Arbitrary-arity Tree Automata and QCTL

François Laroussinie   Nicolas Markey

October 28, 2024

Abstract

We introduce a new class of automata (which we coin EU-automata) running on infinite trees of arbitrary (finite) arity. We develop and study several algorithms to perform classical operations (union, intersection, complement, projection, alternation removal) for those automata, and precisely characterise their complexities. We also develop algorithms for solving membership and emptiness for the languages of trees accepted by EU-automata.

We then use EU-automata to obtain several algorithmic and expressiveness results for the temporal logic QCTL (which extends CTL with quantification over atomic propositions) and for MSO. On the one hand, we obtain decision procedures with optimal complexity for QCTL satisfiability and model checking; on the other hand, we obtain an algorithm for translating any QCTL formula with k quantifier alternations to formulas with at most one quantifier alternation, at the expense of a $(k+1)$-exponential blow-up in the size of the formulas. Using the same techniques, we prove that any MSO formula can be translated into a formula with at most four quantifier alternations (and only two second-order-quantifier alternations), again with a $(k+1)$-exponential blow-up in the size of the formula.

1   Introduction

Logics and automata.   The very tight links between logics and automata on infinite words and trees date back to the early 1960's with the seminal works of Büchi, Elgot, Trakhtenbrot, McNaughton and Rabin [Büc62, Elg61, Tra62, McN66, Rab69]. These early results were mainly concerned with the Monadic Second-Order Logic (MSO), and have been further extended to many other logical formalisms such as modal, temporal and fix-point logics [SVW85, VW86b, BVW94, JW95, Wil01]. Those tight links are embodied as translations back and forth between various logical languages and corresponding classes of automata; translations from logics to automata have allowed to derive efficient algorithms for satisfiability or model checking on the one hand [VW86a, EJ91, BVW94]; with additional translations from automata to logics, we get effective ways for proving expressiveness or succinctness results for some of those logics [Wal96, Wil99, LMS02, KV03, Zan12]. In this paper, we investigate such links between Quantified CTL (QCTL) [Kup95, KMTV00, Fre01, DLM12] and symmetric tree automata [Wal96, Wil99, KV03], and derive algorithmic and expressiveness results for QCTL and its fragments.

QCTL. QCTL extends the classical temporal logic CTL with quantification on atomic propositions. For instance, formula $\exists p.\phi$, where ϕ is a CTL formula, states that there exists a labelling of the model under scrutiny with atomic proposition p under which ϕ holds. QCTL is (much) more expressive than CTL: as an example, formula

$$\exists p.\ (\mathbf{EF}(\phi \wedge p) \wedge \mathbf{EF}(\phi \wedge \neg p))$$

expresses the fact that there are at least two reachable states where ϕ holds. The extension of CTL with *existential* quantification was first studied in [ES84, Kup95]: contrary to CTL, the resulting logic (only allowing formulas in prenex form), which we call EQ^{1}CTL hereafter, is sensitive to unwinding and duplication of transitions; the semantics thus depends on whether the extra labelling refers to the Kripke structure under scrutiny, or on its computation tree. Our sample formula above expresses that there are at least two *different* reachable control states satisfying ϕ in the former case (which we call the *structure semantics*), while it only requires that two different paths lead to some ϕ-states (possibly two copies of the same control state) in the latter semantics (called the tree semantics hereafter).

Universal quantification on atomic propositions can also be added: AQ^{1}CTL is the logic obtained from CTL by adding universal quantification (in prenex form). Mixing existential and universal quantification defines an infinite hierarchy of temporal logics, which we name EQ^{k}CTL and AQ^{k}CTL, where k is the number of quantifier alternations allowed in formulas (still assuming prenex form). QCTL allows unrestricted use of both existential and universal quantifications, and thus contains EQ^{k}CTL and AQ^{k}CTL for all $k \geq 0$. It turns out that QCTL is as expressive as MSO [LM14].

In this paper, we present several results for QCTL with the tree semantics. In particular, we show that any QCTL formula with k quantifier alternations can be translated in EQ^{2}CTL, with a $(k+1)$-exponential blow-up in the size of the formula. Such a result is known to exist also in MSO on trees [Rab69, Tho97]: any MSO formula can be expressed with two alternations of *second-order* quantifiers. While MSO is known to be as expressive as QCTL, this does not directly entail our result because *first-order* quantifiers in MSO involve extra propositional quantifiers when translated in QCTL. The key point of our results is the introduction of a new class of tree automata that are particularly well-suited for characterising models of a QCTL formula, but also of QCTL* or MSO.

Tree automata. We use (top-down[1]) tree-automata techniques to study QCTL. Several results already exist on this topic [ES84, Kup95, LM14], but they all rely on fixed-arity tree automata.

The limitation has several drawbacks. When dealing with model checking, it implies that the compilation of the formula being checked into a tree automaton depends on the (size of the) structure under scrutiny. In particular, it cannot be used directly for evaluating the *program complexity* of QCTL model checking, as it requires bounding the size of the structures that the automaton can handle. An indirect solution to this problem is given in [LM14], by replacing nodes

---

[1]There are several families of tree automata: top-down tree automata explore (finite or infinite) trees starting from the root; bottom-up tree automata explore finite trees from the leaves up to the root; tree-walking automata are a kind of two-way automata for trees. We refer to [CDG$^+$08, Boj08] for more details.

of arbitrary (finite) arity with binary-tree gadgets. A similar problem occurs when dealing with satisfiability: one has to use additional results to ensure that looking for a structure with bounded size is sufficient. More importantly, when deriving expressiveness results, using fixed-arity tree automata again restricts the results to trees or structures with bounded branching.

In order to handle trees of arbitrary branching degree, tree automata must have a *symbolic* way of expressing transitions, with a finite representation that can cope with any arity. We highlight two existing approaches:

- Janin and Walukiewicz introduce MSO-automata [JW95, Wal96], in which transitions are defined as first-order formulas: quantification is over the successors of the current node, and predicates indicate in which states of the automaton those successors must be explored. These automata are shown to be as expressive as MSO, and several expressiveness results have been obtained from this construction [JW95, Wal96, Wal02, Zan12]. However, to the best of our knowledge, the exact complexity of the operations for manipulating those automata has not been studied, so that only qualitative expressiveness results can be obtained, and no bounds on the size and complexity of the translations can be derived without a more careful study.

- Wilke introduces $\{\Box, \Diamond\}$-automata [Wil99], which are alternating tree automata with $\Box q$ and $\Diamond q$ as basic blocs for expressing transitions: the former requires that all successors be explored in state q, while the latter asks that some successor be explored in state q. Any CTL formula can be turned into an equivalent $\{\Box, \Diamond\}$-automata of linear size; this is used to prove that the extension CTL^+ of CTL is exponentially more succinct that CTL. However, $\{\Box, \Diamond\}$-automata are not expressive enough to capture MSO or QCTL.

Our contribution. In this paper, we define a new class of arbitrary-arity alternating tree automata, develop effective operations for their manipulation, and study the complexity of those operations and the size of the resulting automata. Instead of using pairs (k, q) in the transition function to specify that the k-th successor of the current node has to be accepted by the automaton in state q, transitions of our automata are defined with pairs $\langle E; U \rangle$, where E is a multiset of states that have to occur among the set of states involved in the exploration of the successors of the current node, while U is a set of states indicating which states are allowed for exploring successor nodes that are not explored by states of E. For example, $\langle E = \{\!\{q, q, q'\}\!\}; U = \{q''\} \rangle$ requires the presence of at least three successors nodes; two successors will be explored in state q, one in state q', and the remaining ones (if any) in state q''. We name those automata EU-automata[2].

It is not hard to prove that such automata are closed under conjunction and disjunction, thanks to alternation. Closure under negation is harder to prove: while $\Box$ and $\Diamond$ are dual to each other, which provides an easy complementation procedure for $\{\Box, \Diamond\}$-automata, there is no obvious way of expressing the

---

[2]In [KV03], Kupferman and Vardi define another variant of arbitrary-arity alternating tree automata in which transitions are based on pairs (U, E). Those automata are equivalent to Wilke's $\{\Box, \Diamond\}$-automata. We give an overview of those automata in our Section 2.6 on related work.

negation of EU-pairs in terms of EU-pairs. We develop such a translation, and obtain an exponential complementation procedure for EU-automata.

Non-alternating EU-automata are also closed under projection, which is the operation we need to encode quantification over atomic propositions of QCTL, and first- and second-order quantification in MSO. Finally we prove that any alternating EU-automaton can be turned into an equivalent non-alternating EU-automaton. For this operation, we adapt the simulation procedure developed in [Wal96, Zan12] to our setting, and evaluate its exact complexity.

Putting all the pieces together, we prove that any QCTL formula φ can be turned into an equivalent EU-automaton $\mathcal{A}_\varphi$. The size of the automaton is k-exponential in the size of φ, where k is the number of quantifier alternations in φ. This construction then yields optimal algorithms for model-checking and satisfiability for QCTL. Conversely, we prove that acceptance by any EU-automaton can be expressed as an $\mathsf{EQ^2CTL}$ formula. We obtain similar results for MSO. Therefore EU-automata, QCTL (and $\mathsf{EQ^2CTL}$), and MSO (even when restricted to two second-order quantifier alternations) all characterise exactly the same tree languages.

2 Definitions

2.1 Sets and multisets

Let $\mathcal{S}$ be a countable set. A *multiset* over $\mathcal{S}$ is a mapping $\mu\colon \mathcal{S} \to \mathbb{N}$. Sets are seen as special cases of multisets taking values in $\{0,1\}$. We use double-brace notation to distinguish between sets and multisets: $\{a,a,a\}$ is the same as the set $\{a\}$ with one element, while $\{\!\{a,a,a\}\!\}$ is the three-element multiset $a \mapsto 3$. The *empty multiset* is the multiset mapping all elements of $\mathcal{S}$ to zero; we denote it with $\varnothing$.

The *support of a multiset* μ is the set $\mathsf{supp}(\mu) = \{s \in \mathcal{S} \mid \mu(s) > 0\}$. We write $s \in \mu$ for $s \in \mathsf{supp}(\mu)$. The *size* $|\mu|$ of μ is the sum $\sum_{s \in \mathcal{S}} \mu(s)$; the multiset μ is finite whenever $|\mu|$ is. For two multisets μ and μ', we write $\mu \sqsubseteq \mu'$, and say that μ is a submultiset of μ', whenever $\mu(s) \leq \mu'(s)$ for all $s \in \mathcal{S}$. This defines a partial ordering over multisets. We write $\mu \sqsubset \mu'$ when $\mu \sqsubseteq \mu'$ and $\mu \neq \mu'$. We define the following operations on multisets:

$$\mu \uplus \mu'\colon s \in \mathcal{S} \mapsto \mu(s) + \mu'(s) \quad \text{and} \quad \mu' \setminus \mu\colon s \in \mathcal{S} \mapsto \max(0, \mu'(s) - \mu(s))$$

Fix a second countable set $\mathcal{S}'$. For any $c = (s,s') \in \mathcal{S} \times \mathcal{S}'$, we define $\mathsf{proj}_1(c) = s$ and $\mathsf{proj}_2(c) = s'$.

2.2 Markings

Let $\mathcal{S}$ and $\mathcal{S}'$ be two countable sets. A *marking* of $\mathcal{S}'$ by $\mathcal{S}$ is a mapping $\nu\colon \mathcal{S}' \to 2^{\mathcal{S}} \setminus \{\varnothing\}$ decorating each element of $\mathcal{S}'$ with a (non-empty) subset of $\mathcal{S}$. A marking ν is a *submarking* of a marking ν', denoted $\nu \sqsubseteq \nu'$, whenever $\nu(s') \subseteq \nu'(s')$ for all $s' \in \mathcal{S}'$.

A marking ν is *unitary* when $|\nu(s')| = 1$ for all $s' \in \mathcal{S}'$; unitary markings can be seen as mappings from $\mathcal{S}'$ to $\mathcal{S}$. For a unitary marking ν and a subset T of $\mathcal{S}'$, we write $\nu(T)$ for the multiset μ over $\mathcal{S}$ defined as $\mu(s) = \#\{t \in T \mid$

$\nu(t) = s\}$, which we may also write as $\{\!\{\nu(t) \mid t \in T\}\!\}$. We write $\mathsf{img}(\nu)$ for the multiset $\nu(\mathcal{S}')$.

2.3 Words and trees

Let Σ be a finite set. A *word* over Σ (or Σ-*word*) is a sequence $w = (w_i)_{0 \leq i < k}$ of elements of Σ, with $k \in \mathbb{N} \cup \{+\infty\}$. The *length* (or *size*) of w, denoted with $|w|$, is k. We write Σ^* for the set of finite words over Σ, and Σ^∞ for the set of infinite words over Σ. We write ε for the *empty word* (the only word of size 0).

For a word $w = (w_i)_{0 \leq i < k}$ of length $k \in \mathbb{N} \cup \{+\infty\}$, and an integer $j \in \mathbb{N}$ such that $0 \leq j \leq k$, the *prefix* of w of length j (also referred to as its j-th prefix) is the word $w_{\langle j \rangle} = (w_i)_{0 \leq i < j}$. For $0 \leq j < k+1$, the j-th *suffix* of w is the word $w_{[j\rangle} = (w_{j+i})_{0 \leq i < k-j}$. Given a finite word w and a (possibly infinite) word w', their concatenation $w \cdot w'$ is the word x whose $|w|$-th prefix is w and whose $|w|$-th suffix is w'. We identify words of length 1 with their constituent letter, and write $\mathsf{first}(w)$ for the first prefix of w, and, in case w is finite, $\mathsf{last}(w)$ for its $(|w| - 1)$-th suffix.

Let $\mathcal{D}$ be a finite set. A *tree structure* over $\mathcal{D}$ (or $\mathcal{D}$-*tree*) is a subset $t \subseteq \mathcal{D}^*$ that is closed under prefix. In particular, any non-empty tree contains the empty word ε, which is called its *root* (and sometimes denoted with ε_t when we need to distinguish between the roots of different trees). The elements of a tree are called nodes. A node m in t is a *successor* of a node n if $m = n \cdot d$ for some $d \in \mathcal{D}$. In that case, n is the (unique) predecessor of m. We write $\mathsf{succ}(n)$ for the set of successors of node n. Notice that in a $\mathcal{D}$-tree, any node may have at most $|\mathcal{D}|$ successors; this integer $|\mathcal{D}|$ is the *arity* of the tree; as a special case, a *binary tree* is a $\mathcal{D}$-tree with $|\mathcal{D}| = 2$. Notice that not all nodes have to have $|\mathcal{D}|$ successors in a tree of arity $|\mathcal{D}|$. In particular, any tree may contain *leaves*, which are node with no successors. A *branch* of a tree is a (finite of infinite) sequence $b = (n_i)_{0 \leq i < k}$ of nodes of the tree such that $n_0 = \varepsilon$, n_{i+1} is a successor of n_i for all $0 \leq i < k-1$, and if k is finite, n_{k-1} is a leaf. The value of $k \in \mathbb{N} \cup \{+\infty\}$ is the length of b, denoted with $|b|$. A tree is finite when it contains only finitely many nodes. By Kőnig's lemma, since $\mathcal{D}$ is finite, a tree is finite if, and only if, all its branches are finite.

A Σ-*labelled* $\mathcal{D}$-*tree* is a pair $\mathcal{T} = (t, l)$ where t is a $\mathcal{D}$-tree and $l \colon t \to \Sigma$ labels each node of t with a letter in Σ. With any branch $b = (n_i)_{0 \leq i < k}$ of t in a Σ-labelled $\mathcal{D}$-tree $\mathcal{T} = (t, l)$, we associate its *word* $w(b)$ over Σ as the word $(w_i)_{0 \leq i < k}$ defined as $w_i = l(n_i)$ for all $0 \leq i < k$.

Let AP be a finite set of atomic propositions. A *Kripke structure* over AP is a tuple $\mathcal{K} = (V, E, \ell)$ where V is a finite set of vertices, $E \subseteq V \times V$ is a set of edges (requiring that for any $v \in V$, there exists $v' \in V$ s.t. $(v, v') \in E$), and $\ell \colon V \to 2^{\mathsf{AP}}$ is a labelling function.

A *path* in a Kripke structure is a finite or infinite word w over V such that $(w_i, w_{i+1}) \in E$ for all $i < |w|$. We write $\mathsf{Path}^*_{\mathcal{K}}$ for the sets of finite paths of $\mathcal{K}$. Given a vertex $v \in V$, the *computation tree* of $\mathcal{K}$ from v is the 2^{AP}-labelled V-tree $\mathcal{T}_{\mathcal{K},v} = (T_{\mathcal{K},v}, \hat{\ell})$ with $T_{\mathcal{K},v} = \{w \in V^* \mid v \cdot w \in \mathsf{Path}^*_{\mathcal{K}}\}$ and $\hat{\ell}(v \cdot w) = \ell(\mathsf{last}(v \cdot w))$. Notice that two nodes w and w' of $\mathcal{T}_{\mathcal{K},v}$ for which $\mathsf{last}(w) = \mathsf{last}(w')$ give rise to the same subtrees. A tree is said *regular* when it corresponds to the computation tree of some finite Kripke structure.

2.4 Automata over trees of arbitrary arity

In this section, we introduce our automata running over trees of arbitrary arity. The core element of their transition functions are EU-pairs and EU-constraints.

Let $\mathcal{S}$ be a countable set. An *EU-pair* over $\mathcal{S}$ is a pair $\langle E;U \rangle \in \mathbb{N}^{\mathcal{S}} \times 2^{\mathcal{S}}$, where E is a multiset over $\mathcal{S}$ and U is a subset of $\mathcal{S}$. A multiset μ over $\mathcal{S}$ satisfies the EU-pair $\langle E;U \rangle$, denoted $\mu \models \langle E;U \rangle$, whenever $E \sqsubseteq \mu$ and $\mathsf{supp}(\mu \setminus E) \subseteq U$. We write $\mathsf{EU}(\mathcal{S}) = \mathbb{N}^{\mathcal{S}} \times 2^{\mathcal{S}}$ for the set of EU-pairs over $\mathcal{S}$.

Example 1. *Consider a set* $\mathcal{S} = \{q_1, q_2, q_3, q_4\}$. *The EU-pair* $\langle q_1 \mapsto 3, q_2 \mapsto 1; \{q_1, q_3\} \rangle$ *characterises all multisets containing* at least *three occurrences of* q_1, *exactly* one occurrence of q_2, *an arbitrary number of occurrences of* q_3, *and no occurrences of* q_4. ◁

For a finite set $\mathcal{B}$ of boolean variables, we write $\mathsf{PBF}(\mathcal{B})$ for the set of *positive boolean combinations* over $\mathcal{B}$:

$$\mathsf{PBF}(\mathcal{B}) \ni \phi ::= \top \mid \bot \mid v \mid \phi \wedge \phi \mid \phi \vee \phi$$

where v ranges over $\mathcal{B}$. The set of *disjunctions* over $\mathcal{B}$ is the subset of $\mathsf{PBF}(\mathcal{B})$ defined as

$$\mathsf{DBF}(\mathcal{B}) \ni \phi ::= \top \mid \bot \mid v \mid \phi \vee \phi$$

where again v ranges over $\mathcal{B}$. That a subset $V \subseteq \mathcal{B}$ satisfies a formula ϕ of $\mathsf{PBF}(\mathcal{B})$, denoted $V \models \phi$, is defined inductively in the natural way: it is always true when ϕ is $\top$ and always false when ϕ is $\bot$, and

$$V \models v \iff v \in V$$
$$V \models \phi_1 \vee \phi_2 \iff V \models \phi_1 \text{ or } V \models \phi_2$$
$$V \models \phi_1 \wedge \phi_2 \iff V \models \phi_1 \text{ and } V \models \phi_2$$

An *EU-constraint* is a positive boolean formula over EU-pairs. We can now define our class of automata:

Definition 1. *Let* Σ *be a finite alphabet. An* alternating EU tree automaton *(AEUTA for short) over* Σ *is a 4-tuple* $\mathcal{A} = (Q, q_{init}, \delta, \Omega)$ *with*

- Q *is a finite set of states, and* $q_{init} \in Q$ *is the initial state;*

- $\delta \colon Q \times \Sigma \to \mathsf{PBF}(\mathsf{EU}(Q))$ *is the set of transitions;*

- $\Omega \colon Q^* \cup Q^{\infty} \to \{0, 1\}$ *is an acceptance condition.*

An **AEUTA** *is* non-alternating *(and is thus an* EU tree automaton, EUTA *for short) if* δ *takes values in* $\mathsf{DBF}(\mathsf{EU}(Q))$.

The following relation will be the central relation for defining the semantics of AEUTA: it will be used to lift the satisfaction relation of EU-constraints to execution trees.

Definition 2. *Let* $\mathcal{S}$ *and* $\mathcal{S}'$ *be two countable sets. Let* $\langle E;U \rangle$ *be an EU-pair over* $\mathcal{S}$, *and* ν *be a marking of* $\mathcal{S}'$ *by* $\mathcal{S}$. *Then* ν *satisfies* $\langle E;U \rangle$, *denoted* $\nu \Vvdash \langle E;U \rangle$, *if there exists a unitary marking* $\nu' \sqsubseteq \nu$, *such that* $\mathsf{img}(\nu') \models \langle E;U \rangle$.
This definition extends to EU-constraints inductively as follows:

- $\nu \models \phi_1 \vee \phi_2$ *if, and only if, $\nu \models \phi_1$ or $\nu \models \phi_2$;*

- $\nu \models \phi_1 \wedge \phi_2$ *if, and only if, $\nu \models \phi_1$ and $\nu \models \phi_2$.*

Notice that $\nu \models \phi$ is *not* equivalent to having an unitary marking $\nu' \sqsubseteq \nu$ satisfy ϕ, since different submarkings ν' may be needed for different EU-pairs. However, the equivalence holds if ϕ is a disjunction of EU-pairs, since in that case a single EU-pair has to be fulfilled.

We can now define the notion of execution tree of an AEUTA:

Definition 3. *Let $\mathcal{A} = (Q, q_{init}, \delta, \Omega)$ be an AEUTA over Σ and $\mathcal{T} = (t, l)$ be a Σ-labelled $\mathcal{D}$-tree, for some finite set $\mathcal{D}$. An* execution tree *of $\mathcal{A}$ over $\mathcal{T}$ is a $(t \times Q)$-labelled $(\mathcal{D} \times Q)$-tree $\mathcal{U} = (u, \ell)$ such that*

- *the root ε_u of u is labelled with ε_t and q_{init} (formally, $\ell(\varepsilon_u) = (\varepsilon_t, q_{init})$);*

- *any non-root node $n_u = (d_i, q_i)_{0 \le i < |n_u|}$ of u is labelled with $\ell(n_u) = ((d_i)_{0 \le i < |n_u|}, q_{|n_u|-1})$;*

- *for any node n_u of the form $(d_i, q_i)_{0 \le i < |n_u|}$ of u with $\ell(n_u) = (m_t, q)$, letting ν_{n_u} be the marking of $\mathsf{succ}(m_t)$ by Q such that $\nu_{n_u}(m_t \cdot d) = \{q' \in Q \mid n_u \cdot (d, q') \in \mathsf{succ}(n_u)\}$, we have $\nu_{n_u} \models \delta(q, l(m_t))$. We name this marking ν_{n_u} the* marking of $\mathsf{succ}(m_t)$ induced by $\mathcal{U}$.*

The tree $\mathcal{T}$ is accepted *by $\mathcal{A}$ if there exists an execution tree $\mathcal{U}$ of $\mathcal{A}$ over $\mathcal{T}$ such that all branches are* accepting, *i.e., for any branch $b = (b_i)_{0 \le i < |b|}$ in $\mathcal{U}$, it holds $\Omega((\mathsf{proj}_2(\ell(b_i)))_{0 \le i < |b|}) = 1$. Such an execution tree is said to be* accepting. *The* language *of $\mathcal{A}$, denoted by $\mathcal{L}(\mathcal{A})$, is the set of all trees accepted by $\mathcal{A}$. The tree $\mathcal{T}$ is* rejected *if it is not accepted, i.e., if there are no accepting execution trees of $\mathcal{A}$ over $\mathcal{T}$.*

Notice that if a marking ν satisfies some EU-constraint, then any marking ν' containing ν also does. Similarly, any execution tree can be extended with extra subtrees, provided that all their branches are accepting. An execution tree is said *minimal* if it does not contain dispensable subtrees. We may always consider that the execution trees we consider are minimal.

Let us illustrate execution trees with an example:

Example 2. *Consider a node m of some input tree $\mathcal{T}$, with three successors $m \cdot d_1$, $m \cdot d_2$ and $m \cdot d_3$ (as depicted on Fig. 1); assume that this node m is labelled with some letter σ. Consider an AEUTA $\mathcal{A}$ visiting node m in state q, giving rise to a node n in an execution tree $\mathcal{U}$ with $\ell(n) = (m, q)$. The successors of n in the execution tree give rise to the marking ν_n such that $\nu_n(m \cdot d_1) = \{q_1, q_3\}$, $\nu_n(m \cdot d_2) = \{q_2, q_4\}$, and $\nu_n(m \cdot d_3) = \{q_1, q_4\}$, as depicted to the left of Fig. 1.*

Assume that $\delta(q, \sigma)$ is satisfied by the following set W of EU-pairs:

$$W = \left\{ \underbrace{\langle q_1 \mapsto 2; \{q_2\}\rangle}_{\langle E_1; U_1 \rangle}, \underbrace{\langle q_1 \mapsto 1, q_2 \mapsto 1, q_3 \mapsto 1; \{q_4\}\rangle}_{\langle E_2; U_2 \rangle}, \underbrace{\langle q_3 \mapsto 1; \{q_4\}\rangle}_{\langle E_3; U_3 \rangle} \right\}.$$

Figure 1 displays a possible set of successors $\mathsf{succ}(n)$ of n in the execution tree $\mathcal{U}$. Using the following three submarkings ν_i of ν_n, we are able to fulfill all three
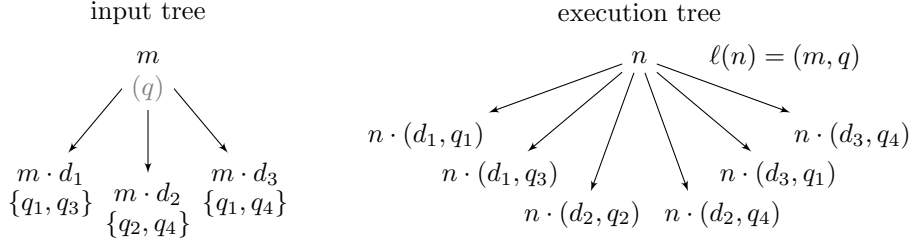
input tree   execution tree

Figure 1: Example of a transition of the automaton when exploring node m of an input tree $\mathcal{T}$ in state q: node m has three successors $m \cdot d_1$, $m \cdot d_2$ and $m \cdot d_3$ in $\mathcal{T}$; if the automaton explores node $m \cdot d_1$ in states q_1 and q_3, node $m \cdot d_2$ in states q_2 and q_4, and node $m \cdot d_3$ in states q_1 and q_4 (as in the partial execution tree on the right), we get the marking of the successors of node m as given on the left, from which we can extract submarkings satisfying the EU-constraint W of Example 2.

EU-pairs of W:

$$\nu_1 : n \cdot d_1 \mapsto q_1 \qquad \nu_2 : n \cdot d_1 \mapsto q_3 \qquad \nu_3 : n \cdot d_1 \mapsto q_3$$
$$n \cdot d_2 \mapsto q_2 \qquad\qquad n \cdot d_2 \mapsto q_2 \qquad\qquad n \cdot d_2 \mapsto q_4$$
$$n \cdot d_3 \mapsto q_1 \qquad\qquad n \cdot d_3 \mapsto q_1 \qquad\qquad n \cdot d_3 \mapsto q_4 \qquad \triangleleft$$

Notice that $\mathcal{D}$ is not constrained by the definition of AEUTAs, so that AEUTAs may accept trees of arbitrary (finite) arity. However, the multisets in the "existential part" of EU-pairs can be used to impose a lower bound on the number of successors for the EU-pair to be satisfied, and an upper bound can be imposed by letting the universal part be empty. We develop such examples in Section 2.5 below.

Remark 1. *We do not define a notion of being* deterministic *for EUTA: even if the transition function δ returns a single EU-pair for each pair (q, σ), there may be many different valid execution trees, since there may be different ways of satisfying a single EU-pair.*

Notice that for (non-alternating) EUTAs, the definition of execution trees can be simplified: in non-alternating automata, $\delta(q, \sigma)$ is a disjunction of EU-pairs, and a single unitary marking, hence a single state for each successor node, is sufficient to fulfill it.

It follows that, for each input tree $\mathcal{T} = (t, l)$ accepted by an EUTA, there is a accepting execution tree of the form $\mathcal{U} = (t, \ell)$, having the same underlying tree structure t. Any node n of $\mathcal{U}$ then has $\ell(n) = (n, q)$ for some $q \in Q$. ◀

In the sequel, we mainly consider *parity acceptance*: Ω can then be defined through a mapping $\omega : Q \to \mathbb{N}$. The integer $\omega(q)$ is called the *priority* of state q. Each branch $b = (b_i)_{0 \le i < |b|}$ in $\mathcal{U}$ thus gives rise through ω to a sequence of priorities $(\omega(\mathsf{proj}_2(\ell(b_i))))_{0 \le i < |b|}$. Parity acceptance for such a sequence is defined as follows: for an infinite branch b, let $\omega_{\min}(b)$ be the least priority appearing infinitely many times along b; then the infinite branch b is accepting if, and only if, $\omega_{\min}(b)$ is even. Notice that we impose no conditions on finite branches (*i.e.*, any finite branch is accepting). AEUTA (resp. EUTA) equipped with a parity acceptance condition are called *AEUPTA* (resp. *EUPTA*).

8

The *size of an AEUPTA*, denoted with $|\mathcal{A}|$, is a 5-tuple $(|Q|, |\delta|_{\mathsf{Bool}}, |\delta|_{\mathsf{E}}, |\delta|_{\mathsf{U}}, |\omega|)$, where $|\delta|_{\mathsf{Bool}}$ is the maximum size of the boolean formulas in δ, $|\delta|_{\mathsf{E}}$ (resp. $|\delta|_{\mathsf{U}}$) is the size of the largest existential part E (resp. universal part U) in some EU-pair in δ, and $|\omega| = |\{\omega(q) \mid q \in Q\}|$ is the number of priorities used in the automaton. Note that contrary to classical, fixed-arity tree automata, we explicitly consider the size of the transition function δ in the size of an AEUPTA; this is motivated by the fact that EU-constraints may succinctly encode very complex transitions, regardless of the size of Q.

In the sequel, we say that the size of an AEUPTA is at most $(s_Q, s_B, s_E, s_U, s_\omega)$ when $|Q| \leq s_Q$, $|\delta|_{\mathsf{Bool}} \leq s_B$, $|\delta|_{\mathsf{E}} \leq s_E$, $|\delta|_{\mathsf{U}} \leq s_U$ and $|\omega| \leq s_\omega$. The fact that we use five different parameters in the size of AEUPTAs will allow us to have more precise bounds on the complexities of our operations for manipulating them.

Remark 2. *We can easily modify any EU-automaton in such a way that every EU-pair it involves uses only a singleton or the empty set as its universal part (i.e., with $|\delta|_U \leq 1$). For this, it suffices to replace any general EU-pair $\langle E; U \rangle$ with $|U| > 1$ by $\langle E; \{q_U\} \rangle$, where q_U is a fresh state s.t. $\delta(q_U, \sigma) = \bigvee_{q \in U} \delta(q, \sigma)$ and $\omega(q_U)$ is the maximal priority used in the automaton. The size of the resulting automaton is then bounded by $(|Q| \, (1 + |\delta|_{Bool} \cdot |\Sigma|), |\delta|_{Bool} \cdot |\delta|_U, |\delta|_E, 1, |\omega|)$. Moreover, as we will see in the sequel, all the constructions we develop for manipulating AEUPTAs preserve this property of having only singleton or empty sets as the universal part of EU-pairs.* ◀

2.5 Examples

We illustrate our definitions with a few examples. Before presenting examples of EU-automata, we begin with examples of (special cases of) EU-constraints.

First, the EU-pair $\langle \varnothing; \varnothing \rangle$ characterises leaves of the input tree: indeed, in order to have $\nu_{n_u} \Vmodels \langle \varnothing; \varnothing \rangle$ (using the notations of Def. 3), the marking ν_{n_u} must contain a unitary submarking with empty image, hence its domain must be empty.

Positive boolean formulas also allow the special cases of $\top$ and $\bot$. Formula $\top$ does not impose any constraints on the node n_u of the execution tree where it is evaluated: this node can be a leaf even if the corresponding node m_t in the input tree is not a leaf. The resulting (finite) branch in the execution tree is accepting. On the other hand, no nodes of any execution tree can satisfy $\bot$, and such a transition can only lead to rejection.

Finally, EU-automata may include a special state $q_\top$ for which $\delta(q_\top, \sigma) = \langle \varnothing; \{q_\top\} \rangle$ for any $\sigma \in \Sigma$. When considering parity acceptance, this is equivalent to having $\delta(q_\top, \sigma) = \top$ if $\omega(q_\top)$ is even. Similarly, we could have a sink state $q_\bot$ with odd priority, which would reject any tree where it appears.

We now present some examples of EU-automata.

Example 3. *Consider the EUTA $\mathcal{A}$ over $\Sigma = \{a\}$ with $Q = \{q_{init}\}$, and*

$$\delta(q_{init}, a) = \langle q_{init} \mapsto 2; \varnothing \rangle$$

and where Ω accepts all branches (for example, Ω is a parity condition with $\omega(q_{init}) = 0$). This automaton accepts a single tree, namely the $\{a\}$-labelled

binary tree in which every node has exactly two successors. Letting

$$\delta(q_{init}, a) = \langle q_{init} \mapsto 2; \varnothing \rangle \vee \langle \varnothing; \varnothing \rangle$$

would accept all binary trees possibly containing finite branches (i.e., each node has either 0 or 2 successors) if $\omega(q_{init}) = 0$, but it would accept only finite binary trees if $\omega(q_{init}) = 1$. ◁

Example 4. *Parity automata on words can be seen as AEUPTA running on trees of arity 1. Formally, a non-alternating parity word automaton (PWA for short) over Σ is a 4-tuple $\mathcal{B} = (Q, q_{init}, \delta, \omega)$ where $\delta\colon Q \times \Sigma \to DBF(Q)$; an execution of $\mathcal{B}$ over a Σ-word $w = (w_i)_{0 \leq i < |w|}$ is a Q-word $s = (s_i)_{0 \leq i < |w|+1}$ such that $s_0 = q_{init}$ and $s_{i+1} \in \delta(s_i, w_i)$ for all $0 \leq i < |w|$. The word w is accepted by $\mathcal{B}$ if some execution of $\mathcal{B}$ on w is accepted by ω.*

Given a PWA $\mathcal{B} = (Q, q_{init}, \delta, \omega)$, we can easily build an EUPTA accepting all trees having at least one branch b whose word $w(b)$ is accepted by $\mathcal{B}$ [KSV06]: we let $\mathcal{A} = (Q \cup \{q_\top\}, q_{init}, \delta', \omega)$, where for all $\sigma \in \Sigma$, we define $\delta(q_\top, \sigma) = \top$, and for all $q \neq q_\top$, we let $\delta'(q, \sigma) = \bigvee_{q' \in \delta(q, \sigma)} \langle q' \mapsto 1; \{q_\top\} \rangle$; It is easily seen that $\mathcal{A}$ precisely accepts those trees containing at least one branch whose word is accepted by $\mathcal{B}$: automaton $\mathcal{A}$ can mimic the behaviour of $\mathcal{B}$ along that branch, and accept the rest of the tree. This entails the following result:

Proposition 4. *Let $\mathcal{B} = (Q, q_{init}, \delta, \omega)$ be a PWA. There exists a EUPTA $\mathcal{A}$ that accepts exactly all trees containing at least one branch b whose word $w(b)$ is accepted by $\mathcal{B}$. The size of $\mathcal{A}$ is $(|Q| + 1, |\delta|, 1, 1, |\omega|)$.*

Now assume that we want to build an AEUPTA accepting all trees in which all branches are accepted by $\mathcal{B}$. The construction above cannot easily be adapted: a natural attempt consists in letting $\delta''(q, \sigma) = \langle \varnothing; \delta(q, \sigma) \rangle$, thereby allowing to choose a different state with which to explore each successor node; however, this is not correct, because the automaton would have to make the same non-deterministic choices on the common prefix of two different branches. This approach works if $\mathcal{B}$ is required to be deterministic (i.e., if $\delta(q, \sigma)$ is a singleton for all $q \in Q$ and all $\sigma \in \Sigma$), and more generally, if it is history-deterministic [KSV06, BL23]. In the deterministic case:

Proposition 5. *Let $\mathcal{B}$ be a deterministic PWA. There exists a EUPTA $\mathcal{A}$ that accepts exactly all trees in which the word $w(b)$ of any branch b is accepted by $\mathcal{B}$. The size of $\mathcal{A}$ is $(|Q|, |\delta|, 0, 1, |\omega|)$.* ◁

2.6 Related formalisms

Several classes of tree automata have been defined in the literature.

***Fixed-arity tree automata* [Rab69, Tho90, Löd21].** Many papers on tree automata assume trees of fixed arity. The transition functions are then defined as positive Boolean combinations of atoms of the form (d, q), such an atom specifying that the d-successor of the current node has to be visited by the automaton in the state q. This in particular requires to have transition functions that depend on the arity in the input tree. Such tree automata clearly have a different expressive power compared to AEUTAs, since one can distinguish the first and second successors in a binary tree, by using directions in the transition

functions; on the other hand, AEUTA can accept trees of arbitrary arity. Note also that AEUTA are often much more succinct, for example by allowing constraints of the form $\langle q \mapsto k; \{q_\top\} \rangle$ that require to enumerate all possible subset of k successors (i.e., directions) in the fixed-arity setting.

Finally, note also that with alternating fixed-arity tree automata, one can use formula $(d,q) \wedge (d,q')$ in transitions to have the d-successor visited by both q and q'. Such a formula is not directly possible in the syntax of AEUTA transitions, because two EU-pairs cannot assign some specific state to some specific successor; however, exploring a single successor in two different states can be achieved by using an extra state $r_{q \wedge q'}$, whose transition function is defined as the conjunction of those of q and q'.

Amorphous tree automata [**BG93**]. Amorphous tree automata are, to our knowledge, the first class of tree automata that can handle trees of arbitrary, varying branching degree. Transitions in an amorphous tree automaton are defined through a *stretch* function, which takes as input an integer d (the arity of the node being visited) and an identifier and returning a d-tuple of states indicating, for each successor node, the state of the automaton in which it will be explored. Amorphous tree automata also have a kind of alternation mechanism, which allows to explore (copies of) the same node of the input tree in different states.

When used for encoding CTL, amorphous tree automata only rely on two stretch functions: (roughly) one that amounts to visiting all successors in the same state q, and one that amounts to visiting all successors in the same state q *but* one of them, which is explored in state q'. The stretch functions can thus be assumed to be fixed, so that any CTL formula ϕ can be turned into an equivalent amorphous tree automata of size linear in $|\phi|$.

$\{\Box, \Diamond\}$***-automata*** [**Wil99**]. In [Wil99], a different class of alternating automata running on tree of arbitrary branching degree is introduced: there, the transition function is defined as a positive boolean combination over $\{\Box, \Diamond\} \times Q$, where $(\Box, q)$ requires that the execution explores all successors of the current node in state q, and $(\Diamond, q)$ requires the execution to explore one successor node in state q. These automata run over Kripke structures in [Wil99], but their semantics could equivalently be defined over trees of arbitrary arity (which is a special case), as in our setting. Then $(\Box, q)$ would correspond to constraint $(\varnothing; \{q\})$ in our formalism ("*explore all successors of current node in state q*"), while $(\Diamond, q)$ would correspond to $(\{q \mapsto 1\}, \{q_\top\})$ ("*select one successor and explore it in state q*"). Clearly the class of $\{\Box, \Diamond\}$-automata corresponds exactly to the subclass of AEUTA where the transition function uses only EU-pairs of the form $\langle q \mapsto 1; \{q_\top\} \rangle$ or $\langle \varnothing; \{q\} \rangle$.

Lemma 6. $\{\Box, \Diamond\}$-*automata are less expressive than EU-automata.*

Proof. Clearly, $\{\Box, \Diamond\}$-automata can be turned into EU-automata. The converse is not true, in particular because $\{\Box, \Diamond\}$-automata cannot impose an upper bound on the number of successors of a node: if a tree $\mathcal{T}$ is accepted by some $\{\Box, \Diamond\}$-automaton, then the tree obtained from $\mathcal{T}$ by duplicating one branch is also accepted. □

$\{\Box, \Diamond\}$-automata (on trees) impose no restrictions on the arity of the trees they take as input, and do not distinguish between the different successors of any node of the input tree. As such, they are often named *symmetric* tree automata. Other variants of symmetric tree automata have been studied in the literature.

***Symmetric Büchi tree automata* [KV03].** In [KV03], a class of *symmetric* non-deterministic[3] tree automata is defined. The general aim of this class is to have a non-alternating equivalent to $\{\Box, \Diamond\}$-automata of [Wil99]. Symmetric Büchi tree automata (symNBT for short) are automata obtained from $\{\Box, \Diamond\}$-automata by applying a powerset construction: the set Q of states of a symNBT is of the form 2^S (where S is the set of states of the $\{\Box, \Diamond\}$-automata, whose elements are coined *micro-states* to distinguish them with the *macro-states* of Q), and transitions return sets of pairs $[U; E] \in 2^S \times 2^S$, where $[U; E]$ intuitively means $\Box U \wedge \Diamond E$: all micro-states in U must be present in *all* macro-states visiting the successors of the current node, and all micro-states in E must be present in *some* macro-state visiting the successors of the current node.

Any $\{\Box, \Diamond\}$-automaton can be turned into a symNBT [KV03]. In a sense, our Theorem 25 is an extension of this result to a richer class of tree automata.

MSO-*automata* [JW95, Wal96, Wal02, BB02, JL04, Zan12] In [JW95, Wal96, Wal02, BB02, JL04, Zan12], arbitrary-arity tree automata are defined where the transition functions are defined using first-order formulas over the successor nodes; these automata are named MSO-automata. In MSO-automata, transitions are given as first-order formulas, with quantification over the successors of the current node and predicates corresponding to states of the automaton. For instance, formula $\exists x.\ q(x)$ corresponds to $\Diamond q$ in $\{\Box, \Diamond\}$-automata: some successor must be visited by the automaton in state q.

Using first-order logic in transition function provides great flexibility. However, in order to facilitate the manipulation of those automata, the first-order formulas defining the transitions can be turned into a disjunction of formulas in *basic form*:

$$\exists (x_i)_{1 \le i \le k}\, diff((x_i)_{1 \le i \le k}) \wedge \bigwedge_{1 \le i \le k} q_i(x_i) \wedge \forall y.\ y \notin (x_i)_{1 \le i \le k} \Rightarrow \bigvee_{q \in U} q(y).$$

Such formula can be seen to correspond to our $\langle E; U \rangle$ formulas, so that MSO-automata have the same expressive power as our EU-automata. However, the transformation of a first-order formula into such a disjunction is based on Ehrenfeucht-Fraïssé games, and is not explicited in [Wal02, BB02, Zan12]. In the sequel, we develop operations (union, intersection, projection, complementation and simulation) with explicit constructions and precise evaluation of the size of the resulting automata, which cannot be directly obtained from the current results about MSO-automata.

3   Game-based semantics

The acceptance of a tree by an AEUPTA can be expressed as the existence of a winning strategy in a two-player turn-based parity game. We first briefly recall

---

[3]With our terminology, we would name them non-alternating.

the definition of parity games, and then explain how they can be used to encode the semantics of alternating tree automata.

3.1 Parity games

A *two-player turn-based parity game* is a 4-tuple $\mathcal{G} = (Y_0, Y_1, R, \theta)$ where Y_0 and Y_1 are disjoint sets of states and, writing $Y = Y_0 \cup Y_1$, $R \subseteq Y^2$ is a set of transitions, and $\theta \colon Y \to \mathbb{N}$ assigns a priority to each state of the game.

In such a game, two players (which we name Player 0 and Player 1) select transitions so as to form a *path* in the graph (Y, R): from some $y \in Y_i$ (with $i \in \{0, 1\}$), Player i selects a transition $(y, y') \in R$, and the game proceeds to y'. A path is *maximal* if it is infinite, or if its last state has no outgoing transitions. A finite maximal path is *winning* for Player 0 if, and only if, its last state belongs to Y_1 (the blocked player loses). For an infinite path π, we write $\theta_{\min}(\pi)$ for the least integer k s.t. there are infinitely many $i \geq 0$ with $\theta(\pi(i)) = k$. The infinite path π is *winning* for Player 0 if, and only if, $\theta_{\min}(\pi)$ is even; otherwise it is *winning* for Player 1.

A *strategy* for Player i in a parity game is a partial function $\alpha_i \colon Y^* \times Y_i \to R$ such that for any finite path $\pi \cdot y$ with $y \in Y_i$, the function α_i is defined at $\pi \cdot y$ if, and only if, there exists y' such that $(y, y') \in R$; in that case, we must have $(y, \alpha_i(\pi \cdot y)) \in R$. A strategy is *memoryless* if for any two paths $\pi \cdot y$ and $\pi' \cdot y$, it holds $\alpha_i(\pi \cdot y) = \alpha_i(\pi' \cdot y)$.

A path $\pi = (y_j)_{0 \leq j \leq |\pi|}$ is *compatible* with a strategy α_i of Player i if for any $0 \leq j < |\pi| - 1$, if $y_j \in Y_i$, then $y_{j+1} = \alpha_i((y_k)_{0 \leq k \leq j})$. A strategy α_i is *winning* for Player i from y if all maximal paths starting from y that are compatible with α_i are winning for Player i.

The following classical property about infinite parity games will be useful in the sequel:

Proposition 7 ([Zie98])**.** *Two-player turn-based parity games are positionally determined: from any state of such games, one of the two players has a memoryless winning strategy.*

3.2 *Game semantics* for tree automata

We now explain how to define a two-player turn-based parity game $\mathcal{G}_{\mathcal{A},\mathcal{T}} = (Y_0, Y_1, R, \theta)$ encoding the acceptance of a tree $\mathcal{T} = (t, l)$ by an AEUPTA $\mathcal{A} = (Q, q_{init}, \delta, \omega)$.

States of $\mathcal{G}_{\mathcal{A},\mathcal{T}}$ are of three kinds: the *main states* of the game are of the form (n, q) where $n \in t$ and $q \in Q$; the *auxiliary states* are of the form (n, q, φ) where φ is a subformula of $\delta(q, l(n))$, and (n, q, ν_n) where ν_n is a unitary marking of $\mathsf{succ}(n)$ with states of $\mathcal{A}$. It remains to partition them into Y_0 and Y_1. The set Y_0 contains the states of the following form:

- the main states (n, q):

- the state $(n, q, \bot)$;

- the states of the form $(n, q, \bigvee_{1 \leq i \leq k} \phi_i)$ (with $k > 1$);

- the states of the form $(n, q, \langle E; U \rangle)$, which correspond to positions where Player 0 has to assign states to the successors of n so as to satisfy $\langle E; U \rangle$;
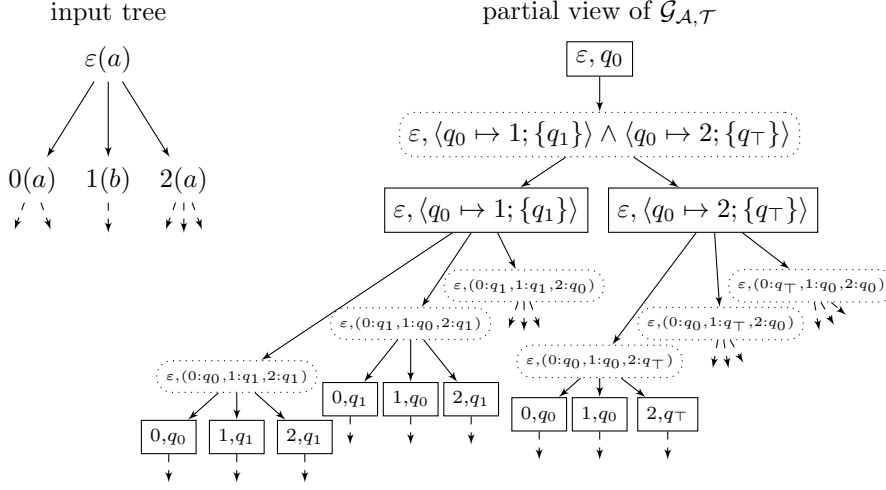
input tree   partial view of $\mathcal{G}_{\mathcal{A},\mathcal{T}}$

Figure 2: Example of game $\mathcal{G}_{\mathcal{A},\mathcal{T}}$

All other states belong to Y_1 (*i.e.*, states of the form $(n, q, \top)$, $(n, q, \bigwedge_{1 \leq i \leq k} \phi_i)$, and (n, q, ν_n)).

Transitions are defined as follows:

- for any state (n, q) with $n \in t$ and $q \in Q$, there is a single transition from (n, q) to $(n, \delta(q, l(n)))$;

- from $(n, q, \bigvee_{1 \leq i \leq k} \phi_i)$ and $(n, q, \bigwedge_{1 \leq i \leq k} \phi_i)$, there are transitions to (n, q, ϕ_i), for each $1 \leq i \leq k$;

- from $(n, q, \langle E; U \rangle)$, for each unitary marking ν_n of $\mathsf{succ}(n)$ with states in Q such that $\nu_n \models \langle E; U \rangle$, there is a transition to (n, q, ν_n). Notice that for ν_n to fulfill $\langle E; U \rangle$, we must have $\mathsf{supp}(\mathsf{img}(\nu_n)) \subseteq \mathsf{supp}(E) \cup U$.

- from (n, q, ν_n) where ν_n is a unitary marking of $\mathsf{succ}(n)$ with states of Q, for each $n_i \in \mathsf{succ}(n)$, there is a transition to $(n_i, \nu_n(n_i))$.

Finally, priorities are defined as $\theta(n, q) = \omega(q)$ for the main states, and to $|\omega| + 1$ for all other states of the game. Since there are infinitely many main states along infinite runs, only priorities of the main states are useful.

Example 5. *We consider an automaton with $\delta(q_0, a) = \langle q_0 \mapsto 1; \{q_1\} \rangle \wedge \langle q_0 \mapsto 2; \{q_\top\} \rangle$, $\delta(q_0, b) = \top$, and $\delta(q_1, a) = \delta(q_1, b) = \langle q_\top \mapsto 1; \varnothing \rangle$. Figure 2 shows the first levels of an input tree and the beginning of the corresponding parity game for this transition function (dotted nodes belong to Player 1). To simplify the figure, we have omitted q_0 in states of the form $(\varepsilon, q_0, \varphi)$ and $(\varepsilon, q_0, \nu_n)$.* ◁

The resulting parity game encodes the acceptance of a tree by an AEUPTA:

Proposition 8. *The Σ-labelled $\mathcal{D}$-tree $\mathcal{T}$ is accepted by the AEUPTA $\mathcal{A}$ if, and only if, Player 0 has a winning strategy from state $(\varepsilon_t, q_{init})$ in the associated parity game $\mathcal{G}_{\mathcal{A},\mathcal{T}}$.*

14

Proof. We prove a slightly stronger result: Player 0 has a winning strategy from a main state (n, q) in $\mathcal{G}_{\mathcal{A}, \mathcal{T}}$ if, and only if, the input tree rooted at n is accepted by the automaton $\mathcal{A}$ with q considered as the initial state (*i.e.*, there exists an accepting execution subtree rooted at a node labelled with (n, q)).

First assume that Player 0 has a winning strategy α_0: by pruning all subtrees that are not selected by α_0, and removing non-main states, we get a tree which we can prove is an accepting execution tree: Boolean operators of the transition function are handled correctly, and from the states of the form $(n, q, \langle E; U \rangle)$, strategy α_0 selects a valid way of exploring the successor nodes, so as to fulfill the EU-constraint, therefore the marking of $\mathsf{succ}(n)$ induced by this tree satisfies $(n, \delta(q, l(n)))$. By definition of the priorities of the states of $\mathcal{G}_{\mathcal{A}, \mathcal{T}}$, all infinite branches are accepting since infinite paths are winning for Player 0. Finite branches in the tree may only originates from auxiliary states of the form $(n, q, \bot)$ and $(n, q, \top)$, and of the form (n, q, ν_n) for which n has no successor nodes. The first two cases are correctly handled by construction of the game (the blocked player loses). Since states of the form (n, q, ν_n) belong to Player 1, they are winning for Player 0 in case n has no successor nodes, *i.e.*, in case ν_n is the marking of the empty set; but there may only be transitions from $(n, q, \langle E; U \rangle)$ to (n, q, ν_n) with the empty marking when $\langle E; U \rangle$ is of the form $\langle \varnothing; U \rangle$: a finite branch of the execution tree ending in a node labelled $(n, \langle \varnothing; U \rangle$ is indeed accepting.

The converse is similar: given an accepting execution tree, we can build a winning strategy for Player 0. For this, it suffices to check which parts of disjunctive formulas in transitions are satisfied, and how EU-constraints are fulfilled. The winningness for infinite and finite paths then again corresponds to the acceptance status of the corresponding branches of the execution tree. $\qquad\square$

Note that the subgames issued from (n, q, φ) and (n, q', φ) (resp. from (n, q, ν_n) and (n, q', ν_n)) are isomorphic and admit the same wining memoryless strategies. Therefore we do not distinguish them in the following and consider only nodes of the form (n, φ) or (n, ν_n).

When $\mathcal{T}$ is regular and corresponds to the execution tree of some Kripke structure $\mathcal{K} = (V, E, \ell)$, we can build a *finite* game $\mathcal{G}_{\mathcal{A}, \mathcal{K}} = (Y_0, Y_1, R, \theta)$ defined exactly as above, but replacing nodes n of t with vertices v of V. The sizes of Y and of the transition relation R are both in $O(|V| \cdot (|Q| \cdot (1 + |\delta|_{\mathsf{Bool}}) + |Q|^{\mathsf{arity}(\mathcal{K})}))$, hence in $O(|V| \cdot (|Q| \cdot |\delta|_{\mathsf{Bool}} + |Q|^{|V|}))$.

Moreover it is worth noticing that the complexity blow-up in the size of the game is due to the treatment of EU-constraints. For automata using only simple constraints of the form $\langle q \mapsto 1; \{q_\top\} \rangle$ or $\langle \varnothing; \{q\} \rangle$ (which correspond to $\{\Box, \Diamond\}$-automata), the number of unitary markings involved in reachable states of the form (n, ν_n) in $\mathcal{G}_{\mathcal{A}, \mathcal{K}}$ is $O(|V| \cdot |Q|)$, so that the sizes of Y and R are in $O(|V| \cdot |Q| \cdot (|\delta|_{\mathsf{Bool}} + |V|))$.

4 Operations on AEUTAs

This section is the main technical part of our paper: we develop algorithms for performing various operations on AEUPTAs (namely union and intersection, projection, complementation and alternation removal), and carefully study the

size of the AEUPTAs we obtain. Table 1 gathers our results, giving the size of the resulting AEUPTAs depending on the size of the AEUPTAs given in input. The rest of this section gives detailed algorithms, explanation of their correctness, and justifications for the sizes of the resulting automata.

intersection union (Thm 9)	$|Q_\cap|, |Q_\cup| \leq |Q| + |Q'| + 1$
	$|\delta_\cap|_{\mathsf{Bool}}, |\delta_\cup|_{\mathsf{Bool}} \leq |\delta|_{\mathsf{Bool}} + |\delta'|_{\mathsf{Bool}} + 1$
	$|\delta_\cap|_{\mathsf{E}}, |\delta_\cup|_{\mathsf{E}} \leq \max(|\delta|_{\mathsf{E}}, |\delta'|_{\mathsf{E}})$
	$|\delta_\cap|_{\mathsf{U}}, |\delta_\cup|_{\mathsf{U}} \leq \max(|\delta|_{\mathsf{U}}, |\delta'|_{\mathsf{U}})$
	$|\omega_\cap|, |\omega_\cup| \leq \max(|\omega|, |\omega'|) + 1$
projection (Thm 10)	$|Q_{\mathsf{proj}}| \leq |Q|$
	$|\delta_{\mathsf{proj}}|_{\mathsf{Bool}} \leq |\Sigma'| \cdot |\delta|_{\mathsf{Bool}}$
	$|\delta_{\mathsf{proj}}|_{\mathsf{E}} \leq |\delta|_{\mathsf{E}}$
	$|\delta_{\mathsf{proj}}|_{\mathsf{U}} \leq |\delta|_{\mathsf{U}}$
	$|\omega_{\mathsf{proj}}| \leq |\omega|$
complement (Thm 17)	$|Q^c| \in O(|Q| \cdot |\delta|_{\mathsf{Bool}} \cdot |\Sigma| \cdot |\delta|_{\mathsf{E}} \cdot 2^{|\delta|_{\mathsf{E}}})$
	$|\delta^c|_{\mathsf{Bool}} \in O(|Q| \cdot |\delta|_{\mathsf{Bool}} \cdot |\delta|_{\mathsf{E}}^3 \cdot 4^{|\delta|_{\mathsf{E}}})$
	$|\delta^c|_{\mathsf{E}} \leq |\delta|_{\mathsf{E}} + 1$
	$|\delta^c|_{\mathsf{U}} \leq \max(|\delta|_{\mathsf{U}}, 1)$
	$|\omega^c| \leq |\omega| + 1$
simulation (Thm 25)	$|Q^s| \in 2^{O(|Q|^2 \cdot \log(|Q|))}$
	$|\delta^s|_{\mathsf{Bool}} \in (2|\delta|_{\mathsf{E}} \cdot |\delta|_{\mathsf{U}})^{O(|Q|^2 \cdot |\delta|_{\mathsf{Bool}}^2 \cdot |\delta|_{\mathsf{E}})}$
	$|\delta^s|_{\mathsf{E}} \leq |Q| \cdot |\delta|_{\mathsf{Bool}} \cdot |\delta|_{\mathsf{E}}$
	$|\delta^s|_{\mathsf{U}} \leq |\delta|_{\mathsf{U}}^{|Q| \cdot |\delta|_{\mathsf{Bool}}}$
	$|\omega^s| \leq 2(|Q| \cdot |\omega| + 1)$

Table 1: Bounds on the size of the automata obtained by our algorithms

4.1 Union and intersection

Union and intersection are straightforward for AEUTAs, thanks to alternation.

Theorem 9. *Let $\mathcal{A} = (Q, q_{init}, \delta, \omega)$ and $\mathcal{A}' = (Q', q'_{init}, \delta', \omega')$ be two AEUPTAs. There exist AEUPTAs $\mathcal{A}_\cup$ and $\mathcal{A}_\cap$, respectively accepting the union and the intersection of the language of $\mathcal{A}$ and $\mathcal{A}'$, and having size at most $(|Q| + |Q'| + 1, |\delta|_{Bool} + |\delta'|_{Bool} + 1, \max(|\delta|_E, |\delta'|_E), \max(|\delta|_U, |\delta'|_U), \max(|\omega|, |\omega'|) + 1)$.*

Proof. We consider intersection (the case of union is similar): automaton $\mathcal{A}_\cap$ is defined as $(Q'', q''_{init}, \delta'', \omega'')$ with

- $Q'' = Q \cup Q' \cup \{q''_{init}\}$ (assuming w.l.o.g. that all three states are pairwise disjoint);

- $\delta''(q''_{init}, \sigma) = \delta(q_{init}, \sigma) \wedge \delta'(q'_{init}, \sigma)$, and δ'' coincides with δ on $Q \times \Sigma$ and with δ' on $Q' \times \Sigma$;

- ω'' coincides with ω on Q and with ω' on Q'; its value in q''_{init} is irrelevant since q''_{init} will be visited only once. Notice that, using straightforward arguments, we may assume that $\omega(Q)$ and $\omega'(Q')$ are subintervals of $[\![0; |\omega|]\!]$ and $[\![0; |\omega'|]\!]$, so that $\omega''(Q'')$ is a subinterval of $[\![0; \max(|\omega|, |\omega'|)]\!]$.

16

The correctness of this construction is not hard to prove, using the game semantics. Consider a tree $\mathcal{T} = (t, l)$ accepted by $\mathcal{A}_\cap$. By Prop. 8, Player 0 has a winning strategy from $(\varepsilon_t, q''_{init})$ in the corresponding game $\mathcal{G}_{\mathcal{A}_\cap, \mathcal{T}}$. Since $\delta''(q''_{init}, l(\varepsilon_t)) = \delta(q_{init}, l(\varepsilon_t)) \wedge \delta'(q'_{init}, l(\varepsilon_t))$, there is a unique transition from $(\varepsilon_t, q''_{init})$ to the Player-1 state $(\varepsilon_t, \delta''(q''_{init}, l(\varepsilon_t)))$; from there, Player 1 can decide to move either to $(\varepsilon_t, \delta(q_{init}, l(\varepsilon_t)))$ or to $(\varepsilon_t, \delta'(q'_{init}, l(\varepsilon_t)))$. Since Player 0 has a winning strategy from $(\varepsilon_t, q''_{init})$, she also has winning strategies from both $(\varepsilon_t, \delta(q_{init}, l(\varepsilon_t)))$ and $(\varepsilon_t, \delta'(q'_{init}, l(\varepsilon_t)))$ in $\mathcal{G}_{\mathcal{A}_\cap, \mathcal{T}}$. Since δ'' coincides with δ on $Q \times \Sigma$ and with δ' on $Q' \times \Sigma$, Player 0 has winning strategies from $(\varepsilon_t, q_{init})$ in $\mathcal{G}_{\mathcal{A}, \mathcal{T}}$ and from $(\varepsilon_t, q'_{init})$ in $\mathcal{G}_{\mathcal{A}', \mathcal{T}}$. Hence $\mathcal{T}$ is accepted by both $\mathcal{A}$ and $\mathcal{A}'$.

The converse implication follows the same arguments. $\square$

Remark 3. *Note that if the minimum priority of ω and ω' are equal, then the number of priorities in $\mathcal{A}_\cup$ and $\mathcal{A}_\cap$ can be bounded by $\max(|\omega|, |\omega'|)$ instead of $\max(|\omega|, |\omega'|) + 1$.* ◄

4.2 Projection

Given an AEUPTA $\mathcal{A}$ over alphabet $\Sigma_1 \times \Sigma_2$, *projection* consists in building another AEUPTA $\mathcal{A}_1$, over alphabet Σ_1, accepting all Σ_1-labelled trees whose labelling can be extended on $\Sigma_1 \times \Sigma_2$ to make the tree accepted by $\mathcal{A}$. This is a classical construction, and it can be performed easily on non-alternating automata [MS85].

Formally, two $\Sigma_1 \times \Sigma_2$-labelled trees $\mathcal{T} = (t, l)$ and $\mathcal{T}' = (t', l')$ are said Σ_1-*equivalent*, denoted $\mathcal{T} \equiv_{\Sigma_1} \mathcal{T}'$, whenever $t = t'$ and for any node n of these trees, it holds $\mathsf{proj}_1(l(n)) = \mathsf{proj}_1(l'(n))$. Σ_2-equivalence is defined analogously.

Theorem 10. *Let $\mathcal{A} = (Q, q_{init}, \delta, \omega)$ be an EUPTA over $\Sigma = \Sigma_1 \times \Sigma_2$. For each $i \in \{1, 2\}$, we can build an EUPTA $\mathcal{A}_i$ over Σ such that, for any Σ-labelled tree $\mathcal{T}$, it holds: $\mathcal{T} \in \mathcal{L}(\mathcal{A}_i)$ if, and only if, there is a Σ-labelled tree $\mathcal{T}'$ in $\mathcal{L}(\mathcal{A})$ such that $\mathcal{T} \equiv_{\Sigma_i} \mathcal{T}'$. The size of $\mathcal{A}_i$ is at most $(|Q|, |\Sigma_{3-i}| \cdot |\delta|_{Bool}, |\delta|_E, |\delta|_U, |\omega|)$.*

Proof. We define $\mathcal{A}_1$ over Σ as $(Q, q_{init}, \delta_1, \omega)$ with:

$$\delta_1(q, (\sigma_1, \sigma_2)) = \bigvee_{\sigma'_2 \in \Sigma_2} \delta(q, (\sigma_1, \sigma'_2)).$$

Take a tree $\mathcal{T} = (t, l)$ accepted by $\mathcal{A}_1$, and pick an accepting execution tree $\mathcal{U} = (t, \ell)$ (automaton $\mathcal{A}_1$ is non-alternating, so we can assume that the execution tree has the same structure as the input tree). Consider any node n of t, labelled with $l(n) = (\sigma_1, \sigma_2)$ in $\mathcal{T}$ and with $\ell(n) = (n, q)$ in $\mathcal{U}$. By definition of δ_1, there exists $\sigma_2^n \in \Sigma_2$ such that the successors of n in $\mathcal{U}$ satisfy $\delta(q, (\sigma_1, \sigma_2^n))$. This holds for all nodes of t, meaning that each node n can be relabelled with (σ_1, σ_2^n) in such a way that this new tree $\mathcal{T}'$ is accepted by $\mathcal{A}$.

Conversely, assume that tree $\mathcal{T}$ admits an Σ_1-equivalent tree $\mathcal{T}'$ that is accepted by the EUPTA $\mathcal{A}$, and take an accepting execution tree $\mathcal{U}$. By construction of δ_1, this execution tree is also an accepting execution tree for $\mathcal{A}_1$ on $\mathcal{T}$. $\square$

Remark 4. *This construction does not directly extend to alternating automata: indeed, let $\Sigma_1 = \{a_1\}$ and $\Sigma_2 = \{a_2, a'_2\}$, and consider the AEUPTA $\mathcal{A} = (Q, q_{init}, \delta, \omega)$ on $\Sigma_1 \times \Sigma_2$ with:*

- $Q = \{q_{init}, q_1, q_2\}$,

- *The transition function is defined as follows (for any $\sigma \in \Sigma_1 \times \Sigma_2$):*

$$\delta(q_{init}, (a_1, a_2)) = \delta(q_{init}, (a_1, a'_2)) = \langle q_1 \mapsto 1; \varnothing \rangle \wedge \langle q_2 \mapsto 1; \varnothing \rangle$$
$$\delta(q_1, (a_1, a_2)) = \delta(q_2, (a_1, a'_2)) = \top$$
$$\delta(q_1, (a_1, a'_2)) = \delta(q_2, (a_1, a_2)) = \bot$$

We let $\omega(q) = 0$ for all states. Now, for a tree to be accepted, its root must be labelled with (a_1, a_2) or (a_1, a'_2) and have a single successor node. That node will be explored both in state q_1 and q_2, and for any σ, either $\delta(q_1, \sigma)$ or $\delta(q_2, \sigma)$ is equal to $\bot$. It follows that $\mathcal{L}(\mathcal{A}) = \varnothing$, hence also $\mathcal{L}(\mathcal{A}_i) = \varnothing$ for $i \in \{1, 2\}$.

Now, applying our construction to this automaton on the first component provides us with an automaton $\mathcal{A}_1$ with the following transition function:

$$\delta_1(q_{init}, \sigma) = \langle q_1 \mapsto 1; \varnothing \rangle \wedge \langle q_2 \mapsto 1; \varnothing \rangle \qquad\qquad \forall \sigma \in \Sigma_1 \times \Sigma_2$$
$$\delta_1(q_1, \sigma) = \delta(q_1, (a_1, a_2)) \vee \delta(q_1, (a_1, a'_2)) = \top \vee \bot = \top$$
$$\delta_1(q_2, \sigma) = \delta(q_2, (a_1, a_2)) \vee \delta(q_2, (a_1, a'_2)) = \bot \vee \top = \top$$

This automaton accepts any tree on $\Sigma_1 \times \Sigma_2$ whose root has a single successor. This shows that Theorem 10 does not hold for alternating automata.

For similar reasons, this theorem does not extend to universal *projection, where $\mathcal{A}_i$ would accept the trees for which* all Σ_i-*equivalent trees would be accepted by $\mathcal{A}$.* ◀

4.3 Complementation

Complementation is the operation of building an automaton accepting the complement of the language accepted by some given automaton. It is usually easy for alternating parity automata: it suffices to dualise the transition function (swapping disjunctions and conjunctions) and shifting the priorities. Such a construction is given in [Kir02] for $\{\Box, \Diamond\}$-automata: in that setting, $\Box$ and $\Diamond$ are dual to each other, and the construction is straightforward. The same is true for MSO-automata [Wal02, Zan12]. For our AEUTA however, we need to express the negation of any EU-pair $\langle E; U \rangle$ as an EU-constraints.

The question then is to characterise nodes that *fail to satisfy* an EU-pair $\langle E; U \rangle$. There can be two reasons for this, which we develop in the sequel:

1. either we cannot find $|E|$ successors of the current node n to associate with the $|E|$ states of the existential part,

2. or for every way to satisfy E with nodes in $\mathsf{succ}(n)$, there remain successors that are accepted by no states in U. Equivalently, there is no way to satisfy the existential part E with (at least) all nodes that are accepted by no states in the universal part U. This includes as a special case the situations where we have more than $|E|$ successors that are accepted by no states in U.

Failing to satisfy the existential part of $\langle E; U \rangle$. We first address the former situation, which is easier and already contains most of the technicalities we need for solving the general case.

Fix an AEUPTA $\mathcal{A} = (Q, q_{init}, \delta, \omega)$ over Σ. We assume w.l.o.g. that $\mathcal{A}$ has a state $q_\top$ from which all trees are accepted. For a state $q \in Q$, we write $\mathcal{A}_q$ for the AEUPTA (Q, q, δ, ω), obtained from $\mathcal{A}$ by taking q as the initial state. For a multiset E over Q, we define the AEUPTA $\mathcal{A}_{\downarrow E} = (Q \cup \{q_E\}, q_E, \delta_E, \omega_E)$ such that

- q_E is a new state, not in Q;

- for any $\sigma \in \Sigma$, $\delta_E(q_E, \sigma) = \langle E; \{q_\top\} \rangle$ and $\delta_E(q, \sigma) = \delta(q, \sigma)$ for all $q \neq q_E$;

- the priority function ω_E coincides with ω on Q, and $\omega_E(q_E) = 0$.

Automaton $\mathcal{A}_{\downarrow E}$ visits some of the successors of the root in each of the states of E. Notice that if $E = \varnothing$, then $\mathcal{A}_{\downarrow E}$ accepts any tree. Notice also that the automaton $\mathcal{A}_q$ defined above is *not* the same automaton as $\mathcal{A}_{\downarrow E}$ with $E = \{q\}$: automaton $\mathcal{A}_q$ starts exploring the root $\varepsilon_\mathcal{T}$ of its input tree in state q, whereas $\mathcal{A}_{\downarrow \{q\}}$ explores the subtree rooted at some successor node of $\varepsilon_\mathcal{T}$ in state q.

Our aim in this part is to compute an AEUPTA $\mathcal{C}_E$ such that $\mathcal{L}(\mathcal{C}_E)$ is the complement of $\mathcal{L}(\mathcal{A}_{\downarrow E})$. This construction is the main ingredient for complementing an EU-automaton. It is based on the notion of blocking pairs:

Proposition 11. *Let $\mathcal{T}$ be an input tree. If $\mathcal{T}$ is not accepted by $\mathcal{A}_{\downarrow E}$, then there exists a submultiset $F \sqsubseteq E$ and a state $g \in E \setminus F$ such that $\mathcal{T}$ is accepted by $\mathcal{A}_{\downarrow F}$ and is rejected by $\mathcal{A}_{\downarrow F \uplus \{g\}}$.*

Definition 12. *A pair (F, g) satisfying the conditions of Prop. 11 is called a blocking pair for $\mathcal{T}$ and E.*

Proof of Prop. 11. The proof is by induction on $|E|$: the result holds vacuously if E is empty, and it is trivial if $|E| = 1$. It then suffices to observe that for any $g \in E$, either $(E \setminus \{g\}, g)$ is a blocking pair for $\mathcal{T}$, or such a blocking pair can be found in $E \setminus \{g\}$ (by induction). □

We now focus on *minimal* blocking pairs for $\mathcal{T}$, *i.e.*, pairs (F, g) such that $(F \setminus \{g'\}, g)$ is not a blocking pair of $\mathcal{T}$, for any $g' \in F$. We will prove that if (F, g) is a minimal blocking pair for $\mathcal{T}$, then in any accepting execution tree $\mathcal{U}$ of $\mathcal{A}_{\downarrow F}$ on $\mathcal{T}$, the subtrees rooted at the successors of the root $\varepsilon_\mathcal{T}$ that are not used to fulfill F can be accepted by no states in $\mathsf{supp}(F) \cup \{g\}$. In the sequel, given a tree $\mathcal{T}$ and a node y, we write $\mathcal{T}_y$ for the subtree of $\mathcal{T}$ rooted at y.

Proposition 13. *Let (F, g) be a blocking pair for some tree $\mathcal{T}$. Let $\mathcal{U} = (u, \ell)$ be a minimal accepting execution tree of $\mathcal{A}_{\downarrow F}$ on $\mathcal{T}$, ν be the corresponding unitary marking of $\mathsf{succ}(\varepsilon_t)$ by Q, satisfying $\nu \models \langle F; \{q_\top\} \rangle$. If there exists a node y in $\mathsf{succ}(\varepsilon_t)$ such that (1) $\nu(y) = q_\top$ and (2) $\mathcal{T}_y \in \mathcal{L}(\mathcal{A}_{g'})$ for $g' \in F$, then $(F \setminus \{g'\}, g)$ is a blocking pair for $\mathcal{T}$.*

Proof. Assume that such a node y exists. The execution tree $\mathcal{U}$ witnesses the fact that $\mathcal{T}$ is accepted by $\mathcal{A}_{\downarrow F \setminus \{g'\}}$. If $(F \setminus \{g'\}, g)$ were not a blocking pair for $\mathcal{T}$, then $\mathcal{T}$ would be accepted by $\mathcal{A}_{\downarrow F \setminus \{g'\} \uplus \{g\}}$. In that case, let $\mathcal{U}' = (u', \ell')$ be a minimal accepting execution tree of $\mathcal{A}_{\downarrow F \setminus \{g'\} \uplus \{g\}}$ over $\mathcal{T}$, and ν' be a unitary

marking of $\mathsf{succ}(\varepsilon_t)$ by Q induced by $\mathcal{U}'$, which satisfies $\nu' \models \langle F \setminus \{g'\} \uplus \{g\};$ $\{q_\top\}\rangle$. Let N and N' be subsets of $\mathsf{succ}(\varepsilon_t)$ such that $\nu(N) = F$ (hence $y \notin N$) and $\nu'(N') = F \setminus \{g'\} \uplus \{g\}$. We pick $\mathcal{U}'$, N and N' so as to maximize the size of $N \cap N'$.

Then $|N| = |N'|$, but $N \nsubseteq N'$: indeed, if $N \subseteq N'$, then $N = N'$, and $y \notin N'$; then $\mathcal{U}'$ could be extended with an execution tree of $\mathcal{A}_{g'}$ on the subtree $\mathcal{T}_y$ entailing that $\mathcal{T}$ would be accepted by $\mathcal{A}_{\downarrow F \uplus \{g\}}$. Pick $n \in N \setminus N'$, and a corresponding node $x \in \mathsf{succ}(\varepsilon_u)$ be such that $\ell(x) = (n, q)$ for some $q \in F$; then $q \neq g'$, as otherwise $\mathcal{U}'$ could again be extended into an accepting execution tree of $\mathcal{A}_{\downarrow F \uplus \{g\}}$ over $\mathcal{T}$. Since $\nu'(N') = F \setminus \{g'\} \uplus \{g\}$, there must exist a node n' in $N' \setminus N$ and a corresponding node $x' \in \mathsf{succ}(\varepsilon_{u'})$ with $\ell(x') = (n', q)$ (if this were not the case, then $\nu(N)$ would contain more copies of q than $\nu'(N')$ does). Replacing the subtree rooted at x' in $\mathcal{U}'$ with the subtree rooted at x in $\mathcal{U}$, we get a minimal accepting execution tree $\mathcal{U}'' = (u'', \ell'')$ of $\mathcal{A}_{\downarrow F \setminus \{g'\} \uplus \{g\}}$ over $\mathcal{T}$ with induced unitary marking ν'', and a set $N'' \subseteq \mathsf{succ}(\varepsilon_t)$ such that $\nu''(N'') = F \setminus \{g'\} \uplus \{g\}$, but for which $|N \cap N''|$ is larger than $|N \cap N'|$, contradicting our choice of $\mathcal{U}'$, N and N'. $\qquad\square$

Now when some tree $\mathcal{T}$ is not accepted by $\mathcal{A}_{\downarrow E}$, we can ensure the existence of a blocking pair (F, g) such that for any node n in $\mathsf{succ}(\varepsilon_\mathcal{T})$ that is not involved in the satisfaction of F, the subtree rooted at n is not accepted by any automaton $\mathcal{A}_{g'}$, for $g' \in F \uplus \{g\}$. Formally, we have:

Proposition 14. *If $\mathcal{T}$ is not accepted by $\mathcal{A}_{\downarrow E}$, then there exists a blocking pair (F, g) such that for any minimal accepting execution tree $\mathcal{U}$ of $\mathcal{A}_{\downarrow F}$ over $\mathcal{T}$, writing ν for the unitary marking of $\mathsf{succ}(\varepsilon_t)$ induced by $\mathcal{U}$, it holds: for any $y \in \mathsf{succ}(\varepsilon_t)$ with $\nu(y) = q_\top$, the subtree $\mathcal{T}_y$ is rejected by $\mathcal{A}_{g'}$ for any $g' \in F \uplus \{g\}$.*

Proof. Consider a blocking pair (F, g) for $\mathcal{T}$ where $F \sqsubseteq E$ is minimal (for inclusion). Let $\mathcal{U}$ be a minimal accepting execution tree of $\mathcal{A}_{\downarrow F}$ over $\mathcal{T}$, ν be the induced unitary marking of $\mathsf{succ}(\varepsilon_t)$, and $N \subseteq \mathsf{succ}(\varepsilon_t)$ such that $\nu(N) = F$.

For any node $y \in \mathsf{succ}(\varepsilon_t) \setminus N$, if the subtree $\mathcal{T}_y$ were accepted by $\mathcal{A}_g$, then $\mathcal{T}$ would be accepted by $\mathcal{A}_{\downarrow F \uplus \{g\}}$. Similarly, if the subtree $\mathcal{T}_y$ were accepted by $\mathcal{A}_{g'}$, for any $g' \in F$, then by Prop. 13, $(F \setminus \{g'\}, g)$ would be a blocking pair, contradicting minimality of (F, g). $\qquad\square$

Applying this result, we will build the complement automaton $\mathcal{C}_E$ of $\mathcal{A}_{\downarrow E}$ by checking the existence of a blocking pair satisfying the conditions of Prop. 14: the transition from the initial state will be a disjunction, over all pairs (F, g), of EU-pairs $\langle F; \{(\mathsf{supp}(\overline{F}) \cup \{\overline{g}\}, \wedge)\}\rangle$, where $(\mathsf{supp}(\overline{F}) \cup \{\overline{g}\}, \wedge)$ denotes a new state accepting any tree that does not belong to $\mathcal{L}(\mathcal{A}_{g'})$, for any $g' \in \mathsf{supp}(F) \cup \{g\}$. This is expressed by the following EU-constraint:

$$\Phi_E = \bigvee_{F \sqsubseteq E} \bigvee_{g \in E \setminus F} \langle F; \{(\mathsf{supp}(\overline{F}) \cup \{\overline{g}\}, \wedge)\}\rangle.$$

Notice that for the special case where E is empty, we end up with an empty disjunction, which is equivalent to false. This is coherent with the fact that $\mathcal{A}_\varnothing$ accepts any tree. Full definitions and correctness proofs will be given after we have explained how we handle the general case of failing to satisfy $\langle E; U\rangle$.

Failing to satisfy $\langle E; U \rangle$. Given a tree $\mathcal{T}$, we call *direct subtree* of $\mathcal{T}$ any subtree of $\mathcal{T}$ rooted in a successor of the root of $\mathcal{T}$. We have the following characterization of trees not accepted by $\mathcal{A}_{\downarrow\langle E;U\rangle}$:

Proposition 15. *If $\mathcal{T}$ is not accepted by $\mathcal{A}_{\downarrow\langle E;U\rangle}$, then*

- *either it has at least $|E| + 1$ direct subtrees accepted by no automata $\mathcal{A}_u$ for all $u \in U$,*

- *or there exists $0 \leq k \leq |E|$ such that it has at least k direct subtrees accepted by no automata $\mathcal{A}_u$ for any $u \in U$, and it does not contain $|E|$ direct subtrees witnessing the fact that $\mathcal{T}$ is accepted by $\mathcal{A}_{\downarrow E}$ and of which k subtrees are accepted by no automata $\mathcal{A}_u$ for any $u \in U$.*

Proof. We first rule out the case where E is empty: in that case, $\mathcal{T}$ is not accepted by $\mathcal{A}_{\downarrow\langle E;U\rangle}$ if, and only if, at least one of its direct subtrees is accepted by no automata $\mathcal{A}_u$, for any $u \in U$.

Now assume that E is not empty; write l for the size of E, and m for the number of direct subtrees of $\mathcal{T}$. That $\mathcal{T}$ is accepted by $\mathcal{A}_{\downarrow\langle E;U\rangle}$ means that there are l direct subtrees each accepted by some automaton $\mathcal{A}_e$ for e ranging over E, and that the remaining $m - l$ direct subtrees are accepted by some automaton $\mathcal{A}_u$, for some $u \in U$. This means that for any k, if there are at least k direct subtrees not accepted by $\mathcal{A}_u$ for any $u \in U$, then there are l direct subtrees, each accepted by some automaton $\mathcal{A}_e$ for e ranging over E, and of which at least k are accepted by no automaton $\mathcal{A}_u$ for any $u \in U$; this simply expresses the fact that those direct subtrees accepted by no automata $\mathcal{A}_u$ for any $u \in U$ must be used to fulfill the E-part of $\langle E; U \rangle$.

By duality, we get that $\mathcal{T}$ is not accepted by $\mathcal{A}_{\downarrow\langle E;U\rangle}$ if, and only if, there exists some integer k such that there are at least k direct subtrees accepted by no automata $\mathcal{A}_u$ for any $u \in U$, and there do not exist l direct subtrees, each accepted by some automaton $\mathcal{A}_e$ for e ranging over E, and of which at least k are accepted by no automata $\mathcal{A}_u$ for any $u \in U$. $\qquad\square$

Construction of the complement automaton. We now define the complement automaton $\mathcal{A}^c$ of $\mathcal{A} = (Q, q_{init}, \delta, \omega)$; we let $\mathcal{A}^c = (Q^c, q_{init}^c, \delta^c, \omega^c)$ be such that:

- Q^c is a subset of $D \cup (2^D \times \{\wedge\}) \cup (2^{(2^D \times \{\vee\})} \times \{\wedge\})$, where $D = Q \cup \overline{Q}$, and $\overline{Q} = \{\overline{q} \mid q \in Q\}$ is a set of fresh states. This defines an operator $\overline{\cdot}$ over Q, which we extend to states $\overline{x}$ of $\overline{Q}$ by letting $\overline{\overline{x}} = x$, to subsets X of D by letting $\overline{X} = \{\overline{x} \mid x \in X\}$, to states $(P, \wedge)$ of $2^D \times \{\wedge\}$ by letting $\overline{(P, \wedge)} = (\overline{P}, \vee)$. We finally extend it to $\mathsf{PBF}(Q^c)$ by letting

$$\overline{\top} = \bot \qquad \overline{\bot} = \top \qquad \overline{\psi \wedge \phi} = \overline{\psi} \vee \overline{\phi} \qquad \overline{\psi \vee \phi} = \overline{\psi} \wedge \overline{\phi}.$$

Intuitively, from a state $q \in Q$, automaton $\mathcal{A}^c$ will accept the same language as from the same state in $\mathcal{A}$, while from a state $\overline{q} \in \overline{Q}$, it will accept its complement. The language accepted by $\mathcal{A}$ from a state $(P, \wedge) \in 2^D \times \{\wedge\}$ will be the intersection of all languages accepted by $\mathcal{A}$ from all states in P, whereas the language accepted from states of the form $(\{(P_i, \vee) \mid 1 \leq i \leq k\}, \wedge) \in 2^{(2^D \times \{\vee\})} \times \{\wedge\}$ will be the intersection (over i) of the unions of the languages accepted from all states in P_i. Notice that the indications of $\wedge$ and $\vee$ are mainly used for the sake of clarity.

- accordingly, $q_{init}^c = \overline{q_{init}}$, as we want $\mathcal{A}^c$ to accept the complement of the language accepted by $\mathcal{A}$ from q_{init};

- δ^c is defined as follows, for $q \in Q$, $P \in 2^D \cup 2^{2^D}$, and $\sigma \in \Sigma$:

$$\delta^c(q, \sigma) = \delta(q, \sigma) \qquad\qquad \delta^c(\overline{q}, \sigma) = \overline{\delta(q, \sigma)}$$

$$\delta^c((P, \wedge), \sigma) = \bigwedge_{r \in P} \delta^c(r, \sigma) \qquad \delta^c((P, \vee), \sigma) = \bigvee_{r \in P} \delta^c(r, \sigma)$$

  and, following the developments above:

$$\overline{\langle E; U \rangle} = \langle (\overline{U}, \wedge) \mapsto |E| + 1; \{q_\top\} \rangle \vee$$
$$\bigvee_{0 \leq k \leq |E|} \left(\langle (\overline{U}, \wedge) \mapsto k; \{q_\top\} \rangle \wedge \bigwedge_{\substack{m \sqsubseteq E \\ |m| = k}} \Phi_{(E \setminus m) \uplus m^{\overline{U}}} \right)$$

  where $m^{\overline{U}}$ is the multiset defined by $\{(\{x\} \cup \overline{U}, \wedge) \mapsto m(x)\}_{x \in \mathsf{supp}(m)}$, and Φ_{G_m} (defined as above) characterises the failure to satisfy $\langle G_m; \{q_\top\} \rangle$, thus the existence of a minimal blocking pair:

$$\Phi_{G_m} \stackrel{\text{def}}{=} \bigvee_{F \sqsubseteq G_m} \bigvee_{g \in G_m \setminus F} \langle F; \{(\mathsf{supp}(\overline{F}) \cup \{\overline{g}\}, \wedge)\} \rangle.$$

  Notice also that F may contain states of the form $(P, \wedge)$ with $P \subseteq D$, since Φ is used as $\Phi_{(E \setminus m) \uplus m^{\overline{U}}}$; then $\overline{F}$ contains states of the form $\overline{(P, \wedge)}$, which we rewrite as $(\overline{P}, \vee)$. This way, δ^c is defined for any state of Q^c and only involves states of Q^c.

  Finally, observe that, as claimed in Remark 2, this construction only introduces new EU-pairs of the form $\langle E; \{u\} \rangle$.

- priorities are defined as follows

$$\omega^c(q) = \omega(q) \qquad \omega^c(\overline{q}) = \omega(q) + 1$$
$$\omega^c(P, op) = \max\{\omega(q) + 1 \mid q \in Q\} \quad \text{for all } (P, op) \in Q^c$$

Examples. We illustrate our construction on simple examples. We begin with $\overline{q_\top}$: we have $\delta^c(\overline{q_\top}, \sigma) = \overline{\langle \varnothing; \{q_\top\} \rangle}$; this gives rise to $\langle (\{q_\top\}, \wedge) \mapsto 1; \{q_\top\} \rangle \vee \Phi_\varnothing$. Formula $\Phi_\varnothing$ is equivalent to false, so that we end up with $\langle \overline{q_\top} \mapsto 1; \{q_\top\} \rangle$. Moreover, $\omega^c(\overline{q_\top}) = 1$, so that as soon as state $\overline{q_\top}$ appear in some execution tree, that execution tree cannot be accepting.

We now consider the simple EU-pairs of the form $\overline{\langle q \mapsto 1; \{q_\top\} \rangle}$ and $\overline{\langle \varnothing; \{q\} \rangle}$, which corresponds to transitions $\Diamond q$ and $\Box q$ in $\{\Box, \Diamond\}$-automata [Wil99].

The EU-constraint $\overline{\langle q \mapsto 1; \{q_\top\} \rangle}$ gives rise to the disjunction of
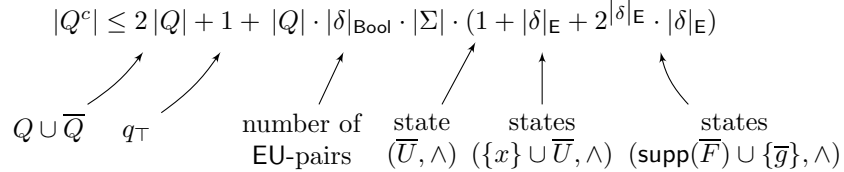
- $\langle (\{\overline{q_\top}\}, \wedge) \mapsto 2; \{q_\top\} \rangle$, which, according to the previous example, can never give rise to an accepting execution tree;

- two formulas of the form $\langle \overline{q_\top} \mapsto k; \{q_\top\} \rangle \wedge \bigwedge_{m \sqsubseteq \{q\}; |m| = k} \Phi_{(\{q\} \setminus m) \uplus m^{\overline{U}}}$, for $k \in \{0, 1\}$. The first part of this formula cannot appear in an accepting execution tree if $k = 1$, so this case simplifies to $\Phi_{\{q\}}$, which in turn simplifies to $\langle \varnothing; \{\overline{q}\} \rangle$. This corresponds to formula $\Box \overline{q}$, which intuitively means that all direct subtrees must be rejected by $\mathcal{A}_q$, as expected.

Similarly, $\overline{\langle\varnothing;\{q\}\rangle}$ is the disjunction of $\langle\overline{q}\mapsto 1;\{q_\top\}\rangle$ and $\langle\overline{q}\mapsto 0;\{q_\top\}\rangle\wedge\Phi_\varnothing$, which as already seen above cannot result in an accepting execution tree. Hence $\overline{\langle\varnothing;\{q\}\rangle}$ is equivalent to $\langle\overline{q}\mapsto 1;\{q_\top\}\rangle$, which corresponds to $\Diamond\overline{q}$.

Size of $\mathcal{A}^c$. We now evaluate the size of the complement automaton:

- *(over)approximating the number of states:* all states of $\mathcal{A}^c$ are in $D\cup(2^D\times\{\wedge\})\cup(2^{(2^D\times\{\vee\})}\times\{\wedge\})$, but not all states of this set are reachable. The reachable states are either in D, or they appear in $\overline{\langle E;U\rangle}$ for some EU-pair $\langle E;U\rangle$ in δ.

  Take an EU-pair $\langle E;U\rangle$ in δ. Besides $(\overline{U},\wedge)$ and $q_\top$, $\overline{\langle E;U\rangle}$ contains a formula of the form Φ_{G_m}, where $G_m=(E\setminus m)\uplus m^{\overline{U}}$, for each submultiset m of E. Formula Φ_{G_m} involves states of G_m, which are either in E (hence already in Q) or of the form $(\{x\}\cup\overline{U},\wedge)$ for $x\in E$, and states of the form $(\mathsf{supp}(\overline{F})\cup\{\overline{g}\},\wedge)$ where F is a submultiset of G_m and $g\in G_m\setminus F$. In the end, using $|Q|\cdot|\Sigma|\cdot|\delta|_{\mathsf{Bool}}$ as an upper bound on the number of EU-pairs,

$$|Q^c|\leq 2\,|Q|+1+|Q|\cdot|\delta|_{\mathsf{Bool}}\cdot|\Sigma|\cdot(1+|\delta|_{\mathsf{E}}+2^{|\delta|_{\mathsf{E}}}\cdot|\delta|_{\mathsf{E}})$$

$$Q\cup\overline{Q}\quad q_\top\qquad\begin{array}{c}\text{number of}\\\text{EU-pairs}\end{array}\quad\begin{array}{c}\text{state}\\(\overline{U},\wedge)\end{array}\quad\begin{array}{c}\text{states}\\(\{x\}\cup\overline{U},\wedge)\end{array}\quad\begin{array}{c}\text{states}\\(\mathsf{supp}(\overline{F})\cup\{\overline{g}\},\wedge)\end{array}$$

  Hence $|Q^c|\leq 1+|Q|\cdot(2+|\Sigma|\cdot|\delta|_{\mathsf{Bool}}\cdot|\delta|_{\mathsf{E}}\cdot(1+2^{|\delta|_{\mathsf{E}}}))$.

- *bounding the size of transition function:* in order to evaluate the size of the transition function, we group the transitions according to the type of their source states. We write δ_Q^c, $\delta_{\overline{Q}}^c$, $\delta_{(\overline{U},\wedge)}^c$, $\delta_{(\mathsf{supp}(\overline{F})\cup\{\overline{g}\},\wedge)}^c$, and δ_F^c for the five categories, which we describe and bound below:

  - δ_Q^c is the set of transitions from states in Q. We have $|\delta_Q^c|_{\mathsf{Bool}}=|\delta|_{\mathsf{Bool}}$, and $|\delta_Q^c|_{\mathsf{E}}=|\delta|_{\mathsf{E}}$ and $|\delta_Q^c|_{\mathsf{U}}=|\delta|_{\mathsf{U}}$, since $\delta^c(q,\sigma)=\delta(q,\sigma)$;

  - $\delta_{\overline{Q}}^c$ contains all transitions from states in $\overline{Q}$; writing $\delta(q,\sigma)=\psi(\langle E_i;U_i\rangle_i)$, then $\delta^c(\overline{q},\sigma)=\overline{\psi}(\overline{\langle E_i;U_i\rangle}_i)$. Hence the boolean size $|\delta_{\overline{Q}}^c|_{\mathsf{Bool}}$ of $\delta_{\overline{Q}}^c$ is at most $|\delta|_{\mathsf{Bool}}\cdot(1+|\delta|_{\mathsf{E}}\cdot(1+|\delta|_{\mathsf{E}}\cdot 2^{2|\delta|_{\mathsf{E}}}))$. Moreover, we have $|\delta_{\overline{Q}}^c|_{\mathsf{E}}\leq|\delta|_{\mathsf{E}}+1$ and $|\delta_{\overline{Q}}^c|_{\mathsf{U}}\leq\max(|\delta|_{\mathsf{U}},1)$.

  - transitions in $\delta_{(\overline{U},\wedge)}^c$ originate from states of the form $(\overline{U},\wedge)$. Since $U\subseteq Q$, those transitions are conjunctions of at most $|Q|$ transitions in $\delta_{\overline{Q}}^c$. Hence $|\delta_{(\overline{U},\wedge)}^c|_{\mathsf{Bool}}\leq|Q|\cdot|\delta_{\overline{Q}}^c|_{\mathsf{Bool}}$, and $|\delta_{(\overline{U},\wedge)}^c|_{\mathsf{E}}\leq|\delta_{\overline{Q}}^c|_{\mathsf{E}}$ and $|\delta_{(\overline{U},\wedge)}^c|_{\mathsf{U}}\leq|\delta_{\overline{Q}}^c|_{\mathsf{U}}$.

  - transitions in $\delta_{(\mathsf{supp}(\overline{F})\cup\{\overline{g}\},\wedge)}^c$ originate from states of the form $(\mathsf{supp}(\overline{F})\cup\{\overline{g}\},\wedge)$. From each such state, transitions are conjunctions of at most $|Q|+1$ subformulas, which in the worst case can be disjunctions of at most $|\delta|_{\mathsf{E}}$ transitions from states in $\overline{Q}$ (corresponding to states belonging to $2^{(2^D\times\{\vee\})}\times\{\wedge\}$). Hence we have $|\delta_{(\mathsf{supp}(\overline{F})\cup\{\overline{g}\},\wedge)}^c|_{\mathsf{Bool}}\leq(|Q|+1)\cdot|\delta|_{\mathsf{E}}\cdot|\delta_{\overline{Q}}^c|_{\mathsf{Bool}}$, and $|\delta_{(\mathsf{supp}(\overline{F})\cup\{\overline{g}\},\wedge)}^c|_{\mathsf{E}}$ and $|\delta_{(\mathsf{supp}(\overline{F})\cup\{\overline{g}\},\wedge)}^c|_{\mathsf{U}}$ are bounded by $|\delta_{\overline{Q}}^c|_{\mathsf{E}}$ and $|\delta_{\overline{Q}}^c|_{\mathsf{U}}$, respectively.

– finally, δ_F^c contains the transitions from states in $F \sqsubset G_m$ in Φ_{G_m}. Since Φ_{G_m} is used with G_m of the form $E \setminus m \uplus m^{\overline{U}}$, the transition in δ_F^c are either transitions from states in E, or transitions from states of the form $(\{x\} \cup \overline{U}, \wedge)$. The former case corresponds to transitions already in δ; the latter case gives rise to conjunctions of at most $|\delta|_{\mathsf{U}}$ transitions in $\delta_{\overline{Q}}^c$ and one transition in δ. Thus $|\delta_F^c|_{\mathsf{Bool}} \leq |\delta|_{\mathsf{U}} \cdot |\delta_{\overline{Q}}^c|_{\mathsf{Bool}} + |\delta|_{\mathsf{Bool}}$, and $|\delta_F^c|_{\mathsf{E}} \leq \max(|\delta|_{\mathsf{E}}, |\delta_{\overline{Q}}^c|_{\mathsf{E}})$ and $|\delta_F^c|_{\mathsf{U}} \leq \max(|\delta|_{\mathsf{U}}, |\tilde{\delta}_{\overline{Q}}^c|_{\mathsf{U}})$.

We end up with $|\delta^c|_{\mathsf{Bool}} \leq 4 \cdot (1+|Q|) \cdot |\delta|_{\mathsf{Bool}} \cdot |\delta|_{\mathsf{E}}^3 \cdot 2^{2|\delta|_{\mathsf{E}}}$ and $|\delta^c|_{\mathsf{E}} \leq |\delta|_{\mathsf{E}} + 1$ and $|\delta^c|_{\mathsf{U}} \leq \max(|\delta|_{\mathsf{U}}, 1)$.

Note that the complementation operation over alternating fixed-arity tree automata does not induce such a complexity blow-up: the construction is performed by using the dual of the transition function and by incrementing the priority of all states. This is an important difference compared to our construction, which is due to the expressiveness (and succinctness) of AEUTA: for example, we can easily express that there are at least k successors that are accepted by some state q with the constraint $\langle q \mapsto k; \{q_\top\}\rangle$; this cannot be expressed with $\{\Box, \Diamond\}$-automata, and it requires a much more complex formula in fixed-arity tree automata (for each arity d, we have to consider all possible subsets of successors of size k).

Correctness proof. We can now state and prove the correctness of the construction:

Proposition 16. *For any Σ-labelled $\mathcal{D}$-tree $\mathcal{T} = (t, l)$, any node n in t and any state q in Q, we have:*

$$\mathcal{T}_n \notin \mathcal{L}(\mathcal{A}, q) \quad \Leftrightarrow \quad \mathcal{T}_n \in \mathcal{L}(\mathcal{A}^c, \overline{q})$$

Proof. We use the game semantics in order to prove both implications. In order to avoid confusions, we name Player 0 and Player 1 the players in $\mathcal{G}_{\mathcal{A}, \mathcal{T}}$, and Player 0^c and Player 1^c the players in $\mathcal{G}_{\mathcal{A}^c, \mathcal{T}}$.

We take a tree $\mathcal{T} = (t, l)$, a node $n \in t$, and a state q of $\mathcal{A}$ such that $\mathcal{T}_n$ is not accepted by $\mathcal{A}_q$. We will prove that $\mathcal{T}_n$ is accepted by $\mathcal{A}_{\overline{q}}^c$.

By Prop. 7 and 8, Player 1 has a memoryless winning strategy σ_1 from (n, q) in $\mathcal{G}_{\mathcal{A}, \mathcal{T}}$. Using this strategy, we build a memoryless strategy σ_0^c for Player 0^c from $(n, \overline{q})$ in $\mathcal{G}_{\mathcal{A}^c, \mathcal{T}}$, which we then prove is winning. By Prop. 8, this entails our result.

From vertex $(n, \overline{q})$, there is only one edge, to $(n, \delta^c(\overline{q}, l(n)))$. By definition, $\delta^c(\overline{q}, l(n))$ is of the form $\overline{\phi}$, where $\phi = \delta(q, l(n))$. Moreover, Player 1 wins from vertex (n, ϕ). Using the memoryless strategy σ_1 of Player 1 from (n, ϕ), we define the strategy of Player 0^c from $(n, \overline{\phi})$, by considering the different possible forms of ϕ:

- if $\phi = \bot$, then $(n, \delta^c(\overline{q}, l(n)))$ is the state $(n, \top)$, which belongs to Player 1^c (and has no outgoing transitions). Notice that we cannot have $\phi = \top$ since Player 1 does not have a winning strategy from $(n, \top)$;

- if $\phi = \bigvee_i \phi_i$, then $\overline{\phi}$ is a conjunction, so that $(n, \overline{\phi})$ belongs to Player 1^c;

- if $\phi = \bigwedge_i \phi_i$, then $\overline{\phi}$ is a disjunction, and $(n, \overline{\phi})$ belongs to Player 0^c. Her strategy consists in following Player 1's move from (n, ϕ) in $\mathcal{G}_{\mathcal{A}, \mathcal{T}}$: letting $(n', s) = \sigma_1(n, \phi)$, we define $\sigma_0^c(n, \overline{\phi}) = (n', \overline{s})$;

- if $\phi = \langle E; U \rangle$: as Player 1 has a winning strategy from (n, ϕ), we know that for *any* possible move (chosen by Player 0) of the form (n, ν) where ν is a marking of $\mathsf{succ}(n)$ by states in $\mathsf{supp}(E) \cup U$ satisfying $\mathsf{img}(\nu) \models \langle E; U \rangle$, Player 1 can choose a winning state $(n', \nu(n'))$. This means that there is no way to satisfy $\langle E; U \rangle$ with the successors of n, and by Prop. 15, this is due to one (or both) of the following cases:

  - there are (at least) $|E| + 1$ nodes in $\mathsf{succ}(n)$ rejected by all states in U. Let X be such a set of $|E| + 1$ nodes. Then Player 0^c may choose the first term of the disjunction in the definition of $\overline{\langle E; U \rangle}$ and then select the move leading to (n, ν_c) with $\nu_c(n') = (\overline{U}, \wedge)$ if $n' \in X$, and $\nu_c(n') = q_\top$ if $n' \notin X$. From (n, ν_c), any successor (chosen by Player 1^c) will be of the form $(n', \overline{q})$ with $n' \in X$ and $q \in U$, or $(n', q_\top)$ with $m \notin X$;

  - there exists an integer k, with $0 \leq k \leq |E|$, such that there are at least k successors of n that are rejected by all states in U, and there is no way to satisfy E with any set of successors containing at least k nodes rejected by all states in U. Take such an integer k, and a set X of k successors of n rejected by all states in U. Then Player 0^c can choose the term corresponding to k in the second part of the formula defining $\overline{\langle E; U \rangle}$. This term is a conjunction, so that Player 1^c decides which successor to move to:

    * if Player 1^c decides to move to state $(n, \langle (\overline{U}, \wedge) \mapsto k; \{q_\top\} \rangle)$, then Player 0^c moves to (n, ν) where ν is the marking mapping the k nodes in X to $(\overline{U}, \wedge)$), and the other nodes to $q_\top$.
    * if Player 1^c chooses $(n, \Phi_{(E \setminus m) \uplus m^{\overline{U}}})$ for some set m of k states, then Player 0^c can choose a blocking pair (F, g) in the subsequent disjunctive states. By Prop. 11, such a blocking pair exists, since it is not possible to satisfy $E \setminus m \uplus m^{\overline{U}}$ (because it is not possible to satisfy E with k successors of n rejected by all states of U, in particular with the k successors in m).
    By Prop. 14, Player 0^c can choose a minimal blocking pair (F, g) such that all successor nodes not used to satisfy F are rejected by all states in $\mathsf{supp}(F) \cup \{g\}$. Player 0^c can then choose a unitary marking ν accordingly.

We apply the same construction from any state of the form $(n', \overline{q}')$. The resulting memoryless strategy for Player 0^c is winning: first observe that any infinite play from $(n, \overline{q})$ following this strategy visits states of the form (n, s) with $s \in Q \cup \overline{Q} \cup \{q_\top\}$ infinitely many times. By definition of ω^c, the other states have no influence on the acceptance status of infinite plays.

Finally, every execution in the game either ends in the game $\mathcal{G}_{\mathcal{A}, \mathcal{T}}$ in a configuration winning for Player 0^c, or it stays in the configurations $(n, \overline{q})$ whose parity is winning for Player 0^c as its corresponding play in the game $\mathcal{G}_{\mathcal{A}, \mathcal{T}}$ are winning for Player 1.

Conversely, we now show how to define a winning strategy for Player 1 in the game $\mathcal{G}_{\mathcal{A},\mathcal{T}}$ from a memoryless winning strategy for Player 0^c in $\mathcal{G}_{\mathcal{A}^c,\mathcal{T}}$. The main case remains the correspondence between $(n, \langle E; U \rangle)$ and $(n, \overline{\langle E; U \rangle})$. As Player 0^c has a winning strategy, she can choose some term in the disjunction $\overline{\langle E; U \rangle}$; there are two cases:

- if Player 0^c chooses the move leading to $(n, \langle (\overline{U}, \wedge) \mapsto |E| + 1; \{q_\top\} \rangle)$, and then a move (n, ν_c) where ν_c maps $|E| + 1$ nodes in $\mathsf{succ}(n)$ to the state $(\overline{U}, \wedge)$: we then let Y be this set of nodes. Note that Player 0^c has a winning strategy from any node $(y, \overline{q})$ for any $q \in U$ and $y \in Y$. Now consider the node $(n, \langle E; U \rangle)$ in $\mathcal{G}_{\mathcal{A},\mathcal{T}}$. Every move of Player 0 leads to a node of the form (n, ν) where ν is a mapping from $\mathsf{succ}(n)$ to $\mathsf{supp}(E) \cup U$ in order to satisfy the constraint $\langle E; U \rangle$. For at least one $y \in Y$, we have $\nu(y) \in U$ (only k nodes are used to fulfil E); then Player 1 can select the move to $(y, \nu(y))$ in order to keep simulating the (winning) strategy of Player 0^c from $(y, \overline{\nu(y)})$.

- if Player 0^c chooses a term of the disjunction corresponding to some k (with $0 \leq k \leq |E|$): then Player 0^c has a winning strategy both from $(n, \langle (\overline{U}, \wedge) \mapsto k; \{q_\top\} \rangle)$ and from any $(n, \Phi_{(E \setminus m) \uplus m \overline{U}})$ with $m \sqsubseteq E$ with $|m| = k$. Let Y be the nodes in $\mathsf{succ}(n)$ that Player 0^c can select in order to satisfy the first part of $\langle (\overline{U}, \wedge) \mapsto k; \{q_\top\} \rangle$. Now consider a move of Player 0 from $(n, \langle E; U \rangle)$ in $\mathcal{G}_{\mathcal{A},\mathcal{T}}$ leading to some (n, ν). If a node $y \in Y$ is associated with a state in U and Player 1 will have a winning strategy because Player 0^c is winning from $(y, \overline{\nu(y)})$. Otherwise all nodes in Y are associated with states in $\mathsf{supp}(E)$, but then we know that Player 0^c has a winning strategy from $(n, \Phi_{(E \setminus m) \uplus m \overline{U}})$ with $m = \uplus_{y \in Y} \nu(y)$ and this ensures that the constraint $\langle E; U \rangle$ is not satisfied from (n, ν) in $\mathcal{G}_{\mathcal{A},\mathcal{T}}$.

With arguments similar to the previous case, we can prove that the resulting strategy is winning for Player 1. $\qquad\square$

To summarise our results:

Theorem 17. *Given an* AEUPTA $\mathcal{A} = (Q, q_{init}, \delta, \omega)$, *we can build an* AE-UPTA $\mathcal{A}^c$ *recognising the complement of* $\mathcal{L}(\mathcal{A})$, *with size at most* $(1 + |Q| \cdot (2 + |\Sigma| \cdot |\delta|_{Bool} \cdot |\delta|_E \cdot (1 + 2^{|\delta|_E})), 4 \cdot (1 + |Q|) \cdot |\delta|_{Bool} \cdot |\delta|_E^3 \cdot 4^{|\delta|_E}, |\delta|_E + 1, \max(|\delta|_U, 1), |\omega| + 1)$.

4.4 Alternation removal (a.k.a. simulation)

Building a non-alternating automaton equivalent to a given alternating automaton is an important construction, e.g. in order to perform projection, or for algorithmic purposes. In this section, we present an *alternation-removal* (a.k.a. *simulation*) algorithm, based on ideas developed in [Wal02, Zan12] for MSO-automata.

For the rest of this section, we fix an AEUPTA $\mathcal{A} = (Q, q_{init}, \delta, \omega)$. Intuitively, (conjunctive) alternation consists in exploring each subtree in several states of the automaton. In order to remove alternation, we follow the classical approach of building a kind of powerset automaton which, instead of visiting a single node in different states $q_{i_1}, ..., q_{i_k}$ of Q, explores that node in a single *macro-state*, corresponding to the union of all states $q_{i_1}, ..., q_{i_k}$.

We illustrate this construction in Example 6, where we show why we need as a first step to keep track of the *origin* of each state (of Q) appearing in a macro-state, in order to be able to evaluate the acceptance condition.

Example 6. *Let $\Sigma = \{a, b\}$. Consider an* **AEUPTA** *$\mathcal{A}$ with an initial state q_{init} with $\omega(q_{init}) = 1$, and a state q_1 with $\omega(q_1) = 0$, and*

$$\delta(q_{init}, a) = \langle q_{init} \mapsto 1; \varnothing \rangle \wedge \langle q_1 \mapsto 1; \varnothing \rangle$$
$$\delta(q_{init}, b) = \langle q_1 \mapsto 1; \varnothing \rangle$$
$$\delta(q_1, a) = \delta(q_1, b) = \langle q_{init} \mapsto 1; \varnothing \rangle.$$

Notice that this automaton only accepts trees with a single branch (i.e., words), because all the constraints are of the form $\langle q \mapsto 1; \varnothing \rangle$. It is easily seen that the word $a^3 \cdot b^\omega$ is accepted, while a^ω is not. If we perform a simple powerset construction, the sequence of sets of states along the (unique) computation for both words is $\{q_{init}\} \cdot \{q_{init}, q_1\}^\omega$. This does not keep enough information to decide if a run is accepting.

Now, if each state is paired with its ancestor (arbitrarily pairing the initial state with itself), then the sequence of sets of pairs of states visited along $a \cdot b^\omega$ is $\{(q_{init}, q_{init})\} \cdot \{(q_{init}, q_{init}), (q_{init}, q_1)\} \cdot \{(q_1, q_{init}), (q_1, q_1)\}^\omega$, while along a^ω it is $\{(q_{init}, q_{init})\} \cdot \{(q_{init}, q_{init}), (q_{init}, q_1)\} \cdot \{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_1)\}^\omega$. In the latter sequence, we can detect the presence of an infinite branch looping in q_{init}. Figure 3 illustrates this difference. ◁

Our construction follows this intuition. It consists in four steps, represented in Fig. 4: the first step just consists in pairing states with their predecessors, as we just illustrated; the second step builds an (alternating) powerset automaton, involving a new satisfaction relation $\models\!\!\models$ and a new acceptance condition; the third step is our main step, where we (inductively) remove conjunctions from the transition function, until it is non-alternating, which allows us to come back to our original satisfaction relation $\models$; the fourth step turns the acceptance condition back into a parity condition, by taking a product with an auxiliary parity word automaton $\mathcal{M}_\omega$ enforcing acceptance along each branch.

4.4.1 Keeping track of ancestor states

In this section, we modify automaton $\mathcal{A}$ so as to store, in each state, its ancestor state. For any state $q' \in Q$, we define the mapping $\phi_{q'} \colon Q \to Q^2$ as $\phi_{q'}(q) = (q', q)$, and extend it to (multi-)sets of states, EU-pairs and EU-constraints in the natural way.

We then define the **AEUPTA** $\mathcal{P} = (Q^2, (q_{init}, q_{init}), \gamma, \omega')$ with $\gamma((q, q'), \sigma) = \phi_{q'}(\delta(q', \sigma))$, and $\omega'(q, q') = \omega(q')$. Intuitively, state (q, q') in $\mathcal{P}$ corresponds to state q' in $\mathcal{A}$, with the extra information that this state originates from state q. Notice that both γ and ω' only depend on the second state of the pair (q, q').

Proposition 18. *The languages of $\mathcal{A}$ and $\mathcal{P}$ are equal. Moreover, if a tree $\mathcal{T}$ is accepted by $\mathcal{P}$, then there is an accepting execution tree $\mathcal{U} = (u, \ell)$ of $\mathcal{P}$ on $\mathcal{T}$ in which the subtrees rooted at any two nodes n_u and n'_u for which $\ell(n_u) = (m, (q', q))$ and $\ell(n'_u) = (m, (q'', q))$ are equal. The size of $\mathcal{P}$ is $(|Q|^2, |\delta|_{Bool}, |\delta|_E, |\delta|_U, |\omega|)$.*

input tree | execution tree | powerset of states | powerset of pairs of states

q_{init}

a

q_{init} q_1 $\quad$ $\{q_{init},q_1\}$ $\quad$ $\{(q_{init},q_{init})\}$

a

q_{init} q_1 q_{init} $\quad$ $\{q_{init},q_1\}$ $\quad$ $\{(q_{init},q_{init}),(q_{init},q_1)\}$

a

q_{init} q_1 q_{init} q_{init} q_1 $\quad$ $\{q_{init},q_1\}$ $\quad$ $\{(q_{init},q_{init}),(q_{init},q_1),(q_1,q_{init})\}$

b

q_1 q_{init} q_1 q_1 q_{init} $\quad$ $\{q_{init},q_1\}$ $\quad$ $\{(q_{init},q_{init}),(q_{init},q_1),(q_1,q_{init})\}$

b

$\{q_{init},q_1\}$ $\quad$ $\{(q_{init},q_1),(q_1,q_{init})\}$

q_{init} $\quad$ q_{init} $\quad$ $\{(q_{init},q_{init})\}$

a

q_{init} q_1 $\quad$ $\{q_{init},q_1\}$ $\quad$ $\{(q_{init},q_{init}),(q_{init},q_1)\}$

a

q_{init} q_1 q_{init} $\quad$ $\{q_{init},q_1\}$ $\quad$ $\{(q_{init},q_{init}),(q_{init},q_1),(q_1,q_{init})\}$

a

q_{init} q_1 q_{init} q_{init} q_1 $\quad$ $\{q_{init},q_1\}$ $\quad$ $\{(q_{init},q_{init}),(q_{init},q_1),(q_1,q_{init})\}$
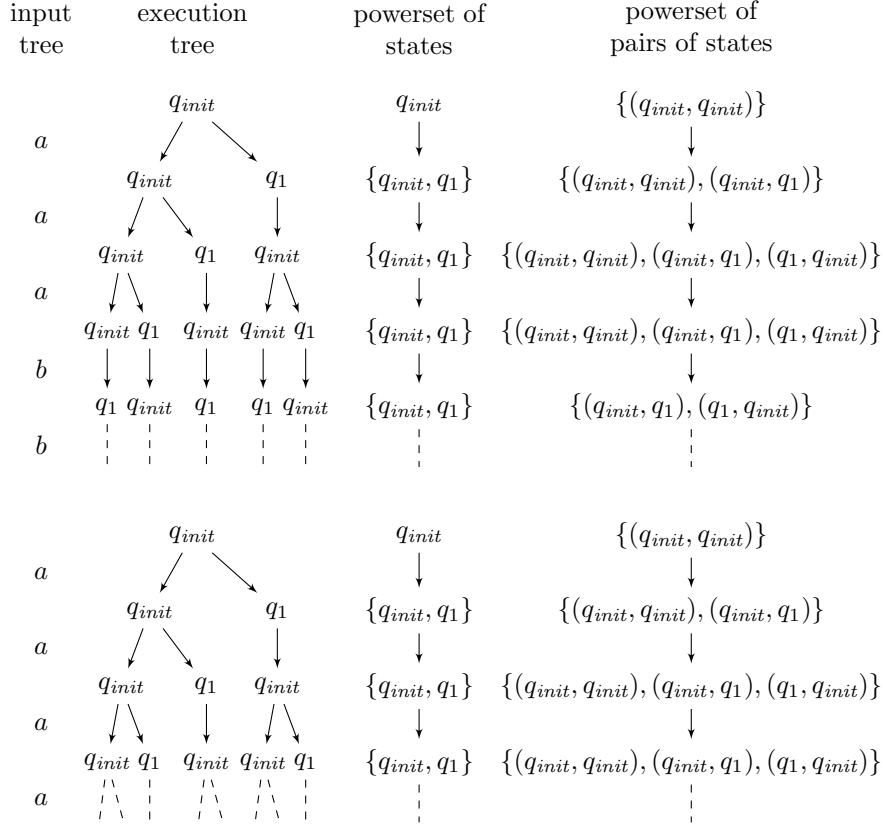
a

Figure 3: Runs of the automaton $\mathcal{A}$ of Example 6 on $a^3 \cdot b^\omega$ and on a^ω

Proof. Take a tree $\mathcal{T}$. Assuming that $\mathcal{T} \in \mathcal{L}(\mathcal{A})$, take an accepting execution tree $\mathcal{U} = (u, \ell)$ of $\mathcal{A}$ on $\mathcal{T}$. With $\mathcal{U}$, we associate another tree $\dot{\mathcal{U}} = (u, \dot{\ell})$ with the same structure, and with labelling function $\dot{\ell}$ defined as follows: for the root, $\dot{\ell}(\varepsilon_u) = (\varepsilon_t, (q_{init}, q_{init}))$ and, for any non-root node $n_u = (d_i, q_i)_{0 \leq i < |n_u|}$ (hence having $\ell(n_u) = ((d_i)_{0 \leq i < |n_u|}, q_{|n_u|-1})$),

$$\dot{\ell}(n_u) = ((d_i)_{0 \leq i < |n_u|}, (q_{|n_u|-2}, q_{|n_u|-1}))),$$

where $q_{|n_u|-2}$ is q_{init} when $|n_u| = 1$.

It should be clear that $\dot{\mathcal{U}}$ is an accepting execution tree of $\mathcal{P}$ on $\mathcal{T}$, since the only difference is the addition of the previous state of the automaton in the labelling, which has no impact on the transition function nor on the satisfaction of the acceptance condition.

Conversely, given an accepting execution tree $\dot{\mathcal{U}} = (\dot{u}, \dot{\ell})$ of $\mathcal{P}$ on $\mathcal{T}$ witnessing the fact that $\mathcal{T}$ is accepted by $\mathcal{P}$, we obtain an accepting execution tree of $\mathcal{A}$ on $\mathcal{T}$ by simply erasing the first item of the second component of the labelling function. Again, it is easily seen that this defines an accepting execution tree of $\mathcal{A}$ on $\mathcal{T}$.

Finally, if a tree $\mathcal{T}$ is accepted by $\mathcal{P}$, then it is accepted by $\mathcal{A}$, and there exists an accepting execution tree $\mathcal{U}$ of $\mathcal{A}$ on $\mathcal{T}$ such that any two nodes of $\mathcal{U}$

28

$\mathcal{A}$: original AEUPTA, using EU-constraints over Q, and relation $\models$

$\mathcal{P}$: AEUPTA, using EU-constraints over $Q \times Q$, and relation $\models$

$\mathcal{Q}$: AEUTA, using EU-constraints over $2^{Q \times Q}$, and relation $\Vdash$

$\mathcal{R}$: EUTA, using EU-constraints over $2^{Q \times Q}$, and relation $\Vdash$

$\mathcal{N}$: EUPTA, using EU-constraints over $Q_\omega \times 2^{Q \times Q}$, and relation $\models$

Figure 4: Sequence of transformations for the simulation construction

carrying the same labels are roots of the same subtrees. The result follows.

$\square$

4.4.2 Building the powerset automaton

In this section, we perform our powerset construction: we build an (alternating) AEUTA $\mathcal{Q}$ whose states are sets of states of $\mathcal{P}$. This requires two important changes in our setting: we will use a modified notion of satisfaction of EU-pairs, based on sets of states, and we will use a new acceptance condition. Notice that we do *not* remove alternation here: this will be done in the next section, and will allow us to come back to our original notion of satisfaction of EU-pairs.

From the AEUPTA $\mathcal{P} = (Q^2, (q_{init}, q_{init}), \gamma, \omega')$, we build the powerset AEUTA $\mathcal{Q} = (K, \{(q_{init}, q_{init})\}, \beta, \Omega_\omega)$ by letting:

- $K = 2^{Q^2}$ contains all the sets of states of $\mathcal{P}$, hence all the sets of pairs of states of $\mathcal{A}$;

- $\beta\big(\{(q_i, q_i') \mid 1 \leq i \leq k\}, \sigma\big) = \bigwedge_{1 \leq i \leq k} \gamma^s((q_i, q_i'), \sigma)$, where $\gamma^s((q_i, q_i'), \sigma)$ is obtained from $\gamma((q_i, q_i'), \sigma)$ by replacing each pair of states (q, q') by the singleton $\{(q, q')\}$. For the time being, this powerset automaton still is alternating.

  Notice that if we keep our definition of execution trees, then $\mathcal{P}$ and $\mathcal{Q}$ would have the same behaviours (and only *singleton states* of $\mathcal{Q}$ would be used). Below, we introduce a new notion of execution trees, which uses the same tree structure as the input tree, and gathers all states of $\mathcal{P}$ visiting a given node of the input tree into a single state of $\mathcal{Q}$ visiting that node.

- the acceptance condition Ω_ω for $\mathcal{Q}$ will be based on ω (and ω'), but it is *not* a parity acceptance condition. We define it formally below.

We call $\mathcal{A}$-*powerset AEUTA* any automaton of the form $(K, \{(q_{init}, q_{init})\}, \beta', \Omega_\omega)$, which only differs from $\mathcal{Q}$ in its transition function $\beta' \colon K \times \Sigma \to \mathsf{PBF}(\mathsf{EU}(K))$.

We now define our new notion of execution trees for $\mathcal{A}$-powerset AEUTA, based on a new notion of satisfaction for EU-pairs. This is based on identifying markings of $\mathcal{S}'$ by $\mathcal{S}$ as unitary markings of $\mathcal{S}'$ by $2^{\mathcal{S}}$.

Definition 19. *Let $\mathcal{S}$ and $\mathcal{S}'$ be two sets, $\langle E; U \rangle$ be an EU-pair over $2^{\mathcal{S}}$, and ν be a marking of $\mathcal{S}'$ by $\mathcal{S}$, seen as a unitary marking of $\mathcal{S}'$ by $2^{\mathcal{S}}$. Then ν set-satisfies $\langle E; U \rangle$, denoted $\nu \Vmodels \langle E; U \rangle$, if there exists a submarking $\nu' \sqsubseteq \nu$, seen as a unitary marking of $\mathcal{S}'$ by $2^{\mathcal{S}}$, such that $\mathsf{img}(\nu') \models \langle E; U \rangle$.*

This relation is extended to positive boolean combinations of EU-pairs in the same way as for $\models$. With this new satisfaction relation, we define a new notion of execution trees for $\mathcal{A}$-powerset automata, whose structure is the same as that of the input tree:

Definition 20. *Let $\mathcal{O} = (K, \{(q_{init}, q_{init})\}, \beta', \Omega_\omega)$ be an $\mathcal{A}$-powerset AEUTA over Σ, and $\mathcal{T} = (t, l)$ be a Σ-labelled $\mathcal{D}$-tree, for some finite set $\mathcal{D}$. An execution tree of $\mathcal{O}$ over $\mathcal{T}$ is a K-labelled $\mathcal{D}$-tree $\mathcal{U} = (u, \ell)$ such that $u = t$ and*

- *the root ε_u in $\mathcal{U}$ is labelled with $\{(q_{init}, q_{init})\}$;*

- *for any node $n_u = (d_i)_{0 \leq i < |n_u|}$ of u (which we can identify with the corresponding node $m_t = (d_i)_{0 \leq i < |m_t|}$ of the input tree), letting ν_{n_u} be the marking of $\mathsf{succ}(m_t)$ by $2^{Q \times Q}$ such that $\nu_{n_u}(m_t \cdot d) = \ell(n_u \cdot d)$, we have $\nu_{n_u} \Vmodels \beta'(\ell(n_u), l(m_t))$.*

Whether such an execution tree is accepting is defined as follows: consider an infinite branch $b = (n_{t_i})_{0 \leq i < \infty}$ of the execution tree of $\mathcal{O}$ on $\mathcal{T}$, with $\ell(n_{t_i}) = \{(q_{i,j}, q'_{i,j}) \mid 0 \leq j \leq z_i\}$ for each $i \in \mathbb{N}$. A sequence $(r_i)_{0 \leq i \leq k}$ of states of $\mathcal{A}$ is said to appear in branch b if for each $i \in \mathbb{N}$, there exists an index $0 \leq j \leq z_i$ such that $(r_{i-1}, r_i) = (q_{i,j}, q'_{i,j})$. Branch b is accepting if all the sequences $(r_i)_{0 \leq i < k}$ that appear in that branch satisfy the parity condition ω of $\mathcal{A}$; the execution tree is accepting if all its branches are.

Example 6 (contd)**.** *Consider again the automaton $\mathcal{A}$ of Example 6, and write $\mathcal{Q}$ for the $\mathcal{A}$-powerset automaton obtained from $\mathcal{A}$ by applying the transformation above. The (one-branch) trees to the right of Fig. 3 are execution trees of $\mathcal{Q}$ on (one-branch) input trees $a^3 \cdot b^\omega$ and a^ω.*

For instance, consider the second node of the execution tree of $\mathcal{Q}$ on $a^3 \cdot b^\omega$, labelled with the state $s = \{(q_{init}, q_{init}), (q_{init}, q_1)\}$ of $\mathcal{Q}$. By construction of $\mathcal{Q}$, we have

$$\beta(s, a) = \left(\langle \{(q_{init}, q_{init})\} \mapsto 1; \varnothing \rangle \wedge \langle \{(q_{init}, q_1)\} \mapsto 1; \varnothing \rangle \right) \wedge$$
$$\langle \{(q_1, q_{init})\} \mapsto 1; \varnothing \rangle,$$

where the first line corresponds to $\delta(q_{init}, a)$ and the second line corresponds to $\delta(q_1, a)$. And the third node of the execution tree indeed set-satisfies $\beta(s, a)$.

On input $a^3 \cdot b^\omega$, the only branch of the execution tree is

$$\{(q_{init}, q_{init})\}\{(q_{init}, q_{init}), (q_{init}, q_1)\}(\{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\})$$
$$(\{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\})(\{(q_{init}, q_1), (q_1, q_{init})\})^\omega$$

There are five sequences of Q^ω appearing in this branch: $q_{init}^3 \cdot (q_{init} \cdot q_1)^\omega$, $q_{init}^2 \cdot (q_{init} \cdot q_1)^\omega$, $q_{init} \cdot (q_{init} \cdot q_1)^\omega$, $q_{init} \cdot q_1 \cdot q_{init} \cdot (q_{init} \cdot q_1)^\omega$, and $(q_{init} \cdot q_1)^\omega$. All five of them are accepting w.r.t. the parity condition of $\mathcal{A}$ ($\omega(q_{init}) = 1$ and $\omega(q_1) = 0$), and thus this execution tree on $a^3 \cdot b^\omega$ is accepting.

On the other hand, on input a^ω, the unique branch of the execution tree is:

$$\{(q_{init}, q_{init})\}\{(q_{init}, q_{init}), (q_{init}, q_1)\}(\{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\})^\omega.$$

The sequences of Q^ω appearing in this branch are of the form $(q_{init}^+ \cdot q_1)^+ \cdot q_{init}^\omega$ and $(q_{init}^+ \cdot q_1)^\omega$; sequences of the former form contain only finitely many occurrences of q_1, so that this branch is not accepting. ◁

Proposition 21. *A Σ-labelled $\mathcal{D}$-tree $\mathcal{T}$ is accepted by $\mathcal{P}$ if, and only if, it is accepted by the $\mathcal{A}$-powerset AEUTA $\mathcal{Q}$. The size of $\mathcal{Q}$ is $(2^{|Q|^2}, |Q|^2 \cdot |\delta|_{Bool}, |\delta|_E, |\delta|_U, -)$ (remember that $\mathcal{Q}$ is not a parity automaton).*

Proof. Assuming that $\mathcal{T} = (t, l)$ is accepted by $\mathcal{P}$, take an accepting execution tree $\mathcal{U} = (u, \ell)$ of $\mathcal{P}$ on $\mathcal{T}$. By Prop. 18, we may assume that any two subtrees rooted at any two nodes n_u and n'_u of $\mathcal{U}$ such that $\ell(n_u) = (m, (q', q))$ and $\ell(n'_u) = (m, (q'', q))$ are the same.

Consider the tree $\mathcal{U}' = (t, \ell')$ having the same tree structure as $\mathcal{T}$ and with $\ell'(n_t) = \{(q, q') \in Q \times Q \mid \exists n_u \in u.\ \ell(n_u) = (n_t, (q, q'))\}$. Notice that the transition function is satisfied at each node: consider a node n_t whose labelling by ℓ' is a set of pairs (q, q'). Then by definition, the successors of n_t collect all the pairs (r, r') used to satisfy every EU-pair required by γ for the labels of the form $(n_t, (r, r'))$, which allows to set-satisfy β from n_t. It follows that $\mathcal{U}'$ is an execution tree of $\mathcal{Q}$ on $\mathcal{T}$.

We now prove that $\mathcal{U}'$ is accepting: take a branch $b = (n_{t_i})_{i \in \mathbb{N}}$ of $\mathcal{U}'$, with $\ell'(n_{t_i}) = \{(q_{i,j}, q'_{i,j}) \mid 0 \le j \le z_i\}$ for each $i \in \mathbb{N}$. Take a sequence $(r_i)_{i \in \mathbb{N}}$ of states that appears in b. We claim that $(n_i, (r_{i-1}, r_i))_{i \in \mathbb{N}}$, with $r_{-1} = q_{init}$, is a branch of $\mathcal{U}$. If this were not the case, consider the first index i_0 such that $(n_i, (r_{i-1}, r_i))_{i \le i_0}$ is a prefix of a branch of $\mathcal{U}$, and $(n_i, (r_{i-1}, r_i))_{i \le i_0+1}$ is not.

By the definition of $\mathcal{U}'$, we know that $\ell'(n_{t_{i_0}}) \ni (r_{i_0-1}, r_{i_0})$, and then there exists a node n_u in $\mathcal{U}$ s.t. $\ell(n_u) = (n_{t_{i_0}}, (r_{i_0-1}, r_{i_0}))$. Moreover we have that there exists $d \in \mathcal{D}$ such that $\ell'(n_{t_{i_0} \cdot d}) \ni (r_{i_0}, r_{i_0+1})$ and then there exists a node n'_u in $\mathcal{U}$ s.t. $\ell(n'_u) = (n_{t_{i_0} \cdot d}, (r_{i_0}, r_{i_0+1}))$.

The predecessor of n'_u in $\mathcal{U}$ is then labelled by some $(n_{t_{i_0}}, (s, r_{i_0}))$, and by Prop. 18, we can assume that the subtrees rooted from this node and from n_u are the same: this entails that $(n_i, (r_{i-1}, r_i))_{i \le i_0+1}$ is a prefix of a branch of $\mathcal{U}$. Therefore $\mathcal{U}'$ is accepting, and $\mathcal{T} = (t, l)$ is accepted by $\mathcal{Q}$.

Conversely, assume that $\mathcal{T} = (t, l)$ is accepted by $\mathcal{Q}$ and consider an accepting execution tree $\mathcal{U}' = (t, \ell')$ of $\mathcal{Q}$ on $\mathcal{T}$. From $\mathcal{U}'$, we build a $(t \times Q^2)$-labelled $(\mathcal{D} \times Q^2)$-tree $\mathcal{U} = (u, \ell)$ level-by-level, in such a way that it is an accepting execution tree of $\mathcal{P}$ on $\mathcal{T}$. During the inductive construction, we will maintain the invariant that for any node n_t at depth i in $\mathcal{T}$, $\ell'(n_t)$ is exactly the set of pairs (q, q') occurring in a labelling of $\mathcal{U}$-nodes at depth i of the form $(n_t, (q, q'))$.

First we define the labelling of the root: $\ell(\varepsilon_u) = (\varepsilon_t, (q_{init}, q_{init}))$. The invariant property clearly holds true at level 0.

Now consider a previously-defined node n_u of $\mathcal{U}$ labelled with $(n_t, (q, q'))$. Then by the invariant, we have $\ell'(n_t) \ni (q, q')$, and by definition of $\mathcal{Q}$, all its successors $\{n_t \cdot d \in \mathcal{T} \mid d \in \mathcal{D}\}$ are labelled by ℓ' with a set of pairs of the form (q', r) that satisfy the γ-function. We precisely add successors to n_u in order to get exactly the same labels $(n_t \cdot d, (q', r))$ for all $n_t \cdot d$ in t. This maintains the invariant and the transition function γ is locally satisfied by the definition of $\mathcal{U}$.
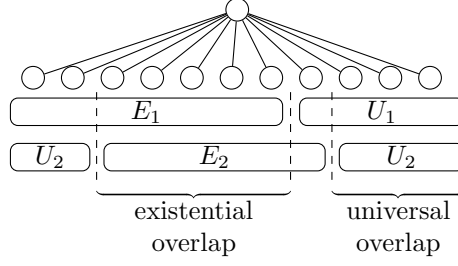
Figure 5: Representation of the overlaps in an execution tree when satisfying a conjunction of two EU-constraints $\langle E_1; U_1 \rangle$ and $\langle E_2; U_2 \rangle$.

Now we can easily see that this execution tree is accepting: consider a branch b of $\mathcal{U}$; its labelling describes a sequence $(r_i)_{i \in \mathbb{N}}$ that also appears in the corresponding branch in $\mathcal{U}'$. $\qquad\square$

4.4.3 Removing conjunctions

We now remove conjunctions from the transition function β of $\mathcal{Q}$. As a first step, we turn each formula $\beta(P, \sigma)$ in disjunctive normal form. We can bound the number of different EU-pairs appearing in any given $\beta(P, \sigma)$ by $|Q| \cdot |\delta|_{\mathsf{Bool}}$: indeed, while it is built as a conjunction of up to Q^2 transition formulas, any two pairs (q', q) and (q'', q) give rise to the same EU-pairs. It follows that $\beta(P, \sigma)$ can be written as the disjunction of at most $2^{|Q| \cdot |\delta|_{\mathsf{Bool}}}$ conjunctions of at most $|Q| \cdot |\delta|_{\mathsf{Bool}}$ EU-pairs.

We now turn those conjunctions into disjunctions. We proceed inductively, by replacing any conjunction $\langle E_1; U_1 \rangle \wedge \langle E_2; U_2 \rangle$ of two EU-pairs over $2^{Q \times Q}$ with an "equivalent" disjunction of EU-pairs over $2^{Q \times Q}$ (in the sense that the transformation preserves the language of the automaton).

Write m_1, n_1, m_2 and n_2 for the sizes of E_1, U_1, E_2 and U_2, respectively. The disjunction we build ranges over the possible ways the "existential" and "universal" parts of the EU-pairs overlap (see Fig. 5). For each combination, we write an EU-pair whose existential part contains the "existential" overlaps and the two "mixed" overlaps, and whose universal part handles the "universal" overlap.

The disjunction of EU-pairs can then be written as follows:

$$C(\langle E_1; U_1 \rangle, \langle E_2; U_2 \rangle) = \bigvee_{\substack{J_1 \sqsubseteq E_1, J_2 \sqsubseteq E_2 \\ |J_1| = |J_2|}} \bigvee_{\substack{\tau \text{ permutation} \\ \text{of } [1; |J_1|]}} \bigvee_{\substack{g_1 \colon E_1 \setminus J_1 \to U_2 \\ g_2 \colon E_2 \setminus J_2 \to U_1}}$$

$$\left(E' = \biguplus \begin{array}{l} \{\!\{ j_k^1 \cup j_{\tau(k)}^2 \mid 1 \le k \le |J_1| \}\!\} \\ \{\!\{ e_k^1 \cup g_1(e_k^1) \mid 1 \le k \le n_1 - |J_1| \}\!\} \\ \{\!\{ g_2(e_k^2) \cup e_k^2 \mid 1 \le k \le n_2 - |J_2| \}\!\} \end{array} \; ; U' = U_1 \otimes U_2 \right)$$

where we use the notations

$$J_1 = \{\!\{j_k^1 \mid 1 \le k \le o\}\!\} \qquad\qquad J_2 = \{\!\{j_k^2 \mid 1 \le k \le o\}\!\}$$
$$E_1 \setminus J_1 = \{\!\{e_k^1 \mid 1 \le k \le n_1 - o\}\!\} \qquad E_2 \setminus J_2 = \{\!\{e_k^2 \mid 1 \le k \le n_2 - o\}\!\}$$
$$U_1 \otimes U_2 = \{u_1 \cup u_2 \mid u_1 \in U_1, u_2 \in U_2\}.$$

Note that the sizes of existential and universal parts of any EU-pair in $C(\langle E_1; U_1\rangle, \langle E_2; U_2\rangle)$ are bounded by $n_1 + n_2$ and $m_1 \cdot m_2$, respectively.

Remark 5. *In case U_1 is empty (the case of U_2 would of course be symmetric), then the only possible overlaps are between E_1 and $\langle E_2; U_2\rangle$. This is reflected in our formula by considering that, when U_1 is empty, there exist no functions $g_2 \colon E_2 \setminus J_2 \to U_1$ when $E_2 \setminus J_2$ is not empty, while there is a single one when $E_2 \setminus J_2$ is empty. In other terms, if U_1 is empty, we must have $J_2 = E_2$. Notice also that if U_1 is empty, then $U_1 \otimes U_2$ also is.* ◄

Example 7. *Consider an AEUTA $\mathcal{A}$ with $Q = \{q_i \mid 0 \le i \le 4\}$, and assume that the transition function for q_1 and q_2 is as follows:*

$$\delta(q_1, \sigma) = \langle q_1 \mapsto 2; \{q_3\}\rangle \vee \langle q_2 \mapsto 2; \{q_2, q_3\}\rangle$$
$$\delta(q_2, \sigma) = \langle q_3 \mapsto 1; \{q_1, q_4\}\rangle$$

Now assume that after building the corresponding automata $\mathcal{P}$ and $\mathcal{Q}$, we have to deal with the state $\{\binom{q_3}{q_1}, \binom{q_4}{q_2}\}$. We get the following formula (where, for the sake of readability, brackets are omitted for singleton sets):

$$\beta\left(\{\binom{q_3}{q_1}, \binom{q_4}{q_2}\}, \sigma\right) = \left(\langle\binom{q_1}{q_1} \mapsto 2; \{\binom{q_1}{q_3}\}\rangle \vee \langle\binom{q_1}{q_2} \mapsto 2; \{\binom{q_1}{q_2}, \binom{q_1}{q_3}\}\rangle\right) \wedge$$
$$\langle\binom{q_2}{q_3} \mapsto 1; \{\binom{q_2}{q_1}, \binom{q_2}{q_4}\}\rangle$$

Turning this into disjunctive normal form gives

$$\beta\left(\{\binom{q_3}{q_1}, \binom{q_4}{q_2}\}, \sigma\right) = \left((\langle\binom{q_1}{q_1} \mapsto 2; \{\binom{q_1}{q_3}\}\rangle \wedge \langle\binom{q_2}{q_3} \mapsto 1; \{\binom{q_2}{q_1}, \binom{q_2}{q_4}\}\rangle\right) \vee$$
$$\left(\langle\binom{q_1}{q_2} \mapsto 2; \{\binom{q_1}{q_2}, \binom{q_1}{q_3}\}\rangle) \wedge \langle\binom{q_2}{q_3} \mapsto 1; \{\binom{q_2}{q_1}, \binom{q_2}{q_4}\}\rangle\right)$$

Consider the first disjunct of this formula:

$$\langle\binom{q_1}{q_1} \mapsto 2; \{\binom{q_1}{q_3}\}\rangle \wedge \langle\binom{q_2}{q_3} \mapsto 1; \{\binom{q_2}{q_1}, \binom{q_2}{q_4}\}\rangle,$$

and write E_1 for the multiset $\binom{q_1}{q_1} \mapsto 2$, E_2 for $\binom{q_2}{q_3} \mapsto 1$, $U_1 = \{\binom{q_1}{q_3}\}$ and $U_2 = \{\binom{q_2}{q_1}, \binom{q_2}{q_4}\}$.

- *In case E_1 and E_2 do not overlap (i.e., for $J_1 = J_2 = \varnothing$), the state $\binom{q_2}{q_3}$ of E_2 will be paired with the only state $\binom{q_1}{q_3}$ of U_1, and the two occurrences of $\binom{q_1}{q_1}$ required to fulfill E_1 may be paired with one of the states $\binom{q_2}{q_1}$ and $\binom{q_2}{q_4}$ of U_2. For this case we obtain a disjunction of three EU-pairs, each having $U' = \{\{\binom{q_1}{q_3}, \binom{q_2}{q_1}\}, \{\binom{q_1}{q_3}, \binom{q_2}{q_4}\}\}$ as their second component, and having the following multisets as their first component:*

$$E_1' = \{\binom{q_1}{q_1}, \binom{q_2}{q_1}\} \mapsto 2, \{\binom{q_2}{q_3}, \binom{q_1}{q_3}\} \mapsto 1$$
$$E_2' = \{\binom{q_1}{q_1}, \binom{q_2}{q_1}\} \mapsto 1, \{\binom{q_1}{q_1}, \binom{q_2}{q_4}\} \mapsto 1, \{\binom{q_2}{q_3}, \binom{q_1}{q_3}\} \mapsto 1$$
$$E_3' = \{\binom{q_1}{q_1}, \binom{q_2}{q_4}\} \mapsto 2, \{\binom{q_2}{q_3}, \binom{q_1}{q_3}\} \mapsto 1.$$

- *otherwise, one of the two occurrences of $\binom{q_1}{q_1}$ required by E_1 will overlap with the state $\binom{q_2}{q_3}$ of E_2, the other occurrence of $\binom{q_1}{q_1}$ being paired with an element of U_2. We get a disjunction of two EU-pairs, again having U' as their second component, and having the following multisets as their first component:*

$$E_4' = \{\binom{q_1}{q_1}, \binom{q_2}{q_3}\} \mapsto 1, \{\binom{q_1}{q_1}, \binom{q_2}{q_1}\} \mapsto 1$$
$$E_5' = \{\binom{q_1}{q_1}, \binom{q_2}{q_3}\} \mapsto 1, \{\binom{q_1}{q_1}, \binom{q_2}{q_4}\} \mapsto 1.$$

For this example, the resulting formula then is a disjunction of 5 EU-constraints.

$\lhd$

In order to prove correctness of this construction, we establish a correspondence between a conjunction $\langle E_1; U_1 \rangle \wedge \langle E_2; U_2 \rangle$ and its resulting formula $C(\langle E_1; U_1 \rangle, \langle E_2; U_2 \rangle)$:

Lemma 22. *Let $\mathcal{S}$ and $\mathcal{S}'$ be two finite sets, and $\langle E_1; U_1 \rangle \wedge \langle E_2; U_2 \rangle$ be a conjunction of two EU-pairs on $2^{\mathcal{S}}$. For any unitary marking ν of $\mathcal{S}'$ by $2^{\mathcal{S}}$, it holds $\nu \models\!\mid \langle E_1; U_1 \rangle \wedge \langle E_2; U_2 \rangle$ if, and only if, $\nu \models\!\mid C(\langle E_1; U_1 \rangle, \langle E_2; U_2 \rangle)$.*

Proof. Assume that $\nu \models\!\mid \langle E_1; U_1 \rangle \wedge \langle E_2; U_2 \rangle$. Then there exists two unitary submarkings ν_1 and ν_2 such that $\nu_1 \models\!\mid \langle E_1; U_1 \rangle$ and $\nu_2 \models\!\mid \langle E_2; U_2 \rangle$. We let $\nu_1 \uplus \nu_2$ be the marking such that $\nu_1 \uplus \nu_2(s') = \nu_1(s') \cup \nu_2(s')$. Then $\nu_1 \uplus \nu_2$ is a unitary submarking of ν by $2^{\mathcal{S}}$; we now prove that $\nu_1 \uplus \nu_2 \models\!\mid C(\langle E_1; U_1 \rangle, \langle E_2; U_2 \rangle)$.

For $i \in \{1, 2\}$, $\nu_i \models\!\mid \langle E_i; U_i \rangle$ means that $\mathsf{img}(\nu_i) \models \langle E_i; U_i \rangle$, which in turn means that there is a subset $S_i' \subseteq \mathcal{S}'$ such that $\nu_i(S_i') = E_i$ and $\mathsf{supp}(\nu_i(\mathcal{S}' \setminus S_i')) \subseteq U_i$. We let $O = S_1' \cap S_2'$ be the overlap between S_1' and S_2', $o = |O|$ be the size of this overlap, and $\{\!\{j_k^i \mid 1 \leq k \leq o\}\!\} = J_i = \nu_i(O)$ be (multiset) images of O by ν_i. Then the multiset $\nu_1 \uplus \nu_2(O)$ corresponds to $\{\!\{j_k^1 \cup j_{\tau(k)}^2 \mid 1 \leq k \leq o\}\!\}$ for some permutation τ. Similarly, letting $H_i = S_i' \cap (\mathcal{S}' \setminus S_{3-i}')$, the elements of $\nu_1 \uplus \nu_2(H_i)$ are unions of one set of E_i and one set of U_{3-i}, of the form $\{\!\{e_k^i \cup g_i(e_k^i) \mid 1 \leq k \leq |E_i| - o\}\!\}$ for some functions $g_i \colon E_i \setminus J_i \to U_{3-i}$. Finally, any $s' \in (\mathcal{S}' \setminus S_1) \cap (\mathcal{S}' \setminus S_2)$, $\nu_1 \uplus \nu_2(s')$ is the union of two sets in U_1 and U_2, respectively. This shows that $\nu_1 \uplus \nu_2 \models\!\mid C(\langle E_1; U_1 \rangle, \langle E_2; U_2 \rangle)$.

The converse direction is similar: assuming $\nu \models\!\mid C(\langle E_1; U_1 \rangle, \langle E_2; U_2 \rangle)$, we pick a unitary submarking ν' such that

$$\mathsf{img}(\nu') \models \left(E' = \biguplus \begin{array}{l} \{\!\{j_k^1 \cup j_{\tau(k)}^2 \mid 1 \leq k \leq |J_1|\}\!\} \\ \{\!\{e_k^1 \cup g_1(e_k^1) \mid 1 \leq k \leq n_1 - |J_1|\}\!\} \\ \{\!\{g_2(e_k^2) \cup e_k^2 \mid 1 \leq k \leq n_2 - |J_1|\}\!\} \end{array} ; U' = U_1 \otimes U_2 \right)$$

for some $J_1 = \{\!\{j_k^i \mid 1 \leq k \leq o\}\!\} \sqsubseteq E_1$ and $J_2 = \{\!\{j_k^i \mid 1 \leq k \leq o\}\!\} \sqsubseteq E_2$ of the same size, some permutation τ of $[1; o]$, and some functions $g_1 \colon E_1 \setminus J_1 \to U_2$ and $g_2 \colon E_2 \setminus J_2 \to U_1$. We fix three disjoint subsets O, H_1 and H_2 of $\mathcal{S}'$ such that

$$\nu'(O) = \{\!\{j_k^1 \cup j_{\tau(k)}^2 \mid 1 \leq k \leq |J_1|\}\!\}$$
$$\nu'(H_1) = \{\!\{e_k^1 \cup g_1(e_k^1) \mid 1 \leq k \leq |E_1| - |J_1|\}\!\}$$
$$\nu'(H_2) = \{\!\{g_2(e_k^2) \cup e_k^2 \mid 1 \leq k \leq |E_2| - |J_1|\}\!\}$$
$$\mathsf{supp}(\nu'(\mathcal{S}' \setminus (O \cup H_1 \cup H_2))) \subseteq U_1 \otimes U_2.$$

It should be clear that from ν' (hence also from ν), we can extract two submarkings ν_1 and ν_2 such that $\nu_1(O \cup H_1) = E_1$ and $\mathsf{supp}(\nu_1(\mathcal{S}' \setminus (O \cup H_1))) \subseteq U_1$, and $\nu_2(O \cup H_2) = E_2$ and $\mathsf{supp}(\nu_2(\mathcal{S}' \setminus (O \cup H_2))) \subseteq U_2$. This proves that $\nu \Vmodels \langle E_1; U_1 \rangle \wedge \langle E_2; U_2 \rangle$. $\qquad\square$

As a consequence, replacing $\langle E_1; U_1 \rangle \wedge \langle E_2; U_2 \rangle$ with $C(\langle E_1; U_1 \rangle, \langle E_2; U_2 \rangle)$ in the transition function of an $\mathcal{A}$-powerset AEUTA does not change the execution trees. Let $\mathcal{R}$ be an $\mathcal{A}$-powerset AEUTA obtained from $\mathcal{Q}$ by replacing all conjunctions $\langle E_1; U_1 \rangle \wedge \langle E_2; U_2 \rangle$ (in no specific order). Then $\mathcal{R}$ is non-alternating, and by Lemma 22:

Proposition 23. *The languages accepted by the two $\mathcal{A}$-powerset AEUTAs $\mathcal{Q}$ and $\mathcal{R}$ are equal.*

Moreover, since $\mathcal{R}$ is non-alternating, it only has to visit each node of the input tree in one of the states given by the transition function, so that both notions of execution trees (with $\models$ and $\Vmodels$) coincide.

Size of $\mathcal{R}$. We now evaluate the size of $\mathcal{R} = (2^{Q \times Q}, \{(q_{init}, q_{init})\}, \zeta, \Omega_\omega)$. In order to evaluate the size of the transition function ζ of $\mathcal{R}$, we first focus on the size of $C(\langle E_1; U_1 \rangle, \langle E_2; U_2 \rangle)$. For this, we define the size of an EU-pair $\langle E; U \rangle$ as the pair $(|E|, |U|)$. In the following, a set of EU-pairs is said to have a size at most (n, m) if all its EU-pairs have existential parts of size at most n and universal parts of size at most m.

Consider a conjunction $\langle E_1; U_1 \rangle \wedge \langle E_2; U_2 \rangle$ of two EU-pairs of size (n_1, m_1) and (n_2, m_2), respectively, and assume w.l.o.g. that $n_1 \leq n_2$. Then the formula $C(\langle E_1; U_1 \rangle, \langle E_2; U_2 \rangle)$ is a disjunction of N EU-pairs of size at most $(n_1 + n_2, m_1 \cdot m_2)$, with:

$$N \leq \sum_{l=0}^{n_1} \binom{n_1}{l} \cdot \binom{n_2}{l} \cdot l! \cdot m_2^{n_1 - l} \cdot m_1^{n2 - l}.$$

This formula follows from the definition of $C(\langle E_1; U_1 \rangle, \langle E_2; U_2 \rangle)$: there is one EU-pair for every possible size l of the overlap of the existential parts (which cannot exceed n_1), every submultisets $J_1 \sqsubseteq E_1$ and $J_2 \sqsubseteq E_2$ of size l, every bijection from J_1 to J_2 (so as to consider any possible combination in the overlap), and every combination outside the overlap, between the remaining states of the existential parts and the states of the universal parts.

This number N can then be overapproximated as follows:

$$\begin{aligned}
N &\leq \left(\sum_{l=0}^{n_1} \binom{n_1}{l} \cdot \frac{n_2!}{(n_2 - l)! \cdot l!} \cdot l! \right) \cdot m_2^{n_1} \cdot m_1^{n2} \\
&\leq \left(\sum_{l=0}^{n_1} \binom{n_1}{l} \cdot n_2^l \right) \cdot m_2^{n_1} \cdot m_1^{n2} \\
&= (n_2 + 1)^{n_1} \cdot m_2^{n_1} \cdot m_1^{n2}
\end{aligned}$$

We now prove that any conjunction of k EU-pairs of sizes at most (n, m) can be turned into disjunctions of at most $((k-1)! \cdot (n+1)^{k-1} \cdot m^{k^2})^n$ EU-pairs, each of size at most $(k \cdot n, m^k)$. According to our computation above, this holds true for $k = 2$.

Consider a conjunction $\mathcal{C}$ of $k + 1$ such EU-pairs, assuming that the result holds for up to k EU-pairs. Then the conjunction of the first k EU-pairs can be turned into a disjunction of at most $((k - 1)! \cdot (n + 1)^{k-1} \cdot m^{k^2})^n$ EU-pairs of size at most $(k \cdot n, m^k)$. By distributing the $(k + 1)$-th conjunction over this disjunction, we obtain an expression of $\mathcal{C}$ as the disjunction of at most $((k - 1)! \cdot (n + 1)^{k-1} \cdot m^{k^2})^n$ conjunctions of two EU-pairs, of sizes at most $(k \cdot n, m^k)$ and (n, m) respectively.

We apply our construction to each conjunction of two EU-pairs. Each such conjunction is then replaced with the disjunction of (at most) $(kn + 1)^n \cdot m^{kn} \cdot (m^k)^n$ EU-pairs of size at most $((k + 1) \cdot n, m^{k+1})$. In the end, we obtain an expression of $\mathcal{C}$ as a disjunction of at most M EU-pairs with

$$
\begin{aligned}
M &\leq \left((k - 1)! \cdot (n + 1)^{k-1} \cdot m^{k^2} \right)^n \cdot (kn + 1)^n \cdot m^{kn} \cdot (m^k)^n \\
&\leq \left(k! \cdot (n + 1)^k \cdot m^{(k+1)^2} \right)^n.
\end{aligned}
$$

We now evaluate the size of the transition function ζ: as explained at the beginning of the present section, ζ is obtained from the transition function β of $\mathcal{Q}$ by first putting each formula $\beta(P, \sigma)$ into disjunctive normal form, as the disjunction of at most $2^{|Q| \cdot |\delta|_{\mathsf{Bool}}}$ conjunctions of at most $|Q| \cdot |\delta|_{\mathsf{Bool}}$ EU-pairs, with EU-pairs of size at most $(|\delta|_{\mathsf{E}}, |\delta|_{\mathsf{U}})$.

Applying our formula above, we get a disjunctive expression for $\beta(P, \sigma)$ involving at most

$$
2^{|Q| \cdot |\delta|_{\mathsf{Bool}}} \cdot \left((|Q| \cdot |\delta|_{\mathsf{Bool}} - 1)! \cdot (|\delta|_{\mathsf{E}} + 1)^{|Q| \cdot |\delta|_{\mathsf{Bool}} - 1} \cdot (|\delta|_{\mathsf{U}})^{(|Q| \cdot |\delta|_{\mathsf{Bool}})^2} \right)^{|\delta|_{\mathsf{E}}}
$$

EU-pairs of size at most $(|Q| \cdot |\delta|_{\mathsf{Bool}} \cdot |\delta|_{\mathsf{E}}, (|\delta|_{\mathsf{U}})^{|Q| \cdot |\delta|_{\mathsf{Bool}}})$. In the end:

Proposition 24. *The languages of the original AEUPTA $\mathcal{A}$ and of the resulting EUTA $\mathcal{R}$ are the same. The size of $\mathcal{R}$ is at most[4] $(2^{|Q|^2}, 2^{\mathfrak{K}} \cdot ((\mathfrak{K} - 1)! \cdot (|\delta|_E + 1)^{\mathfrak{K}-1} \cdot (|\delta|_U)^{\mathfrak{K}^2})^{|\delta|_E}, \mathfrak{K} \cdot |\delta|_E, (|\delta|_U)^{\mathfrak{K}}, -)$ where $\mathfrak{K} = |Q| \cdot |\delta|_{Bool}$.*

Example 6 (contd). *Consider again Example 6. We describe the corresponding automaton $\mathcal{R}$. The initial state still is $\{(q_{init}, q_{init})\}$. In the following, we use $\star$ to represent any of the two states q_{init} and q_1. The previous construction provides the following transition function:*

$$
\begin{aligned}
\zeta(\{(\star, q_{init})\}, a) &= \langle \{(q_{init}, q_{init}), (q_{init}, q_1)\} \mapsto 1; \varnothing \rangle \\
\zeta(\{(\star, q_1)\}, a) &= \langle \{(q_1, q_1)\} \mapsto 1; \varnothing \rangle \\
\zeta(\{(\star, q_{init})\}, b) &= \langle \{(q_{init}, q_1)\} \mapsto 1; \varnothing \rangle \\
\zeta(\{(\star, q_1)\}, b) &= \langle \{(q_1, q_{init})\} \mapsto 1; \varnothing \rangle \\
\zeta(\{(\star, q_{init}), (\star, q_1)\}, a) &= \langle \{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_1)\} \mapsto 1; \varnothing \rangle \\
\zeta(\{(\star, q_{init}), (\star, q_1)\}, b) &= \langle \{(q_{init}, q_1), (q_1, q_{init})\} \mapsto 1; \varnothing \rangle
\end{aligned}
$$

Note that the transition function from the state $\{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_1)\}$ is the same as the one from $\{(\star, q_{init}), (\star, q_1)\}$. Note also that this transition function does not involve any disjunction because the U-parts of the transition function of $\mathcal{A}$ all are empty.

---

[4] We omit the size of the acceptance condition of $\mathcal{R}$ here as it is not a parity condition.

We then obtain the execution trees of $\mathcal{R}$ over the 1-branch trees a^ω and $a^3 \cdot b^\omega$, as depicted on Fig. 6. They (fortunately) correspond to the ones depicted to the right of Fig. 3. The execution tree on the left is not accepting: the branch q_{init}^ω
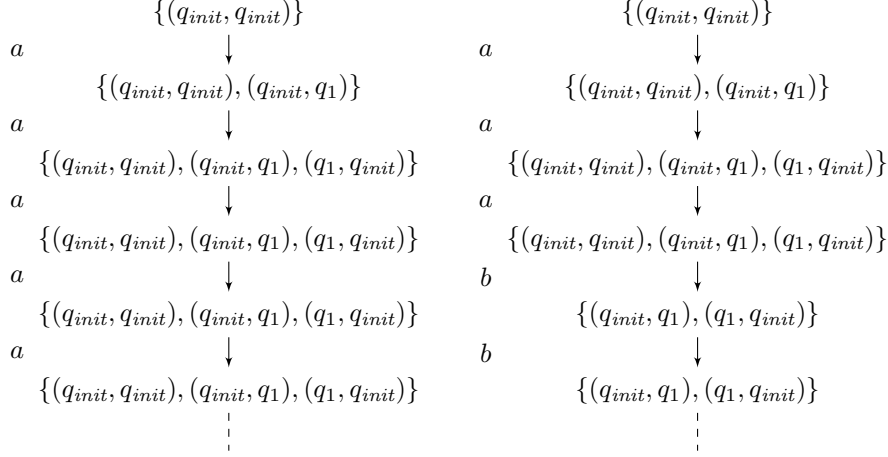

$$\{(q_{init}, q_{init})\} \qquad\qquad\qquad \{(q_{init}, q_{init})\}$$

$a \qquad\qquad \downarrow \qquad\qquad\qquad a \qquad\qquad \downarrow$

$$\{(q_{init}, q_{init}), (q_{init}, q_1)\} \qquad\qquad \{(q_{init}, q_{init}), (q_{init}, q_1)\}$$

$a \qquad\qquad \downarrow \qquad\qquad\qquad a \qquad\qquad \downarrow$

$$\{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\} \qquad \{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\}$$

$a \qquad\qquad \downarrow \qquad\qquad\qquad a \qquad\qquad \downarrow$

$$\{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\} \qquad \{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\}$$

$a \qquad\qquad \downarrow \qquad\qquad\qquad b \qquad\qquad \downarrow$

$$\{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\} \qquad \{(q_{init}, q_1), (q_1, q_{init})\}$$

$a \qquad\qquad \downarrow \qquad\qquad\qquad b \qquad\qquad \downarrow$

$$\{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\} \qquad \{(q_{init}, q_1), (q_1, q_{init})\}$$

Figure 6: Execution tree of $\mathcal{R}$ on a^ω and $a^3 \cdot b^\omega$

does not satisfies the parity condition. But in the execution tree on the right, states q_{init} and q_1 alternate along any sequence appearing in the unique branch, which ensures that the execution tree is accepting. ◁

4.4.4 Adapting the acceptance condition

It remains to turn the acceptance condition of $\mathcal{R}$ into a parity condition. The transformation is the same as in [Zan12]: we first build a non-deterministic parity word automaton $\mathcal{W}$ accepting all words on the alphabet $2^{Q \times Q}$ that contain an infinite sequence of states $(r_i)_{i \in \mathbb{N}}$ of $\mathcal{A}$ (in the sense of Def. 20) *not satisfying* the parity condition of $\mathcal{A}$; we then turn it into a deterministic parity word automaton, take its complement, and run it in parallel with $\mathcal{R}$.

Let $\mathcal{W} = (Q \cup \{q'_{init}\}, q'_{init}, \pi, \omega')$ where q'_{init} is a new state not in Q, $\pi(q'_{init}, L) = \bigvee_{(q', q) \in L} q$ and $\pi(q', L) = \bigvee_{(q', q) \in L} q$, and $\omega'(q) = \omega(q) + 1$ (the value of $\omega'(q'_{init})$ can be set arbitrarily since that state is visited only once). Intuitively, this automaton guesses a sequence of states $(r_i)_{i \in \mathbb{N}}$ contained in the input word $(L_i)_{i \in \mathbb{N}}$ on alphabet $2^{Q \times Q}$ and the parity condition of $\mathcal{W}$ ensures that this sequence does not satisfy the parity condition of $\mathcal{A}$. Note that the number of priorities remains unchanged and equal to $|\omega|$.

From $\mathcal{W}$, we can build an equivalent *deterministic* parity word automaton $\mathcal{W}_d$. For this, we first turn $\mathcal{W}$ into a non-deterministic Büchi automaton $\mathcal{W}'$ with at most $|Q| \cdot |\omega| + 1$ states. This can be achieved by considering several copies of $\mathcal{W}$: an initial one, with no accepting states, and for each even integer p less than or equal to $|\omega'|$, one copy of $\mathcal{W}$ involving only those states with priority larger than or equal to p, with exactly those states of priority p being accepting (for the Büchi condition).

We then apply the construction of [Pit07] to get a deterministic parity word automaton with $2 \cdot (|Q| \cdot |\omega| + 1)^{|Q| \cdot |\omega| + 1} \cdot (|Q| \cdot |\omega| + 1)!$ states and at most

$2 \cdot (|Q| \cdot |\omega| + 1)$ priorities. The number of states can then be bounded by $2^{1+2(|Q| \cdot |\omega|+1) \cdot \log(|Q| \cdot |\omega|+1)}$.

It remains to complement $\mathcal{W}_d$, in order to get an automaton $\mathcal{M}$ that recognises precisely all input words containing only sequences of states satisfying the parity condition of $\mathcal{A}$: these are precisely the branches that $\mathcal{R}$ has to accept. Complementing $\mathcal{W}_d$ is easy, as it consists in incrementing its priorities by 1 (leaving the number of priorities unchanged), and the resulting automaton (namely $\mathcal{M}$) is still deterministic. Let $\omega_{\mathcal{M}}$ be the priority function of $\mathcal{M}$. The number of states of $\mathcal{M}$ is then bounded by $2^{1+2(|Q| \cdot |\omega|+1) \cdot \log(|Q| \cdot |\omega|+1)}$ and the number of priorities is at most $2 \cdot (|Q| \cdot |\omega| + 1)$.

We can then run $\mathcal{R}$ and $\mathcal{M}$ in parallel, thereby obtaining a non-alternating EUPTA $\mathcal{N}$ accepting the same language as $\mathcal{A}$. The construction is as follows: the states of $\mathcal{N}$ are pairs $(\overline{q}, q)$ where $\overline{q}$ is a state of $\mathcal{R}$ and q is a state of $\mathcal{M}$. The transition function τ of $\mathcal{N}$ is defined as follows: take a state $(\overline{q}, q)$ and a letter σ, and write $\zeta(\overline{q}, \sigma) = \bigvee_{i \in I} \langle E_i; U_i \rangle$ for some set I. Then $\tau((\overline{q}, q), \sigma) = \bigvee_{i \in I} \langle E_i'; U_i' \rangle$, where each state $(\overline{q}')$ in E_i and in U_i is replaced with $(\overline{q}', \delta_{\mathcal{M}}(q, \overline{q}'))$, so that when $\mathcal{R}$ explores some successor node in state $\overline{q}'$, the state of $\mathcal{M}$ is updated accordingly. By letting the priority of $(\overline{q}, q)$ be that of q in $\mathcal{M}$, we make $\mathcal{M}$ keep track of whether all the sequences of state of $\mathcal{A}$ that appear along each branch of the execution tree of $\mathcal{R}$ are accepting for the parity condition of $\mathcal{A}$.

The size of the state space of the product automaton $\mathcal{R} \times \mathcal{M}$ is $2^{|Q|^2} \cdot 2^{1+2(|Q| \cdot |\omega|+1) \cdot \log(|Q| \cdot |\omega|+1)}$. The sizes of the transition function and of the EU-constraints are the same as those of $\mathcal{R}$, and the number of priorities is the same as for $\mathcal{M}$.

Summarising our results:

Theorem 25. *Given an AEUPTA $\mathcal{A} = (Q, q_{init}, \delta, \omega)$, we can build an EUPTA $\mathcal{N}$ recognising the same language. The size of $\mathcal{N}$ can be bounded by $(2^{1+|Q|^2+2(|Q| \cdot |\omega|+1) \cdot \log(|Q| \cdot |\omega|+1)}, 2^{\mathfrak{K}} \cdot ((\mathfrak{K}-1)! \cdot (|\delta|_E + 1)^{\mathfrak{K}-1} \cdot (|\delta|_U)^{\mathfrak{K}^2})^{|\delta|_E}, \mathfrak{K} \cdot |\delta|_E, |\delta|_U{}^{\mathfrak{K}}, 2(|Q| \cdot |\omega| + 1))$, where $\mathfrak{K} = |Q| \cdot |\delta|_{Bool}$.*

Example 6 (contd)**.** *We apply the approach above to the automaton of Example 6. We build a deterministic automaton $\mathcal{M} = (Q', q'_{init}, \pi', \omega')$ which recognises the infinite words over the alphabet $2^{Q \times Q}$ that satisfy the parity condition ω (in the sense of Def. 20). Note that we build $\mathcal{M}$ directly here, without using $\mathcal{W}$, because the parity condition we consider turns out to be equivalent to a Büchi condition, which makes the construction simpler. The states of $\mathcal{M}$ are pairs (s, s') with $s, s' \subseteq Q$: $s \cup s'$ is the set of all possible last states of sequences of states that appear in the word $w \in (2^{Q \times Q})^*$ that has been read by $\mathcal{M}$; the states in s' have recently visited an accepting state (with priority 0), while those in s have not. More formally, when reading a letter $\sigma \in 2^{Q \times Q}$, the transition function updates (s, s') into (t, t') by transforming each state q of s or s' into a state q' in t or t', for each $(q, q') \in \sigma$. All accepting states q' are placed in t'; non-accepting states q' are placed in t if they originate from a state q in s, or if s was empty, otherwise they are placed in t'. States of the form $(\varnothing, s')$ are accepting. This corresponds to the classical procedure to transform an alternating Büchi automaton into a non-deterministic one.*

In our case of Example 6, we even get a deterministic automaton:

- *$Q' = 2^{\{q_{init}\}} \times 2^{\{q_{init}, q_1\}}$ and $q'_{init} = (\{q_{init}\}, \varnothing)$;*

- $\omega'((\varnothing, s')) = 0$ and $\omega'((\{q_{init}\}, s')) = 1$ for any $s' \subseteq Q$;

- Given $s \in 2^Q$ and $\sigma \in 2^{Q \times Q}$, we use $\sigma(s)$ to denote $\{q' \mid \exists (q, q') \in \sigma \land q \in s\}$; we then define the transitions as follows:

$$\pi((\varnothing, s'), \sigma) = (\sigma(s') \cap \{q_{init}\}, \sigma(s') \setminus \{q_{init}\})$$
$$\pi((\{q_{init}\}, s'), \sigma) = (\sigma(\{q_{init}\}) \cap \{q_{init}\}, \sigma(\{q_{init}\}) \setminus \{q_{init}\} \cup \sigma(s'))$$

Now we can define $\mathcal{N}$ as the product $\mathcal{R} \times \mathcal{M}$. The states are pairs $(\bar{q}, q')$ wit $\bar{q} \in 2^{Q \times Q}$ and $q' \in 2^{\{q_{init}\}} \times 2^Q$. The initial state is $(\{\binom{q_{init}}{q_{init}}\}, (\{q_{init}\}, \varnothing))$. We simplify the transition function by removing rejecting states and we present only the reachable part (from the initial state) of the relation. We get:

$$\zeta((\{\binom{\star}{q_{init}}\}, (\{q_{init}\}, \varnothing)), a) = \langle(\{\binom{q_{init}}{q_{init}}, \binom{q_{init}}{q_1}\}, (\{q_{init}\}, \{q_1\})) \mapsto 1; \varnothing\rangle$$

$$\zeta((\{\binom{\star}{q_{init}}\}, (\{q_{init}\}, \varnothing)), b) = \langle(\{\binom{q_{init}}{q_1}\}, (\varnothing, \{q_1\})) \mapsto 1; \varnothing\rangle$$

$$\zeta((\{\binom{\star}{q_1}\}, (\varnothing, \{q_1\})), a) = \langle(\{\binom{q_1}{q_1}\}, (\varnothing, \{q_1\})) \mapsto 1; \varnothing\rangle$$

$$\zeta((\{\binom{\star}{q_1}\}, (\varnothing, \{q_1\})), b) = \langle(\{\binom{q_1}{q_{init}}\}, (\{q_{init}\}, \varnothing)) \mapsto 1; \varnothing\rangle$$

$$\zeta((\{\binom{\star}{q_{init}}, \binom{\star}{q_1}\}, (\{q_{init}\}, \{q_1\})), a) = \langle(\{\binom{q_{init}}{q_{init}}, \binom{q_{init}}{q_1}, \binom{q_1}{q_1}\}, (\{q_{init}\}, \{q_1\})) \mapsto 1; \varnothing\rangle$$

$$\zeta((\{\binom{\star}{q_{init}}, \binom{\star}{q_1}\}, (\{q_{init}\}, \{q_1\})), b) = \langle(\{\binom{q_{init}}{q_1}, \binom{q_1}{q_{init}}\}, (\varnothing, \{q_{init}, q_1\})) \mapsto 1; \varnothing\rangle$$

$$\zeta((\{\binom{\star}{q_{init}}, \binom{\star}{q_1}\}, (\varnothing, \{q_{init}, q_1\})), a) = \langle(\{\binom{q_{init}}{q_{init}}, \binom{q_{init}}{q_1}, \binom{q_1}{q_1}\}, (\{q_{init}\}, \{q_1\})) \mapsto 1; \varnothing\rangle$$

$$\zeta((\{\binom{\star}{q_{init}}, \binom{\star}{q_1}\}, (\varnothing, \{q_{init}, q_1\})), a) = \langle(\{\binom{q_{init}}{q_1}, \binom{q_1}{q_{init}}\}, (\{q_{init}\}, \{q_1\})) \mapsto 1; \varnothing\rangle$$

It remains to compare the execution trees of $\mathcal{N}$ over the word a^ω (which does not belong to the language of $\mathcal{A}$) and the word $a^3 \cdot b^\omega$ (which is accepted by $\mathcal{A}$). Figure 7 displays both execution trees; we observe that the execution tree on the left is not accepting (assuming that it continues reading a^ω), while the execution tree on the right is accepting (if it continues reading b^ω). ◁

5   Algorithms for **AEUTAs**

Given some AEUTA $\mathcal{A}$, we are interested in two decision procedures: deciding whether a regular tree belongs to $\mathcal{L}(\mathcal{A})$ and deciding whether $\mathcal{L}(\mathcal{A}) = \varnothing$. Both consist in building a parity game and deciding whether Player 0 has a winning strategy. For this we use the following results:

Proposition 26 ([Löd21]). *Solving a finite parity game can be done in time $O(n^d)$ or in time $n^{O(\log(d))}$, where n is the number of states of the game, and d is the number of priorities.*

Considering that d is usually small, we will mainly use the former result (namely $O(n^d)$) in the sequel.

$$
\begin{array}{cc}
\{(q_{init}, q_{init})\} & \{(q_{init}, q_{init})\} \\
(\{q_{init}\}, \varnothing) & (\{q_{init}\}, \varnothing) \\
a \quad \downarrow & a \quad \downarrow \\
\{(q_{init}, q_{init}), (q_{init}, q_1)\} & \{(q_{init}, q_{init}), (q_{init}, q_1)\} \\
(\{q_{init}\}, \{q_1\}) & (\{q_{init}\}, \{q_1\}) \\
a \quad \downarrow & a \quad \downarrow \\
\{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\} & \{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\} \\
(\{q_{init}\}, \{q_1\}) & (\{q_{init}\}, \{q_1\}) \\
a \quad \downarrow & a \quad \downarrow \\
\{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\} & \{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\} \\
(\{q_{init}\}, \{q_1\}) & (\{q_{init}\}, \{q_1\}) \\
a \quad \downarrow & b \quad \downarrow \\
\{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\} & \{(q_{init}, q_1), (q_1, q_{init})\} \\
(\{q_{init}\}, \{q_1\}) & (\varnothing, \{q_{init}, q_1\}) \\
a \quad \downarrow & b \quad \downarrow \\
\{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\} & \{(q_{init}, q_1), (q_1, q_{init})\} \\
(\{q_{init}\}, \{q_1\}) & (\{q_{init}\}, \{q_1\}) \\
a \quad \downarrow & b \quad \downarrow \\
\{(q_{init}, q_{init}), (q_{init}, q_1), (q_1, q_{init})\} & \{(q_{init}, q_1), (q_1, q_{init})\} \\
(\{q_{init}\}, \{q_1\}) & (\varnothing, \{q_{init}, q_1\}) \\
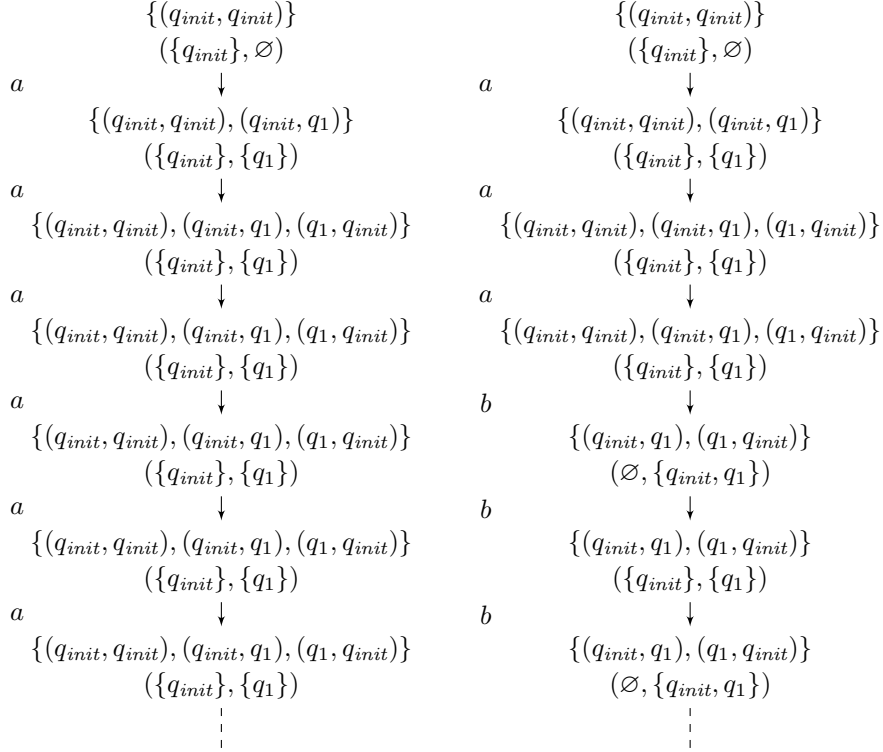\vdots & \vdots
\end{array}
$$

Figure 7: Execution tree of $\mathcal{N} = \mathcal{R} \times \mathcal{M}$ on a^ω and $a^3 \cdot b^\omega$: each state of $\mathcal{N}$ is a pair made of one state of $\mathcal{R}$ (which is a set of pairs of states of $\mathcal{A}$) and one state of $\mathcal{M}$ (which is a pair of sets of states of $\mathcal{A}$)

5.1 Membership checking

Let $\mathcal{K} = (V, E, \ell)$ be a Kripke structure. Deciding whether $\mathcal{T}_\mathcal{K} \in \mathcal{L}(\mathcal{A})$ is equivalent to deciding whether Player 0 has a winning strategy in the parity game $\mathcal{G}_{\mathcal{A}, \mathcal{K}}$ defined in Section 3.2. Remember that the number of states of $\mathcal{G}_{\mathcal{A}, \mathcal{K}}$ is in $O(|V| \cdot (|Q| \cdot |\delta|_{\mathsf{Bool}} + |Q|^{|V|}))$ and the number of priorities is the same as $\mathcal{A}$. Using Prop. 26:

Theorem 27. *Deciding whether a regular tree $\mathcal{T}_K$, defined by a finite Kripke structure $\mathcal{K} = (V, E, \ell)$, is accepted by an AEUPTA $\mathcal{A} = (Q, q_0, \delta, \omega)$ can be performed in time $O((|V| \cdot (|Q| \cdot |\delta|_{Bool} + |Q|^{|V|}))^{|\omega|})$.*

5.2 Emptiness checking

We now address emptiness checking. As for classical tree automata, we consider a non-alternating automaton $\mathcal{A}$, and transform it into a parity game $\mathcal{G}_\mathcal{A}$ such that Player 0 has winning strategy in $\mathcal{G}_\mathcal{A}$ if, and only if, $\mathcal{L}(\mathcal{A}) \neq \varnothing$. We obtain the following:

Theorem 28. *Let $\mathcal{A} = (Q, q_0, \delta, \omega)$ be an EUPTA. Checking emptiness of $\mathcal{L}(\mathcal{A})$ can be performed in time $O((|Q| \cdot (1 + |\delta|_{Bool} \cdot |\Sigma|))^{|\omega|})$.*

Proof. First we build the (non-alternating) automaton $\widehat{\mathcal{A}} = (Q, q_0, \delta', \omega)$ obtained from $\mathcal{A}$ by replacing any EU-constraint $\langle E; U \rangle$ in δ with $\langle \mathsf{supp}(E); \varnothing \rangle$. Observe that $\mathcal{L}(\mathcal{A}) = \varnothing$ if, and only if, $\mathcal{L}(\widehat{\mathcal{A}}) = \varnothing$: first, any tree accepted by $\mathcal{A}$ is also accepted by $\widehat{\mathcal{A}}$, as the latter is less constrained; on the other hand, if a tree $\mathcal{T}$ is accepted by $\widehat{\mathcal{A}}$, then by possibly duplicating subtrees in $\mathcal{T}$ (so as to address the fact that the existential parts E in $\mathcal{A}$ may contain several copies of the states of $\mathsf{supp}(E)$) and in an accepting execution tree of $\widehat{\mathcal{A}}$ on $\mathcal{T}$, we can obtain an accepting execution tree of $\mathcal{A}$ on $\mathcal{T}$.

From $\widehat{\mathcal{A}} = (Q, q_0, \delta', \omega)$, we build the parity game $\mathcal{G}_{\mathcal{A}} = (Y = Y_0 \cup Y_1, R, \theta)$ where:

- Y_0 is Q and Y_1 is the set of all EU-pair $\langle E; \varnothing \rangle$ appearing in some $\delta'(q, \sigma)$.

- R contains two kinds of edges: first, it contains an edge $(q, \langle E; \varnothing \rangle)$ if, and only if, $\langle E; \varnothing \rangle$ occurs in the disjunction $\delta'(q, \sigma)$ for some $\sigma \in \Sigma$; second, it contains an edge $(\langle E; \varnothing \rangle, q)$ if, and only if, $q \in E$.

- $\omega'(q) = \omega(q)$ for all $q \in Q$, and $\omega'(\langle E; \varnothing \rangle)$ is set to the maximum value of ω on Q (in order to have no effect on the acceptance of the play).

In this game, in state q, Player 0 has to select an EU-pair $\langle E; \varnothing \rangle$ in $\delta'(q, \sigma)$ for some $\sigma \in \Sigma$, and from a node $\langle E; \varnothing \rangle$, Player 1 may select any state in E to continue the play. This way, Player 0 builds a Σ-labelled tree step-by-step.

It remains to show that there exists some Σ-labelled tree $\mathcal{T}$ accepted by $\widehat{\mathcal{A}}$ if, and only if, Player 0 has a winning strategy in $\mathcal{G}_{\mathcal{A}}$ from q_0. The proof is based on a direct correspondence between the accepting execution tree for $\mathcal{T}$ and a winning strategy for Player 0. Indeed consider an accepting execution tree for $\mathcal{T}$: this tree has the same structure as $\mathcal{T}$, and it associates with every node n of $\mathcal{T}$ a state $q \in Q$ such that the successor nodes of n satisfy some $\langle E; \varnothing \rangle \in \delta'(q, \sigma)$. This pair $\langle E; \varnothing \rangle$ is precisely the move Player 0 should select to win from q (and whatever the choice of Player 1 in E, Player 0 will be able to continue to select winning moves).

Conversely given a winning strategy for Player 0 from a state q, one can build an accepted tree level-by-level: any move to some $\langle E; \varnothing \rangle$ corresponds to some (possibly several) letter $\sigma \in \Sigma$ s.t. $\langle E; \varnothing \rangle \in \delta(q, \sigma)$); this fixes the arity of the current node to $|E|$, and associates its $|E|$ successors with the states in E.

The number of states of $\mathcal{G}_{\mathcal{A}}$ is bounded by $|Q| + |Q| \cdot |\delta|_{\mathsf{Bool}} \cdot |\Sigma|$; the number of priorities is $|\omega|$. By using Prop. 26, we get a decision procedure in $O((|Q| \cdot (1 + |\delta|_{\mathsf{Bool}} \cdot |\Sigma|))^{|\omega|})$. $\qquad\square$

Emptiness checking for AEUPTA can then be decided by first using the simulation theorem to get a non-alternating automaton (with an exponential blow-up). Then:

Corollary 29. *Let* $\mathcal{A} = (Q, q_0, \delta, \omega)$ *be an AEUPTA. Checking emptiness of* $\mathcal{L}(\mathcal{A})$ *can be performed in time* $2^{O(|Q|^3 \cdot |\omega| \cdot (\log |Q| + |\delta|_{\mathsf{Bool}}{}^2 \cdot |\delta|_E \cdot \log |\delta|_E) + |Q| \cdot |\omega| \cdot \log |\Sigma|)}$.

6 Application to QCTL

QCTL extends the temporal logic CTL with quantifications over atomic propositions. In this section, we establish a tight link between QCTL and EU-automata, from which we obtain expressiveness and algorithmic results for QCTL.

6.1 Syntax and (tree) semantics

Definition 30. *The syntax of* QCTL *over a finite set* AP *of atomic propositions is defined by the following grammar:*

$$QCTL \ni \phi, \psi ::= q \mid \neg\phi \mid \phi \vee \psi \mid \mathbf{EX}\phi \mid \mathbf{E}\phi\mathbf{U}\psi \mid \mathbf{A}\phi\mathbf{U}\psi \mid \exists p.\ \phi$$

where q and p range over AP. *The* size *of a formula $\phi \in$ QCTL, denoted $|\phi|$, is the number of steps needed to build ϕ.* CTL *is the restriction of* QCTL *in which the rule $\exists p.\ \phi$ is not allowed.*

QCTL formulas are evaluated over infinite trees (usually computation trees of finite Kripke structures). It is worth noticing that there exist several semantics for QCTL in the literature [Kup95, Fre06, LM14] and this leads to important differences in term of complexity or expressiveness. Here we consider the so-called *tree semantics*: given a QCTL formula ϕ, a 2^{AP}-labelled $\mathcal{D}$-tree $\mathcal{T} = (t, l)$, and a node n, we write $\mathcal{T}, n \models \phi$ to denote that ϕ holds at node n in $\mathcal{T}$, which is defined inductively as follows:

$$
\begin{aligned}
\mathcal{T}, n &\models p & \text{iff} \quad & p \in l(n) \\
\mathcal{T}, n &\models \neg\phi & \text{iff} \quad & \mathcal{T}, n \not\models \phi \\
\mathcal{T}, n &\models \phi \vee \psi & \text{iff} \quad & \mathcal{T}, n \models \phi \text{ or } \mathcal{T}, n \models \psi \\
\mathcal{T}, n &\models \mathbf{EX}\phi & \text{iff} \quad & \exists d \in \mathcal{D} \text{ s.t. } n \cdot d \in t \text{ and } \mathcal{T}, n \cdot d \models \phi \\
\mathcal{T}, n &\models \mathbf{E}\phi\mathbf{U}\psi & \text{iff} \quad & \exists w \in \mathcal{D}^{\omega} \text{ s.t. } n \cdot w \text{ is an infinite branch in } T \text{ and} \\
& & & \exists i \geq 0 \text{ s.t. } \mathcal{T}, n \cdot w_{\langle i \rangle} \models \psi \text{ and } \forall 0 \leq j < i.\ \mathcal{T}, n \cdot w_{\langle j \rangle} \models \phi \\
\mathcal{T}, n &\models \mathbf{A}\phi\mathbf{U}\psi & \text{iff} \quad & \forall w \in \mathcal{D}^{\omega}. \text{ if } n \cdot w \text{ is an infinite branch in } T, \text{ then } \exists i \geq 0. \\
& & & \mathcal{T}, n \cdot w_{\langle i \rangle} \models \psi \text{ and } \forall 0 \leq j < i \text{ we have } \mathcal{T}, n \cdot w_{\langle j \rangle} \models \phi \\
\mathcal{T}, n &\models \exists p.\ \phi & \text{iff} \quad & \exists \mathcal{T}' \equiv_{2^{\mathsf{AP} \setminus \{p\}}} \mathcal{T} \text{ s.t. } \mathcal{T}', n \models \phi,
\end{aligned}
$$

where, following the definition given at the beginning of Sect. 4.2, $\mathcal{T}' \equiv_{2^{\mathsf{AP} \setminus \{p\}}} \mathcal{T}$ means that $\mathcal{T}$ and $\mathcal{T}'$ are identical except for the labelling with atomic proposition p: formula $\exists p.\ \phi$ intuitively means that it is possible to modify the labelling of $\mathcal{T}$ for proposition p in such a way that ϕ holds.

Finally, for a Kripke structure $\mathcal{K}$ and one of its states v, we write $\mathcal{K}, v \models \phi_s$ whenever $\mathcal{T}_{\mathcal{K},v}, \varepsilon_{\mathcal{T}_{\mathcal{K},v}} \models \phi_s$. We say that two formulas ϕ_1 and ϕ_2 are *equivalent* (denoted $\phi_1 \equiv \phi_2$) when their truth value are equal for every 2^{AP}-labelled tree.

In the sequel, we use standard abbreviations such as $\top \stackrel{\text{def}}{=} p \vee \neg p$, $\bot \stackrel{\text{def}}{=} \neg\top$, $\mathbf{EF}\phi \stackrel{\text{def}}{=} \mathbf{E}\top\mathbf{U}\phi$, $\mathbf{AF}\phi \stackrel{\text{def}}{=} \mathbf{A}\top\mathbf{U}\phi$, $\mathbf{AG}\phi \stackrel{\text{def}}{=} \neg\mathbf{EF}\neg\phi$, $\mathbf{EG}\phi \stackrel{\text{def}}{=} \neg\mathbf{AF}\neg\phi$ and $\forall p.\ \phi_s \stackrel{\text{def}}{=} \neg\exists p.\ \neg\phi_s$.

Quantification over atomic propositions increases the expressive power of CTL. For example, it allows us to count the number of successors, as illustrated by formula

$$\mathbf{E}_1\mathbf{X}\phi = \mathbf{EX}\phi \wedge \neg\exists p.\ \big(\mathbf{EX}(p \wedge \phi) \wedge \mathbf{EX}(\neg p \wedge \phi)\big)$$

where we assume that p does not appear in ϕ. This formula states that there is exactly one successor satisfying ϕ: the first part of the formula enforces the presence of at least one successor satisfying ϕ, and if there were two of them, then labelling only one of them with p would falsify the second part of the formula. It is well-known[HM85] that CTL cannot express such properties.

Generalising the idea above, QCTL can also (succinctly) express that a node has at most 2^k successors: this is achieved by requiring the existence of a labelling with k atomic propositions $(p_i)_{1 \leq i \leq k}$ in such a way that no two successors

have the same labelling:

$$\chi_k = \exists (p_i)_{1 \le i \le k}. \neg \Big(\exists q, q'. \ \mathbf{E}_1 \mathbf{X} q \wedge \mathbf{E}_1 \mathbf{X} q' \wedge \neg \mathbf{E} \mathbf{X}(q \wedge q') \wedge$$
$$\bigwedge_{1 \le i \le k} \mathbf{E} \mathbf{X}(q \wedge p_i) \Leftrightarrow \mathbf{E} \mathbf{X}(q' \wedge p_i) \Big).$$

The negation of this formula expresses the existence of *at least* $2^k + 1$ successors. Using an extra atomic proposition for isolating a single node, we can get a formula of size linear in k expressing the existence of exactly 2^k successors. Similar ideas can be used to succinctly express the existence of 2^k successors satisfying some formula ϕ.

In order to give more precise results about QCTL, we introduce restricted fragments, depending on the nesting of quantifiers. Given two QCTL formulas ϕ and $(\psi_i)_i$, and atomic propositions $(p_i)_i$ that appear free in ϕ (*i.e.*, not as quantified propositions), we write $\phi[(p_i \to \psi_i)_i]$ (or $\phi[(\psi_i)_i]$ when $(p_i)_i$ are understood from the context) for the formula obtained from ϕ by replacing each occurrence of p_i with ψ_i. Given two sublogics L_1 and L_2 of QCTL, we write $L_1[L_2] = \{\phi[(\psi_i)_i] \mid \phi \in L_1, \ (\psi_i)_i \in L_2\}$. We then inductively define the following fragments:

- EQ^{0}CTL and AQ^{0}CTL correspond to CTL, and for $k > 0$, EQ^{k}CTL is the set of formulas of the form $\exists p_1. \exists p_2 \ldots \exists p_n. \phi$ for $\phi \in$ AQ^{k-1}CTL, and AQ^{k}CTL is the set of formulas of the form $\forall p_1. \forall p_2 \ldots \forall p_n. \phi$ for $\phi \in$ EQ^{k-1}CTL,

- Q^{0}CTL is CTL, Q^{1}CTL = CTL[EQ^{1}CTL], and for $k > 1$, Q^{k}CTL is the logic Q^{1}CTL[Q^{k-1}CTL].

Hence formulas in EQ^{k}CTL and AQ^{k}CTL are in prenex form, and involve $k - 1$ quantifier alternations (respectively starting with existential and universal quantifiers); on the other hand, Q^{k}CTL counts the maximal number of nested blocks of quantifiers, allowing boolean and CTL operators between blocks. An easy induction shows that EQ^{k}CTL and AQ^{k}CTL are syntactic fragments of Q^{k}CTL. As examples, it can be seen that formula $\mathbf{E}_1 \mathbf{X} \phi$ (for $\phi \in$ CTL) is in AQ^{1}CTL, and that formula χ_k is in Q^{3}CTL, but can easily be rewritten as a formula in EQ^{3}CTL.

6.2 From QCTL to AEUTA

6.2.1 From CTL to tree automata

Any CTL formula ϕ can be turned into an AEUPTA $\mathcal{A}_\phi$ accepting exactly the trees where ϕ holds. One of the first such constructions is given in [BVW94, KVW00]; it is based on fixed-arity tree automata, but the construction has then been extended to arbitrary-arity trees using $\{\Box, \Diamond\}$-automata [Wil99] (see Sect. 2.6).

Here we adapt the construction of [Wil99] to our EU-constraints in the transitions of the automaton. We assume w.l.o.g. that negations in ϕ may only appear at the level of atomic propositions; transforming a formula in such a negation-normal form may at most double the size of the formula. The automaton $\mathcal{A}_\phi = (Q_\phi, q_{init}, \delta_\phi, \omega_\phi)$ can then be defined as follows:

- Q_ϕ is the set of state subformulas (including $\top$) of ϕ. In order to avoid confusion, for each subformula ψ, we write $[\psi]$ for the associated state in Q_ϕ.

- the initial state q_{init} is $[\phi]$,

- given $[\psi] \in Q_\phi$ and $\sigma \in 2^{\mathsf{AP}}$, we define $\delta_\phi([\psi], \sigma)$ inductively as follows:

$$\delta_\phi([\top], \sigma) = \top \qquad\qquad \delta_\phi([\bot], \sigma) = \bot$$

$$\delta_\phi([P], \sigma) = \begin{cases} \top & \text{if } P \in \sigma \\ \bot & \text{otherwise} \end{cases} \qquad \delta_\phi([\neg P], \sigma) = \begin{cases} \bot & \text{if } P \notin \sigma \\ \top & \text{otherwise} \end{cases}$$

$$\delta_\phi([\psi_1 \wedge \psi_2], \sigma) = \delta_\phi([\psi_1], \sigma) \wedge \delta_\phi([\psi_2], \sigma)$$

$$\delta_\phi([\psi_1 \vee \psi_2], \sigma) = \delta_\phi([\psi_1], \sigma) \vee \delta_\phi([\psi_2], \sigma)$$

$$\delta_\phi([\mathbf{EX}\psi], \sigma) = \langle [\psi] \mapsto 1; \{[\top]\} \rangle$$

$$\delta_\phi([\mathbf{AX}\psi], \sigma) = \langle \varnothing; \{[\psi]\} \rangle$$

$$\delta_\phi([\mathbf{E}\psi_1 \mathbf{U} \psi_2], \sigma) = \delta_\phi([\psi_2], \sigma) \vee \left(\delta_\phi([\psi_1], \sigma) \wedge \langle [\mathbf{E}\psi_1 \mathbf{U}\psi_2] \mapsto 1; \{[\top]\} \rangle \right)$$

$$\delta_\phi([\mathbf{E}\psi_1 \mathbf{W} \psi_2], \sigma) = \delta_\phi([\psi_2], \sigma) \vee \left(\delta_\phi([\psi_1], \sigma) \wedge \langle [\mathbf{E}\psi_1 \mathbf{W}\psi_2] \mapsto 1; \{[\top]\} \rangle \right)$$

$$\delta_\phi([\mathbf{A}\psi_1 \mathbf{U} \psi_2], \sigma) = \delta_\phi([\psi_2], \sigma) \vee \left(\delta_\phi([\psi_1], \sigma) \wedge \langle \varnothing; \{[\mathbf{A}\psi_1 \mathbf{U}\psi_2]\} \rangle \right)$$

$$\delta_\phi([\mathbf{A}\psi_1 \mathbf{W} \psi_2], \sigma) = \delta_\phi([\psi_2], \sigma) \vee \left(\delta_\phi([\psi_1], \sigma) \wedge \langle \varnothing; \{[\mathbf{A}\psi_1 \mathbf{W}\psi_2]\} \rangle \right)$$

- the acceptance condition is a parity condition defined through the following priority function:

$$\omega_\phi([\mathbf{E}\psi_1 \mathbf{U}\psi_2]) = \omega_\phi([\mathbf{A}\psi_1 \mathbf{U}\psi_2]) = 1$$
$$\omega_\phi([\mathbf{E}\psi_1 \mathbf{W}\psi_2]) = \omega_\phi([\mathbf{A}\psi_1 \mathbf{W}\psi_2]) = 0$$

The priority for all other states $[\psi]$ is irrelevant as they can only appear finitely many times along a branch.

Then we have the following theorem, proved by Wilke for his construction with $\{\Box, \Diamond\}$-automata:

Theorem 31 ([Wil99]). *For any CTL formula ϕ, there exists an AEUTA $\mathcal{A}_\phi$ accepting exactly the trees in which ϕ holds. This automaton has size $(O(|\phi|), O(|\phi|), 1, 1, 2)$.*

Remark 6. *Consider a CTL formula ϕ and a Kripke structure $\mathcal{K}$. Using automaton $\mathcal{A}_\phi$ (and the fact that it contains only EU-constraints with $\max(|\delta|_E, |\delta|_U) \leq 1$), we can build a parity game $\mathcal{G}_{\mathcal{A}, \mathcal{K}}$ whose size is in $O(|\mathcal{K}| \cdot |\phi|)$ and that can be solved (using Prop. 26) in time $O\big((|K| \cdot |\phi|)^2\big)$. This is not the optimal complexity for CTL model-checking; however, this can be improved by using the fact that the automaton $\mathcal{A}_\varphi$ is weak [MSS86, BVW94, VW08]. This provides*

us with a weak *game* $\mathcal{G}_{\mathcal{A},\mathcal{K}}$, *which allows us to get an algorithm running in* $O(|\mathcal{K}| \cdot |\phi|)$, *thereby recovering the classical complexity for* CTL *model checking.*
◀

Remark 7. *In Theorem 31, we can observe that the boolean size of the transition function of* $\mathcal{A}_\phi$ *is linear in* $|\phi|$. *This is a direct consequence of its definition. While it is constant for the transition formulas from states of the form* $[\mathbf{E}\psi_1\mathbf{U}\psi_2]$, *it might be linear e.g. for subformulas of the form* $\psi = \psi_1 \wedge \ldots \psi_p$: *the transition function* $\delta_\phi([\psi],\sigma)$ *is then defined as the conjunction of every* $\delta_\phi([\psi_i],\sigma)$, *hence it has linear size.*
◀

6.2.2 A tree-automata construction for QCTL formulas

Combining the construction for CTL and the operations over AEUPTA allows us to extend the automata construction to QCTL formulas. The crucial point is the handling of quantifications.

Consider a QCTL formula Φ where the negations can only be followed by atomic propositions or $\exists p.\psi$ subformulas, and a subformula ϕ of Φ of the form $\exists p.\psi$, assuming that we have built an AEUPTA $\mathcal{A}_\psi$ for ψ. If $\mathcal{A}_\psi$ is non-alternating, we can use the projection operation on $\mathcal{A}_\psi$ (see Sect. 4.2) and get an AEUPTA for ϕ. Otherwise, we have to first turn $\mathcal{A}_\psi$ into a non-alternating automaton (with an exponential blow-up in the size of the automaton) before using projection. Note that we can handle in one step a block of existential quantifiers of the form $\exists p_1 \cdots \exists p_n.\psi$ (hence with a single exponential blow-up). On top of this, there may be negations in front of existential quantifiers, which may require complementing the automaton.

Therefore the size of the resulting automaton will drastically depend on the number of such nested blocks of existential quantifiers in the QCTL formula. Thus complexity results are stated for formulas in $\mathsf{Q}^k\mathsf{CTL}$, $\mathsf{EQ}^k\mathsf{CTL}$ and $\mathsf{AQ}^k\mathsf{CTL}$.

In the following, we write $k\text{-}\mathsf{exp}(n)$ to denote the family of sets of functions of one variable n defined inductively as follows: $0\text{-}\mathsf{exp}(n)$ is the set of functions bounded by a polynomial in n, and $(k+1)\text{-}\mathsf{exp}(n)$ contains all functions f such that $f \in O(2^g)$ with $g \in k\text{-}\mathsf{exp}(n)$.

We can now formally state the result as follows:

Theorem 32. *Given a* $\mathsf{Q}^k\mathsf{CTL}$ *formula* ϕ *over* AP *with* $k > 0$, *we can construct a* AEUPTA $\mathcal{A}_\phi$ *over* 2^{AP} *accepting exactly the trees satisfying* ϕ. *The automaton* $\mathcal{A}_\phi$ *has size* $(k\text{-}\mathsf{exp}(|\phi|), k\text{-}\mathsf{exp}(|\phi|), (k-1)\text{-}\mathsf{exp}(|\phi|), 1, (k-1)\text{-}\mathsf{exp}(|\phi|))$. *If additionally* ϕ *is in* $\mathsf{EQ}^k\mathsf{CTL}$, *then* $\mathcal{A}_\phi$ *is non-alternating.*

Proof. We proceed by induction over k. We prove the result for formulas in $\mathsf{Q}^k\mathsf{CTL}$, showing along the way the property for formulas in $\mathsf{EQ}^k\mathsf{CTL}$.

- if $\phi \in \mathsf{Q}^1\mathsf{CTL}$, then ϕ is of the form $\Phi[(\psi_i)_{1 \le i \le m}]$ where Φ is a CTL formula and $(\psi_i)_{1 \le i \le m}$ are $\mathsf{EQ}^1\mathsf{CTL}$ formulas. We handle each ψ_i separately. Assume that $\psi_i = \exists p_1^i \ldots \exists p_{l_i}^i. \psi_i'$ with $\psi_i' \in \mathsf{CTL}$. From Theorem 31, one can build an AEUPTA $\mathcal{A}_{\psi_i'}$ recognising the trees satisfying ψ_i'; moreover, $|\mathcal{A}_{\psi_i'}|$ is bounded by $(O(|\psi_i'|), O(|\psi_i'|), 1, 1, 2)$. We then apply the constructions of Section 4, and the results summarised in Table 1:

  - we can remove alternation and get an equivalent EUPTA $\mathcal{N}_{\psi_i'}$ whose size is bounded by $(2^{O(|\psi_i'|^2 \cdot \log(|\psi_i'|))}, 2^{O(|\psi_i'|^4)}, O(|\psi_i'|^2), 1, O(|\psi_i'|))$;

45

- applying projection (to $\mathcal{N}_{\psi_i'}$ and for atomic propositions p_1^i to $p_{l_i}^i$), we get an EUPTA $\mathcal{B}_{\psi_i}$ recognising the models of ψ_i, whose size is bounded by $(2^{O(|\psi_i'|^2 \cdot \log(|\psi_i'|))}, 2^{l_i} \cdot 2^{O(|\psi_i'|^4)}, O(|\psi_i'|^2), 1, O(|\psi_i'|))$, hence also by $(2^{O(|\psi_i|^2 \cdot \log(|\psi_i|))}, 2^{O(|\psi_i|^4)}, O(|\psi_i|^2), 1, O(|\psi_i|))$; This shows that for a formula ψ_i in EQ^{1}CTL, we can build a non-alternating EUPTA of size $(1\text{-exp}(|\phi|), 1\text{-exp}(|\phi|), 0\text{-exp}(|\phi|), 1, 0\text{-exp}(|\phi|))$

- in case formula ψ_i is in the (direct) scope of a negation operator, we complement $\mathcal{B}_{\psi_i}$ into an AEUPTA $\mathcal{C}_{\overline{\psi_i}}$ whose size is bounded by $(|\Sigma| \cdot 2^{O(|\psi_i|^4)}, 2^{O(|\psi_i|^4)}, O(|\psi_i|^2), 1, O(|\psi_i|))$ where $|\Sigma|$ can be bounded by $2^{|\phi|}$;

- we build the final AEUPTA $\mathcal{A}_\phi = (Q, q_{init}, \delta, \omega)$ using the construction of Theorem 31, in which we add the following rule to deal with subformulas ψ_i in EQ^{1}CTL (or their negations):

$$\delta(\psi_i, \sigma) = \delta_{\psi_i}(q_{init}^{\psi_i}, \sigma) \qquad \delta(\neg\psi_i, \sigma) = \delta_{\overline{\psi}_i}(q_{init}^{\overline{\psi}_i}, \sigma),$$

where $q_{init}^{\psi_i}$ and $q_{init}^{\overline{\psi}_i}$ are the initial states of $\mathcal{B}_{\psi_i}$ and $\mathcal{C}_{\psi_i}$, respectively. Our automaton has size at most $(2^{O(|\phi|^4)}, 2^{O(|\phi|^4)}, O(|\phi|^2), 1, O(|\phi|))$ This proves the base case of our result.

- if $\phi \in$ Q^{k}CTL with $k > 1$, we show that for every ϕ-subformula $\psi \in$ Q$^{k'}$CTL with $1 \le k' \le k$, the AEUPTA $\mathcal{A}_\psi$ has size $(k'\text{-exp}(|\phi|), k'\text{-exp}(|\phi|), (k'-1)\text{-exp}(|\phi|), 1, (k'-1)\text{-exp}(|\phi|))$. Note that the size of the automaton depends on $|\phi|$ (and not on $|\psi|$) because the complement operation provides an automaton whose size depends on $|\Sigma|$, which can only be bounded by $2^{|\phi|}$. We prove this result by induction over k'. The base case is similar to the previous case with $k = 1$.

Now consider $k' > 1$. Assume $\psi \in$ Q$^{k'}$CTL is of the form $\Psi[(\psi_i)_{1 \le i \le m}]$, where Ψ is a CTL formula and each ψ_i is of the form $\exists p_1^i \ldots \exists p_{l_i}^i. \psi_i'$ with $\psi_i' \in$ Q$^{k'-1}$CTL.

From the induction hypothesis, we can build AEUPTAs $\mathcal{A}_{\psi_i'} = (Q_i, q_{init_i}, \delta_i, \omega_i)$ recognising the trees satisfying ψ_i' for all i, and whose size $(s_i, b_i, e_i, u_i, p_i)$ is bounded by $((k'-1)\text{-exp}(|\phi|), (k'-1)\text{-exp}(|\phi|), (k'-2)\text{-exp}(|\phi|), 1, (k'-2)\text{-exp}(|\phi|))$.

Applying the simulation theorem provides us with EUPTAs $\mathcal{N}_{\psi_i'}$, each accepting the same language as $\mathcal{A}_{\psi_i'}$, whose sizes are at most $(2^{O(s_i^2 \cdot \log(s_i))}, (2 \cdot e_i)^{O(s_i^2 \cdot b_i^2 \cdot e_i)}, s_i \cdot b_i \cdot e_i, 1, 2 \cdot (s_i \cdot p_i + 1))$. After projection, we obtain EUPTAs $\mathcal{B}_{\psi_i}$ for formulas ψ_i whose sizes are bounded by $(2^{O(s_i^2 \cdot \log(s_i))}, 2^{l_i} \cdot (2 \cdot e_i)^{O(s_i^2 \cdot b_i^2 \cdot e_i)}, O(s_i \cdot b_i \cdot e_i), 1, 2 \cdot (s_i \cdot p_i + 1))$.

In case ψ_i is used negatively in ϕ, we compute the complement $\mathcal{C}_{\overline{\psi_i}}$ of $\mathcal{B}_{\psi_i}$, which is an AEUPTA of size at most $(O(|\Sigma| \cdot 2^{O(s_i^2 \cdot \log(s_i))} \cdot 2^{s_i^2 \cdot b_i^2 \cdot e_i \cdot \log(e_i)}) \cdot 2^{l_i}, 2^{O(s_i^2 \cdot \log(s_i))} \cdot 2^{s_i^2 \cdot b_i^2 \cdot e_i \cdot \log(e_i)} \cdot 2^{l_i}, O(s_i \cdot b_i \cdot e_i), 1, 2 \cdot s_i \cdot p_i + 3)$. Again $|\Sigma|$ can be bounded by $2^{|\phi|}$.

Finally, as for the base case, we apply the construction of Theorem 31 to end up with an AEUPTA for ψ, whose size is bounded by $(k'\text{-exp}(|\phi|), k'\text{-exp}(|\phi|), (k'-1)\text{-exp}(|\phi|), 1, (k'-1)\text{-exp}(|\phi|))$.

This result applies to ϕ itself: we get that the size of $\mathcal{A}_\phi$ is bounded by $(k\text{-}\exp(|\phi|), k\text{-}\exp(|\phi|), (k-1)\text{-}\exp(|\phi|), 1, (k-1)\text{-}\exp(|\phi|))$, and if ϕ belongs to $\mathsf{EQ}^k\mathsf{CTL}$, we get a non-alternating automaton as the last step is the projection operation (on a non-alternating automaton). This concludes our proof. □

Model-checking and satisfiability for QCTL can be solved using EU-automata. Using Corollary 29, Theorem 27 and the hardness results of [LM14], we get:

Theorem 33. *The satisfiability problem for $Q^k CTL$, $AQ^k CTL$ and $EQ^{k+1} CTL$ is $(k+1)$-EXPTIME-complete. The model-checking problem for $Q^k CTL$, $AQ^k CTL$ and $EQ^k CTL$ is k-EXPTIME-complete.*

6.2.3 Extension to QCTL*

The automata construction for QCTL can be extended to $QCTL^*$, the extension of CTL* with quantifications over atomic propositions. In CTL^*, we distinguish between state formulas ϕ_s, interpreted over states, and path formulas ϕ_p, interpreted over infinite paths. Informally, a state formula corresponds to some Boolean combination of atomic propositions and formulas of the form $\mathbf{E}\phi_p$ and $\mathbf{A}\phi_p$ (*i.e.* path formulas prefixed by some path quantifier), and path formulas are defined as LTL formulas with state formulas appearing in place of atomic propositions. The logic QCTL* extends CTL* by allowing formulas of the form $\exists p.\ \phi_s$ as state formulas[5].

As for QCTL, we can define several fragments of QCTL*: $Q^k CTL^*$ contains formulas in which the maximum number of nested blocks of quantifiers is at most k. The construction of $\mathcal{A}_\phi$ for $\phi \in$ QCTL follows the same steps as for QCTL; the main difference is that we have to consider formulas of the form $\mathbf{E}\phi_p$ where ϕ_p is an LTL formula: in that case, we have to first build a word automaton to capture ϕ_p, and then use Proposition 4 to derive a tree automaton for $\mathbf{E}\phi_p$. The complexity is then higher. We start with this following result for CTL*:

Proposition 34. *Given a CTL^* formula ϕ over AP, we can construct an AE-UPTA $\mathcal{A}_\phi$ over 2^{AP} accepting exactly the trees satisfying ϕ. The automaton $\mathcal{A}_\phi$ has size $(1\text{-}exp(|\phi|), 1\text{-}exp(|\phi|), 2, 1, 3)$.*

Proof. The key point is the treatment of formulas of the form $\mathbf{E}\phi_p$ with $\phi_p \in$ LTL. In that case, we build a PWA $\mathcal{A}_{\phi_p}$ in a standard way, whose size is exponential in $|\phi|_p$. Then by applying Proposition 4, we get an EUPTA corresponding to the formula $\mathbf{E}\phi_p$ whose size is $(1\text{-}\exp(|\phi|), 1\text{-}\exp(|\phi|), 1, 1, 2)$. In case of $\mathbf{A}\phi_p$ formula, we just have to add a complementation step (Theorem 17) and we get an AEUPTA whose size is $(1\text{-}\exp(|\phi|), 1\text{-}\exp(|\phi|), 2, 1, 3)$.

Now consider a CTL* formula: we apply the previous construction to every subformula $\mathbf{E}\phi_p$ starting with the innermost subformulas. Finally we get an AEUPTA that *combines* the different automata and it provides an AEUPTA whose size is $(1\text{-}\exp(|\phi|), 1\text{-}\exp(|\phi|), 2, 1, 3)$. □

We can now state the construction for QCTL*:

---

[5]This precision is important: allowing such quantifications inside a path formula changes the expressiveness of the logics [LM14].

Theorem 35. *Given a $Q^k CTL^*$ formula ϕ over AP with $k > 0$, we can construct an AEUPTA $\mathcal{A}_\phi$ over 2^{AP} accepting exactly the trees satisfying ϕ. The automaton $\mathcal{A}_\phi$ has size $((k+1)\text{-}exp(|\phi|), (k+1)\text{-}exp(|\phi|), k\text{-}exp(|\phi|), 1, k\text{-}exp(|\phi|))$.*

Proof. • Consider $\phi \in Q^1 CTL^*$. Assume that negations occur only before existential quantifications over propositions. Thus ϕ can be seen as a formula $\Phi[(\psi_i)_{1 \leq i \leq m}]$ where Φ is a CTL^* formula and $(\psi_i)_{1 \leq i \leq m}$ are $EQ^1 CTL^*$ formulas, with $\psi_i = \exists p_1^i \ldots \exists p_{l_i}^i.\ \psi_i'$ with $\psi_i' \in CTL^*$.

We first build an AEUPTA $\mathcal{A}_{\psi_i'}$ as explained in Prop. 34. Then we can transform each of these automata into a non-alternating automaton $\mathcal{N}_{\psi_i'}$ whose size is in $(2\text{-}exp(|\psi|_i), 2\text{-}exp(|\psi|_i), 1\text{-}exp(|\psi|_i), 1, 1\text{-}exp(|\psi|_i))$ (see Theorem 25). We can then apply the projection operation over these automata to deal with the quantification $\exists p_1^i \ldots \exists p_{l_i}^i$. Then we get the (non-alternating) automata $\mathcal{B}_{\psi_i}$.

A complementation procedure is possibly applied (when a negation precedes the corresponding existential quantification in Φ). In that case, we obtain an alternating automaton $\mathcal{C}_{\overline{\psi}_i}$ whose size admits the same bounds as above. Finally it remains to consider the CTL^* context Φ, which corresponds to an AEUPTA of size at most $(1\text{-}exp(|\Phi|), 1\text{-}exp(|\Phi|), 2, 1, 3)$; combined with automata $\mathcal{C}_{\overline{\psi}_i}$s and $\mathcal{B}_{\psi_i}$s, we get an AEUPTA whose size is in $(2\text{-}exp(|\phi|), 2\text{-}exp(|\phi|), 1\text{-}exp(|\phi|), 1, 1\text{-}exp(|\phi|))$.

• Consider $\phi \in Q^{k+1} CTL^*$. Here again we assume that negations occur only before existential quantifications over propositions. As in the construction for QCTL formula, we show that for every ϕ-subformula $\psi \in Q^{k'} CTL^*$ with $1 \leq k' \leq k$, the automaton AEUPTA $\mathcal{A}_\psi$ has size $((k'+1)\text{-}exp(|\phi|), (k'+1)\text{-}exp(|\phi|), k'\text{-}exp(|\phi|), 1, k'\text{-}exp(|\phi|))$. We prove it by induction over k'. The result holds for $k' = 1$ (similar to the previous case). Now assume $1 < k' \leq k$. Consider a ϕ-subformula $\psi \in Q^{k'} CTL^*$. Then ψ is of the form $\Psi[(\psi_i)_{1 \leq i \leq m}]$ where Ψ is a CTL^* formula and every (ψ_i) is of the form $\exists p_1^i \ldots \exists p_{l_i}^i.\ \psi_i'$ with $\psi_i' \in Q^{k'-1} CTL^*$.

By induction hypothesis, we can build an AEUPTA $\mathcal{A}_{\psi_i'}$ for each formula ψ_i, whose size is bounded by $(k'\text{-}exp(|\phi|), k'\text{-}exp(|\phi|), (k'-1)\text{-}exp(|\phi|), 1, (k'-1)\text{-}exp(|\phi|))$. Applying the simulation theorem (Theorem 25), we get an EUPTA whose size is bounded by $((k'+1)\text{-}exp(|\phi|), (k'+1)\text{-}exp(|\phi|), k'\text{-}exp(|\phi|), 1, k'\text{-}exp(|\phi|))$ We can then apply the projection operation, possibly followed by a complementation operation which provides an automaton whose size is still bounded by $((k'+1)\text{-}exp(|\phi|), (k'+1)\text{-}exp(|\phi|), k'\text{-}exp(|\phi|), 1, k'\text{-}exp(|\phi|))$. Finally, it remains to incorporate the CTL^* context Ψ, which we perform as in the base case; we finally get an AEUPTA whose size is in $((k'+1)\text{-}exp(|\phi|), (k'+1)\text{-}exp(|\phi|), k'\text{-}exp(|\phi|), 1, k'\text{-}exp(|\phi|))$. This concludes the proof of the inductive step of the intermediary result. And we can deduce that the automaton for ϕ is therefore in $((k+1)\text{-}exp(|\phi|), (k+1)\text{-}exp(|\phi|), k\text{-}exp(|\phi|), 1, k\text{-}exp(|\phi|))$ which concludes the proof. $\square$

Note also that this construction provides a non-alternating automaton if ϕ belongs to $EQ^k CTL^*$.

As a direct consequence, we get decision procedures for Model-checking and

satisfiability for QCTL* based on EU-automata; again, lower complexity bounds are obtained from [LM14]:

Theorem 36. *The satisfiability problem for $Q^k CTL^*$, $AQ^k CTL^*$, and $EQ^{k+1} CTL^*$ is $(k+2)$-EXPTIME-complete. The model-checking problem for $Q^k CTL^*$, $AQ^k CTL^*$, and $EQ^k CTL^*$ is $(k+1)$-EXPTIME-complete.*

6.3 From AEUPTA to QCTL

In this section, we use EU tree automata to derive expressiveness results: we turn an EUPTA $\mathcal{A} = (Q, q_0, \delta, \omega)$ over 2^{AP} into an equivalent (over all 2^{AP}-labelled trees) QCTL formula $\Phi_{\mathcal{A}}$. Remember that for EUPTA, the transition function $\delta(q, \sigma)$ is a disjunction of EU-pairs.

Theorem 37. *For any $\mathcal{A}$ be an EUPTA over 2^{AP}, we can build an $EQ^2 CTL$ formula $\Phi_{\mathcal{A}}$ such that, for any 2^{AP}-labelled tree $\mathcal{T}$, it holds $\mathcal{T} \in \mathcal{L}(\mathcal{A})$ if, and only if, $\mathcal{T}, \epsilon \models \Phi_{\mathcal{A}}$. The size of the formula $\Phi_{\mathcal{A}}$ is in $O(|Q| \cdot |\omega| + |Q| \cdot 2^{|\mathsf{AP}|} \cdot |\mathsf{AP}| \cdot |\delta|_{Bool} \cdot |\delta|_E \cdot (|\delta|_E + |\delta|_U))$*

Proof. Let $Q = \{q_i \mid 0 \leq i \leq n = |Q| - 1\}$. In $\Phi_{\mathcal{A}}$, we use the set of fresh quantified atomic propositions $\{b, q_0, \ldots, q_n, p_1, \ldots, p_{|\delta|_E}, p\}$ in order to express the existence of an accepting execution tree of the automaton: propositions in $\{q_0, ..., q_n\}$ will be used to label each node of the input tree with the name of the state visiting that node in the execution tree, while propositions in $\{p_1, ..., p_{|\delta|_E}\}$ are used to distinguish between the successors involved in the verification of the E-part of EU-constraints; proposition b is used for expressing the acceptance condition and proposition p is used to ensure that no node has more than one successor labelled with the same p_i[6]. Our formula $\Phi_{\mathcal{A}}$ reads as follows:

$$\exists q_0 \ldots \exists q_n. \ \exists p_1 \ldots \exists p_{|\delta|_E}. \ \forall b. \ \forall p. \ (\Phi_p \wedge \widetilde{\Phi}_{\mathcal{A}}).$$

In this formula, Φ_p will be used to state that no nodes are labelled with several p_is and no nodes have more than one successor labelled with p_i (for any i), while $\widetilde{\Phi}_{\mathcal{A}}$ will enforce that the labelled tree describes an accepting execution tree of $\mathcal{A}$ on $\mathcal{T}$. Formally, Φ_p is defined as the following formula (remember that p is quantified universally):

$$\Phi_p = \mathbf{AG} \bigwedge_{1 \leq i \leq |\delta|_E} \left[(p_i \Rightarrow \bigwedge_{j \neq i} \neg p_j) \wedge \left(\mathbf{AX}(p_i \Rightarrow p) \vee \mathbf{AX}(p_i \Rightarrow \neg p) \right) \right]$$

Formula $\widetilde{\Phi}_{\mathcal{A}}$ is defined as follows:

$$\widetilde{\Phi}_{\mathcal{A}} = q_0 \wedge \bigwedge_{i=0}^{n} \mathbf{AG} \left[q_i \Rightarrow \left(\neg \lambda_{Q \setminus \{q_i\}} \wedge \bigvee_{P \subseteq \mathsf{AP}} \left(\Gamma_P \wedge \Psi_{\delta(q,P)} \right) \right) \right] \wedge \Psi_{\omega}.$$

In this formula, for any set S, formula λ_S is the propositional formula $\bigvee_{q \in S} q$, and for any $P \subseteq \mathsf{AP}$, formula Γ_P is the propositional formula $\bigwedge_{p \in P} p \wedge \bigwedge_{p' \in \mathsf{AP} \setminus P} \neg p'$ (note that the size of Γ_P is in $O(|\mathsf{AP}|)$). We now define formula $\Psi_{\delta(q,P)}$, which

---

[6]For the sake of clarity, we use two distinct propositions b and p, but we could have used the same one.

encodes the satisfaction of the transition function, and formula Ψ_ω, which states that any infinite branch (of the tree) labelled with b satisfies the parity condition.

For the former, we write (remember that $\mathcal{A}$ is non-alternating):

$$\Psi_{\delta(q,P)} = \begin{cases} \top & \text{if } \delta(q,P) = \top \\ \displaystyle\bigvee_{\langle E;U\rangle \in \delta(q,P)} \Psi_{\langle E;U\rangle} & \text{otherwise,} \end{cases}$$

where $\Psi_{\langle E;U\rangle}$ encodes the constraint $\langle E;U\rangle$ of $\mathcal{A}$: writing E as the multiset $\{\!\{E_1, ..., E_{|E|}\}\!\}$, where each E_j belongs to Q, we let:

$$\Psi_{\langle E;U\rangle} = \bigwedge_{j=1}^{|E|} \Big[\mathbf{EX}(p_j \wedge E_j) \wedge \mathbf{AX}\Big((\bigwedge_{j=1}^{|E|} \neg p_j) \Rightarrow \bigvee_{q \in U} q \Big) \Big]$$

Remember that thanks to formula Φ_p, we have ensured that no nodes can be labelled with several p_is and no nodes can have several successors labelled with the same p_j; thus formula $\Psi_{\langle E;U\rangle}$ ensures that all states in E label distinct successors, and nodes with no p_j are labelled with some proposition q corresponding to a state in U.

Formula Ψ_ω expresses the fact that in any infinite subtree labelled with b, there exists no infinite branches where the smallest priority appearing infinitely many times is odd. This can be characterised as follows:

$$\Psi_\omega = \Big(b \wedge \mathbf{AG}(b \Rightarrow \mathbf{EX}b) \ \wedge \ \mathbf{AG}((\neg b) \Rightarrow \mathbf{AX}(\neg b)) \Big) \ \Rightarrow$$
$$\neg \bigvee_{\substack{0 \le d \le |\omega| \\ \text{s.t. } d \text{ odd}}} \Big[\mathbf{AG}\,\mathbf{AF}\Big(b \Rightarrow (\alpha_{=d}) \Big) \wedge \mathbf{EF}(\mathbf{EG}(b \wedge \alpha_{\ge d})) \Big] \Big),$$

where $\alpha_{=d}$ is the formula $\bigvee_{\omega(q_i)=d} q_i$ characterising (atomic propositions corresponding to) states having priority d, and $\alpha_{\ge d}$ is the formula $\bigvee_{\omega(q_i)\ge d} q_i$, identifying states with priorities greater than (or equal to) d. Note that the formula to the left of the implication holds true if, and only if, proposition b labels exactly an infinite subtree from the current node (subformula $(b \Rightarrow \mathbf{EX}b)$ ensures infiniteness, and subformula $(\neg b \Rightarrow \mathbf{AX}\neg b)$ ensures that every b-node is reachable from the current node via a b-path). We will show below why Ψ_ω ensures the satisfaction of the parity condition.

The size of $\Psi_{\delta(q,P)}$ is in $O(|\delta|_{\mathsf{Bool}} \cdot |\delta|_{\mathsf{E}} \cdot (|\delta|_{\mathsf{E}} + |\delta|_{\mathsf{U}})$. The size of Ψ_ω is in $O(|Q| \cdot |\omega|)$. The size of Φ_p is in $O(|\delta|_{\mathsf{E}}^2)$. In the end, we get that the size of $\Phi_{\mathcal{A}}$ is in $O(|\omega| \cdot |Q| + |Q| \cdot 2^{|\mathsf{AP}|} \cdot |\mathsf{AP}| \cdot |\delta|_{\mathsf{Bool}} \cdot |\delta|_{\mathsf{E}} \cdot (|\delta|_{\mathsf{E}} + |\delta|_{\mathsf{U}}))$. Finally, it is easily seen that $\Phi_{\mathcal{A}}$ belongs to $\mathsf{EQ^2CTL}$.

We now prove:

Lemma 38. *Let $\mathcal{A} = (Q, q_0, \delta, \omega)$ be an EUPTA and $\Phi_{\mathcal{A}}$ be the EQ^{2}CTL formula defined above. For any 2^{AP}-labelled tree $\mathcal{T}$, it holds: $\mathcal{T} \in \mathcal{L}(\mathcal{A})$ if, and only if, $\mathcal{T}, \varepsilon \models \Phi_{\mathcal{A}}$.*

Proof. Consider a 2^{AP}-labelled tree $\mathcal{T} = (t, l)$ and assume $\mathcal{T} \in \mathcal{L}(\mathcal{A})$. As $\mathcal{A}$ is non-alternating, there exists an accepting execution tree $\mathcal{U} = (t, \ell)$ of $\mathcal{A}$ with

the same structure t as $\mathcal{T}$. Then any node n of $\mathcal{U}$ is such that $\ell(n) = (n, q)$ for some $q \in Q$.

We aim at showing that $\mathcal{T}, \varepsilon \models \Phi_{\mathcal{A}}$. First, we can label $\mathcal{T}$ with q_is exactly as it is done in $\mathcal{U}$ with ℓ. For each proposition p_j, we proceed as follows: consider a node n labelled with (n, q); the transition function $\delta(q, l(n))$ applies successfully over the subtree rooted at n (as $\mathcal{T} \in \mathcal{L}(\mathcal{A})$). If $\delta(q, l(n)) = \top$, then $\Psi_{\delta(q, l(n))}$ is trivially satisfied. Otherwise there is some EU-pair $\langle E; U \rangle \in \delta(q, l(n))$ that is satisfied from n (since the automaton is non-alternating) and there exist $|E|$ successors of n that satisfy E: writing E as the multiset $\{\!\{E_1, \ldots, E_{|E|}\}\!\}$, we can associate a fixed successor with every E_j. This provides the labelling for proposition p_j at this level (we know that $|E| \leq |\delta|_\mathsf{E}$). All the successors that are not labelled with some p_j with $1 \leq j \leq |E|$, have to be accepted by a state in U. Note also that Φ_p is satisfied by $\mathcal{T}, \varepsilon$.

It remains to verify that Ψ_ω is satisfied. As $\mathcal{U}$ is an accepting execution tree, every infinite branch satisfies the parity condition. Consider an infinite subtree labelled with b, and assume that there exists some odd priority d that appears infinitely many times along every branch of the b-subtree, and such that along one of these branches, eventually all priorities are greater than or equal to d (that is, d is the least priority along that branch); this clearly implies that this branch violates the parity condition, which contradicts our initial assumption.

In conclusion, the chosen labelling makes the formula $\widetilde{\Phi_{\mathcal{A}}}$ hold true.

Conversely, assume $\mathcal{T}, \varepsilon \models \Phi_{\mathcal{A}}$. Consider a labelling for propositions q_is and p_js such that $\widetilde{\Phi_{\mathcal{A}}} \wedge \Phi_p$ holds true at the root for any valuation of p and b. This labelling associates exactly one state of the automaton with every node (thanks to subformulas λ_-). Moreover, for every node x labelled with proposition q, at least one subformula $\Psi_{\langle E; U \rangle}$ allowing to satisfy $\delta(q, l(x))$ is fulfilled (or $\delta(q, l(x)) = \top$, and the result is ensured). When a formula $\Psi_{\langle E; U \rangle}$ holds true at a node x, then there exist $|E|$ distinct successors of x that are labelled with the states in E (they are distinct thanks to subformula Φ_p), and any other successor is labelled with some state in U. Therefore the satisfaction of transitions of the automaton is locally ensured.

Finally, if Ψ_ω is satisfied at the root of $\mathcal{T}$, then for any infinite branch, it is possible to label it with b and then the formula on the right-hand side of the implication states that for any odd priority d, either it appears a finite number of times along the selected branch, or it is not the smallest priority along the branch: this ensures that the smallest infinitely-repeated priority is even, and then the branch satisfies the parity condition. Therefore all branches are accepting, and the tree $\mathcal{T}$ belongs to $\mathcal{L}(\mathcal{A})$. $\qquad \square\square$

Combined with the simulation theorem (Theorem 25), the previous result provides the following corollary for alternating automata:

Corollary 39. *For any $\mathcal{A}$ be an AEUPTA, we can build an exponential-size EQ^2CTL formula $\Phi_{\mathcal{A}}$ such that, for any 2^{AP}-labelled tree $\mathcal{T}$, it holds $\mathcal{T} \in \mathcal{L}(\mathcal{A})$ if, and only if, $\mathcal{T}, \varepsilon \models \Phi_{\mathcal{A}}$.*

This result combined with the automata construction of Section 6.2.2 allows us to prove important properties about the expressive power of QCTL.

6.4 Results about **QCTL** expressiveness

A logic $\mathcal{L}$ is said to be *at least as expressive* as a logic $\mathcal{L}'$ over a class $\mathcal{M}$ of models, which we denote by $\mathcal{L} \succeq_\mathcal{M} \mathcal{L}'$ (omitting to mention $\mathcal{M}$ if it is clear from the context), whenever for any formula $\phi' \in \mathcal{L}'$, there is a $\phi \in \mathcal{L}$ such that ϕ and ϕ' are equivalent over $\mathcal{M}$. Both logics $\mathcal{L}$ and $\mathcal{L}'$ are *equally expressive*, denoted $\mathcal{L} \cong_\mathcal{M} \mathcal{L}'$, when $\mathcal{L} \succeq \mathcal{L}'$ and $\mathcal{L}' \succeq \mathcal{L}$; finally, $\mathcal{L}$ is *strictly more expressive* than $\mathcal{L}'$, written $\mathcal{L} \succ_\mathcal{M} \mathcal{L}'$, if $\mathcal{L} \succeq \mathcal{L}'$ and $\mathcal{L}' \not\succeq \mathcal{L}$. We use $\mathcal{L} \sqcap \mathcal{L}'$ to denote the fragment of $\mathcal{L} \cup \mathcal{L}'$ containing formulas for which there are equivalent formulas in both $\mathcal{L}$ and $\mathcal{L}'$.

Combining the construction of Section 6.2.2 from QCTL formulas into AEUPTA and the construction of the previous section from automata to QCTL allows us to prove that in term of expressive power the hierarchy $\mathsf{Q^kCTL}$ collapses at level 2:

Theorem 40. *QCTL^*, QCTL, $\mathsf{EQ^2CTL}$ and $\mathsf{AQ^2CTL}$ are equally expressive.*

Proof. Given a QCTL^* formula Φ, one can build an AEUPTA $\mathcal{A}_\Phi$ which recognises the 2^{AP_Φ}-labelled trees satisfying Φ, where AP_Φ denotes the set of atomic propositions occurring in Φ. This automaton $\mathcal{A}_\Phi$ can then be transformed into a non-alternating EUPTA $\mathcal{N}_\Phi$, from which we can build a formula $\Phi_\mathcal{N}$ belonging to $\mathsf{EQ^2CTL}$. By construction, we have $\Phi \equiv \Phi_\mathcal{N}$ (over any 2^{AP}-labelled tree). The same holds for $\neg\Phi$, and the negation of the resulting $\mathsf{EQ^2CTL}$ formula belongs to $\mathsf{AQ^2CTL}$ and is equivalent to Φ. $\square$

Translating from QCTL to $\mathsf{EQ^2CTL}$ induces a complexity blow-up. Indeed given a $\mathsf{Q^kCTL}$ formula Φ, the size of the resulting $\mathsf{EQ^2CTL}$ formula Φ' equivalent to Φ and defined in the previous proof is $(k+1)\text{-}\mathsf{exp}(|\Phi|)$; note that our complexity results about the satisfiability of $\mathsf{EQ^2CTL}$ and $\mathsf{Q^kCTL}$ entail that any translation procedure to get such a $\mathsf{EQ^2CTL}$ formula has time complexity at least $(k-1)\text{-}\mathsf{exp}(|\Phi|)$.

We then have $\mathsf{QCTL}^* \cong \mathsf{QCTL} \cong \mathsf{EQ^2CTL} \sqcap \mathsf{AQ^2CTL}$. But there is a difference between $\mathsf{Q^2CTL}$, $\mathsf{Q^1CTL}$ and CTL. The following theorem summarises our expressiveness results:

Theorem 41. *In terms of their relative expressiveness, the fragments of QCTL satisfy the following relations:*

$$CTL \prec EQ^1CTL \sqcap AQ^1CTL \begin{array}{c} \nearrow EQ^1CTL \searrow \\ \\ \searrow AQ^1CTL \nearrow \end{array} Q^1CTL \prec EQ^2CTL \sqcap AQ^2CTL \cong QCTL^*$$

Proof. We first prove that $\mathsf{EQ^2CTL}$ is strictly more expressive than $\mathsf{Q^1CTL}$. By duality, this extends to $\mathsf{AQ^2CTL}$. We already proved that QCTL, hence also $\mathsf{Q^1CTL}$, can be translated in $\mathsf{EQ^2CTL}$ and in $\mathsf{AQ^2CTL}$. We exhibit an $\mathsf{EQ^2CTL}$ formula that $\mathsf{Q^1CTL}$ cannot express, namely:

$$\lambda = \exists p.\ \forall q.\ [\mathbf{EX}(p \wedge (\mathbf{AX}q \vee \mathbf{AX}\neg q)) \wedge \mathbf{EX}(\neg p \wedge (\mathbf{AX}q \vee \mathbf{AX}\neg q))].$$

It specifies that there exist at least two (immediate) successors whose arity is 1. Consider the trees $\mathcal{T}_k$ and $\mathcal{T}'_k$ depicted at Fig. 8. We prove that λ holds in $\mathcal{T}'_k$, but fails to hold in $\mathcal{T}_k$: in $\mathcal{T}'_k$, take the p-labelling where only r'_1 is labelled with p: then for any q-labelling, r'_1 satisfies $p \wedge (\mathbf{AX}q \vee \mathbf{AX}\neg q)$ and r'_2 satisfies
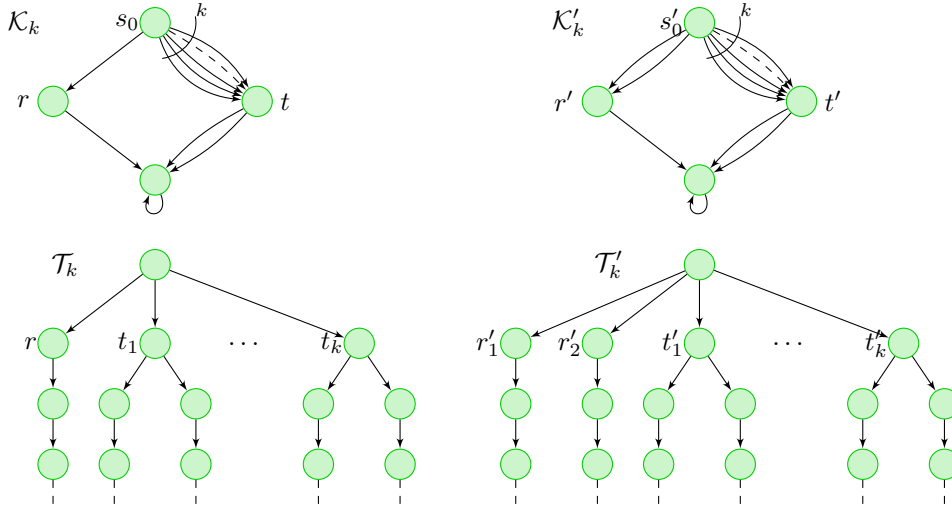
Figure 8: Two (families of) Kripke structures and their computation trees

$\neg p \wedge (\mathbf{AX}q \vee \mathbf{AX}\neg q)$, so that λ holds in $\mathcal{T}'_k$. Now, take any p-labelling of $\mathcal{T}_k$, and the q-labelling in which exactly one of the successors of each node t_i is labelled with q. Then none of the states t_i can be used to satisfy any of the two conjuncts of λ, and r alone can't satisfy both. Hence $\mathcal{T}_k$ does not satisfy λ.

We now prove that $\mathcal{T}_k$ and $\mathcal{T}'_k$ satisfy the same $\mathsf{Q}^1\mathsf{CTL}$ formulas of size at most k. For convenience, we define a binary relation $\mathcal{R}_k$ between states of $\mathcal{T}_k$ and states of $\mathcal{T}'_k$, by letting $\mathcal{R}_k = \{(s_0, s'_0), (r, r'_1), (r, r'_2)\} \cup \{(t_i, t'_j) \mid 1 \leq i, j \leq k\}$. The proof then proceeds in two steps:

- we first prove that all states in relation by $\mathcal{R}_k$ satisfy the same $\mathsf{EQ}^1\mathsf{CTL}$ formulas of size at most k;

- we then prove that s_0 and s'_0 satisfy the same $\mathsf{Q}^1\mathsf{CTL}$ formulas of size at most k.

For the first step: the result is straightforward (as the subtrees are the same) for all pairs of states in $\mathcal{R}_k$ but (s_0, s'_0). Take a formula ϕ in $\mathsf{EQ}^1\mathsf{CTL}$. If ϕ is in CTL, the result is clear as the subtrees are bisimilar. We thus consider the case where $\phi = \exists (p_i)_i. \psi$, where ψ can be written as a boolean combination of at most k CTL formulas of the form $\mathbf{E}\zeta_i$ or $\mathbf{A}\zeta_i$.

Assume $\mathcal{T}_k, s_0 \models \phi$, and consider a labelling ℓ of $\mathcal{T}_k$ with atomic propositions $(p_i)_i$ witnessing this fact. Consider the labelling ℓ' of $\mathcal{T}'_k$ where the subtree under each t'_j is labelled in the same way as ℓ labels the subtree under the corresponding t_j, and the subtrees under r'_1 and r'_2 are labelled in the same way as the subtree under r. The labelled trees $\mathcal{T}_k$ and $\mathcal{T}'_k$ are then bisimilar; since ψ holds in $\mathcal{T}_k$ with labelling ℓ, it also holds in $\mathcal{T}'_k$ with labelling ℓ'.

Conversely, assume that $\mathcal{T}'_k, s_0 \models \phi$, and take a labelling ℓ' witnessing this. Consider a first labelling ℓ of $\mathcal{T}_k$ in which the subtree under each t_j is labelled in the same way as the subtree under t'_j, and the subtree under r is labelled in the same way as the subtree under r_1. All subformulas of ψ of the form $\mathbf{A}\zeta_i$ that hold true at s'_0 in $\mathcal{T}_k$ labelled with ℓ' also hold true at s_0 in $\mathcal{T}_k$ with ℓ, since

53

the paths in the latter are paths in the former. Similarly, all subformulas of ψ of the form $\mathbf{E}\zeta_i$ that hold true at s'_0 in $\mathcal{T}'_k$ under ℓ' also hold true at s_0 in $\mathcal{T}_k$ under ℓ, except for those that are witnessed by the path through r'_2, which is the only path that has no counterpart in $\mathcal{T}_k$ under labelling ℓ. However, since there are at most k such subformulas, at least one path in $\mathcal{T}_k$, say one going through t_k, is not used to fulfill any of the ζ_i subformulas. We then update the labelling ℓ by labelling *both* branches under t_k in the same way as ℓ' labels the subtree under r'_2. This way, the subtrees under r'_2 and under t_k are bisimilar, hence ℓ now fulfill all the required subformulas, so that ϕ also holds in s_0.

The second step of the proof is easy: take a formula ϕ in $\mathsf{Q}^1\mathsf{CTL}$ of size at most k. By definition of $\mathsf{Q}^1\mathsf{CTL}$, it can be written as $\phi[(\psi_i)_i]$, where ψ_i are $\mathsf{EQ}^1\mathsf{CTL}$ formulas. We label the nodes of $\mathcal{T}_k$ and $\mathcal{T}'_k$ with new atomic propositions $(p_i)_i$, in such a way that node n is labelled with p_i if, and only if, it satisfies ψ_i. Since all ψ_i have size at most k, thanks to the result of the first step, the resulting labelled trees are bisimilar. Hence they satisfy the same CTL formulas, in particular they both do or both don't satisfy ϕ, which concludes our proof.

It remains to settle the relative expressiveness of $\mathsf{EQ}^1\mathsf{CTL}$, $\mathsf{AQ}^1\mathsf{CTL}$, and $\mathsf{Q}^1\mathsf{CTL}$. We first show $\mathsf{EQ}^1\mathsf{CTL} \sqcap \mathsf{AQ}^1\mathsf{CTL} \prec \mathsf{EQ}^1\mathsf{CTL}$. For this, it is sufficient to provide an $\mathsf{EQ}^1\mathsf{CTL}$ formula ϕ such that no $\mathsf{AQ}^1\mathsf{CTL}$ formulas are equivalent to φ. Consider $\varphi = \exists p.\,(\mathbf{EX}p \wedge \mathbf{EX}\neg p)$, which characterises all nodes having at least two successors. Now, consider an $\mathsf{AQ}^1\mathsf{CTL}$ formula $\psi = \forall p_1 \ldots p_n.\,\widetilde{\psi}$ with $\widetilde{\psi} \in \mathsf{CTL}$, and assume that ψ is equivalent to φ. Let $\mathcal{K}_1 = (\{s_0, s_1\}, \{(s_0, s_1), (s_1, s_1)\}, \varnothing)$ and $\mathcal{K}_2 = (\{s_0, s_1, s_2\}, \{(s_0, s_1), (s_0, s_2), (s_1, s_1), (s_2, s_2)\}, \varnothing)$ be two Kripke structures such that s_0 in $\mathcal{K}_1$ (resp. s_0 in $\mathcal{K}_2$) has one successor (resp. two successors). Therefore $\mathcal{K}_1, s_0 \not\models \varphi$ and $\mathcal{K}_2, s_0 \models \varphi$.

If ψ is equivalent to φ, then $\mathcal{K}_2, s_0 \models \forall p_1 \ldots p_n.\,\widetilde{\psi}$. Therefore, for any labelling of the tree $\mathcal{T}_{\mathcal{K}_2, s_0}$ with propositions p_1 to p_n, the CTL formula $\widetilde{\psi}$ is satisfied. This is in particular true of the labellings that label both branches of $\mathcal{T}_{\mathcal{K}_2, s_0}$ in the same way; since CTL cannot distinguish between bisimilar structures, we deduce that $\mathcal{K}_1, s_0 \models \forall p_1 \ldots p_n.\,\widetilde{\psi}$; this contradicts the hypothesis that $\varphi \equiv \psi$. Since $\mathsf{EQ}^1\mathsf{CTL} \sqcap \mathsf{AQ}^1\mathsf{CTL}$ is closed under negation, we also get $\mathsf{EQ}^1\mathsf{CTL} \sqcap \mathsf{AQ}^1\mathsf{CTL} \prec \mathsf{AQ}^1\mathsf{CTL}$.

Finally, we notice that $\mathsf{EQ}^1\mathsf{CTL} \sqcap \mathsf{AQ}^1\mathsf{CTL}$ strictly contains CTL: indeed, consider the property $\mathsf{even}(p)$, which characterises all trees in which all nodes at even depth are labelled with some atomic proposition p (and in which all nodes at odd depth may or may not be labelled with p). It is well-known that such a property cannot be expressed in CTL [Wol83]. We now express it in both $\mathsf{EQ}^1\mathsf{CTL}$ and $\mathsf{AQ}^1\mathsf{CTL}$:

- in $\mathsf{EQ}^1\mathsf{CTL}$, we first label all nodes at even depth with a new atomic proposition q, and require that all nodes labelled with q must also be labelled with p:

$$\exists q.\,(q \wedge \mathbf{AG}(q \Leftrightarrow \mathbf{AX}\neg q)) \wedge \mathbf{AG}(q \Rightarrow p);$$

- in $\mathsf{AQ}^1\mathsf{CTL}$, we write that any labelling that labels exactly all nodes at even depth with q (there is a unique such labelling) is such that all nodes labelled with q are also labelled with p:

$$\forall q.\,(q \wedge \mathbf{AG}(q \Leftrightarrow \mathbf{AX}\neg q)) \Rightarrow \mathbf{AG}(q \Rightarrow p). \qquad \square$$

7 Application to MSO

We briefly review *Monadic Second-Order Logic* (MSO) over finite or infinite trees. We use constant monadic predicates P_a for $a \in \mathsf{AP}$ and a relation Edge for the immediate successor relation in a 2^{AP}-labelled tree $\mathcal{T} = (t, l)$.

MSO is built with first-order (or individual) variables for vertices (denoted with lowercase letters $x, y, ...$), and monadic second-order variables for sets of vertices (denoted with uppercase letters $X, Y, ...$). Atomic formulas are of the form $x = y$, $\mathsf{Edge}(x, y)$, $x \in X$, and $\mathsf{P}_a(x)$. General MSO formulas are constructed from atomic formulas using the boolean connectives and the first- and second-order quantifiers $\exists^1$ and $\exists^2$, which we both denote with $\exists$ in the sequel as long as this is not ambiguous. We write $\phi(x_1, ..., x_n, X_1, ..., X_k)$ to state that $x_1, ..., x_n$ and $X_1, ..., X_k$ may appear free (*i.e.*, not within the scope of a quantifier) in ϕ. A closed formula contains no free variables.

We use the standard semantics for MSO: given a tree $\mathcal{T}$, a sequence of nodes s_1 to s_n, and a sequence of sets of nodes S_1 to S_k, we write $\mathcal{T}, s_1, ..., s_n, S_1,$ $..., S_k \models \phi(x_1, ..., x_n, X_1, ..., X_k)$ to indicate that ϕ holds on $\mathcal{T}$ when variables x_1 to x_n in ϕ are replaced with s_1 to s_n, and variables X_1 to X_k are replaced with S_1 to S_k. As an example, the closed formula

$$\forall x. \; \big(P_a(x) \Rightarrow \big[\exists X. \; (x \in X \wedge \forall y. \; (y \in X \Rightarrow \exists z. \; (z \in X \wedge \mathsf{Edge}(y, z))))\big]\big)$$

holds true for any tree in which any node labelled with a belongs to (at least one) infinite branch.

More details about MSO can be found e.g. in [Tho97]. In [LM14], it is proved that MSO and QCTL are equally expressive over trees. This could be used to define translations between MSO and EU-automata, but we prefer direct, more efficient constructions, which we develop below.

7.1 From MSO to AEUPTA

In this section, given a closed formula $\Phi \in \mathsf{MSO}$, we build an AEUTA $\mathcal{A}_\Phi$ such that $\mathcal{L}(\mathcal{A}_\Phi)$ is the set of all trees satisfying Φ. Actually, for any (non-closed) MSO formula $\phi(x_1, ..., x_n, X_1, ..., X_k)$, we build an automaton $\mathcal{A}_\phi$ such that, for any nodes s_1 to s_n and any sets S_1 to S_k, it holds $\mathcal{T}, s_1, ..., s_n, S_1, ..., S_k \models$ $\phi(x_1, ..., x_n, X_1, ..., X_k)$ if, and only if, the $2^{\mathsf{AP} \cup \{x_1, ..., x_n, X_1, ..., X_k\}}$-labelled tree $\mathcal{T}'$, obtained from $\mathcal{T}$ by labelling any node t with x_i if $t = s_i$, and with X_j if $t \in S_j$, belongs to $\mathcal{L}(\mathcal{A}_\phi)$. In the following we consider MSO formulas where the negations can only be followed by an existential quantifier or a atomic formula (of the form $\mathsf{Edge}(x, y)$, $x = y$, $\mathsf{P}_a(x)$ or $x \in X$).

The automaton $\mathcal{A}_\phi$ is built inductively on the structure of ϕ. Handling Boolean connectives $\wedge$ and $\vee$ is done with the corresponding operations of AEUTA. For quantifications $\exists X. \psi$ or $\exists x. \psi$, we use the projection operation (after having built an equivalent non-alternating automaton for $\mathcal{A}_\psi$) exactly as for the QCTL formulas $\exists p. \psi$, and we add a verification step to ensure that every proposition corresponding to some first-order variable labels exactly one node in the tree, we will describe this construction below. First we consider several types of automata to deal with atomic MSO formulas and their negations. In each case, we use a state $q_\top$ that accepts any tree. We have:

- If ϕ is $\mathsf{Edge}(x, y)$: We define the automaton $\mathcal{A}_E = (\{q_E, q'_E\}, q_E, \delta_E, \omega_E)$ as follows:

$$\delta_E(q_E, \sigma) = \begin{cases} \langle q'_E \mapsto 1; \{q_\top\} \rangle & \text{if } x \in \sigma \\ \langle q_E \mapsto 1; \{q_\top\} \rangle & \text{otherwise} \end{cases} \qquad \delta_E(q'_E, \sigma) = \begin{cases} \top & \text{if } y \in \sigma \\ \bot & \text{otherwise} \end{cases}$$

with $\omega_E(q_E) = \omega_E(q'_E) = 1$.

- If ϕ is $\neg\mathsf{Edge}(x, y)$: We define the automaton $\mathcal{A}_{\bar{E}} = (\{q_{\bar{E}}, q'_{\bar{E}}\}, q_{\bar{E}}, \delta_{\bar{E}}, \omega_{\bar{E}})$ as follows:

$$\delta_{\bar{E}}(q_{\bar{E}}, \sigma) = \begin{cases} \langle \varnothing; \{q'_{\bar{E}}\} \rangle & \text{if } x \in \sigma \\ \langle q_{\bar{E}} \mapsto 1; \{q_\top\} \rangle & \text{otherwise} \end{cases} \qquad \delta_{\bar{E}}(q'_{\bar{E}}, \sigma) = \begin{cases} \bot & \text{if } y \in \sigma \\ \top & \text{otherwise} \end{cases}$$

with $\omega_{\bar{E}}(q_{\bar{E}}) = \omega_{\bar{E}}(q'_{\bar{E}}) = 1$.

- If ϕ is $x = y$: We define the automaton $\mathcal{A}_= = (\{q_=\}, q_=, \delta_=, \omega_=)$ as follows:

$$\delta_=(q_=, \sigma) = \begin{cases} \top & \text{if } x, y \in \sigma \\ \bot & \text{if } (x \in \sigma \wedge y \notin \sigma) \vee (x \notin \sigma \wedge y \in \sigma) \\ \langle q_= \mapsto 1; \{q_\top\} \rangle & \text{otherwise} \end{cases}$$

with $\omega_=(q_=) = 1$.

- If ϕ is $\neg(x = y)$: We define the automaton $\mathcal{A}_{\neq} = (\{q_{\neq}\}, q_{\neq}, \delta_{\neq}, \omega_{\neq})$ as follows:

$$\delta_{\neq}(q_{\neq}, \sigma) = \begin{cases} \bot & \text{if } x, y \in \sigma \\ \top & \text{if } (x \in \sigma \wedge y \notin \sigma) \vee (x \notin \sigma \wedge y \in \sigma) \\ \langle q_{\neq} \mapsto 1; \{q_\top\} \rangle & \text{otherwise} \end{cases}$$

with $\omega_{\neq}(q_{\neq}) = 1$.

- if ϕ is $\mathsf{P}_a(x)$: We define the automaton $\mathcal{A}_a = (\{q_a\}, q_a, \delta_a, \omega_a)$ as follows:

$$\delta_a(q_a, \sigma) = \begin{cases} \top & \text{if } x, a \in \sigma \\ \bot & \text{if } x \in \sigma \wedge a \notin \sigma \\ \langle q_a \mapsto 1; \{q_\top\} \rangle & \text{otherwise} \end{cases}$$

with $\omega_a(q_a) = 1$.

- if ϕ is $\neg\mathsf{P}_a(x)$: We define the automaton $\mathcal{A}_{\bar{a}} = (\{q_{\bar{a}}\}, q_{\bar{a}}, \delta_{\bar{a}}, \omega_{\bar{a}})$ as follows:

$$\delta_{\bar{a}}(q_{\bar{a}}, \sigma) = \begin{cases} \top & \text{if } x \in \sigma \wedge a \notin \sigma \\ \bot & \text{if } x, a \in \sigma \\ \langle q_{\bar{a}} \mapsto 1; \{q_\top\} \rangle & \text{otherwise} \end{cases}$$

with $\omega_{\bar{a}}(q_{\bar{a}}) = 1$.

The correctness of the constructions for Edge is stated as follows: given a $2^{\mathsf{AP}\cup\{x,y\}}$-labeled tree $\mathcal{T} = (t,l)$ such that there exists exactly one node $n \in t$ (resp. $n' \in t$) such that $x \in l(n)$ (resp. $y \in l(n')$), we have $\mathcal{T} \in \mathcal{L}(\mathcal{A}_E)$ if, and only if, $\mathcal{T}, n, n' \models \mathsf{Edge}(x,y)$. We proceed in a similar way for the other cases. The correctness proofs are straightforward.

Now we can follow exactly the same steps as for QCTL formulas. Let Φ be an MSO formula without any (first-order or second-order) quantifier. The previous automata constructions can then be composed with union and intersection operations in order to get an AEUPTA whose size is bounded by $(O(|\Phi|), O(|\Phi|), O(1), O(1), O(1))$.

Consider a formula $\Phi = \exists \mathcal{V}.\phi$ where $\mathcal{V}$ is a set of variables $\{x_1, \dots x_m\} \cup \{X_1 \dots X_p\}$ (where every x_j is a first-order variable and every X_j is a second-order variable) and ϕ is an MSO formula without any quantifier. As explained above, one can build an AEUPTA $\mathcal{A}_\phi$ for ϕ. We can also combine $\mathcal{A}_\phi$ with (a conjunction of) automata $\mathcal{A}_{x_j}$, for every $1 \le j \le m$, to ensure that the letter x_j labels exactly one node. Such an automaton $\mathcal{A}_x$ is then defined as $\mathcal{A}_x = (\{q_x, q_{\bar{x}}\}, q_x, \delta_x, \omega_x)$ with:

$$\delta_x(q_x, \sigma) = \begin{cases} \langle \varnothing; \{q_{\bar{x}}\} \rangle & \text{if } x \in \sigma \\ \langle q_x \mapsto 1; \{q_{\bar{x}}\} \rangle & \text{otherwise} \end{cases} \qquad \delta_x(q_{\bar{x}}, \sigma) = \begin{cases} \bot & \text{if } x \in \sigma \\ \langle \varnothing; \{q_{\bar{x}}\} \rangle & \text{otherwise} \end{cases}$$

with $\omega_x(q_x) = 1$ and $\omega_x(q_{\bar{x}}) = 0$.

This provides an AEUPTA $\mathcal{A}_{\phi'}$ whose size is bounded by $(O(|\phi|), O(|\phi|), 1, 1, 2)$ and which recognises precisely the trees satisfying ϕ and where every first-order variable x_i labels exactly one node in the tree. Applying the simulation theorem (Theorem 25), we get an EUPTA whose size is bounded by $(1\text{-}\mathsf{exp}(|\phi|), 1\text{-}\mathsf{exp}(|\phi|), 0\text{-}\mathsf{exp}(|\phi|), 1, 0\text{-}\mathsf{exp}(|\phi|))$. It remains to use the projection to get a (non-alternating) automaton which recognises precisely the infinite trees satisfying the formula $\Phi = \exists \mathcal{V}.\phi$.

As for QCTL and QCTL*, one can extend the previous construction for any MSO formula ϕ. As we have done for QCTL with the definition of $\mathsf{Q}^k\mathsf{CTL}$, $\mathsf{EQ}^k\mathsf{CTL}$ and $\mathsf{AQ}^k\mathsf{CTL}$, we can define a similar notion of *maximal number of quantifier alternation* in an MSO formula. An important point is that we do not distinguish between first-order and second-order quantifiers: both quantifiers are treated in the same way, via the projection operation; in both cases, each quantifier alternation induces an exponential blow-up, due to the simulation step. By proceeding exactly as for QCTL, and with the specific treatment of first-order quantifiers as explained above, we get:

Theorem 42. *Given a closed MSO formula ϕ over AP with at most k quantifier alternations (with $k > 0$), we can construct a AEUPTA $\mathcal{A}_\phi$ over 2^{AP} accepting exactly the trees satisfying ϕ. The automaton $\mathcal{A}_\phi$ has size $(k\text{-}\mathsf{exp}(|\phi|), k\text{-}\mathsf{exp}(|\phi|), (k-1)\text{-}\mathsf{exp}(|\phi|), 1, (k-1)\text{-}\mathsf{exp}(|\phi|))$.*

As a corollary, we get the following results about decision procedures for MSO via AEUPTA construction:

Corollary 43. *Let ϕ be an MSO formula with at most k quantifier alternations. The satisfiability problem for ϕ is $(k+1)$-EXPTIME-complete. The model-checking problem for ϕ over a finite Kripke structure is k-EXPTIME-complete.*

7.2 From **AEUPTA** to **MSO**

Expressing acceptance of some tree $\mathcal{T}$ by some EUPTA $\mathcal{A}$ as an MSO formula is based on the same techniques as the ones we used for QCTL: an existential (second-order) quantification is used to label every node of $\mathcal{T}$ with a unique state of $\mathcal{A}$; the rest of the formula checks that the (non-alternating) transition function is fulfilled locally at any node, and that for every infinite branch (which we encode using second-order quantification), there is a suffix (this again involves second-order quantification[7]) along which all nodes are labelled with states of $\mathcal{A}$ whose priorities are greater than or equal to some even value that occurs infinitely often along the suffix.

Consider an EUPTA $\mathcal{A} = (Q, q_0, \delta, \omega)$ over $\Sigma = 2^{\mathsf{AP}}$. Let Q be $\{q_0, \ldots, q_n\}$ and let D be the set of priorities in ω (in the following, we use $D_{\geq k}$ to denote the subset of priorities greater than or equal to k). Formally, we define $\Phi_{\mathcal{A}}$ as:

$$\Phi_{\mathcal{A}} = \exists Q_0 \ldots Q_n. \exists x_\varepsilon. \left(\neg (\exists y. \mathsf{Edge}(y, x_\varepsilon)) \wedge (x_\varepsilon \in Q_0) \wedge \Phi_\delta \wedge \right.$$
$$\left. \forall B. \left[\mathsf{Br}(B, x_\varepsilon) \Longrightarrow \bigvee_{\substack{d \in D \\ d \text{ even}}} \left(\exists S. \mathsf{Suff}(S, B, d) \wedge \neg \exists S. \mathsf{Suff}(S, B, d+1) \right) \right] \right).$$

In this formula, quantification over Q_0 to Q_n is used to label the nodes of $\mathcal{T}$ with states of $\mathcal{A}$, and quantification over x_ε is used to characterise the root of $\mathcal{T}$; subformula Φ_δ ensures that each node is labelled with exactly one state of $\mathcal{A}$, and that the transition function is satisfied; the second line of the formula encodes the parity acceptance condition, requiring that the minimal repeated priority along any infinite branch is even: $\mathsf{Br}(B, u)$ states that the set of nodes labelled with B forms an infinite branch from u, and $\mathsf{Suff}(S, B, d)$ states that the set of nodes labelled with S is a suffix of the branch B and contains only nodes whose labels in Q have priorities larger than or equal to d.

Consistency w.r.t. the transition function is expressed as follows:

$$\Phi_\delta = \forall x. \left[\bigvee_{0 \leq i \leq n} \left(x \in Q_i \wedge \bigwedge_{j \neq i} x \notin Q_j \wedge \right. \right.$$
$$\left. \left. \bigvee_{R \in 2^{\mathsf{AP}}} \left[\Phi_R(x) \wedge \left(\delta(q_i, R) \overset{?}{=} \top \vee \bigvee_{\langle E_j; U_j \rangle \in \delta(q_i, R)} \Phi_{\langle E_j; U_j \rangle}(x) \right) \right] \right) \right]$$

where the subformula $\delta(q_i, R) \overset{?}{=} \top$ on the second line is just replaced by $\top$ for all R for which the equality holds, and with $\bot$ otherwise. For a subset $R \subseteq \mathsf{AP}$, the formula $\Phi_R(x)$ specifies that x is labelled exactly with the propositions in R: $\Phi_R(x) = \bigwedge_{p \in R} \mathsf{P}_p(x) \wedge \bigwedge_{p \in \mathsf{AP} \setminus R} \neg \mathsf{P}_p(x)$.

The formula $\Phi_{\langle E; U \rangle}(x)$ requires that the successors of the node labeled by x satisfy the EU-pairs $\langle E; U \rangle$ and it is defined as follows:

$$\Phi_{\langle \{\!\!\{ r_1, \ldots, r_k \}\!\!\}; \{s_1, \ldots, s_m\} \rangle}(x) = \exists x_1 \ldots x_k. \left(\bigwedge_{1 \leq i \leq k} \left(\mathsf{Edge}(x, x_i) \wedge \bigwedge_{\substack{1 \leq j \leq k \\ j \neq i}} x_i \neq x_j \wedge \right. \right.$$
$$\left. \left. \mathsf{P}_{r_i}(x_i) \right) \wedge \forall z. \left[\left(\mathsf{Edge}(x, z) \wedge \bigwedge_{1 \leq i \leq k} z \neq x_i \right) \Rightarrow \left(\bigvee_{1 \leq i \leq m} \mathsf{P}_{s_i}(z) \right) \right] \right)$$

---

[7]Remember that our definition of MSO does not include the transitive *successor* relation $<$ in its signature: it only allows the non-transitive *direct successor* relation Edge. This is why we need second-order quantification to quantify over suffixes.

Formula $\mathsf{Br}(B, u)$, stating that the set of nodes labelled with B forms an infinite branch from the node labelled with u, can be written as

$$\mathsf{Br}(B, u) = (u \in B) \wedge \forall x \in B. \ \Big([(x = u) \Leftrightarrow \neg(\exists y \in B. \ \mathsf{Edge}(y, x))] \wedge$$

$$\exists y \in B. \ (\mathsf{Edge}(x, y) \wedge (\forall z.((\mathsf{Edge}(x, z) \wedge z \neq y) \Rightarrow z \notin B)))\Big)$$

Formula $\mathsf{Suff}(S, B, d)$ stating that the set of nodes labelled with S forms a suffix of the branch labelled with B in which all nodes are labelled with states having priorities greater than or equal to d, is expressed as:

$$\mathsf{Suff}(S, B, d) = \big[\forall x \in S. \ x \in B\big] \wedge \exists u \in S. \ \big[\forall x \in B.$$

$$\big(x \in S \Leftrightarrow (x = u \vee \exists y \in S. \ \mathsf{Edge}(y, x))\big)\big] \wedge \Big(\forall z \in S. \bigvee_{\substack{q_i \in Q \ \text{s.t.} \\ \omega(q_i) \in D_{\geq d}}} z \in Q_i \Big)$$

Finally we can observe that the size of $\Phi_{\mathcal{A}}$ is in $O(|Q| \cdot (|Q| + 2^{|\mathsf{AP}|} \cdot (|\mathsf{AP}| + |\delta|_{\mathsf{Bool}} \cdot (|\delta|_{\mathsf{E}} + |\delta|_{\mathsf{U}}))) + |\omega| \cdot |Q|)$, from which we can deduce that $|\Phi|_{\mathcal{A}}$ is in $O(|Q|^2 \cdot 2^{|\mathsf{AP}|} \cdot (|\mathsf{AP}| + |\delta|_{\mathsf{Bool}} \cdot (|\delta|_{\mathsf{E}} + |\delta|_{\mathsf{U}})))$. Formula $\Phi_{\mathcal{A}}$ contains four alternations of (first-order or second-order) quantifiers. Note also that it contains four blocks of second-order quantifiers and the number of second-order alternations is 2.

It follows:

Theorem 44. *Any closed MSO formula ϕ with at most k quantifier alternations (with $k > 0$) can be translated into another, equivalent MSO formula with four alternations of (first-order or second-order) quantifiers and two alternations of second-order quantifiers. The size of the resulting formula can be bounded by $(k + 1)$-exp$(|\phi|)$.*

Proof. From ϕ, we can build an AEUPTA $\mathcal{A}_\phi$ over Σ whose size is in $(k$-exp$(|\phi|)$, k-exp$(|\phi|), (k - 1)$-exp$(|\phi|), 1, (k - 1)$-exp$(|\phi|))$. Here the alphabet Σ is 2^{AP_ϕ} where AP_ϕ is the set of monadic predicates occurring in ϕ (and then $|\mathsf{AP}_\phi| \leq |\phi|$).

Applying the simulation theorem provides us with EUPTAs $\mathcal{N}_\phi$ whose size is $((k + 1)$-exp$(|\phi|), (k + 1)$-exp$(|\phi|), k$-exp$(|\phi|), 1, k$-exp$(|\phi|))$. It remains to build $\Phi_{\mathcal{N}_\phi}$ as above to get the result. $\qquad\square$

Remark 8. *It is claimed in [Tho97] that MSO formulas can be translated into formulas with only one alternation of second-order quantifiers (for binary trees). This is because the signature of MSO in that paper includes the transitive relation $<$ in place of our non-transitive Edge relation.*

If we allow for the use of $<$, we can drop the existential quantification over S, and characterise suffixes of a branch B with only their starting node u (as explained in note 7). ◀

8 Conclusion

We have introduced a new class of symmetric tree automata (AEUPTA) for trees of arbitrary branching degrees. We showed that these automata have exactly the same expressive power as the temporal logics QCTL and QCTL*, and as the logic MSO: given a formula Φ in those formalisms, the set of infinite

trees satisfying Φ can be defined as the language of some automaton $\mathcal{A}_\Phi$, and conversely for any AEUPTA $\mathcal{A}$ one can build a formula $\Phi_{\mathcal{A}}$ whose models are precisely $\mathcal{L}(\mathcal{A})$.

In order to prove those results, we have developed algorithms for manipulating our AEUPTA, and have carefully studied their complexities. This has allowed us to obtain decision procedures for satisfiability and model checking for QCTL* and its fragments whose complexities match the lower-bound established in previous papers [LM14]. It also allowed us to obtain an effective translation from QCTL to EQ^{2}CTL, and similarly, from MSO to its fragment with only two second-order-quantifier alternations.

References

[BB02]   Dietmar Berwanger and Achim Blumensath. The monadic theory of tree-like structures. In Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors, *Automata, Logics, and Infinite Games*, volume 2500 of *Lecture Notes in Computer Science*, chapter 16, pages 285–302. Springer-Verlag, 2002.

[BG93]   Orna Bernholtz and Orna Grumberg. Branching time temporal logic and $\mathcal{A}$mo$r$P$\mathcal{H}$0us tree automata. In Eike Best, editor, *Proceedings of the 4th International Conference on Concurrency Theory (CONCUR'93)*, volume 715 of *Lecture Notes in Computer Science*, pages 262–277. Springer-Verlag, August 1993.

[BL23]   Udi Boker and Karoliina Lehtinen. When a little nondeterminism goes a long way: an introduction to history-determinism. *SIGLOG News*, 35:24–51, January 2023.

[Boj08]   Mikołaj Bojańczyk. Tree-walking automata. In Carlos Martín-Vide, Friedrich Otto, and Henning Fernau, editors, *Revised Papers of the 2nd International Conference on Language and Automata Theory and Applications (LATA'08)*, volume 5196 of *Lecture Notes in Computer Science*, pages 1–2. Springer-Verlag, March 2008.

[Büc62]   Julius R. Büchi. On a decision method in restricted second order arithmetic. In Ernest Nagel, Patrick Suppes, and Alfred Tarski, editors, *Proceedings of the 1960 International Congress on Logic, Methodology and Philosophy of Science (LMPS'60)*, pages 1–11. Stanford University Press, June 1962.

[BVW94]   Orna Bernholtz, Moshe Y. Vardi, and Pierre Wolper. An automata-theoretic approach to branching-time model checking (extended abstract). In David L. Dill, editor, *Proceedings of the 6th International Conference on Computer Aided Verification (CAV'94)*, volume 818 of *Lecture Notes in Computer Science*, pages 142–155. Springer-Verlag, June 1994.

[CDG$^+$08] Hubert Comon, Max Dauchet, Rémi Gilleron, Florent Jacquemard, Denis Lugiez, Christof Löding, Sophie Tison, and Marc Tommasi. Tree automata techniques and applications. Technical Report hal-03367725, H.A.L., 2008.

[DLM12]   Arnaud Da Costa, François Laroussinie, and Nicolas Markey. Quantified CTL: Expressiveness and model checking. In Maciej Koutny and Irek Ulidowski, editors, *Proceedings of the 23rd International Conference on Concurrency Theory (CONCUR'12)*, volume 7454 of *Lecture Notes in Computer Science*, pages 177–192. Springer-Verlag, September 2012.

[EJ91]   E. Allen Emerson and Charanjit S. Jutla. Tree automata, mu-calculus and determinacy. In *Proceedings of the 32nd Annual Symposium on Foundations of Computer Science (FOCS'91)*, pages 368–377. IEEE Comp. Soc. Press, October 1991.

[Elg61]   Calvin C. Elgot. Decision problems of finite automata design and related arithmetics. *Transactions of the American Mathematical Society*, 98(1):21–51, January 1961.

[ES84]   E. Allen Emerson and A. Prasad Sistla. Deciding full branching time logic. *Information and Control*, 61(3):175–201, June 1984.

[Fre01]   Tim French. Decidability of quantified propositional branching time logics. In Markus Stumptner, Dan Corbett, and Mike Brooks, editors, *Proceedings of the 14th Australian Joint Conference on Artificial Intelligence (AJCAI'01)*, volume 2256 of *Lecture Notes in Computer Science*, pages 165–176. Springer-Verlag, December 2001.

[Fre06]   Tim French. *Bisimulation Quantifiers for Modal Logics*. PhD thesis, School of Computer Science & Software Engineering, University of Western Australia, December 2006.

[HM85]   Matthew C. B. Hennessy and Robin Milner. Algebraic laws for non-determinism and concurrency. *Journal of the ACM*, 32(1), January 1985.

[JL04]   David Janin and Giacommo Lenzi. On the relationship between monadic and weak monadic second order logic on arbitrary trees. *Fundamenta Informaticae*, 61(3-4):247–265, 2004.

[JW95]   David Janin and Igor Walukiewicz. Automata for the modal μ-calculus and related results. In Jirí Wiedermann and Petr Hájek, editors, *Proceedings of the 20th International Symposium on Mathematical Foundations of Computer Science (MFCS'95)*, volume 969 of *Lecture Notes in Computer Science*, pages 552–562. Springer-Verlag, August 1995.

[Kir02]   Daniel Kirsten. Alternating tree automata and parity games. In Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors, *Automata, Logics, and Infinite Games*, volume 2500 of *Lecture Notes in Computer Science*, chapter 9, pages 153–167. Springer-Verlag, 2002.

[KMTV00] Orna Kupferman, Parthasarathy Madhusudan, P. S. Thiagarajan, and Moshe Y. Vardi. Open systems in reactive environments: Control and synthesis. In Catuscia Palamidessi, editor, *Proceedings of*

the 11th International Conference on Concurrency Theory (CON-CUR'00), volume 1877 of Lecture Notes in Computer Science, pages 92–107. Springer-Verlag, August 2000.

[KSV06]   Orna Kupferman, Shmuel Safra, and Moshe Y. Vardi. Relating word and tree automata. Annals of Pure and Applied Logic, 138(1-3):126–146, March 2006.

[Kup95]   Orna Kupferman. Augmenting branching temporal logics with existential quantification over atomic propositions. In Pierre Wolper, editor, Proceedings of the 7th International Conference on Computer Aided Verification (CAV'95), volume 939 of Lecture Notes in Computer Science, pages 325–338. Springer-Verlag, July 1995.

[KV03]   Orna Kupferman and Moshe Y. Vardi. $\Pi_2 \cap \Sigma_2 \equiv$ AFMC. In Jos C. M. Baeten, Jan Karel Lenstra, Joachim Parrow, and Gerhard J. Woeginger, editors, Proceedings of the 30th International Colloquium on Automata, Languages and Programming (ICALP'03), volume 2719 of Lecture Notes in Computer Science, pages 697–713. Springer-Verlag, June-July 2003.

[KVW00]   Orna Kupferman, Moshe Y. Vardi, and Pierre Wolper. An automata-theoretic approach to branching-time model-checking. Journal of the ACM, 47(2):312–360, March 2000.

[LM14]   François Laroussinie and Nicolas Markey. Quantified CTL: expressiveness and complexity. Logical Methods in Computer Science, 10(4), December 2014.

[LMS02]   François Laroussinie, Nicolas Markey, and Philippe Schnoebelen. Temporal logic with forgettable past. In Proceedings of the 17th Annual Symposium on Logic in Computer Science (LICS'02), pages 383–392. IEEE Comp. Soc. Press, July 2002.

[Löd21]   Christof Löding. Automata on infinite trees. In Jean-Éric Pin, editor, Handbook of automata theory, volume 1, chapter 8, pages 265–302. EMS Press, 2021.

[McN66]   Robert McNaughton. Testing and generating infinite sequences by a finite automaton. Information and Control, 9(6):521–530, December 1966.

[MS85]   David E. Muller and Paul E. Schupp. Alternating automata on infinite objects, determinacy and Rabin's theorem. In Maurice Nivat and Dominique Perrin, editors, Automata on Infinite Words – École de Printemps d'Informatique Théorique (EPIT'84), volume 192 of Lecture Notes in Computer Science, pages 99–107. Springer-Verlag, 1985.

[MSS86]   David E. Muller, Ahmed Saoudi, and Paul E. Schupp. Alternating automata, the weak monadic theory of the tree, and its complexity. In Laurent Kott, editor, Proceedings of the 13th International Colloquium on Automata, Languages and Programming (ICALP'86),

volume 226 of *Lecture Notes in Computer Science*, pages 275–283. Springer-Verlag, July 1986.

[Pit07]   Nir Piterman. From non-deterministic Büchi and Streett automata to deterministic parity automata. *Logical Methods in Computer Science*, 3(3:5), August 2007.

[Rab69]   Michael O. Rabin. Decidability of second-order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141:1–35, July 1969.

[SVW85]   A. Prasad Sistla, Moshe Y. Vardi, and Pierre Wolper. The complementation problem for Büchi automata, with applications to temporal logic. In Wilfried Brauer, editor, *Proceedings of the 12th International Colloquium on Automata, Languages and Programming (ICALP'85)*, volume 194 of *Lecture Notes in Computer Science*, pages 465–474. Springer-Verlag, July 1985.

[Tho90]   Wolfgang Thomas. Automata on infinite objects. In Jan van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 133–191. Elsevier, 1990.

[Tho97]   Wolfgang Thomas. Languages, automata and logics. In Grzegorz Rozenberg and Arto Salomaa, editors, *Handbook of Formal Languages*, volume 3, pages 389–455. Springer-Verlag, 1997.

[Tra62]   Boris A. Trakhtenbrot. Finite automata and the logic of one-place predicates. *Siberskii Matematicheskii Zhurnal*, 3(1):103–131, 1962.

[VW86a]   Moshe Y. Vardi and Pierre Wolper. An automata-theoretic approach to automatic program verification. In *Proceedings of the 1st Annual Symposium on Logic in Computer Science (LICS'86)*, pages 332–344. IEEE Comp. Soc. Press, June 1986.

[VW86b]   Moshe Y. Vardi and Pierre Wolper. Automata theoretic techniques for modal logics of programs. *Journal of Computer and System Sciences*, 32(2):183–221, April 1986.

[VW08]   Moshe Y. Vardi and Thomas Wilke. Automata: from logics to algorithms. In Jörg Flum, Erich Grädel, and Thomas Wilke, editors, *Logic and Automata: History and Perspectives*, volume 2 of *Texts in Logic and Games*, pages 629–736. Amsterdam University Press, 2008.

[Wal96]   Igor Walukiewicz. Monadic second order logic on tree-like structures. In Claude Puech and Rüdiger Reischuk, editors, *Proceedings of the 13th Symposium on Theoretical Aspects of Computer Science (STACS'96)*, volume 1046 of *Lecture Notes in Computer Science*, pages 401–413. Springer-Verlag, February 1996.

[Wal02]   Igor Walukiewicz. Monadic second order logic on tree-like structures. *Theoretical Computer Science*, 275(1-2):311–346, March 2002.

[Wil99]   Thomas Wilke. CTL$^+$ is exponentially more succinct than CTL. In C. Pandu Rangan, Venkatesh Raman, and R. Ramanujam, editors, *Proceedings of the 19th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'99)*, volume 1738 of *Lecture Notes in Computer Science*, pages 110–121. Springer-Verlag, December 1999.

[Wil01]   Thomas Wilke. Alternating tree automata, parity games, and modal μ-calculus. *Bulletin of the Belgian Mathematical Society – Simon Stevin*, 8(2):359–391, 2001.

[Wol83]   Pierre Wolper. Temporal logic can be more expressive. *Information and Control*, 56(1-2):72–99, 1983.

[Zan12]   Fabio Zanasi. Expressiveness of monadic second-order logics on infinite trees of arbitrary branching degrees. Master's thesis, Amsterdam University, the Netherlands, August 2012.

[Zie98]   Wiesław Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200(1-2):135–183, June 1998.