5 (right). If the play starts in v_1, we observe that $\mathcal{P}_2$ loses by not using the loop in v_2 and going immediately to v_3, as well as by staying infinitely often in v_2. Player $\mathcal{P}_2$ can actually win, but needs to loop at least once (and finitely many times) in v_2 before going to v_3, which cannot be done without memory. For this objective, $\mathcal{P}_2$ can actually play optimally with two states of memory: intuitively, $\mathcal{P}_2$ must keep track of whether the current history is in q_ε or q_a, but there is no point in keeping track of state q_{aa}, as the play is already lost in that state for $\mathcal{P}_2$. ⌟

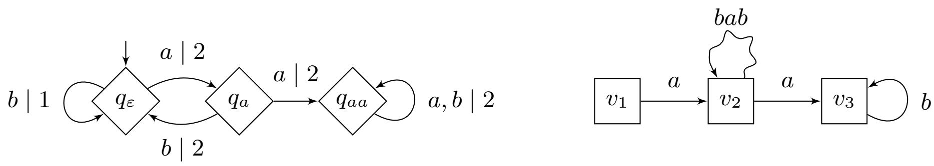

Figure 5 A DPA representing the objective $W = (b^*a)^\omega \cup (C^*aaC^\omega)$ from Example 14 (left), and an arena in which $\mathcal{P}_2$ cannot play optimally with a memoryless strategy (right).

The following questions therefore remain: given an ω-regular objective, what is a *minimal* (i.e., with as few states as possible) memory skeleton sufficient to play optimally for both players? And for a *single* player? A less ambitious (but still open) question would be to understand the memoryless case: how to characterize/decide memoryless determinacy or half-positionality? And would these precise results give us even more information on the representation of ω-regular objectives, perhaps using other acceptance conditions than the parity one? Progress toward these questions, which can be seen as strengthenings of Theorem 9, has been obtained on specific classes of ω-regular objectives. We discuss two of them here.

Memory requirements of Muller conditions. Muller conditions are objectives whose winning words depend solely on the set of colors seen *infinitely often*. They are usually specified by a set $\mathcal{F} \subseteq 2^C$ of sets of colors. The related Muller condition then contains the set of words $w \in C^\omega$ such that the set of colors seen infinitely often by w is a set of $\mathcal{F}$. They are in particular prefix-independent, i.e., their prefix classifier has just one state. A systematic study of the memory requirements of Muller conditions started in the '80s, with first general upper bounds through the later appearance record construction [26, 37], culminating in a complete characterization of their (*chaotic*) memory requirements [18].

More relevant to our chromatic memory considerations, we mention the recent work by Casares [11] that characterizes the chromatic memory requirements of Muller conditions for each individual player. The characterization uses the *Rabin* acceptance condition, which subsumes parity acceptance conditions.

▶ **Theorem 15** ([11, Theorem 27]). *Let W be a Muller condition and $\mathcal{M}$ be a memory skeleton. Structure $\mathcal{M}$ suffices for $\mathcal{P}_1$ for objective W if and only if W is recognized by a deterministic Rabin automaton built on top of $\mathcal{M}$.*

Once again, one direction has been known for some time: Rabin conditions have been known to be half-positional for some time [29] (but unlike parity conditions, their complement may not be). The other, novel direction was obtained thanks to results about the representation of ω-regular objectives [12].

This provides, in the special case of Muller conditions, a characterization of the memory requirements of each player without any blow-up in any direction of the equivalence, and independently of the memory requirements of the other player. It therefore goes beyond the two limitations of Theorem 9 sketched above. As a bonus, this result allows to link the problem of finding a minimal memory skeleton to the problem of minimizing a Rabin automaton, and implies that the related decision problem (given a Muller condition, is there a sufficient memory skeleton with $\leq k$ states for a fixed k?) is NP-complete.

Half-positional deterministic Büchi automata. *Deterministic Büchi automata* (DBAs), unlike their nondeterministic counterparts, only recognize a proper subclass of the ω-regular objectives [44]. They can be seen as a special case of DPAs using only priorities 1 and 2 (the automaton in Figure 5 is also a DBA). The objectives that they recognize are incomparable to Muller conditions. In particular, they recognize some non-prefix-independent objectives, hence with a non-trivial prefix classifier. Currently, their complete memory requirements are not understood – only their half-positionality has been fully characterized. Article [3] gives a characterization of the objectives that are half-positional among those that can be recognized by a DBA. This characterization is a conjunction of three properties that are decidable in polynomial time. The first property is equivalent for this class of objectives to the aforementioned *monotony* property, so we do not discuss it here.

The second property deals with the notion of *progress*: a finite word w_2 is said to be a progress after a finite word w_1 if $w_1 w_2$ has strictly more winning continuations than w_1. We illustrate this property on the objective from Example 14. We take the point of view of $\mathcal{P}_2$, who wants to avoid seeing infinitely many a and avoid seeing a twice in a row at some point. If $w_1 = a$ and $w_2 = bab$, for $\mathcal{P}_2$, w_2 is a progress after w_1: any winning continuation of w_1 is still winning after $w_1 w_2$, and $w_1 w_2 ab^\omega$ is winning, while $w_1 ab^\omega$ is losing. See the link between our choice of words and the arena in Figure 5. The reason that $\mathcal{P}_2$ needs memory here is that although w_2 is a progress after w_1, repeating $w_1 w_2^\omega$ is not winning. Hence, it is useful to play w_2, but an optimal strategy cannot just repeat w_2. We define a *progress-consistent* objective as an objective such that for all progresses w_2 after some w_1, $w_1 w_2^\omega$ is winning. This is necessary for half-positionality.

The third property relates once more to the representation of an ω-regular objective, this time using a DBA. Any deterministic automaton representing an ω-regular objective W needs at least one state per equivalence class of the right congruence $\sim_W$. The celebrated Myhill-Nerode theorem [38] states that to represent *regular* languages (of finite words), we do not need more than one state per equivalence class. This does not hold in general for ω-regular objectives, as was shown with objective $W_2 = C^*(ab)^\omega$ in Example 11. Still, *some* ω-regular objectives admit representations using exactly one state per equivalence class [1]. One example was given in Example 14, which admits a representation as a DBA with three states – one per equivalence class. One of the main technical contribution of [3] is that for an objective recognizable by a DBA to be half-positional, it is necessary that it admits a representation as a DBA with just one state per equivalence class. In such cases, a DBA can be minimized in polynomial time.

▶ **Theorem 16** ([3, Theorem 10]). *Let W be an objective recognized by a DBA. Objective W is half-positional if and only if it is monotone, progress-consistent, and recognized by a DBA built on top of its prefix classifier.*

5 Perspectives

We highlight some of the remaining open paths in the topic of strategy complexity of zero-sum games on graphs. As shown in Section 4, the memory requirements of ω-regular objectives, despite their relevance to logic and synthesis, are not fully mapped out. As discussed in Section 1, there are two relevant memory formalisms for this question: the first one is the chromatic one (which we discussed extensively), and the second one is the *chaotic* memory formalism. A promising research direction about *chaotic* memory requirements was recently given in [13], which proved a bridge between this question and the theory of *good-for-games* automata for Muller conditions. Moreover, it is shown that for some Muller conditions, chaotic memory requirements (expressed as a number of memory *states*) can be exponentially smaller than chromatic ones, at the cost of having to specialize the transitions of the memory structure for each distinct arena.

We also mention a recent tool used to study zero-sum games: *universal graphs*. Universal graphs were originally introduced to design algorithms to decide the winner for some classes of games [15]. Recently, Ohlmann [39] showed an equivalence between the existence of some "good" universal graph for an objective and half-positionality of the objective. Roughly, a universal graph (with the right properties) can be a structural witness of the half-positionality of an objective, and any half-positional objective has a good universal graph. This result has already been used to prove the sufficient condition of Theorem 16 [3] by mechanizing the construction of good universal graphs given a DBA with the right properties. This suggests that using universal graphs may be one path toward understanding half-positionality of ω-regular objectives.

References

1. Dana Angluin and Dana Fisman. Regular ω-languages with an informative right congruence. *Inf. Comput.*, 278:104598, 2021. doi:10.1016/j.ic.2020.104598.
2. 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.
3. Patricia Bouyer, Antonio Casares, Mickael Randour, and Pierre Vandenhove. Half-positional objectives recognized by deterministic Büchi automata. In Bartek Klin, Slawomir Lasota, and Anca Muscholl, editors, *33rd International Conference on Concurrency Theory, CONCUR 2022, September 12-16, 2022, Warsaw, Poland*, volume 243 of *LIPIcs*, pages 20:1–20:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.CONCUR.2022.20.
4. Patricia Bouyer, Stéphane Le Roux, Youssouf Oualhadj, Mickael Randour, and Pierre Vandenhove. Games where you can play optimally with arena-independent finite memory. *Log. Methods Comput. Sci.*, 18(1), 2022. doi:10.46298/lmcs-18(1:11)2022.
5. Patricia Bouyer, Nicolas Markey, Mickael Randour, Kim G. Larsen, and Simon Laursen. Average-energy games. *Acta Inf.*, 55(2):91–127, 2018. doi:10.1007/s00236-016-0274-1.
6. Patricia Bouyer, Youssouf Oualhadj, Mickael Randour, and Pierre Vandenhove. Arena-independent finite-memory determinacy in stochastic games. In Serge Haddad and Daniele Varacca, editors, *32nd International Conference on Concurrency Theory, CONCUR 2021, August 24-27, 2021, Virtual Conference*, volume 203 of *LIPIcs*, pages 26:1–26:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.CONCUR.2021.26.
7. Patricia Bouyer, Mickael Randour, and Pierre Vandenhove. Characterizing omega-regularity through finite-memory determinacy of games on infinite graphs. In Petra Berenbrink and Benjamin Monmege, editors, *39th International Symposium on Theoretical Aspects of Computer Science (STACS 2022)*, volume 219 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 16:1–16:16, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.STACS.2022.16.

8 Romain Brenguier, Lorenzo Clemente, Paul Hunter, Guillermo A. Pérez, Mickael Randour, Jean-François Raskin, Ocan Sankur, and Mathieu Sassolas. Non-zero sum games for reactive synthesis. In Adrian-Horia Dediu, Jan Janousek, Carlos Martín-Vide, and Bianca Truthe, editors, *Language and Automata Theory and Applications – 10th International Conference, LATA 2016, Prague, Czech Republic, March 14-18, 2016, Proceedings*, volume 9618 of *Lecture Notes in Computer Science*, pages 3–23. Springer, 2016. doi:10.1007/978-3-319-30000-9_1.

9 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 Seventh International Symposium on Games, Automata, Logics and Formal Verification, GandALF 2016, Catania, Italy, 14-16 September 2016*, volume 226 of *EPTCS*, pages 135–148, 2016. doi:10.4204/EPTCS.226.10.

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

11 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, *30th EACSL Annual Conference on Computer Science Logic, CSL 2022, February 14-19, 2022, Göttingen, Germany (Virtual Conference)*, volume 216 of *LIPIcs*, pages 12:1–12:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.CSL.2022.12.

12 Antonio Casares, Thomas Colcombet, and Nathanaël Fijalkow. Optimal transformations of games and automata using Muller conditions. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, pages 123:1–123:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.ICALP.2021.123.

13 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 *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, volume 229, pages 117:1–117:20, 2022. doi:10.4230/LIPIcs.ICALP.2022.117.

14 Krishnendu Chatterjee, Mickael Randour, and Jean-François Raskin. Strategy synthesis for multi-dimensional quantitative objectives. *Acta Inf.*, 51(3-4):129–163, 2014. doi:10.1007/s00236-013-0182-6.

15 Thomas Colcombet, Nathanaël Fijalkow, Pawel Gawrychowski, and Pierre Ohlmann. The theory of universal graphs for infinite duration games. *CoRR*, abs/2104.05262, 2021. arXiv:2104.05262.

16 Thomas Colcombet and Damian Niwiński. On the positional determinacy of edge-labeled games. *Theor. Comput. Sci.*, 352(1-3):190–196, 2006. doi:10.1016/j.tcs.2005.10.046.

17 Florent Delgrange, Joost-Pieter Katoen, Tim Quatmann, and Mickael Randour. Simple strategies in multi-objective MDPs. In Armin Biere and David Parker, editors, *Tools and Algorithms for the Construction and Analysis of Systems – 26th International Conference, TACAS 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings, Part I*, volume 12078 of *Lecture Notes in Computer Science*, pages 346–364. Springer, 2020. doi:10.1007/978-3-030-45190-5_19.

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

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

20 E. Allen Emerson and Charanjit S. Jutla. The complexity of tree automata and logics of programs. In *FOCS*, pages 328–337. IEEE Computer Society, 1988. doi:10.1109/SFCS.1988.21949.

21 E. Allen Emerson and Charanjit S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991*, pages 368–377. IEEE Computer Society, 1991. doi:10.1109/SFCS.1991.185392.

22 Nathanaël Fijalkow and Florian Horn. The surprizing complexity of reachability games. *CoRR*, abs/1010.2420, 2010. doi:10.48550/arXiv.1010.2420.

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

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

25 Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002. doi:10.1007/3-540-36387-4.

26 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, May 5-7, 1982, San Francisco, California, USA*, pages 60–65. ACM, 1982. doi:10.1145/800070.802177.

27 Ernst Moritz Hahn, Mateo Perez, Sven Schewe, Fabio Somenzi, Ashutosh Trivedi, and Dominik Wojtczak. Good-for-MDPs automata for probabilistic analysis and reinforcement learning. In Armin Biere and David Parker, editors, *Tools and Algorithms for the Construction and Analysis of Systems – 26th International Conference, TACAS 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings, Part I*, volume 12078 of *Lecture Notes in Computer Science*, pages 306–323. Springer, 2020. doi:10.1007/978-3-030-45190-5_17.

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

29 Nils Klarlund. Progress measures, immediate determinacy, and a subset construction for tree automata. *Ann. Pure Appl. Log.*, 69(2-3):243–268, 1994. doi:10.1016/0168-0072(94)90086-8.

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

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

32 Alexander Kozachinskiy. Infinite separation between general and chromatic memory. *CoRR*, abs/2208.02691, 2022. doi:10.48550/arXiv.2208.02691.

33 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, *38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2018, December 11-13, 2018, Ahmedabad, India*, volume 122 of *LIPIcs*, pages 38:1–38:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.FSTTCS.2018.38.

34 James C. A. Main and Mickael Randour. Different strokes in randomised strategies: Revisiting Kuhn's theorem under finite-memory assumptions. In Bartek Klin, Slawomir Lasota, and Anca Muscholl, editors, *33rd International Conference on Concurrency Theory, CONCUR 2022, September 12-16, 2022, Warsaw, Poland*, volume 243 of *LIPIcs*, pages 22:1–22:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.CONCUR.2022.22.

35 Oded Maler and Ludwig Staiger. On syntactic congruences for omega-languages. *Theor. Comput. Sci.*, 183(1):93–112, 1997. doi:10.1016/S0304-3975(96)00312-X.

36 Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.

37 Robert McNaughton. Infinite games played on finite graphs. *Ann. Pure Appl. Logic*, 65(2):149–184, 1993. doi:10.1016/0168-0072(93)90036-D.

38 A. Nerode. Linear automaton transformations. *Proceedings of the American Mathematical Society*, 9(4):541–544, 1958. doi:10.2307/2033204.

39 Pierre Ohlmann. Characterizing positionality in games of infinite duration over infinite graphs. In Christel Baier and Dana Fisman, editors, *LICS '22: 37th Annual ACM/IEEE Symposium on Logic in Computer Science, Haifa, Israel, August 2 – 5, 2022*, pages 22:1–22:12. ACM, 2022. doi:10.1145/3531130.3532418.

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

41 Mickael Randour. Automated synthesis of reliable and efficient systems through game theory: A case study. In *Proc. of ECCS 2012*, Springer Proceedings in Complexity XVII, pages 731–738. Springer, 2013. doi:10.1007/978-3-319-00395-5_90.

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

43 Ludwig Staiger. Finite-state omega-languages. *J. Comput. Syst. Sci.*, 27(3):434–448, 1983. doi:10.1016/0022-0000(83)90051-X.

44 Klaus Wagner. On ω-regular sets. *Information and control*, 43(2):123–177, 1979.

45 Wieslaw Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theor. Comput. Sci.*, 200(1-2):135–183, 1998. doi:10.1016/S0304-3975(98)00009-7.