

Rabin-Mostowski index problem: a step beyond deterministic automata

Alessandro Facchini, Filip Murlak, and Michał Skrzypczak
University of Warsaw
{facchini, fmurlak, mskrzypczak}@mimuw.edu.pl

Abstract—For a given regular language of infinite trees, one can ask about the minimal number of priorities needed to recognise this language with a non-deterministic or alternating parity automaton. These questions are known as, respectively, the non-deterministic and the alternating Rabin-Mostowski index problems. Whether they can be answered effectively is a long-standing open problem, solved so far only for languages recognisable by deterministic automata (the alternating variant trivialises).

We investigate a wider class of regular languages, recognisable by so-called game automata, which can be seen as the closure of deterministic ones under complementation and composition. Game automata are known to recognise languages arbitrarily high in the alternating Rabin-Mostowski index hierarchy, i.e., the alternating index problem does not trivialise any more.

Our main contribution is that both index problems are decidable for languages recognisable by game automata. Additionally, we show that it is decidable whether a given regular language can be recognised by a game automaton.

I. INTRODUCTION

Finite state automata running over infinite words and infinite binary trees lie at the core of the seminal works of Büchi [1] and Rabin [2]. Known to be equivalent to the monadic second-order (MSO) logic and the modal μ -calculus on both classes of structures, they subsume all standard linear and branching temporal logics. Because of these properties, they constitute fundamental tools in the theory of verification and model-checking, where the model-checking problem is reduced to the non-emptiness problem for automata: a given formula is translated into an automaton recognising its models. From this perspective, a natural question is, which parameter in the definition of an automaton reflects the complexity of the language recognised by it. A naïve approach is to look at the number of states; a more meaningful one is to consider the infinitary behaviour of the automata, captured by the complexity of its acceptance condition.

Out of different acceptance conditions proposed for tree automata, Büchi, Muller, Rabin, Streett, and parity [3], [4], the last one has proved to be the most appropriate, as it enabled unveiling the subtle correspondences between games, automata, and the modal μ -calculus [5], [6]. In a parity automaton, each state is assigned a natural number, called its priority. A sequence of states is said to be accepting if the lowest priority occurring infinitely often is even (min-parity condition). The pair (i, j) consisting of the minimal priority i and the maximal priority j in a given automaton is called its Rabin-Mostowski index. The index of a language is the minimal index of a recognising automaton. Practical importance of this parameter comes from the fact that the best

known algorithms deciding emptiness of (non-deterministic) automata are exponential in the number of priorities.

Given a regular tree language, what is the minimal range of priorities needed to recognise it? The answer to this question depends on which mode of computation is used, i.e., whether the automata are deterministic, non-deterministic, and alternating. While deterministic automata are weaker, non-deterministic and alternating parity automata recognise all regular tree languages. Still, alternating automata often need less priorities than non-deterministic ones. Thus, for each of these three classes there is the respective index problem.

C Index Problem: Given i, j and a regular language L , decide if L is recognised by an automaton in \mathcal{C} of index (i, j) .

The solution of this problem for the most important cases—when \mathcal{C} is the class of non-deterministic or alternating automata—seems still far away. The results of [7]–[9], later extended in [10], show that it is decidable if a given regular tree language can be recognised by a combination of reachability and safety conditions (which corresponds to the Boolean combination of open sets). It is also known that the non-deterministic (min-parity) index problem is decidable for $(i, j) = (1, 2)$, and for $(i, j) = (0, 1)$ if the input language is given by an alternating automaton of index $(1, 2)$ [11]–[13]. The non-deterministic index problem has been reduced to the uniform universality problem for so-called distance-parity automata [14], but decidability of latter problem remains open.

The index problems become easier when we restrict the input to languages recognised by deterministic automata. This is mostly due to the fact that in a deterministic automaton, each sub-automaton can be replaced with any automaton recognising a language of the same index, without influencing the index of the whole language. This observation has been essential in providing a full characterisation of the combinatorial structure of a language L in terms of certain patterns in a deterministic automaton recognising L . This so-called pattern method [15] has been successfully used for solving all three index problems for languages recognised by deterministic automata:

Theorem I.1. For languages recognised by deterministic automata the following problems are decidable:

- 1) the deterministic index problem [16];
- 2) the non-deterministic index problem [17], [18]; and
- 3) the alternating index problem [19].

The pattern method cannot be applied to non-deterministic nor alternating automata; the reason is that both these classes are closed under union and union is not an operation that preserves the index of languages. However, it turns out that

if we avoid closure under union, we can extend the pattern method well beyond deterministic automata, to so-called game automata.

Game automata can be seen as a combination of deterministic and co-deterministic ones. They were introduced in [20] as the largest subclass of alternating tree automata extending the deterministic ones, closed under complementation and composition, and for which the latter operation preserves natural equivalence relations on recognised languages, like the topological equivalence, or having the same index. As game automata recognise the game languages $W_{i,j}$ [21], the alternating index problem does not trivialise, unlike for deterministic deterministic automata.

Here, we extend Theorem I.1(2),(3) and prove the following.

Theorem I.2. *For languages recognised by game automata the following problems are decidable:*

- 1) *the non-deterministic index problem,*
- 2) *the alternating index problem.*

Decidability of the non-deterministic index problem for languages recognised by game automata is obtained via an easy reduction to the non-deterministic index problem for deterministic automata (Section III).

The alternating index problem is solved by providing a recursive procedure computing the alternating index of the language recognised by a given game automaton (Section IV).

Finally, we give an effective characterisation of languages recognised by game automata, within the class of all regular languages (Section V). As the characterisation effectively yields an equivalent game automaton, we obtain procedures computing the alternating and non-deterministic index for a given alternating automaton equivalent to some game automaton.

II. PRELIMINARIES

To simplify the presentation of inductive arguments, all our definitions allow partial objects: trees have leaves, automata have exits (where computation stops) and games have final positions (where the play stops and no player wins). The definitions become standard when restricted to *total* objects: trees without leaves, automata without exits, and games without final positions. We also do not distinguish the initial state of an automaton but treat it as an additional parameter for the recognised language.

A. Trees

For a function f we write $\text{dom}(f)$ for the domain of f and $\text{rg}(f)$ for the range of f . For a finite alphabet A , we denote by PTr_A the set of partial trees over A , i.e., functions $t: \text{dom}(t) \rightarrow A$ from a prefix-closed subset $\text{dom}(t) \subseteq \{L, R\}^*$ to A . By Tr_A we denote the set of *total* trees, i.e., trees t such that $\text{dom}(t) = \{L, R\}^*$. For a direction $d \in \{L, R\}$ by \bar{d} we denote the opposite direction. For $v \in \text{dom}(t)$, $t \upharpoonright_v$ denotes the subtree of t rooted at v . The sequences $u, v \in \{L, R\}^*$ are naturally ordered by the prefix relation: $u \preceq v$ if u is a prefix of v .

A tree that is not total contains *holes*. A *hole* of a tree t is a minimal sequence $h \in \{L, R\}^*$ that does not belong to $\text{dom}(t)$. By $\text{holes}(t) \subseteq \{L, R\}^*$ we denote the set of holes of a tree t . If h is a hole of $t \in \text{PTr}_A$, for $s \in \text{PTr}_A$ we define the partial tree $t[h := s]$ obtained by putting the root of s into the hole h of t .

B. Games

A *parity game* G is a tuple $\langle V = V_\exists \cup V_\forall, v_I, F, E, \Omega \rangle$, where

- V is a countable *arena*;
- $V_\exists, V_\forall \subseteq V$ are positions of the game *belonging*, respectively, to player \exists and player \forall , $V_\exists \cap V_\forall = \emptyset$;
- $v_I \in V$ is the initial position of the game;
- F is a countable set of *final positions*, $F \cap V = \emptyset$;
- $E \subseteq V \times (V \cup F)$ is the transition relation;
- $\Omega: V \rightarrow \{i, \dots, j\}$ is a *priority function*.

We assume that all parity games are finitely branching (for each $v \in V$ there are only finitely many $u \in V \cup F$ such that $(v, u) \in E$), and that there are no dead-ends (for each $v \in V$ there is at least one $u \in V \cup F$ such that $(v, u) \in E$). We also implicitly assume that sets V, F are restricted to elements that are accessible by E from v_I .

A *play* in a parity game G is a finite or infinite sequence π of positions starting from v_I . If π is finite then $\pi = v_I v_1 \dots v_n$ and v_n is required to be a final position (that is $v_n \in F$). In that case v_n is called the *final position of* π . An infinite play π is *winning* for \exists if $\liminf_{n \rightarrow \infty} \Omega(\pi(n))$ is even. Otherwise π is winning for \forall .

A (positional) *strategy* σ for a player $P \in \{\exists, \forall\}$ in a game G is defined as usual, as a function assigning to every P 's position $v \in V_P$ the chosen successor $\sigma(v) \in V \cup F$ such that $(v, \sigma(v)) \in E$. A play π *conforms to* σ if whenever π visits a vertex $v \in V_P$ then the next position of π is $\sigma(v)$. We say that a strategy σ is *winning* for P if every infinite play conforming to σ is winning for P . For a winning strategy σ we define the *guarantee of* σ as the set of all final positions that can be reached in plays conforming to σ . In each parity game one of the players has a (positional) winning strategy [3], [6].

C. Automata

For the purpose of the inductive argument we incorporate into the definition of automata a finite set of *exits*. Therefore, an alternating automaton \mathcal{A} is defined as a tuple $\langle A, Q, F, \delta, \Omega \rangle$, where A is a finite alphabet, Q is a finite set of states, F is a finite set of exits disjoint from Q , $\Omega: Q \rightarrow \mathbb{N}$ is a function assigning to each state of \mathcal{A} its priority, and δ assigns to each pair $(q, a) \in Q \times A$ the transition $b = \delta(q, a)$ built using the grammar

$$b ::= \top \mid \perp \mid (q, d) \mid (f, d) \mid b \vee b \mid b \wedge b$$

for $q \in Q$, $f \in F$, and $d \in \{L, R\}$.

For an alternating automaton \mathcal{A} , a state $q_I \in Q$, and a partial tree $t \in \text{PTr}_A$ we define the game $\mathbf{G}(\mathcal{A}, t, q_I)$ as follows:

- $V = \text{dom}(t) \times (S_\delta \cup Q)$, where S_δ is the set of all subformulae of formulae in $\text{rg}(\delta)$; all positions of the

form $(v, b_1 \vee b_2)$ belong to \exists and the remaining ones to \forall ;

- $F = (\text{holes}(t) \times (Q \cup F)) \cup \text{dom}(t) \times F$;
- $v_I = (\epsilon, q_I)$;
- E contains the following pairs (for all $v \in \text{dom}(t)$):
 - $((v, b), (v, b))$ for $b \in \{\top, \perp\}$,
 - $((v, b), (v, b_i))$ for $b = b_1 \wedge b_2$ or $b = b_1 \vee b_2$,
 - $((v, (q, d)), (vd, q))$ for $d \in \{L, R\}$, $q \in Q \cup F$,
 - $((v, q), (v, \delta(q, t(v))))$ for $q \in Q$;
- $\Omega(v, \top) = 0$, $\Omega(v, \perp) = 1$, $\Omega(v, q) = \Omega(q)$ for $q \in Q$, $v \in \text{dom}(t)$, and for other positions Ω is $\max(\text{rg}(\Omega))$.

An automaton \mathcal{A} is *total* if $F = \emptyset$. A total automaton \mathcal{A} *accepts* a total tree $t \in \text{Tr}_A$ from $q_I \in Q$ if \exists has a winning strategy in $\mathbf{G}(\mathcal{A}, t, q_I)$. By $\text{L}(\mathcal{A}, q_I)$ we denote the set of total trees accepted by a total automaton \mathcal{A} from a state q_I . A total automaton \mathcal{A} *recognises* a language $L \subseteq \text{Tr}_A$ if $\text{L}(\mathcal{A}, q_I) = L$ for some $q_I \in Q$. A state $q \in Q$ is *non-trivial* if $\emptyset \subsetneq \text{L}(\mathcal{A}, q) \subsetneq \text{Tr}_A$. Without loss of generality we implicitly assume that all total automata have only non-trivial states.

The *(Rabin-Mostowski) index* of an automaton \mathcal{A} is the pair (i, j) where i is the minimal and j is the maximal priority of the states of \mathcal{A} . In that case \mathcal{A} is called an (i, j) -automaton.

An automaton \mathcal{A} is *deterministic* if all its transitions are deterministic, i.e., of the form $\top, \perp, (q_d, d)$, or $(q_L, L) \wedge (q_R, R)$, for $d \in \{L, R\}$. Similarly, \mathcal{A} is *non-deterministic* if its transitions are disjunctions of deterministic transitions.

D. Compositionality

Let $\mathcal{A} = \langle A, Q, F, \delta, \Omega \rangle$ be an alternating automaton and $Q' \subseteq Q$ be a set of states. By $\mathcal{A}|_{Q'}$ we denote the *restriction of \mathcal{A} to Q'* obtained by replacing the set of states by Q' , the set of exits by $F \cup (Q - Q')$, the ranking function by $\Omega|_{Q'}$, and the transition function by $\delta|_{Q' \times A}$. We say that \mathcal{B} is a *sub-automaton of \mathcal{A}* (denoted $\mathcal{B} \subseteq \mathcal{A}$) if $\mathcal{B} = \mathcal{A}|_{Q^B}$.

For automata \mathcal{A}, \mathcal{B} over an alphabet A with $Q^A \cap Q^B = \emptyset$, we define the composition $\mathcal{A} \cdot \mathcal{B}$ as the automaton over A , with states $Q = Q^A \cup Q^B$, exits $(F^A \cup F^B) - Q$, transitions $\delta^A \cup \delta^B$, and ranks $\Omega^A \cup \Omega^B$. Note here that some exits of \mathcal{A} may be states of \mathcal{B} and *vice versa*.

Fact II.1. *If \mathcal{A} is an alternating automaton and $Q = Q_1 \cup Q_2$ is a partition of the states of \mathcal{A} then $\mathcal{A}|_{Q_1} \cdot \mathcal{A}|_{Q_2} = \mathcal{A}$.*

E. Game automata

In this work we study the so-called *game automata*, i.e., alternating automata with transitions of the following forms:

$$\top, \perp, (q_d, d), (q_L, L) \vee (q_R, R), (q_L, L) \wedge (q_R, R)$$

for $d \in \{L, R\}$ and $q_L, q_R \in Q \cup F$.

The main similarity between game automata and deterministic automata is that their acceptance can be expressed in terms of *runs*, which are relabellings of input trees induced by transitions. For a total game automaton \mathcal{A} and an initial state q_I , with each partial tree t one can associate the *run*

$$\rho(\mathcal{A}, t, q_I): \text{dom}(t) \cup \text{holes}(t) \rightarrow Q^A \cup \{\top, \perp, *\}$$

such that $\rho(\epsilon) = q_I$ and for all $v \in \text{dom}(t)$, if $\rho(v) = q$, $\delta(q, t(v)) = b_v$, then

- if b_v is $(q_L, L) \vee (q_R, R)$ or $(q_L, L) \wedge (q_R, R)$, then $\rho(vL) = q_L$ and $\rho(vR) = q_R$;
- if $b_v = (q_d, d)$ for some $d \in \{L, R\}$, then $\rho(vd) = q_d$ and $\rho(v\bar{d}) = *$;
- if $b_v = \perp$ then $\rho(vL) = \rho(vR) = \perp$, and dually for \top ;

and if $\rho(v) \in \{\top, \perp, *\}$, then $\rho(vL) = \rho(vR) = *$. Observe that $\rho(v)$ is uniquely determined by the labels of t on the path leading to v .

The run $\rho = \rho(\mathcal{A}, t, q_I)$ is naturally interpreted as a game $\mathbf{G}_\rho(\mathcal{A}, t, q_I)$ with positions $\text{dom}(t) - \rho^{-1}(*)$, final positions $\text{holes}(t)$, where edges follow the child relation and loop on $\rho^{-1}(\{\top, \perp\})$, priority of v is $\Omega^A(\rho(v))$ with $\Omega^A(\perp) = 1$, $\Omega^A(\top) = 0$, and the owner of v is \exists iff $\delta(\rho(v), t(v)) = (q_L, L) \vee (q_R, R)$ for some $q_L, q_R \in Q^A$. Clearly $\mathbf{G}_\rho(\mathcal{A}, t, q_I)$ is equivalent to $\mathbf{G}(\mathcal{A}, t, q_I)$. If t is total, we say that ρ is *accepting*, if \exists has a winning strategy in $\mathbf{G}_\rho(\mathcal{A}, t, q_I)$.

Let $t \in \text{PTr}_A$ be a partial tree and $\rho = \rho(\mathcal{A}, t, q_I)$ be the run of an automaton \mathcal{A} on t . We say that t *resolves \mathcal{A} from $q_I \in Q^A$* if $\rho(h) \neq *$ for each hole h of t and whenever $t \upharpoonright_{vd}$ is the only total tree in $\{t \upharpoonright_{vL}, t \upharpoonright_{vR}\}$, either $\rho(vd) = *$ or vd is losing for the owner of v in $\mathbf{G}_\rho(\mathcal{A}, t, q_I)$.

Fact II.2. *Assume that t resolves \mathcal{A} from q_I and $\rho = \rho(\mathcal{A}, t, q_I)$. If t has a single hole h then $t[v := s] \in \text{L}(\mathcal{A}, q_I)$ iff $s \in \text{L}(\mathcal{A}, \rho(h))$ for all $s \in \text{Tr}_A$.*

If t has two holes h, h' whose closest common ancestor u satisfies $\delta_A(\rho(u), t(u)) = (q_L, L) \wedge (q_R, R)$ for some q_L, q_R , then $t[h := s, h' := s'] \in \text{L}(\mathcal{A}, q_I)$ iff $s \in \text{L}(\mathcal{A}, \rho(h))$ and $s' \in \text{L}(\mathcal{A}, \rho(h'))$ for all s, s' ; dually for $(q_L, L) \vee (q_R, R)$.

III. NON-DETERMINISTIC INDEX PROBLEM

Decidability of the non-deterministic index problem for languages recognised by game automata is an immediate consequence of decidability for deterministic tree languages [18] and the following observation.

Proposition III.1. *For each game automaton \mathcal{A} and a state $q_I^A \in Q^A$ one can effectively construct a deterministic automaton \mathcal{D} with initial state q_I^D , such that $\text{L}(\mathcal{A}, q_I^A)$ is recognised by a non-deterministic automaton of index (i, j) if and only if so is $\text{L}(\mathcal{D}, q_I^D)$.*

Proof: Essentially, \mathcal{D} recognises the set of winning strategies for \exists in games induced by the runs of \mathcal{A} . For two total trees $t \in \text{Tr}_A$, $s \in \text{Tr}_B$ let $t \otimes s \in \text{Tr}_{A \times B}$ be given by $(t \otimes s)(v) = (t(v), s(v))$. Let $W_{\mathcal{A}, q_I}^{\exists}$ be the set of all total trees $t \otimes s$ over the alphabet $A^A \times \{L, R, *\}$ such that s encodes a winning strategy for \exists in the game $\mathbf{G}_\rho(\mathcal{A}, t, q_I)$ in the following sense: if $s(v) \in \{L, R\}$, \exists should choose $v \cdot s(v)$, and $s(v) = *$ means that \exists has no choice in v . It is easy to see that $W_{\mathcal{A}, q_I}^{\exists}$ can be recognised by a deterministic automaton.

Note that

$$\text{L}(\mathcal{A}, q_I^A) = \{t \in \text{Tr}_{A^A} \mid \exists s. t \otimes s \in W_{\mathcal{A}, q_I}^{\exists}\}.$$

Hence, if $W_{\mathcal{A}, q_I^A}^\exists = L(\mathcal{B}, q_I^B)$ for some non-deterministic automaton \mathcal{B} then $L(\mathcal{A}, q_I^A) = L(\mathcal{B}', q_I^B)$, where \mathcal{B}' is the standard projection of \mathcal{B} on the alphabet A^A : for all $q \in Q^A$ and $a \in A^A$, $\delta^{\mathcal{B}'}(q, a) = \delta^{\mathcal{B}}(q, (a, L)) \vee \delta^{\mathcal{B}}(q, (a, R)) \vee \delta^{\mathcal{B}}(q, (a, \star))$. The projection does not influence the index.

For the other direction, the proof is based on the following observation. For $t \in \text{Tr}_{A^A}$ and $s \in \text{Tr}_{\{L, R, \star\}}$ let $t \odot s \in \text{Tr}_{A^A}$ be the tree obtained from t by the following operation: for each v , if $\rho_{t, q_I}(v) = q$, $\delta(q, t(v)) = (q_L, L) \vee (q_R, R)$, and $s(v) = L$, then replace the subtree of t rooted at vR by some fixed regular tree in the complement of $L(\mathcal{A}, q_R)$; dually for $s(v) = R$. (Recall that \mathcal{A} has only non-trivial states, so $L(\mathcal{A}, q_R) \subsetneq \text{Tr}_{A^A}$.) If s encodes a strategy σ_s for \exists in $\mathbf{G}_\rho(\mathcal{A}, t, q_I^A)$, then σ_s is winning if and only if $t \odot s \in L(\mathcal{A}, q_I^A)$. Hence, $t \otimes s \in W_{\mathcal{A}, q_I^A}^\exists$ if and only if s encodes a strategy for \exists in $\mathbf{G}_\rho(\mathcal{A}, t, q_I^A)$ and $t \odot s \in L(\mathcal{A}, q_I^A)$. These conditions can be checked by a non-deterministic automaton of index (i, j) as soon as $L(\mathcal{A}, q_I^A)$ can be recognised by such an automaton. ■

IV. ALTERNATING INDEX PROBLEM

In this section we show that the alternating index problem is decidable for game automata. Let us start with some notation.

Definition IV.1. For $i < j \in \mathbb{N}$, let $\text{RM}(i, j)$ be the class of languages recognised by alternating automata of index (i, j) . Let $\Sigma_j^{\text{RM}} = \text{RM}(0, j)$, $\Pi_j^{\text{RM}} = \text{RM}(1, j + 1)$, and $\Delta_j^{\text{RM}} = \text{RM}(0, j) \cap \text{RM}(1, j + 1)$. The above classes are naturally ordered by inclusion.

The result we prove not only gives decidability of the alternating index problem but also shows that languages recognisable by game automata collapse inside the Δ_i^{RM} classes. To express it precisely we recall the so-called *comp* classes [22] that can be defined in terms of strongly connected components (SCCs) of a graph naturally associated with each alternating automaton.

Definition IV.2. Let \mathcal{A} be an alternating automaton. Let $\text{Gph}(\mathcal{A})$ be the directed edge-labelled graph over the set of vertices Q such that there is an edge $p \xrightarrow{(a, d)} q$ whenever (q, d) occurs in $\delta(p, a)$. Additionally, vertices of $\text{Gph}(\mathcal{A})$ are labelled by values of Ω . We write $p \xrightarrow{w} q$ if there is a path in $\text{Gph}(\mathcal{A})$ whose edge-labels yield the word w .

Definition IV.3. An alternating automaton \mathcal{A} is in $\text{Comp}(i, j)$ if (ignoring edge-labels) each SCC in $\text{Gph}(\mathcal{A})$ has priorities between i and j or between $i + 1$ and $j + 1$.

It follows from the definition that each $\text{Comp}(i, j)$ automaton is a $(i, j + 1)$ automaton, and can be transformed into an equivalent $\text{Comp}(i + 1, j + 2)$ automaton by scaling the priorities. We write Comp_j for the class of languages recognised by $\text{Comp}(0, j)$ automata. We then have

$$\Sigma_j^{\text{RM}} \cup \Pi_j^{\text{RM}} \subseteq \text{Comp}_j \subseteq \Delta_{j+1}^{\text{RM}}.$$

The Comp_0 automata are more widely known as *weak alternating automata*. They recognise exactly those languages

that are definable in weak monadic second-order logic. An important result by Rabin [23] states that the classes Comp_0 and Δ_1^{RM} coincide. However, as shown by Arnold and Santocanale [22], for higher levels the inclusion is strict

$$\text{Comp}_j \subsetneq \Delta_{j+1}^{\text{RM}} \quad \text{for } j > 0,$$

i.e., there are examples of regular languages in Δ_{j+1}^{RM} but not in Comp_j . It turns out that, as a consequence of our characterisation, in the case of languages recognisable by game automata the respective classes Comp_j and Δ_{j+1}^{RM} coincide for all levels.

Theorem IV.1. For each game automaton \mathcal{A} and an initial state q_I , the language $L(\mathcal{A}, q_I)$ belongs to exactly one of the classes: Comp_0 , $\Sigma_i^{\text{RM}} - \Pi_i^{\text{RM}}$, $\Pi_i^{\text{RM}} - \Sigma_i^{\text{RM}}$, or $\text{Comp}_i - (\Sigma_i^{\text{RM}} \cup \Pi_i^{\text{RM}})$, for $i > 0$.

Moreover, it can be effectively decided which class it is and an automaton from this class can be constructed.

The rest of this section is devoted to showing this result. Section IV-A describes a recursive procedure to compute the class of the given language $L(\mathcal{A}, q_I)$, i.e., Σ_i^{RM} , Π_i^{RM} , or Comp_i , depending on which of the possibilities holds. Sections IV-B, IV-C show that the procedure is correct. The estimation of Section IV-B is in fact an effective construction of an automaton from the respective class.

A. The algorithm

Let \mathcal{A} be an alternating automaton of index (i, j) . For $n \in \mathbb{N}$ we denote by $\mathcal{A}^{\geq n}$ the sub-automaton obtained from \mathcal{A} by restricting to states of priority at least n . Observe that the index of $\mathcal{A}^{\geq n}$ is at most (n, j) . A sub-automaton $\mathcal{B} \subseteq \mathcal{A}$ is an *n-component* of \mathcal{A} if $\text{Gph}(\mathcal{B})$ is a strongly connected component of $\text{Gph}(\mathcal{A}^{\geq n})$. We say that \mathcal{B} is *non-trivial* if $\text{Gph}(\mathcal{B})$ contains at least one edge. Our algorithm computes the class of each *n-component* \mathcal{B} of \mathcal{A} , based on the classes of $n + 1$ -components of \mathcal{B} and transitions between them. (We shall see that for *n-components* the class does not depend on the initial state.)

We begin with a simple preprocessing. An automaton \mathcal{A} is *rank-reduced* if for all $n > 0$, each *n-component* of \mathcal{A} is non-trivial and contains a state of priority n .

Lemma IV.1. Each game automaton can be effectively transformed into an equivalent rank-reduced game automaton.

The main algorithm uses three simple notions. An $n + 1$ -component \mathcal{B}_0 of \mathcal{B} is \exists -*branching* if \mathcal{B} contains a transition

$$\delta(p, a) = (q_L, L) \vee (q_R, R)$$

with $p, q_L \in Q^{\mathcal{B}_0}$ or $p, q_R \in Q^{\mathcal{B}_0}$. For \forall replace \vee with \wedge .

For a class K , operations K^\exists and K^\forall are defined as

$$\begin{aligned} (\Sigma_m^{\text{RM}})^\exists &= (\Pi_{m-1}^{\text{RM}})^\exists = (\text{Comp}_{m-1})^\exists = \Sigma_m^{\text{RM}}, \\ (\Pi_m^{\text{RM}})^\forall &= (\Sigma_{m-1}^{\text{RM}})^\forall = (\text{Comp}_{m-1})^\forall = \Pi_m^{\text{RM}}. \end{aligned}$$

We write $\bigvee_{\ell=1}^k K_\ell$ for the largest class among K_1, K_2, \dots, K_ℓ if it exists, or Comp_m if among these classes there are two maximal ones, Σ_m^{RM} and Π_m^{RM} .

Let \mathcal{A} be a rank-reduced game automaton of index (i, j) . The algorithm starts from $n = j$ and proceeds downward. Let \mathcal{B} be an n -component.

- If \mathcal{B} has only states of priority n , set $\text{class}(\mathcal{B}) = \text{Comp}_0$.
- If \mathcal{B} has no states of priority n , it coincides with a single 1-component \mathcal{B}_1 . Set $\text{class}(\mathcal{B}) = \text{class}(\mathcal{B}_1)$.
- Otherwise, assume that n is even (for odd n replace \exists with \forall). Let $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_k$, be the $(n+1)$ -components of \mathcal{B} that are \exists -branching, and let $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_{k'}$ be the ones that are not \exists -branching. We set

$$\text{class}(\mathcal{B}) = \bigvee_{\ell=1}^k \text{class}(\mathcal{B}_\ell)^\exists \vee \bigvee_{\ell=1}^{k'} \text{class}(\mathcal{C}_\ell),$$

Let $\text{class}(\mathcal{A}, q_I) = \bigvee_{\ell=1}^k \text{class}(\mathcal{A}_\ell)$ where $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_k$ are the i -components of \mathcal{A} reachable from q_I in $\text{Gph}(\mathcal{A})$.

B. Upper bounds

In this subsection we show that $L(\mathcal{A}, q_I)$ can be recognised by a $\text{class}(\mathcal{A}, q_I)$ -automaton. The argument will closely follow the recursive algorithm, pushing through an invariant guaranteeing that each n -component \mathcal{B} of \mathcal{A} can be replaced with an “equivalent” $\text{class}(\mathcal{B})$ -automaton. The notion of equivalence for non-total automata is formalised by simulations.

Definition IV.4. An alternating automaton \mathcal{S} simulates a game automaton \mathcal{A} if $F^{\mathcal{S}} \subseteq F^{\mathcal{A}}$ and there exists an embedding $\iota: Q^{\mathcal{A}} \rightarrow Q^{\mathcal{S}}$ (usually $Q^{\mathcal{A}} \subseteq Q^{\mathcal{S}}$) such that for all $t \in \text{Tr}_A$, $q_I^{\mathcal{A}} \in Q^{\mathcal{A}}$, and for each winning strategy σ for player P in $\mathbf{G}(\mathcal{A}, t, q_I^{\mathcal{A}})$ there is a winning strategy $\sigma^{\mathcal{S}}$ for P in $\mathbf{G}(\mathcal{S}, t, \iota(q_I^{\mathcal{A}}))$ such that the guarantee of $\sigma^{\mathcal{S}}$ is contained in the guarantee of σ , and if there is an infinite play conforming to $\sigma^{\mathcal{S}}$ then there is an infinite play conforming to σ .

Note that if \mathcal{A} and \mathcal{S} are total and \mathcal{S} simulates \mathcal{A} then $L(\mathcal{A}, q_I^{\mathcal{A}}) = L(\mathcal{S}, \iota(q_I^{\mathcal{A}}))$.

Lemma IV.2. For each n -component \mathcal{B} of a game automaton \mathcal{A} , \mathcal{B} can be simulated by a $\text{class}(\mathcal{B})$ -automaton.

Proof: Assume that the index of \mathcal{A} is (i, j) . We proceed by induction on $n = j, j-1, \dots, i$. If all states of \mathcal{B} have priority n or all have priority strictly greater than n , the claim is immediate. Let us assume that neither is the case. By symmetry it is enough to give the construction for even n .

Suppose \mathcal{B} has only \exists -branching $n+1$ components, $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_k$. Then $\text{class}(\mathcal{B}) = \bigvee_{\ell} \text{class}(\mathcal{B}_\ell)^\exists = \Sigma_m^{RM}$ for some $m \geq 1$. By the inductive hypothesis we get a $\text{class}(\mathcal{B}_\ell)$ -automaton \mathcal{B}_ℓ^S , simulating \mathcal{B}_ℓ . Since $\Sigma_m^{RM} \geq \text{class}(\mathcal{B}_\ell)^\exists$, \mathcal{B}_ℓ^S can be assumed to be an $(n, n+m)$ -automaton. Hence, we can put

$$\mathcal{B}^S = \mathcal{B} \upharpoonright_{\Omega^{-1}(n)} \cdot \mathcal{B}_1^S \cdot \mathcal{B}_2^S \cdot \dots \cdot \mathcal{B}_k^S$$

to get an $(n, n+m)$ -automaton simulating \mathcal{B} .

Now, assume that \mathcal{B} contains also $n+1$ components $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_{k'}$ that are not \exists -branching. Repeating the construction above would now result in an automaton of index $\bigvee_{\ell} \text{class}(\mathcal{B}_\ell)^\exists \vee \bigvee_{\ell} \text{class}(\mathcal{C}_\ell)^\exists$, potentially higher than

$\text{class}(\mathcal{B}) = \bigvee_{\ell} \text{class}(\mathcal{B}_\ell)^\exists \vee \bigvee_{\ell} \text{class}(\mathcal{C}_\ell)$. Hence, instead of \mathcal{C}_ℓ^S we shall use $\mathcal{C}_\ell^R \cdot \mathcal{C}_\ell^T$, where

- \mathcal{C}_ℓ^T is a copy of \mathcal{C}_ℓ^S with each transition leading to an exit of \mathcal{C}_ℓ^S that is not an exit of \mathcal{B} , replaced with a transition to \top (losing for \forall);
- \mathcal{C}_ℓ^R is \mathcal{C}_ℓ^S with all priorities set to n and additional ε -transitions (which can be eliminated in the usual way): for each state q of \mathcal{C}_ℓ^R allow \forall to decide to stay in q or move to the copy of q in \mathcal{C}_ℓ^T (treated as an exit in \mathcal{C}_ℓ^R).

Thus,

$$\mathcal{B}^S = \mathcal{B} \upharpoonright_{\Omega^{-1}(n)} \cdot \mathcal{B}_1^S \cdot \dots \cdot \mathcal{B}_k^S \cdot \mathcal{C}_1^R \cdot \mathcal{C}_1^T \cdot \dots \cdot \mathcal{C}_{k'}^R \cdot \mathcal{C}_{k'}^T.$$

The composition of automata $\mathcal{B} \upharpoonright_{\Omega^{-1}(n)}$, \mathcal{B}_ℓ^S , \mathcal{C}_ℓ^R gives a $\text{class}(\mathcal{B})$ -automaton (each \mathcal{C}_ℓ^S was replaced with an (n, n) -automaton \mathcal{C}_ℓ^R). This is further composed with $\text{class}(\mathcal{C}_\ell)$ -automata \mathcal{C}_ℓ^T in a loop-less way. Hence, \mathcal{B}^S is a $\text{class}(\mathcal{B})$ -automaton.

Let us see that \mathcal{B}^S simulates \mathcal{B} . Let ι be defined as identity on $\mathcal{B} \upharpoonright_{\Omega^{-1}(n)}$, on $Q^{\mathcal{B}_\ell}$ as the embedding $Q^{\mathcal{B}_\ell} \rightarrow Q^{\mathcal{B}_\ell^S}$, and on $Q^{\mathcal{C}_\ell}$ as the embedding $Q^{\mathcal{C}_\ell} \rightarrow Q^{\mathcal{C}_\ell^R}$. Consider a tree $t \in \text{Tr}_A$, a state $q_I^{\mathcal{B}}$ of \mathcal{B} , and games $\mathbf{G}(\mathcal{B}, t, q_I^{\mathcal{B}})$ and $\mathbf{G}(\mathcal{B}^S, t, \iota(q_I^{\mathcal{B}}))$.

First, consider a strategy σ for \exists in $\mathbf{G}(\mathcal{B}, t, q_I^{\mathcal{B}})$. We decompose this strategy into parts corresponding to the sub-automata \mathcal{B}_ℓ and \mathcal{C}_ℓ , for each part we use the fact that \mathcal{B}_ℓ^S simulates \mathcal{B}_ℓ and \mathcal{C}_ℓ^S simulates \mathcal{C}_ℓ . This gives us a strategy for \exists on parts of $\mathbf{G}(\mathcal{B}^S, t, \iota(q_I^{\mathcal{B}}))$ corresponding to sub-automata \mathcal{B}_ℓ^S , \mathcal{C}_ℓ^R , \mathcal{C}_ℓ^T . Outside of \mathcal{B}_ℓ^S , \mathcal{C}_ℓ^R , and \mathcal{C}_ℓ^T , \exists has the same choices in \mathcal{B}^S as in \mathcal{B} . Therefore, she can make her choices according to σ . This gives a complete strategy σ^S . Now consider any play conforming to σ^S . Such a play either visits infinitely many times a state of priority n in \mathcal{B}^S , and so is winning for \exists , or from some point on it stays in some sub-automaton \mathcal{B}_ℓ^S , \mathcal{C}_ℓ^R or \mathcal{C}_ℓ^T . In this case the play is also winning for \exists , by the assumption on σ and by the fact that all the changes of priorities in \mathcal{C}_ℓ^R 's and transitions in \mathcal{C}_ℓ^T 's are favourable to \exists . By the definition of σ^S , the guarantee of σ^S is contained in the guarantee of σ , and if there is an infinite play conforming to σ^S then there is an infinite play conforming to σ .

For a winning strategy σ for \forall in $\mathbf{G}(\mathcal{B}, t, q_I^{\mathcal{B}})$, we construct a winning strategy σ^S for \forall in $\mathbf{G}(\mathcal{B}^S, t, \iota(q_I^{\mathcal{B}}))$ as follows:

- in positions corresponding to states of priority n in \mathcal{B} the strategy σ^S follows the decisions of σ ;
- in components \mathcal{B}_ℓ^S , \mathcal{C}_ℓ^R , \mathcal{C}_ℓ^T the strategy σ^S simulates σ (using the fact that \mathcal{C}_ℓ^R and \mathcal{C}_ℓ^T have the same states and exits as the automaton \mathcal{C}_ℓ^S that simulates \mathcal{C}_ℓ) with the following exception: \forall immediately moves from \mathcal{C}_ℓ^R to \mathcal{C}_ℓ^T whenever each extension of the current play, conforming to the simulating strategy, stays forever in \mathcal{C}_ℓ^R (possibly reaching an exit that is also an exit of \mathcal{B}^S).

An easy inductive argument shows that

- 1) each position (v, p) with $p \in \mathcal{B} \upharpoonright_{\Omega^{-1}(n)}$ that is reached in some play conforming to σ^S is also reached in some play in $\mathbf{G}(\mathcal{B}, t, q_I^{\mathcal{B}})$ conforming to σ ;
- 2) whenever a play conforming to σ^S enters \mathcal{B}_ℓ^S (resp. \mathcal{C}_ℓ^R) in a position (v, p) , then $p = \iota(q)$ for some $q \in \mathcal{B}_\ell$

(resp. $q \in \mathcal{C}_\ell$) and (v, q) is reached in some play in $\mathbf{G}(\mathcal{B}, t, q_I^{\mathcal{B}})$ conforming to σ .

Consider any play b^S conforming to σ^S .

Assume that b^S is a finite play leading to a final position (v, f) . Unless (v, f) is entered directly from some \mathcal{C}_ℓ^T , by the two observations above (and by the definition of σ^S) it follows that (v, f) can also be reached in some play conforming to σ . Assume that (v, f) is entered directly from some \mathcal{C}_ℓ^T . Let $(w, \iota(q))$ be the last moment when b^S entered \mathcal{C}_ℓ^R (recall that \mathcal{C}_ℓ^T is only entered from \mathcal{C}_ℓ^R). Since σ^S in \mathcal{C}_ℓ^R and \mathcal{C}_ℓ^T mimics the simulating strategy in \mathcal{C}_ℓ^S , the final position (v, f) can be reached in some play in $\mathbf{G}(\mathcal{B}, t, q_I^{\mathcal{B}})$ starting in (w, q) , conforming to σ . By observation 2 it follows that (v, f) is reached in a play conforming to σ and starting in $(\varepsilon, q_I^{\mathcal{B}})$.

The remaining case is when b^S is an infinite play. Should b^S visit infinitely often positions of priority n , by the observation 2 and by the definition of σ^S we would define a play in $\mathbf{G}(\mathcal{B}, t, q_I^{\mathcal{B}})$ conforming to σ that visits infinitely often positions of priority n . This is impossible since σ is winning for \forall . It follows that from some point on b^S stays in some sub-component. If the sub-component is \mathcal{B}_ℓ^S , \forall wins as he is playing with a winning strategy in \mathcal{B}_ℓ^S . The other possibility is that b^S stays forever in $\mathcal{C}_\ell^R \cdot \mathcal{C}_\ell^T$ for some ℓ . Since \mathcal{C}_ℓ is not \exists -branching, in each transition of the form $\delta(p, a) = (q_L, L) \vee (q_R, R)$ with $p \in Q^{\mathcal{B}_\ell}$, at least one of the states q_L, q_R is an exit state in \mathcal{B} , or both are outside of \mathcal{C}_ℓ . Hence, after entering \mathcal{C}_ℓ , σ becomes a single path in \mathcal{C}_ℓ , with all the branchings (choices of \exists) going directly to exits of \mathcal{B} . In general, this path may end in a position belonging to \exists , such that both choices lead outside of \mathcal{C}_ℓ (not necessarily to exits of \mathcal{B} .) In our case the path must stay in \mathcal{C}_ℓ forever: since b^S is infinite and stays forever in $\mathcal{C}_\ell^R \cdot \mathcal{C}_\ell^T$, there is an infinite play conforming to the strategy simulating σ in \mathcal{C}_ℓ^S and, by Definition IV.4, an infinite play conforming to σ in \mathcal{C}_ℓ . Consequently, all exits reachable with σ in \mathcal{C}_ℓ are also exits of \mathcal{B} . Hence, as soon as b^S enters \mathcal{C}_ℓ^R for the last time, σ^S tells \forall to move to \mathcal{C}_ℓ^T where \forall wins all infinite plays. ■

It follows easily that $L(\mathcal{A}, q_I)$ can be recognised by a $\text{class}(\mathcal{A}, q_I)$ -automaton: the automaton can be obtained as a loop-less composition of the $\text{class}(\mathcal{A}_\ell)$ -automata simulating the i -components \mathcal{A}_ℓ of \mathcal{A} reachable from q_I . In other words, the alternating index bounds computed by the algorithm in Section IV-A are correct.

C. Lower bounds

It remains to see that $L(\mathcal{A}, q_I)$ cannot be recognised by an alternating automaton of index lower than $\text{class}(\mathcal{A}, q_I)$. Our proof uses the concept of topological hardness. A classical notion of topological hardness relies on the Borel hierarchy and the projective hierarchy [24], but these notions are not suitable for us, since most regular tree languages live on the same level of these hierarchies: Δ_2^1 . We use a more refined notion based on continuous reductions [25] and so-called game languages [21], [26], [27].

Definition IV.5. For $i < j$ consider the following alphabet

$$A_{i,j} = \{\exists, \forall\} \times \{i, i+1, \dots, j\}.$$

With each $t \in \text{PTr}_{A_{i,j}}$ we associate a parity game \mathbf{G}_t where

- $V = \text{dom}(t)$, $F = \text{holes}(t)$,
- $E = \{(v, vd) \mid v \in \text{dom}(t), d \in \{L, R\}\}$,
- if $t(v) = (P, n)$ then $\Omega(v) = n$ and $v \in V_P$ for $P \in \{\exists, \forall\}$.

Let $W_{i,j}$ be the set of total trees over $A_{i,j}$ such that \exists has a winning strategy in \mathbf{G}_t .

Topological hardness of languages can be compared using continuous reductions. A *continuous reduction* of $L_1 \subseteq X$ to $L_2 \subseteq Y$ is a continuous function $f: X \rightarrow Y$ such that $f^{-1}(L_2) = L_1$. The fact that L_1 can be continuously reduced to L_2 is denoted by $L_1 \leq_W L_2$. On Borel sets, the pre-order \leq_W induces the so-called Wadge hierarchy (see [25]) which greatly refines the Borel hierarchy and has the familiar ladder shape with pairs of mutually dual classes alternating with single self-dual classes. Here, we are interested in the following connection between continuous reductions, languages $W_{i,j}$, and the alternating index hierarchy.

Fact IV.1 ([21], [26], [27]). For all $i < j$,

- 1) $W_{i,j}$ is regular and $W_{i,j} \in \text{RM}(i, j)$,
- 2) $L \leq_W W_{i,j}$ for each $L \in \text{RM}(i, j)$,
- 3) $W_{i,j} \not\leq_W W_{i+1, j+1}$,
- 4) $W_{i,j} \in \Delta_2^1$, $W_{0,1}$ is Σ_1^1 -complete, $W_{1,2}$ is Π_1^1 -complete.

This gives a criterion for proving index lower bounds.

Corollary IV.1. If $W_{i,j} \leq_W L$ then $L \notin \text{RM}(i+1, j+1)$.

In consequence, in order to show that the index bound computed by the algorithm from Section IV-A is tight, it suffices to show that if $\text{RM}(i, j) \leq \text{class}(\mathcal{A}, q_I)$, then $W_{i,j} \leq_W L(\mathcal{A}, q_I)$. We construct the reduction in three steps:

- 1) we show that if the class computed by the algorithm is at least $\text{RM}(i, j)$, then this is witnessed with a certain subgraph in $\text{Gph}(\mathcal{A})$, called (i, j) -edelweiss;
- 2) we introduce intermediate languages $\widehat{W}_{i,j}$, whose internal structure corresponds precisely to (i, j) -edelweisses, and show that $\widehat{W}_{i,j} \leq_W L(\mathcal{A}, q_I)$ if only \mathcal{A} contains an (i, j) -edelweiss reachable from q_I ;
- 3) we prove that $W_{i,j} \leq_W \widehat{W}_{i,j}$.

The combinatorial core of the argument is the last step.

Definition IV.6. We say that in a game automaton \mathcal{B} there is an i -loop rooted in p if there exists a word w such that on the path $p \xrightarrow{w} p$ in $\text{Gph}(\mathcal{B})$ the minimal priority is i .

Automaton \mathcal{B} contains an (i, j) -loop for \exists rooted in p , if there exist states q, q_L, q_R of \mathcal{B} , a letter a , and words w, w_L, w_R such that:

- $\delta(q, a) = (q_L, L) \vee (q_R, R)$;
- $p \xrightarrow{w} q$; $q_L \xrightarrow{w_L} p$; $q_R \xrightarrow{w_R} p$;
- on one of the paths $p \xrightarrow{w(a, L) w_L} p$, $p \xrightarrow{w(a, R) w_R} p$ the minimal priority is i and on the other it is j .

For \forall dually, with \vee replaced with \wedge .

For an even $j > i$, \mathcal{B} contains an (i, j) -edelweiss rooted in p (see Fig. 1) if for some even n it contains

- $(n+k)$ -loops for $k = i, i+1, \dots, j-3$,
- $(n+j-2, n+j-1)$ -loop for \exists , if $i \leq j-2$; and
- $(n+j-1, n+j)$ -loop for \forall ;

all rooted in p . For odd j swap \forall and \exists but keep n even.

Lemma IV.3. Let \mathcal{A} be a game automaton and q_I a state of \mathcal{A} . If $\text{class}(\mathcal{A}, q_I) \geq \text{RM}(i, j)$ then \mathcal{A} contains an (i, j) -edelweiss rooted in a state reachable from q_I .

Proof: Let us first assume that $(i, j) = (0, 1)$. Analysing the algorithm we see that the only case when $\text{class}(\mathcal{A}, q_I)$ jumps to $\text{RM}(0, 1)$ is when for some even n there is an n -component \mathcal{B} in \mathcal{A} , reachable from q_I , and containing states of priority n , such that some $n+1$ component \mathcal{B}_ℓ of \mathcal{B} is \exists -branching in \mathcal{B} , i.e., \mathcal{B} contains a transition of the form

$$\delta(p, a) = (q_L, L) \vee (q_R, R)$$

with $p, q_L \in Q^{\mathcal{B}_\ell}$, $q_R \in Q^{\mathcal{B}}$ (or symmetrically, $p, q_R \in Q^{\mathcal{B}_\ell}$, $q_L \in Q^{\mathcal{B}}$). Since \mathcal{A} is rank-reduced, p is reachable from q_L within \mathcal{B}_ℓ via a state of priority $n+1$, and from q_R within \mathcal{B} via a state of priority n . This gives an $(n, n+1)$ -loop for \exists (a $(0, 1)$ -edelweiss) rooted in a state reachable from q_I . The argument for $(1, 2)$ is entirely dual.

Next, assume that $(i, j) = (0, 2)$. It follows immediately from the algorithm that \mathcal{A} contains an n -component \mathcal{B} (reachable from q_I , containing states of priority n) such that n is even and there exists an \exists -branching $(n+1)$ -component \mathcal{B}_ℓ in \mathcal{B} such that $\text{class}(\mathcal{B}_\ell) = \Pi_1^{RM}$ or $\text{class}(\mathcal{B}_\ell) = \text{Comp}_1$. In either case, $\text{class}(\mathcal{B}_\ell) \geq \text{RM}(1, 2)$ and by the previous case \mathcal{B}_ℓ contains an $(n', n'+1)$ -loop for \forall , for some odd $n' \geq n$. Since \mathcal{A} is rank-reduced, for each state q in \mathcal{B}_ℓ and each r between n and $\Omega(q)$, there is a loop from q to q with the lowest priority r . Hence, the $(n', n'+1)$ -loop can be turned into an $(n+1, n+2)$ -loop. Thus, \mathcal{B}_ℓ contains an $(n+1, n+2)$ -loop for \forall , rooted in a state p . We claim that \mathcal{B} contains an $(n, n+1)$ -loop for \exists , also rooted in p (giving a $(0, 2)$ -edelweiss rooted in p). Indeed, since \mathcal{B}_ℓ is \exists -branching, arguing like for $(0, 1)$, we obtain an $(n, n+1)$ -loop for \exists rooted in a state p' in \mathcal{B}_ℓ . Since \mathcal{B}_ℓ is an $n+1$ -component, there are paths in \mathcal{B}_ℓ from p to p' and back; the lowest priority on these paths is at least $n+1$. Using these paths one easily transforms the $(n, n+1)$ -loop rooted in p' into an $(n, n+1)$ -loop rooted in p .

The inductive step is easy. Suppose that $j-i > 2$. Then, for some even n , \mathcal{A} contains an $(n+i)$ -component \mathcal{B} (reachable from q_I , containing states of priority $n+i$), which has an $(n+i+1)$ -component \mathcal{B}_ℓ such that $\text{class}(\mathcal{B}_\ell) = \text{RM}(i+1, j)$ or $\text{class}(\mathcal{B}_\ell) = \text{Comp}(i+1, j)$. Since for each state p in \mathcal{B}_ℓ , \mathcal{B} contains an $(n+i)$ -loop rooted in p , we can conclude by the inductive hypothesis. ■

Definition IV.7. For $i \leq 2k-2$ consider the alphabet

$$\widehat{A}_{i,2k} = \{i, i+1, \dots, 2k-3, e, a\}.$$

With each $t \in \text{PTr}_{\widehat{A}_{i,2k}}$ we associate a parity game $\widehat{\mathbf{G}}_t$ with positions $\text{dom}(t)$ and final positions $\text{holes}(t)$ such that




Fig. 1. $(0, 4)$ -edelweiss and $(1, 5)$ -edelweiss.

- if $t(v) = a$, then in v player \forall can choose to go to vL or to vR , and $\Omega(vL) = 2k-1$, $\Omega(vR) = 2k$,
- if $t(v) = e$, then in v player \exists can choose to go to vL or to vR , and $\Omega(vL) = 2k-2$, $\Omega(vR) = 2k-1$,
- if $t(v) \in \{i, i+1, \dots, 2k-3\}$, the only move from v is to vL and $\Omega(vL) = t(v)$.

For $i = 2k-1$, let $\widehat{A}_{i,2k} = \{a, \top\}$, and let $\widehat{\mathbf{G}}_t$ be defined like above, except that if $t(v) = \top$ then $\Omega(v) = 2k$ and the only move from v is back to v .

Let $\widehat{W}_{i,2k}$ be the set of all total trees over $A_{i,2k}$ such that \exists has a winning strategy in $\widehat{\mathbf{G}}_t$.

The languages $\widehat{W}_{i,2k+1}$ are defined dually, with e, a and \exists, \forall swapped, and \top replaced with \perp .

Lemma IV.4. If a total game automaton \mathcal{A} contains an (i, j) -edelweiss rooted in a state reachable from an initial state q_I then $\widehat{W}_{i,j} \leq_W L(\mathcal{A}, q_I)$.

Proof: We only give a proof for $(i, j) = (1, 2)$; for other values of (i, j) the argument is entirely analogous. By the definition, \mathcal{A} contains an $(1, 2)$ -loop for \forall , rooted in a state p reachable from q_I . Since \mathcal{A} is a game automaton and has no trivial states, it follows that there exist

- a partial tree t_I resolving \mathcal{A} from q_I , with a single hole v , labelled with p in $\rho(\mathcal{A}, t_I, q_I)$;
- a partial tree t_a resolving \mathcal{A} from p with two holes v_1, v_2 , such that in $\rho(\mathcal{A}, t_a, p)$ both holes are labelled p , the lowest priority on the path from the root to v_i is i , and the closest common ancestor u of v_1 and v_2 is labelled with a state q such that $\delta_{\mathcal{A}}(q, t(u)) = (q_L, L) \wedge (q_R, R)$ for some q_L, q_R ; and
- a total tree $t_{\top} \in L(\mathcal{A}, p)$.

Let us see how to build t_a . The paths $p \xrightarrow{w(a,L)w_L} p$, $p \xrightarrow{w(a,R)w_R} p$ guaranteed by Definition IV.6 give as a partial tree s with a single branching in some node u and two leaves v_1, v_2 , which we replace with holes. For $\rho = \rho(\mathcal{A}, s, p)$, $\rho(v_1) = \rho(v_2) = p$ and $\delta(\rho(u), t(u)) = (q_L, L) \wedge (q_R, R)$. At each hole of s , except v_1 and v_2 , we substitute a total tree such that the run on the resulting tree with two holes resolves \mathcal{A} from p , e.g., if vL is a hole and $\delta(s(v), \rho(v)) = (q', L) \vee (q'', R)$, we substitute at vL any tree that is not in $L(\mathcal{A}, q')$, relying on the assumption that \mathcal{A} has no trivial states.

Let us define the reduction $g: \text{Tr}_{\{a, \top\}} \rightarrow \text{Tr}_{A^A}$. Let $t \in \text{Tr}_{\{a, \top\}}$. For $v \in \text{dom}(t)$, define t_v co-inductively as follows: if $t(v) = \top$, set $t_v = t_{\top}$; if $t(v) = a$, then t_v is obtained by plugging in the holes v_1, v_2 of t_a the trees t_{vL} and t_{vR} . Let $g(t)$ be obtained by plugging t_ε in the hole of t_I . It is easy to

check that g continuously reduces $\widehat{W}_{1,2}$ to $L(\mathcal{A}, qI)$. ■

It remains to see that $W_{i,j} \leq_W \widehat{W}_{i,j}$. For the lowest level we give a separate proof.

Lemma IV.5. $W_{0,1} \leq_W \widehat{W}_{0,1}$ and $W_{1,2} \leq_W \widehat{W}_{1,2}$.

Proof: By the symmetry it is enough to prove the first claim. Let us take $t \in \text{Tr}_{A_{0,1}}$. By König's lemma, Player \exists has a winning strategy in \mathbf{G}_t if and only if she can produce a sequence of finite strategies $\sigma_0, \sigma_1, \sigma_2, \dots$ (viewed as subtrees of t) such that

- 1) σ_0 consists of the root only;
- 2) for each n the strategy σ_{n+1} extends σ_n in such a way that below each leaf of σ_n a non-empty subtree is added, and all leaves of σ_{n+1} have priority 0.

Using this observation we can define the reduction. Let $(\tau_i)_{i \in \mathbb{N}}$ be the list of all finite unlabelled binary trees. Some of these trees naturally induce a strategy for \exists in \mathbf{G}_t . For those we define $t_{\tau_i} \in \text{Tr}_{\{e, \perp\}}$ co-inductively, as follows:

- $t_{\tau_i}(R^j) = e$ for all j ;
- if τ_j induces in \mathbf{G}_t a strategy that is a legal extension of the strategy induced by τ_i in the sense of item 2) above, then the subtree of t_{τ_i} rooted at $R^j L$ is t_{τ_j} ;
- otherwise, all nodes in this subtree are labelled with \perp .

Let $f(t) = t_{\sigma_0}$. By the initial observation, $t_{\sigma_0} \in \widehat{W}_{0,1}$ if and only if \exists has a winning strategy in \mathbf{G}_t . The function f is continuous: to determine the labels in nodes $R^{n_1} L R^{n_2} L \dots R^{n_k}$ and $R^{n_1} L R^{n_2} L \dots R^{n_k} L$ we only need to know the restriction of t to the union of the domains of $\tau_{n_1}, \tau_{n_2}, \dots, \tau_{n_k}$. Hence, f continuously reduces $W_{0,1}$ to $\widehat{W}_{0,1}$. ■

Lemma IV.6. For all i and $j \geq i + 2$, $W_{i,j} \leq_W \widehat{W}_{i,j}$.

Proof: By duality we can assume that $j = 2k$. For $t \in \text{Tr}_{A_{i,2k}}$, let us consider game $\widehat{\mathbf{G}}_t$ defined as follows. The positions are pairs (v, σ) , where v is a node of t , and σ is finite strategy from v for \forall (viewed as a subtree of $t \upharpoonright v$). Initially $v = \varepsilon$ is the root of t and $\sigma = \{\varepsilon\}$. In each round, in position (v, σ) , the players make the following moves:

- \forall extends σ under leaves of priority $2k - 1$ to σ' in such a way that on every path leading from a leaf of σ to a leaf of σ' all nodes have priority $2k$, except the leaf of σ' , which has priority at most $2k - 1$;
- \exists has the following possibilities:
 - select a leaf v' of σ' with priority at most $2k - 2$, and let the next round start with $(v', \{v'\})$, or
 - if σ' has some leaves of priority $2k - 1$, continue with (v, σ') .

A play is won by \exists if she selects a leaf infinitely many times and the least priority of these leaves seen infinitely often is even, or \forall is unable to extend σ in some round. Otherwise, the play is won by \forall .

We claim that player P has a winning strategy in \mathbf{G}_t if and only if P has a winning strategy in $\widehat{\mathbf{G}}_t$.

For a winning strategy σ_{\exists} for \exists in \mathbf{G}_t , let $\tilde{\sigma}_{\exists}$ be the strategy in $\widehat{\mathbf{G}}_t$ in which \exists selects a leaf v' in σ' if and only

if $v' \in \sigma_{\exists}$. Consider an infinite play conforming to $\tilde{\sigma}_{\exists}$. If in the play \exists selects a leaf infinitely many times, she implicitly defines a path in t conforming to σ_{\exists} , and so the play must be winning for \exists . Assume that \exists selects a leaf only finitely many times. Then, \forall produces an infinite sequence of finite strategies $\{v\} = \sigma_0 \subseteq \sigma_1 \subseteq \dots$ in \mathbf{G}_t . Let σ_{∞} be the union of these strategies. Consider the play π in \mathbf{G}_t passing through v and conforming to σ_{∞} and σ_{\exists} . Observe that for each σ_i , the strategy σ_{\exists} must choose some path; hence, either \exists selects a leaf of σ_i , or this path goes via a leaf of priority $2k - 1$. Thus, π is infinite and by the rules of $\widehat{\mathbf{G}}_t$ priorities at most $2k - 1$ are visited infinitely often. Since \exists selects a leaf only finitely many times, priorities strictly smaller than $2k - 1$ are visited finitely many times in π . Hence, π is won by \forall , what contradicts the assumption that σ_{\exists} is winning for \exists .

Now, let σ_{\forall} be a winning strategy for \forall in \mathbf{G}_t . Then, for each $v \in \sigma_{\forall}$ there exists a finite sub-strategy σ' of σ_{\forall} from v such that all internal nodes of σ' have priority $2k$ and leaves have priority at most $2k - 1$. This shows that for each current strategy $\sigma \subseteq \sigma_{\forall}$, \forall is able to produce a legal extension $\sigma' \subseteq \sigma_{\forall}$. Let $\tilde{\sigma}_{\forall}$ be a strategy of \forall in $\widehat{\mathbf{G}}_t$ that extends every given σ by σ' as above. Consider any play conforming to $\tilde{\sigma}_{\forall}$. By the initial observation, the play is infinite, so priorities strictly smaller than $2k$ are visited infinitely often. If \exists selects a leaf only finitely many times, priorities strictly smaller than $2k - 1$ occur only finitely many times and \forall wins. If \exists selects a leaf infinitely many times, then the lowest priority seen infinitely often must be odd, as otherwise \exists would show a losing path in σ_{\forall} . Hence, \forall wins in this case as well.

It remains to encode $\widehat{\mathbf{G}}_t$ as a tree $f(t) \in \text{Tr}_{\widehat{A_{i,2k}}}$ in a continuous manner. The argument is similar to the one in Lemma IV.5. Let $(\tau_n)_{n \in \mathbb{N}}$ be the list of all unlabelled finite trees. For some pairs (v, τ_n) , τ_n induces a strategy in \mathbf{G}_t from the node v . For such (v, τ_n) we define t_{v, τ_n}^{\forall} and t_{v, τ_n}^{\exists} co-inductively, as follows:

- $t_{v, \tau_n}^{\forall}(R^m) = a$ for all m ;
- the subtree of t_{v, τ_n}^{\forall} rooted at $R^m L$ is t_{v, τ_m}^{\exists} if τ_m induces a strategy from v that is a legal extension of τ_n according to the rules of $\widehat{\mathbf{G}}_t$, and otherwise the whole subtree is labelled with e 's (losing choice for \forall);
- $t_{v, \tau_n}^{\exists}(R^m) = e$ for $m = 0, 1, \dots, \ell$, where $v_0, v_1, \dots, v_{\ell}$ are the leaves in the strategy induced by τ_n from v ;
- the subtree of t_{v, τ_n}^{\exists} rooted at $R^{\ell+1}$ is t_{v, τ_n}^{\forall} if the strategy induced by τ_n from v has some leaves of priority $2k - 1$, otherwise the whole subtree is labelled with a 's (losing choice for \exists);
- for $m \leq \ell$, consider the following cases to define the subtree s_m of t_{v, τ_n}^{\exists} rooted at $R^m L$:
 - if $\Omega(v_m) \in \{2k - 1, 2k\}$ then s_m is labelled everywhere with a 's (losing choice for \exists),
 - if $\Omega(v_m) = 2k - 2$ then $s_m = t_{v_m, \{v_m\}}^{\forall}$,
 - if $\Omega(v_m) = r < 2k - 2$ then $s_m(\varepsilon) = r$, the left subtree of s_m is $t_{v_m, \{v_m\}}^{\forall}$, and the right subtree of s_m is labelled with a 's (irrelevant for \mathbf{G}_t).

Let $f(t)$ be $t_{\varepsilon, \{\varepsilon\}}^{\forall}$. Checking that f continuously reduces $W_{i,j}$

to $\widehat{W}_{i,j}$ does not pose any difficulties. ■

V. RECOGNISABILITY BY GAME AUTOMATA

In this section we give an effective characterisation of the class of languages recognised by game automata within the class of all regular languages. The characterisation is inspired by the one for deterministic automata [19], however, due to the alternation of players, the arguments here are more involved.

We begin with a handful of definitions. Let us fix a finite alphabet A . A *trace* is a finite word w over $A \cup \{L, R\}$, with letters from A on even positions, and directions from $\{L, R\}$ on odd positions. If the last symbol of w is a letter, the trace is *labelled*, otherwise it is *unlabelled*. A trace w can be seen as a partial tree $t_w \in \text{PTr}_A$ consisting of a single path: for a labelled trace $w = a_0 d_1 a_1 \dots d_k a_k$, $\text{dom}(t_w) = \{d_1 d_2 \dots d_i \mid i \leq k\}$ and $t_w(d_1 d_2 \dots d_i) = a_i$ for all $i \leq k$. Abusing the notation, we write w instead of $d_1 d_2 \dots d_k$. The tree t_w has two *final holes*, wL and wR , and *side holes* $d_1 d_2 \dots d_{i-1} \bar{d}_i$ for $i \leq k$. For an unlabelled trace $w = a_0 d_1 a_1 \dots d_k$, t_w is defined similarly, but this time it has only one final hole: $d_1 d_2 \dots d_k$. We shall also write w for this hole.

A partial tree $t \in \text{PTr}_A$ is a *realisation* of a trace w if it is obtained from t_w by putting some *total* trees in all the side holes of t_w . If w is an unlabelled trace, t still has a hole w . We write $t(t')$ for the tree obtained by putting t' in the hole w , and $t^{-1}M$ for $\{t' \mid t(t') \in M\}$. Similarly, if w is a labelled trace, we write $t(t_L, t_R)$ for the total tree obtained by putting t_L, t_R in the holes wL and wR , respectively, and we define $t^{-1}M$ as $\{(t_L, t_R) \mid t(t_L, t_R) \in M\}$. Additionally, $a^{-1}M$ stands for $t_a^{-1}M$ for the root-only tree t_a with $t_a(\varepsilon) = a$.

A language Z is *non-trivial* if neither Z nor its complement Z^c is empty. The following notions are semantic counter-parts of states and transitions of game automata.

Definition V.1. A unary profile is $*$ (standing for trivial) or a non-trivial regular tree language Z . A binary profile is $*$, \emptyset , $\text{Tr}_A \times \text{Tr}_A$, or a non-trivial subset of $\text{Tr}_A \times \text{Tr}_A$ in one of the forms $Z_L \times \text{Tr}_A$, $\text{Tr}_A \times Z_R$, $(Z_L \times \text{Tr}_A) \cup (\text{Tr}_A \times Z_R)$, or $Z_L \times Z_R$, for some non-trivial regular tree languages Z_L, Z_R .

We shall see that the binary profiles (except $*$) correspond to transitions of the form \perp , \top , (q_L, L) , (q_R, R) , $(q_L, L) \vee (q_R, R)$, and $(q_L, L) \wedge (q_R, R)$, respectively. As a first step, let us relate traces to profiles.

Definition V.2. A trace w has non-trivial profile Z in a regular language M , if for each realisation t of w either $t^{-1}M$ is trivial or $t^{-1}M = Z$, and for some realisation t_0 , $t_0^{-1}M = Z$; here Z is unary for unlabelled w and binary for labelled w .

An unlabelled trace w has profile $*$ in M if for each realisation t of w , $t^{-1}M$ is trivial. A labelled trace wa has profile $Z \in \{\emptyset, \text{Tr}_A \times \text{Tr}_A\}$ in M if w has a non-trivial profile Z' and $a^{-1}Z' = Z$; if w has profile $*$, so does the trace wa .

Note that each trace, labelled or unlabelled, has at most one profile in M . We write p_M for the partial function assigning profiles to traces. We say that M is *locally game* if each trace

has a profile in M . Equivalently, one could assume that all unlabelled traces have profiles in M .

Let us now examine the connection between profiles, and states and transitions of game automata. Let \mathcal{B} be a total game automaton and let q_I be a state of \mathcal{B} . For a trace w , let $\rho_w = \rho(\mathcal{B}, t_w, q_I)$ be the run over the tree t_w associated with w .

If w is an unlabelled trace, define $p_{\mathcal{B}, q_I}(w) = L(\mathcal{B}, q)$, if $\rho_w(w) = q \in Q^{\mathcal{B}}$; if $\rho_w(w) \notin Q^{\mathcal{B}}$, set $p_{\mathcal{B}, q_I}(w) = *$.

If w is a labelled trace, set $p_{\mathcal{B}, q_I}(w) = *$ if $\rho_w(w) \notin Q^{\mathcal{B}}$; otherwise, let $\rho_w(w) = q$ and $b_w = \delta_{\mathcal{B}}(q, a)$, where a is the last symbol of w , and set $p_{\mathcal{B}, q_I}(w) = L(\mathcal{B}, b_w)$ where $L(\mathcal{B}, b)$ is the *profile of the transition b in \mathcal{B}* , defined as

$$\begin{aligned} \emptyset & \text{ for } b = \perp; \\ \text{Tr}_A \times \text{Tr}_A & \text{ for } b = \top; \\ L(\mathcal{B}, q_L) \times \text{Tr}_A & \text{ for } b = (q_L, L); \\ \text{Tr}_A \times L(\mathcal{B}, q_R) & \text{ for } b = (q_R, R); \\ L(\mathcal{B}, q_L) \times \text{Tr}_A \cup \text{Tr}_A \times L(\mathcal{B}, q_R) & \text{ for } b = (q_L, L) \vee (q_R, R); \\ L(\mathcal{B}, q_L) \times L(\mathcal{B}, q_R) & \text{ for } b = (q_L, L) \wedge (q_R, R). \end{aligned}$$

The following is an easy consequence of Fact II.2.

Lemma V.1. For each trace w , $p_{\mathcal{B}, q_I}(w) = p_{L(\mathcal{B}, q_I)}(w)$.

Corollary V.1. Languages recognised by game automata are locally game.

Being locally game is necessary but not sufficient to be recognisable by a game automaton. In what follows, for a given locally game language M we construct a game automaton \mathcal{G}_M that locally computes the profiles and globally reflects the infinitary aspects of M . We show that M is recognised by a game automaton if and only if it is recognised by \mathcal{G}_M .

We say that a DFA $\mathcal{A} = \langle A, Q, q_I, \delta, F \rangle$ computes a partial function $f: A^* \rightarrow X$ if \mathcal{A} recognises $\text{dom}(f)$ and it comes equipped with a function $\tau^A: F \rightarrow \text{rg}(f)$, such that $\tau^A(\delta(q_I, w)) = f(w)$ for each $w \in \text{dom}(f)$, where $\delta(q, v)$ is the state of \mathcal{A} after reading word v from state q .

Lemma V.2. For each regular tree language M one can effectively construct a DFA \mathcal{A} that computes (a finite representation of) the profile of the trace w in M .

In particular, it is decidable whether M is locally game, and the set Profiles_M of all possible profiles of traces in M is finite and can be computed from M .

The infinitary aspects of M are captured by the notion of *correct* infinite traces. An infinite trace is an infinite word π over $A \cup \{L, R\}$ with letters from A on even positions and directions from $\{L, R\}$ on odd positions. Just like a finite trace, π can be seen as a tree t_π consisting of a single infinite branch which has only side holes. A tree t realises π if it is obtained by plugging total trees in the side holes of t_π .

Assume that M is locally game and let $p_M(w)$ be the profile of w in M . We say that t *resolves M up to π* if t realises π and for each labelled trace w that is a prefix of π , if wL is a prefix of π then

$$\bullet t \upharpoonright_{wR} \notin Z_R \text{ if } p_M(w) = (Z_L \times \text{Tr}_A) \cup (\text{Tr}_A \times Z_R),$$

- $t \upharpoonright_{wR} \in Z_R$ if $p_M(w) = Z_L \times Z_R$,

and symmetrically if wR is a prefix of π . An infinite trace π is M -correct if some tree $t \in M$ resolves M up to π .

If π is M -correct, it must be accepted by each game automaton recognising M in the following sense: a game automaton \mathcal{D} accepts π from q_I if \exists wins the play corresponding to π in the game associated with $\rho = \rho(\mathcal{D}, t_\pi, q_I)$, or $\rho(v) = *$ for some $v \in \text{dom}(t_\pi)$ and $\rho(w) \neq \perp$ for all $w \in \text{dom}(t_\pi)$.

Lemma V.3. *Automaton \mathcal{D} accepts π from q_I if and only if π is $L(\mathcal{D}, q_I)$ -correct.*

Lemma V.4. *For each regular tree language M one can effectively construct a deterministic parity automaton \mathcal{D} recognising the set of M -correct infinite traces.*

We define \mathcal{G}_M as a product of \mathcal{A} and \mathcal{D} , with priorities inherited from \mathcal{D} and the types of transitions (\vee , \wedge , etc.) determined by the type of profile computed by \mathcal{A} . More precisely, for $a \in A$, $(p, q) \in Q^A \times Q^D$, $\tau = \tau^A(\delta^A(p, a))$, define $\delta((p, q), a)$ as

$$\top \text{ if } \tau \in \{*, \text{Tr}_A \times \text{Tr}_A\};$$

$$\perp \text{ if } \tau = \emptyset;$$

$$\beta(p, q, a, L) \text{ if } \tau = Z_L \times \text{Tr}_A;$$

$$\beta(p, q, a, R) \text{ if } \tau = \text{Tr}_A \times Z_R;$$

$$\beta(p, q, a, L) \vee \beta(p, q, a, R) \text{ if } \tau = Z_L \times \text{Tr}_A \cup \text{Tr}_A \times Z_R;$$

$$\beta(p, q, a, L) \wedge \beta(p, q, a, R) \text{ if } \tau = Z_L \times Z_R;$$

where $\beta(p, q, a, d)$ is defined as $((\delta^A(p, ad), \delta^D(q, ad)), d)$. Let $q_M = (q_I^A, q_I^D)$.

Theorem V.1. *A regular language M is recognised by a game automaton iff M is locally game and $L(\mathcal{G}_M, q_M) = M$.*

Proof: Assume that $M = L(\mathcal{B}, q_I^B)$ for some game automaton \mathcal{B} and $q_I^B \in Q^B$. By Corollary V.1, M is locally game. Fix $t \in \text{Tr}_A$ and let $\rho_M = \rho(\mathcal{G}_M, t, q_M)$ and $\rho_B = \rho(\mathcal{B}, t, q_I^B)$. By Lemma V.1, $p_M(w)$ determines the profiles of the corresponding transitions in ρ_B and ρ_M . Hence, the games associated to these runs are isomorphic if the priorities are ignored. Let π be an infinite trace in t . By the construction, \mathcal{G}_M accepts π from q_M iff π is M -correct. By Lemma V.3, π is M -correct iff \mathcal{B} accepts π from q_I^B . It follows that ρ_B is accepting iff ρ_M is accepting. ■

As an immediate corollary we obtain the following.

Theorem V.2. *Given an alternating automaton \mathcal{A} and a state q_I , it is decidable whether $L(\mathcal{A}, q_I)$ is recognised by a game automaton. If so, some game automaton recognising $L(\mathcal{A}, q_I)$ can be effectively constructed from \mathcal{A} and q_I .*

ACKNOWLEDGEMENTS

The authors thank the anonymous referees for helpful comments and for the motivation to improve the presentation of the alternating index algorithm.

The first author was supported by the *Expressiveness of Modal Fixpoint Logics* project realised within the Homing Plus

programme of the Foundation for Polish Science, co-financed by the European Union from the Regional Development Fund within the Operational Programme Innovative Economy (“Grants for Innovation”). The second author was supported by the Polish government grant no. N N206 567840. The last author was supported by Poland’s National Science Centre grant no. DEC-2012/05/N/ST6/03254 and by ERC Starting grant ”SOSNA”.

REFERENCES

- [1] J. R. Büchi, “On a decision method in restricted second-order arithmetic,” in *Proc. 1960 Int. Congr. for Logic, Methodology and Philosophy of Science*, 1962, pp. 1–11.
- [2] M. O. Rabin, “Decidability of second-order theories and automata on infinite trees,” *Trans. of the American Math. Soc.*, vol. 141, pp. 1–35, 1969.
- [3] A. W. Mostowski, “Games with forbidden positions,” University of Gdańsk, Tech. Rep., 1991.
- [4] —, “Regular expressions for infinite trees and a standard form of automata,” in *Symposium on Computation Theory*, 1984, pp. 157–168.
- [5] A. Arnold and D. Niwiński, *Rudiments of μ -calculus*, ser. Studies in Logic and the Foundations of Mathematics. Burlington, MA: Elsevier, 2001.
- [6] A. Emerson and C. Jutla, “Tree automata, μ -calculus, and determinacy,” in *FOCS’91*, 1991, pp. 368–377.
- [7] M. Otto, “Eliminating recursion in the μ -calculus,” in *STACS*, 1999, pp. 531–540.
- [8] R. Küsters and T. Wilke, “Deciding the first level of the μ -calculus alternation hierarchy,” in *FSTTCS*, 2002, pp. 241–252.
- [9] I. Walukiewicz, “Deciding low levels of tree-automata hierarchy,” *Electr. Notes Theor. Comput. Sci.*, vol. 67, pp. 61–75, 2002.
- [10] M. Bojańczyk and T. Place, “Regular languages of infinite trees that are Boolean combinations of open sets,” in *ICALP (2)*, 2012, pp. 104–115.
- [11] M. Vanden Boom, “Weak cost automata over infinite trees,” Ph.D. dissertation, University of Oxford, 2012.
- [12] D. Kuperberg, “Etude de classes de fonctions de coût régulières,” Ph.D. dissertation, Université Paris Diderot, 2012.
- [13] T. Colcombet, D. Kuperberg, C. Löding, and M. Vanden Boom, “Deciding the weak definability of Büchi definable tree languages,” 2013, submitted.
- [14] T. Colcombet and C. Löding, “The non-deterministic Mostowski hierarchy and distance-parity automata,” in *ICALP (2)*, 2008, pp. 398–409.
- [15] F. Murlak, “Effective topological hierarchies of recognizable tree languages,” Ph.D. dissertation, University of Warsaw, 2008.
- [16] D. Niwiński and I. Walukiewicz, “Relating hierarchies of word and tree automata,” in *STACS*, 1998, pp. 320–331.
- [17] T. F. Urbański, “On deciding if deterministic Rabin language is in Büchi class,” in *ICALP*, 2000, pp. 663–674.
- [18] D. Niwiński and I. Walukiewicz, “Deciding nondeterministic hierarchy of deterministic tree automata,” *Electr. Notes Theor. Comput. Sci.*, vol. 123, pp. 195–208, 2005.
- [19] —, “A gap property of deterministic tree languages,” *Theor. Comput. Sci.*, vol. 1, no. 303, pp. 215–231, 2003.
- [20] J. Duparc, A. Facchini, and F. Murlak, “Definable operations on weakly recognizable sets of trees,” in *FSTTCS*, 2011, pp. 363–374.
- [21] A. Arnold, “The μ -calculus alternation-depth hierarchy is strict on binary trees,” *ITA*, vol. 33, no. 4/5, pp. 329–340, 1999.
- [22] A. Arnold and L. Santocanale, “Ambiguous classes in μ -calculi hierarchies,” *TCS*, vol. 333, no. 1–2, pp. 265–296, 2005.
- [23] M. O. Rabin, “Weakly definable relations and special automata,” in *Proceedings of the Symposium on Mathematical Logic and Foundations of Set Theory*. North-Holland, 1970, pp. 1–23.
- [24] A. Kechris, *Classical descriptive set theory*. New York: Springer-Verlag, 1995.
- [25] W. Wadge, “Reducibility and determinateness in the Baire space,” Ph.D. dissertation, University of California, Berkeley, 1983.
- [26] J. Bradfield, “Simplifying the modal μ -calculus alternation hierarchy,” in *STACS*, 1998, pp. 39–49.
- [27] A. Arnold and D. Niwiński, “Continuous separation of game languages,” *Fundamenta Informaticae*, vol. 81, no. 1-3, pp. 19–28, 2007.
- [28] S. Shelah, “The monadic theory of order,” *The Annals of Mathematics*, vol. 102, no. 3, pp. 379–419, 1975.

Proofs for II. Preliminaries

A. Trivial states

Fact A.1. Every game automaton recognising non-trivial language can be effectively transformed into an equivalent game automaton without trivial states.

Proof: Let \mathcal{A} be a game automaton. We just remove trivial states of \mathcal{A} . If q is trivial then in each transition we replace each subterm of the form (q, d) by \perp or \top (depending on whether $L(\mathcal{A}, q) = \emptyset$ or $L(\mathcal{A}, q) = \text{Tr}_{\mathcal{A}}$) and then propagate it up the transition expression using the standard laws: $(\top \wedge b) = b$, $(\perp \wedge b) = \perp$, $(\top \vee b) = \top$, $(\perp \vee b) = b$. After this step the automaton is still a game automaton but does not contain any trivial states. ■

B. Resolving trees

Fact II.2. Assume that t resolves \mathcal{A} from q_I and $\rho = \rho(\mathcal{A}, t, q_I)$ assigns states to all holes of t . If t has a single hole h , then $t[v := s] \in L(\mathcal{A}, q_I)$ iff $s \in L(\mathcal{A}, \rho(h))$ for all s .

If t has two holes h, h' , whose closest common ancestor u satisfies $\delta_{\mathcal{A}}(\rho(u), t(u)) = (q_L, L) \wedge (q_R, R)$ for some q_L, q_R , then $t[h := s, h' := s'] \in L(\mathcal{A}, q_I)$ iff $s \in L(\mathcal{A}, \rho(h))$ and $s' \in L(\mathcal{A}, \rho(h'))$ for all s, s' ; dually for $(q_L, L) \vee (q_R, R)$. □

Proof: Let us prove the first claim. There are two cases.

- One of the players $P \in \{\exists, \forall\}$ has a winning strategy σ in the game associated to ρ such that node v does not belong to σ . In that case, there exists an ancestor u of v such that the player P owns u and σ moves to from u to ud such that ud is not ancestor of v . In that case σ is a winning strategy for P in the subtree under ud , which contradicts the definition of a resolving tree.
- Whenever σ is a winning strategy for a player $P \in \{\exists, \forall\}$ in the game associated to ρ , then the node corresponding to v belongs to σ . Take any total tree s . If $t[v := s] \in L(\mathcal{A}, q_I)$ then \exists has a winning strategy in the game associated with $\rho(\mathcal{A}, t[v := s], q_I)$. In particular, she can win from the position v in this game. Therefore, by the definition, $s \in L(\mathcal{A}, \rho(q))$. If $t[v := s] \notin L(\mathcal{A}, q_I)$ then the property is symmetrical: \forall has a winning strategy and $s \notin L(\mathcal{A}, \rho(v))$.

For the second claim, it follows easily that in this case the trees $t \upharpoonright_{uL}$, $t \upharpoonright_{uR}$ and the tree obtained by putting a hole in t instead of u , resolve \mathcal{A} from q_L , q_R , and q_I , respectively. We obtain the second claim by applying the first claim three times. ■

Proofs for III. Non-deterministic index problem
--

C. Winning strategies

Proposition III.1. For each game automaton \mathcal{A} and state q_I one can construct effectively a deterministic automaton recognising a language $W_{\mathcal{A}, q_I}^{\exists}$ such that $W_{\mathcal{A}, q_I}^{\exists}$ and $L(\mathcal{A}, q_I)$ have the same non-deterministic index.

Proof: Let us define a deterministic automaton \mathcal{D} recognising $W_{\mathcal{A}, q_I}^{\exists}$. It inherits the state-space and the ranking

function from \mathcal{A} and its transitions are modified as follows: for all $q \in Q$, $a \in A$, $d \in \{L, R\}$, if $\delta_{\mathcal{A}}(q, a) = (q_L, L) \vee (q_R, R)$ for some q_L, q_R , then

$$\delta(q, (a, d)) = (q_L, d), \quad \delta(q, (a, \star)) = \perp,$$

otherwise,

$$\delta(q, (a, d)) = \perp, \quad \delta(q, (a, \star)) = \delta_{\mathcal{A}}(q, a).$$

It is easy to check that $L(\mathcal{D}, q_I) = W_{\mathcal{A}, q_I}^{\exists}$.

What remains is to show that if $L(\mathcal{A}, q_I) = L(\mathcal{B}, q_I^B)$ for some non-deterministic automaton \mathcal{B} of index (i, j) , then we can construct a non-deterministic automaton \mathcal{C} of index at most (i, j) recognising $W_{\mathcal{A}, q_I}^{\exists}$. Following the idea described in the body of the paper, the automaton \mathcal{C} simply checks for the input tree $t \otimes s$ if s encodes a strategy for \exists in the parity game associated with ρ_{t, q_I} and if $t \odot s \in L(\mathcal{B}, q_I^B)$. This can be implemented as follows.

By the Rabin's theorem, for each $q \in Q_{\mathcal{A}}$ there exists a regular tree $t_q \notin L(\mathcal{A}, q)$. We define a sequence of regular languages, then we argue that they can be recognised by non-deterministic automata of indices at most (i, j) :

$$\begin{aligned} \text{St} &= \{t \otimes s : s \text{ is a strategy for } \exists \text{ in } \mathbf{G}_{\rho}(\mathcal{A}, t, q_I)\}, \\ \text{StE} &= \{t \otimes s \otimes t' : t \otimes s \in \text{St} \wedge t \odot s = t'\}, \\ \text{StEW} &= \{t \otimes s \otimes t' \in \text{StE} : t' \in L(\mathcal{B}, q_I^B) = L(\mathcal{A}, q_I)\}, \\ \text{StW} &= \{t \otimes s \in \text{St} : t \odot s \in L(\mathcal{A}, q_I)\}. \end{aligned}$$

Where:

- St corresponds to a safety condition that can be verified by a Comp_0 deterministic automaton,
- StE additionally enforces that the respective subtrees equal t_q , as above it can be checked by a Comp_0 deterministic automaton,
- StEW can be recognised by a product of automata recognising StE and \mathcal{B} — the resulting non-deterministic automaton can be constructed in such a way that its index equals (i, j) ,
- StW is obtained as the projection of StEW onto the first two coordinates, as such can also be recognised by a non-deterministic (i, j) -automaton.

What remains to show is the following equation

$$W_{\mathcal{A}, q_I}^{\exists} = \text{StW} \tag{1}$$

First assume that $t \otimes s \in W_{\mathcal{A}, q_I}^{\exists}$. In that case s encodes a winning strategy σ for \exists in $\mathbf{G}_{\rho}(\mathcal{A}, t, q_I)$. Let $t' = t \odot s$ and $D = \text{dom}(\sigma)$ be the set of vertices belonging to σ . Note that if $v \in D$ then $t(v) = t'(v)$, so also $\rho_{t, q_I}(v) = \rho_{t', q_I}(v)$. Therefore, the strategy σ is also winning in $\mathbf{G}_{\rho}(\mathcal{A}, t', q_I)$. So $t' \in L(\mathcal{A}, q_I)$ what implies that $t \otimes s \otimes t' \in \text{StEW}$ and $t \otimes s \in \text{StW}$.

Now assume that $t \otimes s \in \text{StW}$. Let $t' = t \odot s$ and σ be the strategy for \exists in $\mathbf{G}_{\rho}(\mathcal{A}, t, q_I)$ encoded by s . By the definition of StEW we obtain that $t' \in L(\mathcal{A}, q_I)$ so there exists a winning strategy σ' for \exists in $\mathbf{G}_{\rho}(\mathcal{A}, t', q_I)$. Similarly as above, let D (resp. D') be the set of vertices in σ (resp.

σ'). If $D' \not\subseteq D$ then there exists a minimal (w.r.t. the prefix order) vertex $v \in D' - D$. By the definition of $t \odot s$ we obtain that $t' \upharpoonright_v$ is t_q for $q = \rho(\mathcal{A}, t, q_I)(v)$. Therefore, since $t_q \notin L(\mathcal{A}, q)$, so there is no winning strategy for \exists in $\mathbf{G}_\rho(\mathcal{A}, t_q, q)$ and we obtain a contradiction. Therefore $D' \subseteq D$ and for every $v \in D'$ we have $\rho(\mathcal{A}, t, q_I)(v) = \rho(\mathcal{A}, t', q_I)(v)$, so σ' is also a strategy in $\mathbf{G}_\rho(\mathcal{A}, t', q_I)$. Since strategies form an anti-chain with respect to inclusion, so $\sigma = \sigma'$, $t' \in L(\mathcal{A}, q_I)$, and $t \otimes s \in W_{\mathcal{A}, q_I}^\exists$. ■

Proofs for IV. Alternating index problem

D. Rank reduced automata

Lemma IV.1. *Each game automaton can be effectively transformed into an equivalent rank-reduced game automaton.*

Proof: We iteratively decrease ranks in n -components of \mathcal{A} , for $n \geq 1$. As long as there is an n -component that is not rank-reduced, pick any such n -component, if it is trivial, set all its ranks to $n - 1$, if it is non-trivial but does not contain a state of rank n , decrease all its ranks by 2 (this does not influence the recognised language). After finitely many steps the automaton is rank-reduced. Note that no trivial states are introduced. ■

Proofs for V. Recognisability by game automata

E. Unary profiles give binary profiles

The following fact implies in particular, that if all unlabelled traces have profiles in M , so do labelled traces.

Lemma A.1. *Given the profiles of traces w , waL and waR in M , one can effectively compute the profile of w in M .*

Proof: Let us consider a situation where w has a non-trivial profile $K \subseteq \text{Tr}_A$, and waL and waR have profiles K_L , and K_R . Then, by Definition V.2, wa cannot have profile $*$. Let $a^{-1}K = \{(s, t) \in \text{Tr}_A \times \text{Tr}_A \mid a(s, t) \in K\}$. Is easy to see that wa has a profile $Z \subseteq \text{Tr}_A \times \text{Tr}_A$ if and only if $Z = a^{-1}K$. Thus, it remains to check that $a^{-1}K$ is of one of the forms allowed by Definition V.1.

For a set $U \subseteq X \times Y$, we define the *lower section* of U by $x \in X$ as $U_x = \{y \in Y \mid (x, y) \in U\}$, and the *upper section* of U by $y \in Y$ as $U^y = \{x \in X \mid (x, y) \in U\}$.

Since waL and waR have profiles K_L and K_R , respectively, it follows easily that

- each lower section of $a^{-1}K$ is \emptyset , Tr_A , or K_R ; and
- each upper section of $a^{-1}K$ is \emptyset , Tr_A , or K_L .

The following three sets form a partition of Tr_A :

$$\begin{aligned} X_{\text{Tr}_A} &= \{s_L \mid (a^{-1}K)_{s_L} = \text{Tr}_A\}, \\ X_{K_R} &= \{s_L \mid (a^{-1}K)_{s_L} = K_R\}, \\ X_\emptyset &= \{s_L \mid (a^{-1}K)_{s_L} = \emptyset\}, \end{aligned}$$

and $a^{-1}K = X_{\text{Tr}_A} \times \text{Tr}_A \cup X_{K_R} \times K_R \cup X_\emptyset \times \emptyset$.

Fix $s_R \in K_R$ and $s'_R \in \text{Tr}_A - K_R$. It follows that

$$\begin{aligned} (a^{-1}K)^{s_R} &= X_{\text{Tr}_A} \cup X_{K_R}, \\ (a^{-1}K)^{s'_R} &= X_{\text{Tr}_A}. \end{aligned}$$

Hence, by the initial observation on upper sections, X_{Tr_A} is Tr_A , K_L , or \emptyset , and similarly for $X_{\text{Tr}_A} \cup X_{K_R}$.

Assume K_R is non-trivial. Then for some realisation t of waR , $t^{-1}(M) = K_R$. Let s_L be the subtree of t rooted at awL , and let t_0 be obtained by replacing wa in t with a hole. Then, t_0 is a realisation of w . Since $(a^{-1}t_0^{-1}M)_{s_L} = K_R$, the set $t_0^{-1}M$ is non-trivial. Hence, $t_0^{-1}M = K$ and $(a^{-1}K)_{s_L} = K_R$. This means that $X_{K_R} \neq \emptyset$.

We have three cases, depending on the value of X_{Tr_A} .

- If $X_{\text{Tr}_A} = \text{Tr}_A$, then $X_{K_R} = \emptyset$. Contradiction.
- If $X_{\text{Tr}_A} = \emptyset$, then $(a^{-1}K)^{s'_R} = X_{K_R}$. Since X_{K_R} is non-empty, it is either Tr_A or K_L . In the first case, $a^{-1}K = \text{Tr}_A \times K_R$; in the second case, $a^{-1}K = K_L \times K_R$.
- If $X_{\text{Tr}_A} = K_L$, then either $X_{\text{Tr}_A} \cup X_{K_R} = K_L$ or $X_{\text{Tr}_A} \cup X_{K_R} = \text{Tr}_A$. In the first case, $a^{-1}K = K_L \times \text{Tr}_A$; in the second case, $a^{-1}K = K_L \times K_R \cup \text{Tr}_A \times K_L$.

If K_L is non-trivial the situation is analogous.

If K_L and K_R are trivial, then each section of $a^{-1}K$ is either empty or full. The only two possibilities are $a^{-1}K = \text{Tr}_A \times \text{Tr}_A$ and $a^{-1}K = \emptyset$. ■

F. Languages recognised by game automata are locally game

Lemma V.1. *For each total game automaton \mathcal{B} and a state $q_I \in Q^B$, for each trace w ,*

$$p_{\mathcal{B}, q_I}(w) = p_{L(\mathcal{B}, q_I)}(w).$$

Proof: Let $M = L(\mathcal{B}, q_I)$.

First consider the case of an unlabelled trace w . Let $\rho = \rho(\mathcal{B}, t_w, q_I)$ be the run. If $\rho(w) = *$ then by the definition $p_{\mathcal{B}, q_I}(w) = *$. For every realisation t of w the position w is not accessible in the game $\mathbf{G}_\rho(\mathcal{B}, t, q_I)$ so $t^{-1}(M)$ is either \emptyset or Tr_A and by the definition w has profile $*$ in M .

Now let $\rho(w) = q \in Q^B$. In that case $p_{\mathcal{B}, q_I}(w) = L(\mathcal{B}, q)$. Let t be any realisation of w . Observe that either:

- 1) Player P has a winning σ strategy in $\mathbf{G}_\rho(\mathcal{B}, t, q_I)$ such that $w \notin \sigma$. Then $t^{-1}(M)$ is either \emptyset or Tr_A depending on P .
- 2) Every winning strategy σ of P in $\mathbf{G}_\rho(\mathcal{B}, t, q_I)$ contains w . In that case $t^{-1}(M) = L(\mathcal{B}, q)$ since the following conditions are equivalent:
 - a composition $t[w := s]$ belongs to M ,
 - there exists a winning strategy for \exists in the game $\mathbf{G}_\rho(\mathcal{B}, t[w := s], q_I)$,
 - \exists can win $\mathbf{G}_\rho(\mathcal{B}, t[w := s], q_I)$ from w ,
 - \exists has a winning strategy in the game $\mathbf{G}_\rho(\mathcal{B}, s, q)$,
 - $s \in L(\mathcal{B}, q)$.

Recall that there exists a tree t_0 that realises w and resolves \mathcal{B} from q_I — we plug subtrees in the side holes of t_w accordingly to the states assigned by ρ . By Fact II.2 we obtain that $t_0^{-1}M = L(\mathcal{B}, q)$ so t_0 is a witness that w has profile $L(\mathcal{B}, q_I)$.

For the case when w is a labelled trace we use Lemma A.1 — since every unlabelled trace has a profile, we know that every labelled trace also has a profile. It is then easy to verify that the respective equality holds. ■

G. Locally game is not enough

Proposition A.1. *There exists a regular tree language L such that L is locally game but L cannot be recognised by a game automaton.*

Proof: Consider the alphabet $A = \{a, b\}$. Let $t \in \text{Tr}_A$ be a tree. Let us denote $\text{Cut}(a, t)$ as the subtree of t containing those nodes that are accessible by only letters a from the root of t . A total tree $t \in \text{Tr}_A$ is called *thin* if $\text{Cut}(a, t)$ has only countably many infinite branches. Let Thin contain all thin trees. This language is regular by observation that the following conditions are equivalent:

- 1) a given tree $t \in \text{Tr}_A$ is not thin,
- 2) there exists an embedding of the full binary tree $\{L, R\}^*$ into $\text{Cut}(a, t)$.

Note that every trace w has a profile Z_w in Thin :

- if w contains a letter b then $Z_w = *$,
- otherwise either w is labelled and therefore $Z_w = \text{Thin} \times \text{Thin}$ or w is unlabelled and $Z_w = \text{Thin}$.

This means that Thin is locally game.

Assume that Thin is recognised by a game automaton \mathcal{B} . In that case all transitions of \mathcal{B} have profile $\text{Thin} \times \text{Thin}$ (see Lemma V.1 in Section V), so \mathcal{B} is a deterministic automaton. However, a standard argument shows that Thin is not recognisable by any deterministic automaton. ■

H. Effectiveness of locally game property

We prove Lemma V.2 in the following formulation.

Lemma V.2. *Let M be a regular tree language over an alphabet A . There exists a finite automaton that reads a word w over $A \cup \{L, R\}$ and outputs:*

- *NotTrace if w is not a trace;*
- *NoProfile if w is a trace but w has no profile in M ; and*
- *a finite representation of $p_M(w)$ if w is a trace and has a profile in M .*

The proof of this fact can be easily obtained using the composition method [28]. However, to make the paper self-contained, we give a direct reasoning. The crucial observation is that if a tree t' is put in a hole of a tree t , then the only thing that matters for the acceptance of t is *the type of t'* . For the sake of this proof let us fix a regular tree language M recognised by a non-deterministic tree automaton \mathcal{B} from an initial state $q_I \in Q$.

The type of a total tree $t \in \text{Tr}_A$ is defined as follows:

$$\text{tp}(t) = \{q \in Q : t \in \text{L}(\mathcal{B}, q)\} \subseteq Q.$$

The set of types of all total trees is finite and effective, we denote it by $\text{Tp} \subseteq \mathcal{P}(Q)$. For a set $T \subseteq \text{Tp}$, by $\text{L}(T)$ we denote the language of all total trees t such that $\text{tp}(t) \in T$.

Fact A.2. *Let $t_L, t_R, t'_L, t'_R \in \text{Tr}_A$, let q be a state of \mathcal{B} , and t be a tree with two holes. If*

$$(\text{tp}(t_L), \text{tp}(t_R)) = (\text{tp}(t'_L), \text{tp}(t'_R))$$

then

$$t(t_L, t_R) \in \text{L}(\mathcal{B}, q) \Leftrightarrow t(t'_L, t'_R) \in \text{L}(\mathcal{B}, q).$$

In particular, the type $\text{tp}(t(t_L, t_R))$ does not depend on the choice of representatives t_L, t_R .

By the fact above, we can write $t(\tau_L, \tau_R)$ for the type of $t(t_L, t_R)$ for any t_L, t_R with $(\text{tp}(t_L), \text{tp}(t_R)) = (\tau_L, \tau_R)$.

Our aim is to construct a finite automaton \mathcal{A} that reads a finite word $w \in (A \cup \{L, R\})^*$, checks that w is a trace, and computes a representation of $p_M(w)$, provided that w has profile.

First let us fix

$$\begin{aligned} Q_1 &= \mathcal{P}(\text{Tp}), \\ Q_2 &= \{S \subseteq \text{Tp}^2 \mid \text{L}(S) \text{ is a profile}\} \cup \{\emptyset, \text{Tp}^2\}, \\ Q_E &= \{\text{NotTrace}, \text{NoProfile}\}, \\ Q^A &= Q_1 \cup Q_2 \cup Q_E, \\ q_I^A &= \{T \subseteq \text{Tp} : q_I \in T\} \in Q_1. \end{aligned}$$

Our aim is to define the transition function δ^A in such a way that Lemma A.2, given below, is satisfied. First, for every $T \in Q_1$, $S \in Q_2$, $U \in Q^A$, $a \in A$, $d \in \{L, R\}$, and $l \in A \cup \{L, R\}$ we put

- $\delta^A(T, d) = \text{NotTrace}$,
- $\delta^A(S, a) = \text{NotTrace}$,
- $\delta^A(U, l) = U$.

Second, assume that the current state is $T \in Q_1$ and a letter a is given. We define the successive state $\delta^A(T, a) \in Q_2 \cup \{\text{NoProfile}\}$. Let us define the following set of pairs of types

$$S = \{(\tau_L, \tau_R) : a(\tau_L, \tau_R) \in T\}. \quad (2)$$

Note that, given S , we can decide if $\text{L}(S)$ is a profile and, if it is, we define $\delta^A(T, a) = S$. Otherwise we put $\delta^A(T, a) = \text{NoProfile}$.

Third, assume that the current state is $S \in Q_2$ and a direction d is given. We define the successive state $\delta^A(S, d) \in Q_1$. By the symmetry assume that $d = L$. Consider the following cases:

- if $S = T_L \times \text{Tp} \cup \text{Tp} \times T_R$ for some $T_L, T_R \subseteq \text{Tp}$ then $\delta^A(S, d) = T_L$,
- otherwise, $\delta^A(S, d) = \pi_1(S)$ — the projection of S onto the first coordinate.

In the case $d = R$ we consider T_R instead of T_L and the projection onto the second coordinate of S .

Lemma A.2. *Let w be a word and U be the state of \mathcal{A} after reading w . The following conditions hold*

- 1) *if $U \in (Q_1 \cup Q_2) - \{\emptyset, \text{Tp}, \text{Tp}^2\}$ then w is a trace and $\text{L}(U)$ is the profile of w in M ,*
- 2) *if $U \in \{\emptyset, \text{Tp}, \text{Tp}^2\}$ then w is a trace and has profile $*$ in M ,*
- 3) *if $U = \text{NotTrace}$ then w is not a trace,*
- 4) *if $U = \text{NoProfile}$ then w is a trace but has no profile in M .*

Proof: The first three statements follow easily from the definition of a profile. What remains is to show that if $U = \text{NoProfile}$ then the trace w has no profile in M .

Assume on the contrary and consider a minimal counterexample. Notice, that by the definition, such minimal counterexample is a trace of the form wa for some letter a . Assume that wa has profile Z in M . Let T be the state of \mathcal{A} after reading w and let S be the set computed in (2). If $T \in \{\emptyset, \text{Tp}\}$ then w has profile $*$ in M by Item 2. Then wa has also profile $*$ in M and the state of \mathcal{A} after reading wa belongs to $\{\emptyset, \text{Tp}^2\}$. Assume that $T \notin \{\emptyset, \text{Tp}\}$. By Item 1 we obtain that w has profile $L(T)$ in M .

Let t_0 be a realisation of w such that $t_0^{-1}(M) = L(T)$. Then, by the definition of t^{-1} we obtain that

$$(t_0[w := a])^{-1}(M) = a^{-1}(L(T)) = Z. \quad (3)$$

It is enough to show that $L(S) = Z$ and thus $S \in Q_2$ and $R = S \neq \text{NoProfile}$.

Note that the following conditions are equivalent for a pair of trees $t_L, t_R \in \text{Tr}_A$:

- 1) $(t_L, t_R) \in Z$,
- 2) $a(t_L, t_R) \in L(T)$ (by (3)),
- 3) $a(\text{tp}(t_L), \text{tp}(t_R)) \in T$ (by the definition of $a(\tau_L, \tau_R)$),
- 4) $(\text{tp}(t_L), \text{tp}(t_R)) \in S$ (by (2)),
- 5) $t_L, t_R \in L(S)$ (by the definition of $L(S)$).

■

I. Game automata accept correct infinite traces

Lemma V.3. \mathcal{D} accepts π from q_I if and only if π is $L(\mathcal{D}, q_I)$ -correct.

Proof: Let us denote $L(\mathcal{D}, q_I)$ as M' . Let $\pi \in (A \times \{L, R\})^\omega$ be an infinite trace. First assume that π is M' -correct. Let $t' \in M'$ be a tree witnessing it. Recall, that t' is obtained by putting some total trees in holes of t_π . Since $t' \in M'$ so there exists a winning strategy σ for \exists in $\mathbf{G}_\rho(\mathcal{D}, t', q_I^D)$. Since, whenever \exists could make a choice to leave the branch π , the respective subtree in t' is losing for her, so the whole branch π must be contained in the strategy σ . Therefore, \mathcal{D} must accept π from q_I^D .

Now assume that \mathcal{D} accepts π from q_I^D . Let t' be some realisation of π . Consider σ to be the strategy of \exists in $\mathbf{G}_\rho(\mathcal{D}, t', q_I^D)$ defined as follows:

- in all the nodes along π follow this branch,
- whenever \forall selects to go off the branch π , use some winning strategy in the respective subtree (it exists by the construction).

By the definition, σ is a winning strategy for \exists in the game $\mathbf{G}_\rho(\mathcal{D}, t', q_I^D)$. Therefore, $t' \in M'$ so t' is a witness that π is an M' -correct branch. ■

J. Effectiveness of the languages of M -correct infinite traces

Lemma V.4. For every regular tree language M the language of M -correct branches C is ω -regular. Moreover, a deterministic parity automaton recognising C can be constructed from a representation of M .

Proof: Let \mathcal{B} be a non-deterministic automaton recognising the given regular tree language M from a state q_I . We show how to express in monadic second-order logic over

ω the fact, that a given ω -word π is an M -correct branch. By the results by Büchi and McNaughton such a formula φ can be effectively translated into a deterministic parity ω -word automaton \mathcal{D} .

Similarly as in the proof of Lemma V.1, we use compositional tools. In particular, by $\text{Tp} \subseteq \mathcal{P}(Q^B)$ we denote the set of all types of total trees with respect to the automaton \mathcal{B} .

Intuitively, the formula φ guesses the types of the total subtrees that need to be put in the side holes of t_π to obtain a tree t that resolves M up to π . Basing on this guessed types φ can verify that $t \in M$.

Recall that an infinite trace π is defined as a word in the language $(A \cdot \{L, R\})^\omega$. A witness for the existence of t will be an infinite word over the alphabet $A \cup (\{L, R\} \times \text{Tp})$ denoted $\hat{\pi}$ and called *enrichment of π by types*. We require that every even position of $\hat{\pi}$ belongs to A and every odd position belongs to $\{L, R\} \times \text{Tp}$.

$$\hat{\pi} = a_0 \cdot (d_0, \tau_0) \cdot a_1 \cdot (d_1, \tau_1) \cdot \dots$$

If $w = a_0 \cdot d_0 \cdot a_1 \cdot d_1 \cdot \dots \cdot d_{n-1} \cdot a_n$ is a labelled trace that is a prefix of π we say that d_n is the *final direction of w* and τ_n is the *final type of w* .

Let the formula ψ_R express that for every trace $w \prec \pi$ with final direction d , final type τ , and such that S is the state of \mathcal{A} (see Lemma V.1) after reading w , the following conditions are satisfied:

- if $S = S_L \times \text{Tp} \cup \text{Tp} \times S_R$ then $\tau \notin S_d$,
- otherwise $\tau \in \pi_i(S)$ where $i = 1, 2$ for $d = L, R$ respectively.

Note that every $\hat{\pi}$ that is enrichment of π by types induces a total tree t_π over the alphabet A where in the n 'th side hole of π we put some total tree of type τ_n . Note also that $\hat{\pi} \models \psi_R$ if and only if t_π resolves M up to π . In particular the exact subtrees we put into side holes of t_π are irrelevant, we only need to take care of their types.

What remains is to express in MSO logic on $\hat{\pi}$ that $t_\pi \in M$. For this we say that there exists an infinite word ρ coding a run the non-deterministic automaton \mathcal{B} on t_π . Formally a *word coding a run* ρ is defined as a word over the alphabet $Q^B \cup \{*\} \cup D^B$ where D^B is the set of all deterministic transitions appearing in the transitions of \mathcal{B} . The elements of $Q^B \cup \{*\}$ are supposed to appear on even positions of ρ and elements of D^B are supposed to appear on odd positions of ρ :

$$\rho = q_0 \cdot b_0 \cdot q_1 \cdot b_1 \cdot \dots$$

Let the formula ψ_M express the following facts about the combination of words $\hat{\pi} \otimes \rho$ in the language

$$[A \times (Q^B \cup \{*\}) \cdot ((\{L, R\} \times \text{Tp}) \times D^B)]^\omega$$

- the state q_0 equals q_I ,
- for every n the transition b_n is one of the deterministic transitions appearing in $\delta^B(q_n, a_n)$,
- for every n the state assigned to \bar{d}_n by b_n (if any) belongs to τ_n ,

- for every n the state assigned to d_n by b_n (if any) equals q_{n+1} , if there is no such state then $q_{n+1} = *$,
- either from some point on $b_n = *$ or the parity condition is satisfied by the sequence of states q_0, q_1, \dots

Note, that $\hat{\pi} \otimes \rho \models \psi_M$ if and only if ρ encodes an accepting run of \mathcal{B} on t_π that assigns to the n 'th hole of t_π a state belonging to τ_n . Therefore, $\hat{\pi} \otimes \rho \models \psi_M$ if and only if the run encoded by ρ can be extended to an accepting run of \mathcal{B} on t_π^\wedge .

Let φ express for a given infinite trace π that there exists an enrichment of π by types τ_n and an encoding of run ρ such that $\hat{\pi} \models \psi_R$ and $\hat{\pi} \otimes \rho \models \psi_M$. Note that $\pi \models \varphi$ if and only if there exists a tree $t = t_\pi^\wedge \in M$ that realises π and that resolves M up to π . ■