

On the Size of Good-For-Games Rabin Automata and Its Link with the Memory in Muller Games

Antonio Casares  

LaBRI, Université de Bordeaux, France

Thomas Colcombet  

CNRS, IRIF, Université Paris Cité, France

Karoliina Lehtinen  

CNRS, Aix-Marseille Université, Université de Toulon, LIS, France

Abstract

In this paper, we look at good-for-games Rabin automata that recognise a Muller language (a language that is entirely characterised by the set of letters that appear infinitely often in each word). We establish that minimal such automata are exactly of the same size as the minimal memory required for winning Muller games that have this language as their winning condition. We show how to effectively construct such minimal automata. Finally, we establish that these automata can be exponentially more succinct than equivalent deterministic ones, thus proving as a consequence that chromatic memory for winning a Muller game can be exponentially larger than unconstrained memory.

2012 ACM Subject Classification Theory of computation → Automata over infinite objects

Keywords and phrases Infinite duration games, Muller games, Rabin conditions, omega-regular languages, memory in games, good-for-games automata

Digital Object Identifier 10.4230/LIPIcs.ICALP.2022.117

Category Track B: Automata, Logic, Semantics, and Theory of Programming

Related Version *Full Version*: <https://arxiv.org/abs/2204.11333>

Funding *Thomas Colcombet*: ANR Delta and Duall

Acknowledgements We would like to thank Marthe Bonamy and Pierre Charbit for their help with graph theory.

1 Introduction

Games. Games, as considered in this work, are played by two antagonistic players, called the existential and universal players, who move a token around finite edge-coloured directed graphs. When the token lands on a position belonging to one of the players, this player moves it along an outgoing edge onto a new position. At the end of the day, the players have constructed an infinite path, called a play, and the winner is determined based on some language \mathbb{W} of winning infinite sequences of colours, called the winning condition (we call \mathbb{W} -games the games which use the winning condition \mathbb{W}). Solving such games consists of deciding whether the existential player has a winning strategy, i.e. a way to guarantee, whatever the moves of the opponent are, that the play will end up in the winning condition. Solving infinite duration games is at the crux of many algorithms used in verification, synthesis, and automata theory [9, 19, 38, 30]. Difficulties in solving them are both theoretical and practical, and many questions pertaining to game resolution still remain unanswered.



© Antonio Casares, Thomas Colcombet, and Karoliina Lehtinen;
licensed under Creative Commons License CC-BY 4.0

49th International Colloquium on Automata, Languages, and Programming (ICALP 2022).

Editors: Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff;

Article No. 117; pp. 117:1–117:20



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Memory. Several parameters are relevant for solving a game: its size, of course, but also its winning condition and the complexity of winning strategies. A measure of this complexity is the memory used by a strategy. The simplest strategies are those that use no memory (positional strategies): decisions depend exclusively on the current position, and not on the past of the game. A strategy uses a finite amount of memory if the information that we need to retain from the past can be summarized by a finite state machine that processes the sequence of moves played in the game. In this case, the amount of memory used by the strategy is the number of states of this machine. Given a winning condition \mathbb{W} , a fundamental question is what is the minimal quantity m such that if the existential player wins a \mathbb{W} -game, there is a winning strategy using a memory of size m (we call m the memory requirements of \mathbb{W}). In addition to its size, a memory also has structure, which further elucidates the game dynamics. Understanding both the size and the structure of memories for \mathbb{W} is a crucial step to design algorithms for solving \mathbb{W} -games.

Question A: Give a structural description of the optimal memory in \mathbb{W} -games.

Muller conditions. While there is a large zoo of winning conditions in the literature, here we are interested in ω -regular ones (described by finite state automata over infinite words), and, in particular, so called Muller conditions, for which the winner depends only on the colours that are seen infinitely often in the play. Memory requirements for Muller conditions have been studied in depth by Dziembowski, Jurdziński and Walukiewicz [15]. They provide a “formula” for computing the size of the minimal memory sufficient for winning in all games with a given Muller winning condition, based on the Zielonka tree [43], which describes the structure of a Muller condition. The Zielonka tree has also been used to characterise the memory requirements of Muller conditions when randomised strategies are allowed [21] and to provide minimal parity automata recognising a Muller condition [11]. This fundamental structure is also at the heart of our contribution.

Game reductions and good-for-gameness. When confronted with a \mathbb{W} -game, a standard solution is to reduce it to a game with a larger underlying graph, but a simpler winning condition. The typical way to do this (but not the only one) is to perform the composition of the game with a suitable automaton with another acceptance condition \mathbb{W}' that accepts the language \mathbb{W} . The result is an \mathbb{W}' -game which has as size the product of the size of the original game and the size of the automaton. There is a subtlety here: not all automata can be used for this operation. For a non-deterministic automaton, this is in general incorrect, while using a deterministic automaton is always correct. So here, finding a minimal deterministic automaton for a given language improves the complexity of game resolution, and there is a large body of research in this direction (see [11, 26, 31] for Muller conditions, [1, 10, 41, 42] for minimisation of automata, and [32, 39, 36, 37, 40, 34, 31, 14, 29] for determinisation). However, some non-deterministic automata can also be used to perform this reduction. These are called good-for-games automata (GFG) [20, 12].

Some languages are known to be recognised by good-for-games automata that are exponentially more succinct than any equivalent deterministic automaton [27], and several lines of research concerning good-for-games automata are under study (how to decide “good-for-gameness” [3, 27, 4, 6], how expressive is “good-for-gameness” for pushdown automata [28, 18] what are good-for-games quantitative automata [5], etc). However, one key question that has not yet been addressed concerning good-for-games automata is how to design techniques as general as possible for building them. To the best of our knowledge, the only existing result in this direction is a polynomial-time algorithm to minimise co-Büchi GFG automata [1].

Question B: Provide general tools for constructing good-for-games automata.

In this paper, in the context of Muller conditions, we relate these two lines of study, and in particular give partial answers to the general questions A and B. Indeed, we show that the memory needed to win in L -games for a Muller language L coincides with the size of minimal GFG Rabin automata for L , and, in this sense, we give a structural description of the memory for Muller games, thus giving a refined answer to question A in this case. We also provide an optimal way to construct these minimal good-for-games automata, thus answering question B in the context of Muller conditions.

Contributions

1. We show that for all ω -regular languages L , the size (number of states) of a good-for-games Rabin-automaton for L is an upper bound on the memory that the existential player needs to implement winning strategies for L -games. This inequality is straightforward, but had not been stated explicitly prior to this work.
2. We establish that when L is a Muller language, the following two quantities are equal: the least size of a good-for-games Rabin-automaton for L and the least memory required for the existential player in all L -games in which she wins. Furthermore, we provide an efficient way to construct such a minimal automaton from the Zielonka tree of the condition [43]. This automaton can be seen, in a certain way, as a quotient of the minimal deterministic parity automaton for this language, as described in [11].

Let us note that the least amount of memory needed to win a Muller game was described precisely by Dziembowski, Jurdziński and Walukiewicz [15]. We show here that the optimal strategy described in [15] can be implemented in a good-for-games Rabin-automaton. In combination with Item 1, this provides another proof of the upper bound in [15].

3. Finally, we provide a family of Muller languages such that the smallest GFG Rabin automata recognising it are of linear size in the number of letters, while equivalent deterministic Rabin automata grow exponentially. Note that the least size of a deterministic Rabin automaton for a Muller language L is known to coincide with the chromatic memory needed for winning L -games [10] (i.e. a memory that is updated based only on the letters seen, independently of the position in the game). The question of equivalence between chromatic memory and memory was asked by Kopczyński [23, 24], and an arbitrary difference between these two notions was established only recently by Casares [10]. Our new result, which is incomparable, shows that the chromatic memory can grow exponentially in the size of the alphabet, even when the general memory remains linear.

Together these three points develop techniques to solve Muller games in an optimal way by means of good-for-games Rabin automata reductions. The last point shows that an exponential gain can be achieved compared to using classical deterministic Rabin automata. Overall, our contribution supplements our understanding of Muller languages and highlights the – so far unexplored – fundamental role of GFG automata in the equation. Indeed, up to now GFG automata had mainly been studied for their succinctness, expressivity or algorithmic properties. Here, we shed light on a novel dimension of this automata class.

Related work. There is vast amount of literature on the memory requirements of different games. The first results in this direction where the proofs of the positionality of parity conditions and half-positionality of Rabin conditions [16, 22] and the finite-memory determinacy of Muller games [19]. The exact memory requirements of Muller conditions were characterised in [15]. In his PhD Thesis [23, 24], Kopczyński characterises several classes of

conditions that are half-positional, introduces the concept of chromatic memories (memories that are updated based only on colours seen) and provides an algorithm to decide the chromatic memory requirements of a winning condition. Conditions that are positional for both players over all graphs where characterised in [13] and those that are positional over finite graphs in [17]. More recently, these two results have been generalized to finite-memory conditions [8, 7]. The memory requirements have been proved to be different to the chromatic memory requirements in general [10], but conditions that are finite-memory determined are also chromatic-finite-memory determined [25].

Structure of this document. In Section 2, we describe the classical definitions related to our work such as games, automata and good-for-gamesness. In Section 3, we show why good-for-games Rabin automaton can be used as a memory structure for the existential player, the optimality of the construction for Muller conditions, and how to construct the least such automaton. In Section 4, we establish that this construction can be exponentially more succinct than deterministic Rabin automata. Section 5 concludes the paper.

2 Definitions

Notations. $|A|$ denotes the cardinality of a set A , $\mathcal{P}(A)$ its power set and $\mathcal{P}_+(A) = \mathcal{P}(A) \setminus \{\emptyset\}$. For a finite non-empty alphabet Σ , we write Σ^* and Σ^ω for the sets of finite and *infinite words* over Σ , respectively. The empty word is denoted by ε . Given $w = w_0w_1w_2\cdots \in \Sigma^\omega$, we denote $\text{Inf}(w) \subseteq \Sigma$ the set of letters that appear infinitely often in w . We let $w[i..j]$ be the finite word $w_iw_{i+1}\dots w_j$ if $i \leq j$, and ε if $j < i$.

We extend maps $\sigma : A \rightarrow B$ to A^* and A^ω component-wise and we denote these extensions by σ whenever no confusion arises. For a positive rational number $q \in \mathbb{Q}$ we denote by $\lfloor q \rfloor$ the greatest integer $n \in \mathbb{N}$ such that $n \leq q$.

2.1 Games and their memory

Games. We consider turn-based infinite duration games played between the existential and the universal player (referred to as *Exist* and *Univ*) over a directed graph. Formally, a Γ -coloured *game* is a tuple $\mathcal{G} = (V = V_E \uplus V_A, E, x_0, \mathbb{W})$, which consists of a set of vertices V partitioned into Exist's positions V_E and Univ's ones, V_A ; a set of *transitions* (also called edges or moves) $E \subseteq V \times (\Gamma \cup \{\varepsilon\}) \times V$; an initial vertex $x_0 \in V$ and a subset $\mathbb{W} \subseteq \Gamma^\omega$ of winning sequences. We make the assumptions that there is at least one move from every position and that no cycle is labelled exclusively by ε . We will denote by $\gamma : E \rightarrow \Gamma \cup \{\varepsilon\}$ the function that assigns to each edge its colour. We write $\text{Out}(x)$ for the set of outgoing moves from x , that is, $\text{Out}(x) = \{e \in E : e = (x, c, x') \text{ for some } x' \in V, c \in \Gamma \cup \{\varepsilon\}\}$. If Γ is a game using the winning condition \mathbb{W} we call it a \mathbb{W} -*game*.

Each player moves a pebble along an outgoing edge whenever it lands on a position belonging to that player, forming an infinite path $\pi \in E^\omega$ starting in x_0 called a *play*.

We denote $\gamma(\pi)$ sequence of colours labelling π omitting the ε labels (we remark that $\gamma(\pi) \in \Gamma^\omega$, since there are no cycles entirely labelled by ε). The play is *winning* for the existential player if $\gamma(\pi) \in \mathbb{W}$. A *partial play* is a finite path $\pi \in E^*$ in \mathcal{G} starting in x_0 . A *strategy for the existential player* is a function $\sigma : E^* \rightarrow E$ such that if a partial play π ends in a position $x \in V_E$, then $\sigma(\pi) \in \text{Out}(x)$. We say that a play π is *consistent with* the strategy σ if for every partial play π' that is a prefix of π ending in a position controlled by Exist, the next edge in π is $\sigma(\pi')$. The strategy σ is *winning* if every play consistent with σ is winning for the existential player. We say that the game \mathcal{G} is *won by* the existential player

if that player has a winning strategy in \mathcal{G} . A strategy is *positional* if it can be represented by a function $\sigma: V_E \rightarrow E$ (that is, the choice of the next transition only depends on the current position, and not on the history of the path).

Winning conditions. We fix an alphabet Γ .

Muller. A *Muller condition* over the alphabet Γ is given by a family $\mathcal{F} \subseteq \mathcal{P}_+(\Gamma)$. A word $w \in \Gamma^\omega$ satisfies the Muller condition \mathcal{F} if $\text{Inf}(w) \in \mathcal{F}$. The language of the Muller condition $\mathcal{L}_{\mathcal{F}} \subseteq \Gamma^\omega$ contains the ω -words that satisfy \mathcal{F} .

Rabin. A *Rabin condition* over the alphabet Γ is represented by a family of *Rabin pairs* $R = \{(G_1, R_1), \dots, (G_r, R_r)\}$, where $G_i, R_i \subseteq \Gamma$ and $G_i \cap R_i = \emptyset$. The Rabin pair j is said to be *green in c* if $c \in G_j$, to be *red in c* if $c \in R_j$, or to be *orange in c* if none of the previous occur. A word $w \in \Gamma^\omega$ satisfies the Rabin condition R if $\text{Inf}(w) \cap G_j \neq \emptyset$ and $\text{Inf}(w) \cap R_j = \emptyset$ for some index $j \in \{1, \dots, r\}$. Said differently, there is a Rabin pair j which is red in finitely many letters from w , and green for infinitely many letters of w . The language of the Rabin condition $\mathcal{L}_{\mathcal{R}} \subseteq \Gamma^\omega$ contains the ω -words that satisfy \mathcal{R} .

Parity. To define a *parity condition* we suppose that $\Gamma \subseteq \mathbb{N}$. A word $w \in \Gamma^\omega$ satisfies the parity condition if the maximum in $\text{Inf}(w)$ is even. The language of the parity condition contains the ω -words that satisfy it.

We say that a language $L \subseteq \Gamma^\omega$ is a *Muller language* if it is the language of some Muller condition \mathcal{F} . Equivalently, L is a Muller language if it can be described as a boolean combination of atomic propositions of the form “the letter “ a ” appears infinitely often” and their negations. Note that languages of Rabin conditions are languages of Muller conditions, and that languages of parity conditions are languages of Rabin conditions, but the converses do not hold. Given a Muller condition $\mathcal{F} \subseteq \mathcal{P}_+(\Gamma)$ and a subset $C \subseteq \Gamma$, we define the *restriction of \mathcal{F} to C* as the Muller condition over C given by $\mathcal{F}|_C = \{A \subseteq C : A \in \mathcal{F}\}$.

Memory structures. A *memory structure for the game \mathcal{G}* of moves E is a tuple $\mathcal{M} = (M, m_0, \mu, \sigma)$ where M is a set of memory states, $m_0 \in M$ is an *initial memory state*, $\mu: M \times E \rightarrow M$ is an *update function*, and $\sigma: M \times V_E \rightarrow E$ maps each position x owned by the existential player to a move from x . The *size* of \mathcal{M} is the cardinal of M . We extend the function μ to paths by induction: $\mu(m, \varepsilon) = m$, and $\mu(m, \pi e) = \mu(\mu(m, \pi), e)$ for a path $\pi \cdot e \in E^*$. The memory structure \mathcal{M} induces a strategy $\sigma_{\mathcal{M}}: E^* \rightarrow V_E$ given by $\sigma_{\mathcal{M}}(\pi) = \sigma(\mu(m_0, \pi), \text{Last}(\pi))$, where $\text{Last}(\pi)$ is the last position of the partial play π .

We say that \mathcal{M} is a *chromatic memory structure* if there is a function $\mu_c: M \times (\Gamma \cup \{\varepsilon\}) \rightarrow M$ such that $\mu_c(m, \varepsilon) = m$ for every $m \in M$ and $\mu(m, e) = \mu_c(m, \gamma(e))$ for all edges $e \in E$.

The *memory requirements* of a winning condition \mathbb{W} are defined as the least integer n such that if \mathcal{G} is an \mathbb{W} -game won by the existential player, then she has a winning strategy given by a memory of size at most n . We denote this quantity by $\text{mem}(\mathbb{W})$.

We say that a winning condition \mathbb{W} is *exist-positional* (also called *half-positional*) if $\text{mem}(\mathbb{W}) = 1$. Equivalently, \mathbb{W} is exist-positional if Exist has a winning positional strategy whenever Exist has a winning strategy at all.

2.2 Automata and good-for-gameness

Automata. A *non-deterministic automaton* $\mathcal{A} = (Q, \Sigma, Q_0, \Delta, \Gamma, \mathbb{W})$ (or simply an *automaton*) consists of a finite set of states Q , an input alphabet Σ , a non-empty set of initial states $Q_0 \subseteq Q$, a transition relation $\Delta \subseteq Q \times \Sigma \times \Gamma \times Q$ and an acceptance condition $\mathbb{W} \subseteq \Gamma^\omega$. We will

write $\delta: Q \times \Sigma \rightarrow \mathcal{P}(Q)$ for the function $\delta(q, a) = \{q' \in Q : (q, a, c, q') \in \Delta \text{ for some } c \in \Gamma\}$. The *size* of the automaton, $|\mathcal{A}|$, is the number of its states. A *run of the automaton* over a word $a_0a_1a_2 \cdots \in \Sigma^\omega$ is a sequence of transitions of the form:

$$\rho = (q_0, a_0, c_0, q_1)(q_1, a_1, c_1, q_2) \cdots \in \Delta^\omega, \text{ such that } q_0 \in Q_0 \text{ is an initial state.}$$

A run is *accepting* if $c_0c_1c_2 \cdots \in \mathbb{W}$. If an accepting run over a word $w \in \Sigma^\omega$ exists, the automaton *accepts* w . The set of accepted words is the *language accepted by the automaton*, written $L(\mathcal{A})$. The automaton is *deterministic* if Q_0 is a singleton and Δ is such that for all states q and letter $a \in \Sigma$, there exists exactly one transition of the form (q, a, c, q') . In this case, for all words $w \in \Sigma^\omega$ there exists one and exactly one run of the automaton over w .

An automaton using an acceptance condition \mathbb{W} (resp. an acceptance condition of type $X \in \{\text{Muller, Rabin, parity}\}$) is called an \mathbb{W} -*automaton* (resp. X -automaton).

Good-for-gameness. The automaton \mathcal{A} is *good-for-games* (GFG) if there is a resolver for it, consisting of a choice of an initial state $r_0 \in Q_0$ and a function $r: \Sigma^* \times \Sigma \rightarrow \Delta$ such that for all words $w \in L(\mathcal{A})$, the run $t_0t_1 \dots \in \Delta^\omega$, called the *run induced by r* and defined by $t_i = r(w[0..i-1], w[i])$, starts in r_0 and is an accepting run over w . In other words, r should be able to construct an accepting run in \mathcal{A} letter-by-letter with only the knowledge of the word so far, for all words in $L(\mathcal{A})$.

2.3 The Zielonka tree of a Muller condition

A *tree* $T = (N, \sqsubseteq)$ is a nonempty finite set of *nodes* N equipped with an order relation \sqsubseteq called the *ancestor relation* (x is an ancestor of y if $x \sqsubseteq y$), such that (1) there is a minimal node for \sqsubseteq , called *the root*, and (2) the ancestors of an element are totally ordered by \sqsubseteq . The converse relation is the *descendant* relation. Maximal nodes are called *leaves*, and the set of leaves of T is denoted by $Leaves(T)$. Given a node n of a tree T , the *subtree of T rooted at n* is the tree T restricted to the nodes that have n as ancestor. A node n' is a *child* of n if it is a minimal strict descendant of it. The set of children of n is written $Children_T(n)$. The *height* of a tree T is the maximal length of a chain for the ancestor relation. An *A -labelled tree* is a tree T together with a labelling function $\nu: N \rightarrow A$.

► **Definition 1 ([43]).** Let $\mathcal{F} \subseteq \mathcal{P}_+(\Gamma)$ be a Muller condition. A Zielonka tree for \mathcal{F} , denoted $\mathcal{Z}_{\mathcal{F}} = (N, \sqsubseteq, \nu: N \rightarrow \mathcal{P}_+(\Gamma))$ is a $\mathcal{P}_+(\Gamma)$ -labelled tree with nodes partitioned into round nodes and square nodes, $N = N_{\circ} \uplus N_{\square}$ such that:

- The root is labelled Γ .
- If a node is labelled $X \subseteq \Gamma$, with $X \in \mathcal{F}$, then it is a round node, and its children are labelled exactly with the maximal subsets $Y \subseteq X$ such that $Y \notin \mathcal{F}$.
- If a node is labelled $X \subseteq \Gamma$, with $X \notin \mathcal{F}$, then it is a square node, and its children are labelled exactly with the maximal subsets $Y \subseteq X$ such that $Y \in \mathcal{F}$.

We remark that if n is a node of $\mathcal{Z}_{\mathcal{F}}$, then the subtree of $\mathcal{Z}_{\mathcal{F}}$ rooted at n is a Zielonka tree for $\mathcal{F}|_{\nu(n)}$, (the restriction of \mathcal{F} to the label of n).

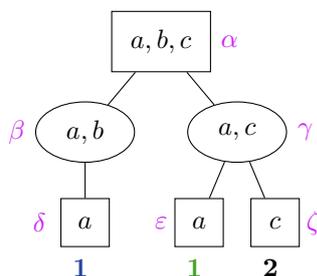
We equip trees with an order in order to navigate in them. An *ordered Zielonka tree* is a Zielonka tree for which the set of children of each node is equipped with a total order (“from left to right”). The function $Next_n$ maps each child of n to its successor for this order, in a cyclic way. For each node $n \in N$ and each leaf l below n we define the set $Jump_n(l)$ containing a leaf l' if there are two children n_1, n_2 of n such that $n_1 \sqsubseteq l, n_2 \sqsubseteq l'$ and $n_2 = Next_n(n_1)$ (we remark that $n_1 = n_2$ if n has only one child). For $n = l$ we define $Jump_n(l) = \{l\}$. That

is, $l' \in \text{Jump}_n(l)$ if we can reach l' by the following procedure: we start at l , we go up the tree until finding the node n , we change to the next branch below n (in a cyclic way) and we re-descend to l' . From now on, we will suppose that all Zielonka trees are ordered, without explicitly mentioning it.

► **Example 2.** We will use the following Muller condition as a running example throughout the paper. Let $\Gamma = \{a, b, c\}$ and let \mathcal{F} be the Muller condition defined by:

$$\mathcal{F} = \{\{a, b\}, \{a, c\}, \{b\}\}.$$

In Figure 1 we show the Zielonka tree for \mathcal{F} . We use Greek letters to name the nodes of the tree, N . We have that $\text{Jump}_\alpha(\delta) = \{\varepsilon, \zeta\}$ and $\text{Jump}_\gamma(\zeta) = \{\varepsilon\}$. The numbering of the branches will be used in Section 3.2.



■ **Figure 1** Zielonka tree $\mathcal{Z}_{\mathcal{F}}$ for $\mathcal{F} = \{\{a, b\}, \{a, c\}, \{b\}\}$.

► **Definition 3** ([15]). Let T be a tree with nodes partitioned into round and square nodes, its memory for the existential player (memory for short), denoted $\text{mem-tree}(T)$, is defined inductively as:

- 1 if T has exactly one node.
- The sum of the memories of the subtrees of T rooted at the children of the root, if the root is round.
- The maximum of the memories of the subtrees of T rooted at the children of the root, if the root is square.

For instance, for the Zielonka tree from Example 2, $\text{mem-tree}(\mathcal{Z}_{\mathcal{F}}) = 2$.

The key result justifying the introduction of this notion is that it characterises precisely the quantity of memory required for winning an \mathcal{F} -game, as shown by the next proposition.

► **Proposition 4** ([15]). For all Muller conditions \mathcal{F} , $\text{mem}(\mathcal{F}) = \text{mem-tree}(\mathcal{Z}_{\mathcal{F}})$.

3 GFG Rabin automata correspond to memory structures for Muller games

In this section, we prove the following result:

► **Theorem 5.** Let L be a Muller language. The memory requirements for L coincide with the size of a minimal GFG Rabin automaton recognising L .

In Section 3.1 we show the first direction: the size of a GFG Rabin automaton for a Muller language L is always an upper bound on the memory required by the existential player on L -games. In Section 3.2, we show how to construct a GFG Rabin automaton from

the Zielonka tree of a Muller condition of size $\text{mem-tree}(\mathcal{Z}_L) = \text{mem}(L)$, completing the equivalence. Moreover, we can build this minimal GFG Rabin automaton in polynomial time given the Zielonka tree of the Muller condition.

3.1 A GFG Rabin automaton induces a memory structure for any game

In this section, we establish Corollary 10 which states that the size of a GFG Rabin automaton \mathcal{A} accepting a Muller language L is an upper bound on the memory required for winning all L -games. Concretely, given an L -game won by the existential player, we are able to construct a memory structure inducing a winning strategy based on \mathcal{A} . The argument is standard: we construct the product game of \mathcal{A} and the L -game, which is a Rabin game in which the existential player enjoys a positional winning strategy; then we use the \mathcal{A} -component of this product as a memory structure for a strategy in the original L -game.

► **Lemma 6.** *Let $\mathcal{A} = (Q, \Sigma, Q_0, \Delta, \Gamma, \mathbb{W})$ be a GFG \mathbb{W} -automaton recognising a language $L \subseteq \Sigma^\omega$, with \mathbb{W} exist-positional. Then if \mathcal{G} is an L -game won by Exist, she can win it using a strategy given by a memory structure $\mathcal{M} = (Q, r_0, \mu, \sigma)$.*

Proof. Let $\mathcal{G} = (V = V_E \uplus V_A, E, x_0, L)$ and $\mathcal{A} = (Q, \Sigma, r_0, \Delta, \Gamma, \mathbb{W})$, where $r_0 \in Q_0$ is the initial state chosen by some resolver. We consider the product \mathbb{W} -game \mathcal{G}' in which:

- Positions are elements in $V' = (V \times Q) \cup (V \times \Gamma \times \mathcal{P}(Q))$. The initial position is (x_0, r_0) .
- Exist's positions are $V'_E = (V_E \times Q) \cup (V \times \Gamma \times \mathcal{P}(Q))$.
- There is an ε -coloured edge from (x, q) to $(x', c, \delta(q, c))$ if $(x, c, x') \in E$, $c \neq \varepsilon$, and to (x', q) if $(x, \varepsilon, x') \in E$.
- There is a c -labelled edge from (x, c, S) to (x, q') for all $q' \in S \subseteq Q$.
- The winning condition is \mathbb{W} .

In short, in this game the players still negotiate a play in \mathcal{G} , but in addition, the existential player must simultaneously build an accepting run on the labelling of this play in \mathcal{A} .

If \mathcal{A} is GFG then whenever the existential player wins in \mathcal{G} , she also wins in \mathcal{G}' [20, Theorem 3] by playing a winning strategy on the \mathcal{G} component of \mathcal{G}' and using the resolver for \mathcal{A} to choose the successor state in the \mathcal{A} component.

This strategy is not necessarily positional. However, since \mathbb{W} is exist-positional, the existential player also has a positional strategy $s : V'_E \rightarrow E'$. We can now build a memory structure $\mathcal{M}_s = (Q, r_0, \mu, \sigma)$ that projects the strategy s onto \mathcal{G} (and updates itself with the projection of s onto \mathcal{A}):

- $\mu(q, e) = q$ if e is an ε -coloured move of \mathcal{G} and
- $\mu(q, (x, c, x')) = q'$ where $q' = s(x', c, \delta(q, c))$ otherwise;
- $\sigma(q, x) = s(x, q)$.

Since s is a winning strategy in \mathcal{G}' , its projection onto \mathcal{G} is a strategy that only agrees with plays with labels in $L(\mathcal{A})$, that is, winning plays. ◀

► **Remark 7.** Note that there is a slight subtlety here: the resolver, which induces a winning strategy in the product game, does not need to be positional. In fact, it might require exponential memory [27, Theorem 1]. Yet, for each L -game, a memory based on \mathcal{A} suffices.

► **Lemma 8** ([22, 43]). *Rabin conditions are exist-positional.*

As a direct consequence we obtain the following Proposition.

► **Proposition 9.** *Let L be a Muller language accepted by a GFG Rabin automaton \mathcal{A} . Then, in every L -game \mathcal{G} won by the existential player, she can win using a strategy given by a memory structure of size $|\mathcal{A}|$.*

► **Corollary 10.** *If \mathcal{A} is a GFG Rabin automaton accepting a Muller language L , then $\text{mem}(L) \leq |\mathcal{A}|$.*

3.2 An optimal construction of a GFG Rabin automaton

So far, we have seen that given a GFG Rabin automaton, it can serve as a memory structure for the games of which it accepts the winning condition. In this section we do the converse: we build a minimal GFG Rabin automaton for a Muller language L of the same size as the minimal memory required to win in L -games.

3.2.1 The construction

► **Proposition 11.** *Let $\mathcal{F} \subseteq \mathcal{P}_+(\Gamma)$ be a Muller condition. There exists a GFG Rabin automaton recognising $\mathcal{L}_{\mathcal{F}}$ of size $\text{mem}(\mathcal{F})$.*

To prove Proposition 11, we build a GFG Rabin automaton $\mathcal{R}_{\mathcal{F}} = (Q, \Gamma, q_0, \Delta, \Gamma', R)$ for $\mathcal{L}_{\mathcal{F}}$ based on the Zielonka tree $\mathcal{Z}_{\mathcal{F}}$, as illustrated in Figure 2. We use a mapping from the leaves of $\mathcal{Z}_{\mathcal{F}}$ to the states of the automaton that guarantees that two leaves of which the last common ancestor is a round node cannot map to the same state. The number of states required to satisfy this condition (\star below) coincides with $\text{mem-tree}(\mathcal{Z}_{\mathcal{F}})$. Then, for each leaf of $\mathcal{Z}_{\mathcal{F}}$ and letter $c \in \Gamma$, we identify its last ancestor n in $\mathcal{Z}_{\mathcal{F}}$ containing c , and, using the Jump_n function (defined in Section 2.3), pick a leaf below the next child of n . We add a c -transition with label n between the states mapped to from these leaves. This way, we can identify a run in the automaton with a promenade through the nodes of the Zielonka tree. If during this promenade a unique minimal node (for \sqsubseteq) is visited infinitely often, it is not difficult to see that the sequence of input colours belongs to \mathcal{F} if and only if the label of this minimal node is an accepting set (it is a round node). We devise a Rabin condition over the set of nodes of the Zielonka tree accepting exactly these sequences of nodes.

We now describe the construction of the automaton $\mathcal{R}_{\mathcal{F}} = (Q, \Gamma, q_0, \Delta, N, R)$ formally, starting from the Zielonka tree $\mathcal{Z}_{\mathcal{F}} = (N, \sqsubseteq, \nu : N \rightarrow \mathcal{P}_+(\Gamma))$, and then proceed to prove its correctness.

States. First, we set $Q = \{1, 2, \dots, \text{mem-tree}(\mathcal{Z}_{\mathcal{F}})\}$ and we label the leaves of $\mathcal{Z}_{\mathcal{F}}$ by a mapping $\eta : \text{Leaves}(\mathcal{Z}_{\mathcal{F}}) \rightarrow \{1, 2, \dots, \text{mem-tree}(\mathcal{Z}_{\mathcal{F}})\}$ verifying the property:

If $n \in \mathcal{Z}_{\mathcal{F}}$ is a round node with children $n_1 \neq n_2$, for any pair of leaves l_1 and l_2 below n_1 and n_2 , respectively, $\eta(l_1) \neq \eta(l_2)$. (\star)

► **Lemma 12.** *For every Zielonka tree $\mathcal{Z}_{\mathcal{F}}$ there is a mapping verifying Property \star of the form $\eta : \text{Leaves}(\mathcal{Z}_{\mathcal{F}}) \rightarrow \{1, 2, \dots, \text{mem-tree}(\mathcal{Z}_{\mathcal{F}})\}$.*

Proof. We prove it by induction in the height of $\mathcal{Z}_{\mathcal{F}}$. Let n_1, \dots, n_k be the children of the root of $\mathcal{Z}_{\mathcal{F}}$. We write \mathcal{F}_i to denote the Muller condition restricted to $\nu(n_i)$ and $\eta_i : \text{Leaves}(\mathcal{Z}_{\mathcal{F}_i}) \rightarrow \{1, 2, \dots, \text{mem-tree}(\mathcal{Z}_{\mathcal{F}_i})\}$ be a labelling verifying Property \star , for $1 \leq i \leq k$. We distinguish two cases according to the shape of the root.

If the root of $\mathcal{Z}_{\mathcal{F}}$ is a square node ($\Gamma \notin \mathcal{F}$), then the mapping $\eta(l) = \eta_i(l)$ if the leaf l belongs to the subtree $\mathcal{Z}_{\mathcal{F}_i}$ verifies Property \star .

117:10 On the Size of Good-For-Games Rabin Automata and the Memory in Muller Games

If the root of $\mathcal{Z}_{\mathcal{F}}$ is a round node ($\Gamma \in \mathcal{F}$), then $\text{mem-tree}(\mathcal{Z}_{\mathcal{F}}) = \sum_{i=1}^k \text{mem-tree}(\mathcal{Z}_{\mathcal{F}_i})$, and we can partition $\{1, 2, \dots, \text{mem-tree}(\mathcal{Z}_{\mathcal{F}})\}$ into disjoint sets C_1, \dots, C_k of size $|C_i| = \text{mem-tree}(\mathcal{Z}_{\mathcal{F}_i})$. We write σ_i for a bijection from $\{1, \dots, \text{mem-tree}(\mathcal{Z}_{\mathcal{F}_i})\}$ to C_i . Then, the mapping $\eta(l) = \sigma_i(\eta_i(l))$, if the leaf l belongs to the subtree $\mathcal{Z}_{\mathcal{F}_i}$ verifies Property \star . ◀

We suppose that the image of the leftmost leaf under η is 1 and choose the initial state q_0 to be 1. In Example 2, the labelling $\eta(\delta) = \eta(\varepsilon) = 1$, $\eta(\zeta) = 2$ verifies Property \star .

Transitions. For each leaf $l \in \text{Leaves}(\mathcal{Z}_{\mathcal{F}})$ and each letter $c \in \Gamma$, we define a c -transition from $\eta(l)$, with an output label from $\Gamma' = N$, as follows: let n be the maximal ancestor of l that contains the letter c in its label and let l' be the leftmost leaf in $\text{Jump}_n(l)$ ¹. Then, $(\eta(l), c, n, \eta(l')) \in \Delta$. That is, if we are in a state $\eta(l)$, when we read the letter $c \in \Gamma$ we can choose to go up in the Zielonka tree from l until visiting a node n with c in its label. We produce the letter n as output, then we change to the next child of n (in a cyclic way) and we descend to the leftmost leaf below it. The destination is the η -label of this leaf².

Following the above definition, we obtain a mapping from transitions in the automaton to $\text{Leaves}(\mathcal{Z}_{\mathcal{F}}) \times \Gamma \times N \times \text{Leaves}(\mathcal{Z}_{\mathcal{F}})$. We say that l and l' are the leaves *corresponding to* the transition (q, a, n, q') if this transition is sent to (l, a, n, l') by this mapping. The node n produced as output is the last common ancestor of l and l' . The automaton obtained in this way might present multiple transitions labelled by the same input letter between two states. We will show in Proposition 18 that duplicated transitions can be removed.

Acceptance condition. We define a Rabin condition over the alphabet $\Gamma' = N$, that is the set of nodes of the Zielonka tree. We define a Rabin pair for each round node of $\mathcal{Z}_{\mathcal{F}}$ (that is, nodes whose label is an accepting set of letters for \mathcal{F}): $R = \{(G_n, R_n)\}_{n \in N_{\circ}}$. Let n be a round node and n' be a general node of $\mathcal{Z}_{\mathcal{F}}$:

$$\begin{cases} n' \in G_n & \text{if } n' = n, \\ n' \in R_n & \text{if } n' \neq n \text{ and } n \text{ is not an ancestor of } n'. \end{cases}$$

That is, for the letter n' , the Rabin pairs corresponding to round ancestors of n' are not affected by it (they are “orange in n' ”). If this node n' is round, then it belongs to $G_{n'}$ (this pair is “green in n' ”). For any other node $n \in N_{\circ}$, we have $n' \in R_n$ (the pair is “red in n' ”).

► **Remark 13.** The construction presented depends on the order of the nodes of the Zielonka tree. However, the size of the resulting automaton is independent of this order.

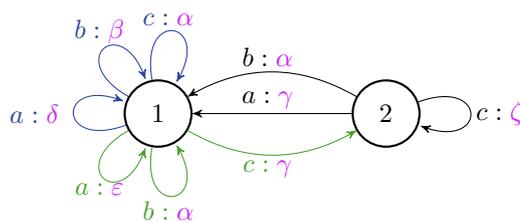
► **Example 14.** Let $\mathcal{F} = \{\{a, b\}, \{a, c\}, \{b\}\}$ be the Muller condition from Example 2. The labelling of the leaves of the Zielonka tree given by $\eta(\delta) = \eta(\varepsilon) = 1$, $\eta(\zeta) = 2$ verifies Property \star . Figure 2 shows the GFG Rabin automaton obtained by following this procedure.

The Rabin condition of this automaton is given by two Rabin pairs (corresponding to the round nodes of the Zielonka tree in Figure 1):

$$\begin{aligned} G_{\beta} &= \{\beta\}, & R_{\beta} &= \{\alpha, \gamma, \varepsilon, \zeta\}, \\ G_{\gamma} &= \{\gamma\}, & R_{\gamma} &= \{\alpha, \beta, \delta\}. \end{aligned}$$

¹ We could add all transitions $\{(\eta(l), n, \eta(l')) : l' \in \text{Jump}_n(l)\}$ to Δ . However, a resolver for the GFG automaton just needs to make use of one of these transitions, so in order to simplify the automaton we make an arbitrary choice (the leftmost leaf).

² We remark that l' is the target of the transition of the Zielonka tree parity automaton from l reading letter “ c ” [11]. See also Section 3.2.3.



■ **Figure 2** The GFG Rabin automaton obtained from the Zielonka tree $\mathcal{Z}_{\mathcal{F}}$.

3.2.2 Proof of correctness

► **Lemma 15.** *Let $w = n_0 n_1 n_2 \dots \in N^\omega$ be an infinite sequence of nodes of the Zielonka tree. The word w satisfies the Rabin condition defined above if and only if there is a unique minimal node for the ancestor relation in $\text{Inf}(w)$ and this minimal node is round (recall that the root is the minimal element in $\mathcal{Z}_{\mathcal{F}}$).*

Proof. Suppose that there is a unique minimal node in $\text{Inf}(w)$, called n , and that n is round. We claim that w is accepted by the Rabin pair (G_n, R_n) . It is clear that $\text{Inf}(w) \cap G_n \neq \emptyset$, because $n \in G_n$. It suffices to show that $\text{Inf}(w) \cap R_n = \emptyset$: By minimality, any other node $n' \in \text{Inf}(w)$ is a descendant of n (equivalently, n is an ancestor of n'), so $n' \notin R_n$.

Conversely, suppose that $w \in N^\omega$ satisfies the Rabin condition. Then, there is some round node $n \in N_{\bigcirc}$ such that $\text{Inf}(w) \cap G_n \neq \emptyset$ and $\text{Inf}(w) \cap R_n = \emptyset$. Since $G_n = \{n\}$, we deduce that $n \in \text{Inf}(w)$. Moreover, as $\text{Inf}(w) \cap R_n = \emptyset$, all nodes in $\text{Inf}(w)$ are descendants of n . We conclude that n is the unique minimal node in $\text{Inf}(w)$, and it is round. ◀

► **Lemma 16.** *The automaton $\mathcal{R}_{\mathcal{F}}$ recognises the language $\mathcal{L}_{\mathcal{F}}$ and is good-for-games.*

Proof. $\mathcal{L}(\mathcal{R}_{\mathcal{F}}) \subseteq \mathcal{L}_{\mathcal{F}}$: Let $u \in \mathcal{L}(\mathcal{R}_{\mathcal{F}})$ and let $w \in N^\omega$ be the sequence of nodes produced as output of an accepting run over u in $\mathcal{R}_{\mathcal{F}}$. By Lemma 15, there is a unique minimal node n for \sqsubseteq appearing infinitely often in w and moreover n is round. Let n_1, \dots, n_k be an enumeration of the children of n (from left to right), with labels $\nu(n_i) \subseteq \Gamma$ (we remark that $\nu(n_i) \notin \mathcal{F}$, for $1 \leq i \leq k$). We will prove that $\text{Inf}(u) \subseteq \nu(n)$ and $\text{Inf}(u) \not\subseteq \nu(n_i)$ for $1 \leq i \leq k$. By definition of the Zielonka tree, as n is round, this implies that $\text{Inf}(u) \in \mathcal{F}$.

Since eventually all nodes produced as output are descendants of n (by minimality), $\text{Inf}(u)$ must be contained in $\nu(n)$ (by definition of the transitions of $\mathcal{R}_{\mathcal{F}}$).

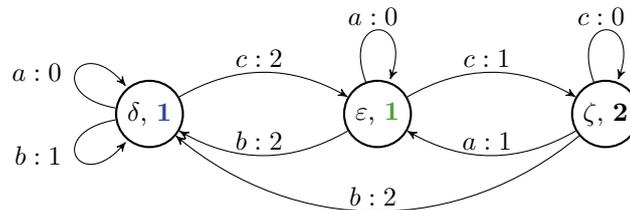
We suppose, towards a contradiction, that $\text{Inf}(u) \subseteq \nu(n_j)$ for some $1 \leq j \leq k$. Let $Q_i = \{\eta(l) : l \text{ is a leaf below } n_i\}$ be the set of states corresponding to leaves under n_i , for $1 \leq i \leq k$. We can suppose that the leaves corresponding to transitions of an accepting run over u are all below n , and therefore, transitions of such a run only visit states in $\bigcup_{i=1}^k Q_i$. Indeed, eventually this is going to be the case, because if some of the leaves l, l' corresponding to a transition (q, a, n', q') are not below n , then n' would not be a descendant of n (since n' is the least common ancestor of l and l'). Also, by Property \star , we have $Q_i \cap Q_j = \emptyset$, for all $i \neq j$. By definition of the transitions of $\mathcal{R}_{\mathcal{F}}$, if $c \in \Gamma$ is a colour in $\nu(n)$ but not in $\nu(n_i)$, all transitions from some state in Q_i reading the colour c go to Q_{i+1} , for $1 \leq i \leq k-1$ (and to Q_1 if $i = k$). Also, if $c \in \nu(n_i)$, transitions from states in Q_i reading c stay in Q_i . We deduce that a run over u will eventually only visit states in Q_j , for some j such that $\text{Inf}(u) \subseteq \nu(n_j)$. However, the only transitions from Q_j that would produce n as output are those corresponding to a colour $c \notin \nu(n_j)$, so the node n is not produced infinitely often, a contradiction.

$\mathcal{L}_{\mathcal{F}} \subseteq \mathcal{L}(\mathcal{R}_{\mathcal{F}})$ and good-for-gameness: We will describe a strategy for a resolver in $\mathcal{R}_{\mathcal{F}}$ using as memory the set of leaves of the Zielonka tree³. It will verify at every step that if the memory is on the leaf l , then $\mathcal{R}_{\mathcal{F}}$ is on the state $\eta(l)$. The initial state of the memory is the leftmost leaf of $\mathcal{Z}_{\mathcal{F}}$. If we are on the memory state l and the letter $c \in \Gamma$ is read, we take the transition $(\eta(l), c, n, \eta(l')) \in \Delta$, where n is the maximal ancestor of l such that $c \in \nu(n)$ and l' is the leftmost leaf in $\text{Jump}_n(l)$; the memory state is updated to l' . Let us suppose that a word $u \in \mathcal{L}_{\mathcal{F}}$ is given as input to the automaton. We will see that the run produced by this strategy is accepting. We can suppose that the only colours appearing in u are those of $\text{Inf}(u)$. Let n be the leftmost \sqsubseteq -maximal node such that $\text{Inf}(u) \subseteq \nu(n)$. Since $\text{Inf}(u) \in \mathcal{F}$, n is a round node. We will prove that the run produced by the resolver above only produces nodes that are descendants of n (including n) infinitely often and that it produces n infinitely often and is therefore accepting. Let n_1, \dots, n_k be the children of n from left to right, and let L_1, \dots, L_k be the (disjoint) sets of leaves below them, respectively. By the definition of the transitions and the strategy, the memory will eventually only consider leaves in $\bigcup_{i=1}^k L_i$, and will produce as output nodes that are descendants of n (including n itself). Also, each time that the memory is in some state in L_i and a colour not in $\nu(n_i)$ is given, a transition leading to some state in $\eta(L_{i+1})$ ($\eta(L_1)$ if $i = k$) producing the node n as output is taken. Since $\text{Inf}(u)$ is not contained in $\nu(n_i)$ for $1 \leq i \leq k$ (by the maximality assumption), this occurs infinitely often. ◀

► **Remark 17.** We have shown that given as input the Zielonka tree of a Muller condition \mathcal{F} we can build in polynomial time a minimal GFG Rabin automaton for $\mathcal{L}_{\mathcal{F}}$. On the other hand, with the same input, it is NP-complete to decide whether there is a deterministic Rabin automaton of size k recognising $\mathcal{L}_{\mathcal{F}}$ [10, Theorem 31]. Therefore, unless $\text{P} = \text{NP}$, there are Muller languages for which minimal deterministic Rabin automata are strictly greater than minimal GFG Rabin automata. We will explicitly show some of these languages in Section 4.

3.2.3 Relation with the Zielonka-tree parity automaton

The Zielonka tree has been previously used to provide a minimal deterministic parity automaton for a Muller condition [11, 33]. The automata states are the leaves of the Zielonka tree, and the transition from a leaf l reading colour c goes to the leftmost leaf in $\text{Jump}_n(l)$, where n is the last ancestor of l containing colour c . For example, Figure 3 shows a parity automaton recognising the Muller condition $\mathcal{F} = \{\{a, b\}, \{a, c\}, \{b\}\}$ from Example 2.



■ **Figure 3** Parity automaton $\mathcal{P}_{\mathcal{F}}$ obtained from the Zielonka tree from Figure 1.

³ This strategy is given by the (deterministic) Zielonka tree parity automaton $\mathcal{P}_{\mathcal{F}}$. It suffices to note that there is a morphism from $\mathcal{P}_{\mathcal{F}}$ to $\mathcal{R}_{\mathcal{F}}$ preserving all edges and the acceptance of loops.

This minimal parity automaton $\mathcal{P}_{\mathcal{F}}$ is closely related to the GFG Rabin automaton $\mathcal{R}_{\mathcal{F}}$ presented in Section 3.2. More precisely, the automaton $\mathcal{R}_{\mathcal{F}}$ can be regarded as a quotient of $\mathcal{P}_{\mathcal{F}}$ given by the numbering $\eta: \text{Leaves}(\mathcal{Z}_{\mathcal{F}}) \rightarrow \{1, \dots, \text{mem-tree}(\mathcal{Z}_{\mathcal{F}})\}$. That is, to obtain $\mathcal{R}_{\mathcal{F}}$ we merge the states $\eta^{-1}(i)$ for $1 \leq i \leq \text{mem-tree}(\mathcal{Z}_{\mathcal{F}})$ and we keep all transitions. Moreover, the strategy for a resolver for $\mathcal{R}_{\mathcal{F}}$ as presented in the proof of Lemma 16 is exactly given by the deterministic automaton $\mathcal{P}_{\mathcal{F}}$. However, we note that in general a parity condition is not sufficient in $\mathcal{R}_{\mathcal{F}}$ to accept $\mathcal{L}_{\mathcal{F}}$ and we need to replace it by a Rabin one.

We observe that the GFG Rabin automaton from Figure 2 is obtained as a quotient of the deterministic parity automaton in Figure 3.

3.2.4 Simplifications and optimisations

Given an automaton $\mathcal{A} = (Q, \Sigma, Q_0, \Delta, \Gamma, \mathbb{W})$ we say that it has *duplicated edges* if there are some pair of states $q, q' \in Q$ and two different transitions between them labelled with the same input letter: $(q, a, \alpha, q'), (q, a, \beta, q') \in \Delta$.

As remarked previously, the construction we have presented provides an automaton potentially having duplicated edges, which can be seen as an undesirable property (even if some automata models such as the HOA format [2] allow them). We show next that we can always derive an equivalent automaton without duplicated edges. Intuitively, in the Rabin case, if we want to merge two transitions having as output letters α and β , we add a fresh letter $(\alpha\beta)$ to label the new transition. For each Rabin pair, this new letter will simulate the best of either α or β depending upon the situation.

► **Proposition 18** (Simplification of automata). *Let \mathcal{A} be a Muller (resp. Rabin) automaton presenting duplicated edges. There exists an equivalent Muller (resp. Rabin) automaton \mathcal{A}' on the same set of states without duplicated edges. Moreover, if \mathcal{A} is GFG, \mathcal{A}' can be chosen GFG. In the Rabin case, the number of Rabin pairs is also preserved.*

Proof. For the Rabin case, let \mathcal{A}' be an automaton that is otherwise as \mathcal{A} except that instead of the transitions Δ of \mathcal{A} it only has one a -transition $q \xrightarrow{a:x} q' \in \Delta'$ (with a fresh colour x per transition) per state-pair q, q' and letter $a \in \Sigma$. That is, $\Delta' = \{(q, a, x_j, q') : (q, a, y, q') \in \Delta \text{ for some } y\}$. The new Rabin condition $\{(G'_1, R'_1), \dots, (G'_r, R'_r)\}$ is defined as follows. For each transition $q \xrightarrow{a:x} q'$:

- $x \in G'_i$ if $q \xrightarrow{a:y} q' \in \Delta$ for some $y \in G_i$ (there is a green transition for the i^{th} pair)
- $x \in R'_i$ if for all $q \xrightarrow{a:y} q' \in \Delta$, $y \in R_i$ (there is no green or orange transition for the i^{th} pair).

We claim that $L(\mathcal{A}') = L(\mathcal{A})$. Indeed, if $u \in L(\mathcal{A})$, as witnessed by some run ρ and a Rabin pair (G_i, R_i) , then the corresponding run ρ' in \mathcal{A}' over u is also accepting with Rabin pair (G'_i, R'_i) : the transitions of $\text{Inf}(\rho) \cap G_i$ induce transitions of $\text{Inf}(\rho') \cap G'_i$ and the fact that $\text{Inf}(\rho) \cap R_i = \emptyset$ guarantees that $\text{Inf}(\rho') \cap R'_i = \emptyset$.

Conversely, if $u \in L(\mathcal{A}')$ as witnessed by a run ρ' and Rabin pair (G'_i, R'_i) , then there is an accepting run ρ over u in \mathcal{A} : such a run can be obtained by choosing for each transition $q \xrightarrow{a:x} q'$ of ρ' where $x \in G'_i$ a transition $q \xrightarrow{a:y} q' \in \Delta$ such that $y \in G_i$, which exists by definition of \mathcal{A}' , for each transition $q \xrightarrow{a:x} q'$ where $x \notin G_i \cup R_i$ a transition $q \xrightarrow{q,y} q' \in \Delta$ such that $y \notin R_i$, which also exists by definition of \mathcal{A}' , and for other transitions $q \xrightarrow{a:x} q'$ (that is, those for which $x \in R'_i$) an arbitrary transition $q \xrightarrow{a:y} q' \in \Delta$. Since ρ' is accepting, we have $\text{Inf}(\rho') \cap G'_i \neq \emptyset$ and $\text{Inf}(\rho) \cap R_i = \emptyset$, that is, ρ is also accepting.

For the Muller case, the argument is even simpler. As above, we consider \mathcal{A}' that is otherwise like \mathcal{A} except that instead of the transitions Δ of \mathcal{A} , it only has one a -transition $q \xrightarrow{a:x} q' \in \Delta'$ (with a fresh colour per transition) per state-pair q, q' and the accepting

condition is defined as follows. A set of transitions T is accepting if and only if for each $t = q \xrightarrow{a:x} q' \in T$ there is a non-empty set $S_t \subseteq \{q \xrightarrow{a:y} q' \in \Delta\}$ such that $\bigcup_{t \in T} S_t$ is accepting in \mathcal{A} . In other words, a set of transitions in \mathcal{A}' is accepting if for each transition we can choose a non-empty subset of the original transitions in \mathcal{A} that form an accepting run in \mathcal{A} .

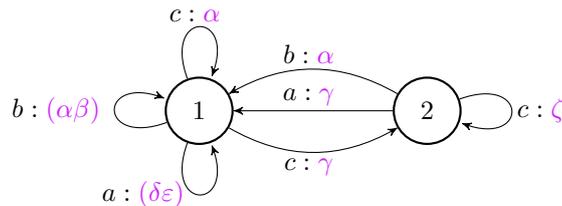
We claim that $L(\mathcal{A}') = L(\mathcal{A})$. Indeed if $u \in L(\mathcal{A})$, as witnessed by some run ρ , the run ρ' that visits the same sequence of states in \mathcal{A}' is accepting as witnessed by the transitions that occur infinitely often in ρ .

Conversely, assume $u \in L(\mathcal{A}')$, as witnessed by a run ρ' and a non-empty subset S_t for each transitions t that occurs infinitely often in ρ' such that $\bigcup_{t \in \text{Inf}(\rho')} S_t$ is accepting in \mathcal{A} . Then there is an accepting run ρ over u in \mathcal{A} that visits the same sequence of states as ρ' and chooses instead of a transition $t \in \text{Inf}(\rho')$ each transition in S_t infinitely often, and otherwise takes an arbitrary transition. The set of transitions ρ visits infinitely often is exactly $\bigcup_{t \in \text{Inf}(\rho')} S_t$, and is therefore accepting.

Finally, observe that in both cases, if \mathcal{A} is GFG, then the automaton \mathcal{A}' without duplicate edges is also GFG since \mathcal{A}' is obtained from \mathcal{A} by merging transitions. Indeed, the resolver r of \mathcal{A} induces a resolver r' for \mathcal{A}' by outputting the unique transition with the same letter and state-pair as r . By the same argument as above, the run induced by r' is accepting if and only if the run induced by r is. \blacktriangleleft

► **Example 19.** The GFG Rabin automaton from Figure 2 has duplicated transitions. In Figure 4 we present an equivalent GFG Rabin automaton without duplicates. For this, we have merged the self-loops in state 1 labelled with a and b respectively. We have added the output letters $(\alpha\beta)$ and $(\delta\varepsilon)$. The new Rabin pairs are given by:

$$\begin{aligned} G'_\beta &= \{\beta, (\alpha\beta)\}, & R'_\beta &= \{\alpha, \gamma, \varepsilon, \zeta\}, \\ G'_\gamma &= \{\gamma\}, & R'_\gamma &= \{\alpha, \beta, (\alpha\beta), \delta\}. \end{aligned}$$



■ **Figure 4** The simplified GFG Rabin automaton.

► **Remark 20 (Optimisation on the number of Rabin pairs).** An important parameter in the study of Rabin automata is the number of Rabin pairs used. The automaton presented in this Section uses a number of Rabin pairs that equals the number of round nodes in the Zielonka tree. This can be improved by using only the round nodes in the Zielonka directed acyclic graph (obtained from the tree by merging nodes with the same labels). However, even this latter option is not always optimal and we conjecture that minimising the number of Rabin pairs without increasing the size of the automaton is NP-complete.

4 GFG Rabin automata recognising Muller conditions can be exponentially more succinct than deterministic ones in number of states

On his PhD Thesis [23, 24], Kopczyński raised the question of whether the general and the chromatic memory requirements of winning conditions always coincide. By Theorem 5 and [10, Theorem 28], in the case of Muller conditions, this question is equivalent to the following:

Is there a Muller language L such that minimal GFG Rabin automata recognising L are strictly smaller than deterministic Rabin automata for L ?

In [10] this question is answered positively. It is shown that for every $n \in \mathbb{N}$ there is a Muller language L_n over an alphabet Γ_n such that a minimal GFG Rabin automaton for it has size 2, but a minimal deterministic Rabin automaton for it has size n . However, the size of the alphabet Γ_n in that example also has size n . A natural question is whether GFG Rabin automata recognising Muller conditions can be exponentially more succinct than deterministic ones, when also taking into account the alphabet size. This is indeed the case:

► **Theorem 21.** *There exists a constant $\alpha > 1$, a sequence of natural numbers $n_1 < n_2 < n_3 \dots$ and a sequence of Muller conditions \mathcal{F}_{n_i} over $\Gamma_{n_i} = \{1, \dots, n_i\}$ such that*

- *a minimal GFG Rabin automaton for $\mathcal{L}_{\mathcal{F}_{n_i}}$ has size $\lfloor n_i/2 \rfloor$,*
- *a minimal deterministic Rabin automaton for $\mathcal{L}_{\mathcal{F}_{n_i}}$ has size at least α^{n_i} .*

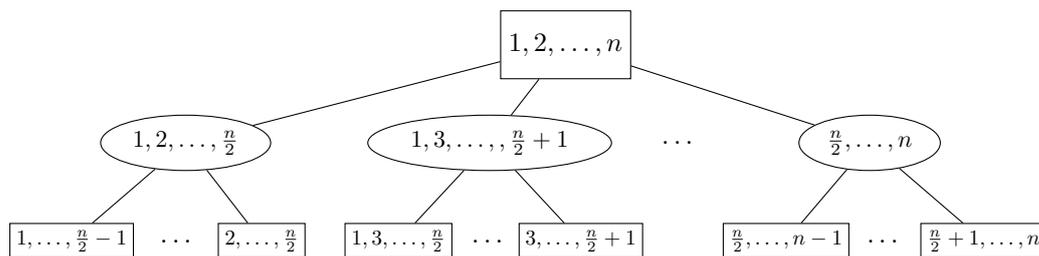
A lower bound for such a constant is $\alpha = 1.116$.

We devote the rest of this Section to proving Theorem 21. In brief, the Muller conditions in question require half the colours to be seen infinitely often. The construction of the small GFG Rabin automaton follows from constructing the Zielonka tree of the condition. For the lower bound on the deterministic Rabin automaton, we reduce the problem to finding a lower bound on the chromatic number of a certain graph, which we finally show to be sufficiently large for a family of our Muller conditions.

Let $n \in \mathbb{N}$. We define the following Muller condition over $\Gamma_n = \{1, \dots, n\}$:

$$\mathcal{F}_n = \{C \subseteq \Gamma_n : |C| = \lfloor n/2 \rfloor\}.$$

The Zielonka tree of \mathcal{F}_n is depicted in Figure 5 (for n even).



■ **Figure 5** Zielonka tree $\mathcal{Z}_{\mathcal{F}_n}$ for $\mathcal{F}_n = \{C \subseteq \Gamma_n : |C| = \lfloor n/2 \rfloor\}$.

Each round node in $\mathcal{Z}_{\mathcal{F}_n}$ has exactly $\lfloor n/2 \rfloor$ children, and therefore $\text{mem-tree}(\mathcal{Z}_{\mathcal{F}_n}) = \lfloor n/2 \rfloor$. Thus, a minimal GFG Rabin automaton recognising $\mathcal{L}_{\mathcal{F}_n}$ has size $\lfloor n/2 \rfloor$ (by Proposition 4 and Theorem 5).

We now give a lower bound for deterministic Rabin automata recognising $\mathcal{L}_{\mathcal{F}_n}$. Our main tool will be Lemma 22, which uses the notion of *cycles*. A *cycle* of an automaton \mathcal{A} is a set of transitions forming a closed path (not necessarily simple). The set of states of a cycle consists

117:16 On the Size of Good-For-Games Rabin Automata and the Memory in Muller Games

of those states that are the source of some transition in it. If \mathcal{A} is a Rabin automaton, we say that a cycle is accepting (resp. rejecting) if the colours c_1, \dots, c_k appearing in its transitions form a word $(c_1 c_2 \dots c_k)^\omega$ that satisfies (resp. does not satisfy) the Rabin condition.

► **Lemma 22** ([11]). *Let \mathcal{A} be a deterministic Rabin automaton. If ℓ_1 and ℓ_2 are two rejecting cycles in \mathcal{A} with some state in common, then the union of ℓ_1 and ℓ_2 is also a rejecting cycle.*

For the following, let \mathcal{A} be a deterministic Rabin automaton recognising $\mathcal{L}_{\mathcal{F}_n}$. For each subset of letters $C \subseteq \Gamma_n$, we define a *final C -Strongly Connected Component* (C -FSCC for short) as a set of states P of \mathcal{A} such that:

- For every pair of states $p, q \in P$, there is a word $w \in C^*$ labelling a path from p to q .
- For every $p \in P$ and $w \in C^*$, the run over w starting in p remains in P .

It is easy to see that for every $C \subseteq \Gamma_n$ there exists some C -FSCC in \mathcal{A} .

► **Lemma 23.** *Let $C_1, C_2 \subseteq \Gamma_n$ such that $|C_i| < \lfloor n/2 \rfloor$, for $i = 1, 2$ and such that $|C_1 \cup C_2| = \lfloor n/2 \rfloor$. If P_1 and P_2 are two C_1 and C_2 -FSCC, respectively, then $P_1 \cap P_2 = \emptyset$.*

Proof. For $i = 1, 2$, let ℓ_i be a cycle visiting all states of P_i and reading exactly the set of letters C_i . By definition of $\mathcal{L}_{\mathcal{F}_n}$, ℓ_i is a rejecting cycle. If P_1 and P_2 had some state in common, we could take the union of the cycles ℓ_1 and ℓ_2 , producing an accepting cycle, which is impossible by Lemma 22. ◀

We associate the following (undirected) graph $\mathcal{G}_{\mathcal{F}_n} = (V_{\mathcal{F}_n}, E_{\mathcal{F}_n})$ to the Muller condition \mathcal{F}_n :

- $V_{\mathcal{F}_n} = \mathcal{P}(\Gamma_n)$.
- There is an edge between two subsets $C_1, C_2 \subseteq \Gamma_n$ if and only if $|C_i| < \lfloor n/2 \rfloor$, for $i = 1, 2$, and $|C_1 \cup C_2| = \lfloor n/2 \rfloor$.

That is, we connect two vertices if they correspond to rejecting sets but taking their union we obtain an accepting set.

We reduce finding lower bounds in the size of deterministic Rabin automata to giving lower bounds for the chromatic number of $\mathcal{G}_{\mathcal{F}_n}$. A *colouring* of an undirected graph $G = (V, E \subseteq V \times V)$ is a mapping $c : V \rightarrow \Lambda$ such that $c(v) = c(v') \Rightarrow (v, v') \notin E$ for every pair of nodes $v, v' \in V$. We say that such a colouring has size $|\Lambda|$. The *chromatic number* of G is the minimal number k such that G has a colouring of size k . We denote it $\chi(G)$.

► **Lemma 24.** *A lower bound for the size of a minimal deterministic Rabin automaton recognising $\mathcal{L}_{\mathcal{F}_n}$ is given by $\chi(\mathcal{G}_{\mathcal{F}_n})$.*

Proof. Let \mathcal{A} be a deterministic Rabin automaton recognising $\mathcal{L}_{\mathcal{F}_n}$ with states Q . We define a colouring c of $\mathcal{G}_{\mathcal{F}_n}$ using Q as colours. For each $C \subseteq \Gamma_n$, we let P_C be a C -FSCC and we pick a state $q_C \in P_C$. We define $c(C) = q_C$. We prove that this is a correct colouring. Suppose that C_1 and C_2 are two vertices in $\mathcal{G}_{\mathcal{F}_n}$ connected by some edge, that is, $|C_i| < \lfloor n/2 \rfloor$ and $|C_1 \cup C_2| = \lfloor n/2 \rfloor$. If $q_{C_1} = q_{C_2}$, it means that $P_{C_1} \cap P_{C_2} \neq \emptyset$, contradicting Lemma 23. ◀

► **Remark 25.** The definition of $\mathcal{G}_{\mathcal{F}_n}$ is not specific to this Muller condition. It can be defined analogously for any other Muller condition and Lemma 24 holds by the same argument.

► **Proposition 26.** *There exists a constant $\alpha > 1$ and a sequence of natural numbers $n_1 < n_2 < n_3 \dots$ such that $\alpha^{n_i} \leq \chi(\mathcal{G}_{\mathcal{F}_{n_i}})$.*

In order to prove Proposition 26 we introduce some further graph-theoretic notions. Let $\mathcal{G} = (V, E)$ be an undirected graph. An *independent set* of \mathcal{G} is a set $S \subseteq V$ such that $(v, v') \notin E$ for every pair of vertices $v, v' \in S$.

► **Lemma 27.** *Let $R \subseteq V$, and let $\mathcal{G}_R = (R, E|_{R \times R})$ be the subgraph of \mathcal{G} induced by R . Then, $\chi(\mathcal{G}) \geq \chi(\mathcal{G}_R)$.*

► **Lemma 28.** *Let m be an upper bound on the size of the independent sets in \mathcal{G} . Then*

$$\chi(\mathcal{G}) \geq \frac{|V|}{m}.$$

Proof. Let $c: V \rightarrow \Lambda$ be a colouring of \mathcal{G} with $|\Lambda| = \chi(\mathcal{G})$. Then, by definition of a colouring, for each $x \in \Lambda$, $c^{-1}(x)$ is an independent set in \mathcal{G} , so $|c^{-1}(x)| \leq m$. Also, $V = \bigcup_{x \in \Lambda} c^{-1}(x)$, so

$$|V| = \sum_{x \in \Lambda} |c^{-1}(x)| \leq \chi(\mathcal{G}) \cdot m. \quad \blacktriangleleft$$

We will find a subgraph of $\mathcal{G}_{\mathcal{F}_n}$ for which we can provide an upper bound on the size of its independent sets. The upper bound is provided by the following theorem (adapted from [35, Theorem 15]).

► **Theorem 29** ([35], Theorem 15). *Let $n > k > 2t$ such that $k - t$ is a prime number. Suppose that \mathcal{B} is a family of subsets of size k of Γ_n such that $|A \cap B| \neq t$ for any pair of subsets $A, B \in \mathcal{B}$. Then,*

$$|\mathcal{B}| \leq \binom{n}{k-t-1}.$$

We conclude this section with the proof of Proposition 26.

Proof of Proposition 26. Let p be a prime number and let $n = 5p$. We will study the subgraph of $\mathcal{G}_{\mathcal{F}_n}$ formed by the subsets of size exactly $k = \lfloor 3n/10 \rfloor$. We denote this subgraph by $H_{n,k}$. Two subsets $A, B \subseteq \Gamma_n$ of size k verify that $|A \cup B| = \lfloor n/2 \rfloor$ if and only if $|A \cap B| = \lfloor n/10 \rfloor$. We set $t = \lfloor n/10 \rfloor$. We get $k - t = p$ so we can apply Theorem 29 and we obtain that any independent set in $H_{n,k}$ has size at most $\binom{n}{\frac{1}{5}n-1}$. By Lemma 28, $\chi(H_{n,k}) \geq \binom{n}{\lfloor \frac{3}{10}n \rfloor} / \binom{n}{\frac{1}{5}n-1}$. By Lemma 27 we know that this lower bound also holds for $\mathcal{G}_{\mathcal{F}_n}$. Using Stirling's approximation we obtain that

$$\chi(\mathcal{G}_{\mathcal{F}_n}) \geq \binom{n}{\lfloor \frac{3}{10}n \rfloor} / \binom{n}{\frac{1}{5}n-1} = \Omega \left(\left(\frac{(1/5)^{1/5} (4/5)^{4/5}}{(3/10)^{3/10} (7/10)^{7/10}} \right)^n \right) = \Omega(1.116^n).$$

To conclude, we take an enumeration of prime numbers, $p_1 < p_2 < \dots$ and we set $n_i = 5p_i$. ◀

► **Remark 30** (Choices of k and t). The choice of $k = \lfloor 3n/10 \rfloor$ and $t = \lfloor n/10 \rfloor$ in the previous proof might appear quite enigmatic. We try to explain them now.

We want to find a number k such that there is not a big family of sets $\{A_i \subseteq \Gamma_n\}$ of size $|A_i| = k$ such that $|A_i \cup A_j| \neq n/2$, and express this fact in terms of $|A_i \cap A_j|$. Since $|A \cup B| = 2k - |A \cap B|$, if we define $t = 2k - n/2$, then $|A \cup B| \neq n/2$ if and only if $|A \cap B| \neq t$, so the value of t will be completely determined by the choice of k . Our objective is to minimise the upper bound given in Theorem 29 (what we do by minimising $k - t$) while making sure that the hypothesis $k > 2t$ is verified. In the boundary of this condition ($k = 2t$) we obtain $k = n/3$, so we express our choices as $k = (1/3 - \varepsilon)n$ and $t = (1/6 - 2\varepsilon)n$. Moreover, $k - t = (1/6 + \varepsilon)n$ has to be a prime number (for infinite n). If $1/6 + \varepsilon = 1/q$ for some $q \in \mathbb{N}$, we would succeed by considering n of the form $q \cdot p$, for p a prime number. We will therefore take $\varepsilon = \frac{6-q}{6q}$, for some $q, 1 \leq q \leq 5$. With the optimal choice, $q = 5$, we obtain $k = 3n/10$, $t = n/10$ and $k - t = n/5$. Since k and t will not be integers for n of the form $5p$ (p a prime number) we are forced to take the integer part in the proof of Proposition 26.

5 Conclusion

We believe that our work is a significant advance in the understanding of the memory needed for winning ω -regular games. In combination with the literature, we can describe the current understanding of Muller languages as follows:

- The least memory necessary for winning all won L -games equals the least number of states of a GFG Rabin automaton for L .
- Computing this quantity can be done in polynomial time for L given by its Zielonka tree.
- The least chromatic memory necessary for winning all won L -games equals the least number of states of a deterministic Rabin automaton for L .
- Computing this quantity is NP-complete for L given by its Zielonka tree.
- The chromatic memory can be arbitrarily larger than the memory. It can be exponential in the size the alphabet, even while the memory remains linear.

This description shows that GFG automata play a key, and, up till now, unexplored role in understanding the complexity of Muller languages and that this role is – in some respect – even more important than that of the more classical deterministic automata.

References

- 1 Bader Abu Radi and Orna Kupferman. Minimizing GFG transition-based automata. In *ICALP*, volume 132, pages 100:1–100:16, 2019. doi:10.4230/LIPIcs.ICALP.2019.100.
- 2 Tomáš Babiak, František Blahoudek, Alexandre Duret-Lutz, Joachim Klein, Jan Křetínský, David Müller, David Parker, and Jan Strejček. The Hanoi omega-automata format. In *CAV*, pages 479–486, 2015.
- 3 Marc Bagnol and Denis Kuperberg. Büchi good-for-games automata are efficiently recognizable. In *FSTTCS*, page 16, 2018.
- 4 Udi Boker, Denis Kuperberg, Karoliina Lehtinen, and Michal Skrzypczak. On succinctness and recognisability of alternating good-for-games automata. *CoRR*, abs/2002.07278, 2020. arXiv:2002.07278.
- 5 Udi Boker and Karoliina Lehtinen. History determinism vs. good for gameness in quantitative automata, 2021. arXiv:2110.14238.
- 6 Udi Boker and Karoliina Lehtinen. Token games and history-deterministic quantitative-automata, 2022. To appear in proceedings of FoSSaCS’22. arXiv:2110.14308.
- 7 Patricia Bouyer, Mickael Randour, and Pierre Vandenhove. Characterizing omega-regularity through finite-memory determinacy of games on infinite graphs. *CoRR*, abs/2110.01276, 2021. arXiv:2110.01276.
- 8 Patricia Bouyer, Stéphane Le Roux, Youssef Oualhadj, Mickael Randour, and Pierre Vandenhove. Games where you can play optimally with arena-independent finite memory. In *CONCUR*, volume 171, pages 24:1–24:22, 2020. doi:10.4230/LIPIcs.CONCUR.2020.24.
- 9 J. Richard Büchi. Using determinacy of games to eliminate quantifiers. In *FCT*, volume 56 of *Lecture Notes in Computer Science*, pages 367–378. Springer, 1977. doi:10.1007/3-540-08442-8_104.
- 10 Antonio Casares. On the minimisation of transition-based Rabin automata and the chromatic memory requirements of Muller conditions. In *CSL*, volume 216, pages 12:1–12:17, 2022. doi:10.4230/LIPIcs.CSL.2022.12.
- 11 Antonio Casares, Thomas Colcombet, and Nathanaël Fijalkow. Optimal transformations of games and automata using Muller conditions. In *ICALP*, volume 198, pages 123:1–123:14, 2021. doi:10.4230/LIPIcs.ICALP.2021.123.
- 12 Thomas Colcombet. The theory of stabilisation monoids and regular cost functions. In *ICALP*, pages 139–150, 2009. doi:10.1007/978-3-642-02930-1_12.

- 13 Thomas Colcombet and Damian Niwiński. On the positional determinacy of edge-labeled games. *Theoretical Computer Science*, 352(1):190–196, 2006. doi:10.1016/j.tcs.2005.10.046.
- 14 Thomas Colcombet and Konrad Zdanowski. A tight lower bound for determinization of transition labeled Büchi automata. In *ICALP*, pages 151–162, 2009. doi:10.1007/978-3-642-02930-1_13.
- 15 Stefan Dziembowski, Marcin Jurdziński, and Igor Walukiewicz. How much memory is needed to win infinite games? In *LICS*, pages 99–110, 1997. doi:10.1109/LICS.1997.614939.
- 16 E. Allen Emerson and Charanjit S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *FOCS*, pages 368–377, 1991. doi:10.1109/SFCS.1991.185392.
- 17 Hugo Gimbert and Wiesław Zielonka. Games where you can play optimally without any memory. In *CONCUR*, volume 3653, pages 428–442, 2005. doi:10.1007/11539452_33.
- 18 Shibashis Guha, Ismaël Jecker, Karoliina Lehtinen, and Martin Zimmermann. A Bit of Nondeterminism Makes Pushdown Automata Expressive and Succinct. In *MFCS*, volume 202, pages 53:1–53:20, 2021. doi:10.4230/LIPIcs.MFCS.2021.53.
- 19 Yuri Gurevich and Leo Harrington. Trees, automata, and games. In *STOC*, pages 60–65, 1982. doi:10.1145/800070.802177.
- 20 Thomas A. Henzinger and Nir Piterman. Solving games without determinization. In *Computer Science Logic*, pages 395–410, 2006.
- 21 Florian Horn. Random fruits on the zielonka tree. In *STACS*, volume 3, pages 541–552, 2009. doi:10.4230/LIPIcs.STACS.2009.1848.
- 22 Nils Klarlund. Progress measures, immediate determinacy, and a subset construction for tree automata. *Annals of Pure and Applied Logic*, 69(2):243–268, 1994. doi:10.1016/0168-0072(94)90086-8.
- 23 Eryk Kopczyński. Half-positional determinacy of infinite games. In *ICALP*, pages 336–347, 2006. doi:10.1007/11787006_29.
- 24 Eryk Kopczyński. *Half-positional Determinacy of Infite Games*. Phd thesis, University of Warsaw, 2008.
- 25 Alexander Kozachinskiy. State complexity of chromatic memory in infinite-duration games. *CoRR*, abs/2201.09297, 2022. arXiv:2201.09297.
- 26 Jan Křetínský, Tobias Meggendorfer, Clara Waldmann, and Maximilian Weininger. Index appearance record for transforming Rabin automata into parity automata. In *TACAS*, pages 443–460, 2017. doi:10.1007/978-3-662-54577-5_26.
- 27 Denis Kuperberg and Michał Skrzypczak. On determinisation of good-for-games automata. In *ICALP*, pages 299–310, 2015. doi:10.1007/978-3-662-47666-6_24.
- 28 Karoliina Lehtinen and Martin Zimmermann. Good-for-games ω -pushdown automata. In *LICS*, pages 689–702, 2020. doi:10.1145/3373718.3394737.
- 29 Christof Löding and Anton Pirogov. Determinization of Büchi automata: Unifying the approaches of Safra and Muller-Schupp. In *ICALP*, pages 120:1–120:13, 2019. doi:10.4230/LIPIcs.ICALP.2019.120.
- 30 Michael Luttenberger, Philipp J. Meyer, and Salomon Sickert. Practical synthesis of reactive systems from LTL specifications via parity games. *Acta Informatica*, pages 3–36, 2020. doi:10.1007/s00236-019-00349-3.
- 31 Christof Löding. Optimal bounds for transformations of ω -automata. In *FSTTCS*, pages 97–109, 1999. doi:10.1007/3-540-46691-6_8.
- 32 Robert McNaughton. Testing and generating infinite sequences by a finite automaton. *Information and control*, 9:521–530, 1966.
- 33 Philipp Meyer and Salomon Sickert. On the optimal and practical conversion of Emerson-Lei automata into parity automata. *Personal Communication*, 2021.
- 34 Max Michel. Complementation is more difficult with automata on infinite words. *CNET, Paris*, 15, 1988.
- 35 Dhruv Mubayi and Vojtech Rödl. Specified intersections. *Transactions of the American Mathematical Society*, 366(1):491–504, 2014. URL: <http://www.jstor.org/stable/23813142>.

- 36 David E. Muller and Paul E. Schupp. Simulating alternating tree automata by nondeterministic automata: New results and new proofs of the theorems of Rabin, McNaughton and Safra. *Theor. Comput. Sci.*, 141(1–2):69–107, 1995. doi:10.1016/0304-3975(94)00214-4.
- 37 Nir Piterman. From nondeterministic Büchi and Streett automata to deterministic parity automata. In *LICS*, pages 255–264, 2006. doi:10.1109/LICS.2006.28.
- 38 Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190, 1989. doi:10.1145/75277.75293.
- 39 Schmuel Safra. On the complexity of ω -automata. In *FOCS*, pages 319–327, 1988. doi:10.1109/SFCS.1988.21948.
- 40 Sven Schewe. Tighter bounds for the determinisation of Büchi automata. In *FoSSaCS*, pages 167–181, 2009. doi:10.1007/978-3-642-00596-1_13.
- 41 Sven Schewe. Beyond hyper-minimisation—minimising DBAs and DPAs is NP-complete. In *FSTTCS*, volume 8, pages 400–411, 2010. doi:10.4230/LIPIcs.FSTTCS.2010.400.
- 42 Sven Schewe. Minimising Good-For-Games automata is NP-complete. In *FSTTCS*, volume 182, pages 56:1–56:13, 2020. doi:10.4230/LIPIcs.FSTTCS.2020.56.
- 43 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.