

The Complexity of the Graded μ -Calculus

Orna Kupferman¹, Ulrike Sattler², Moshe Y. Vardi³

¹ School of Computer Science and Engineering, Hebrew University, Jerusalem, Israel

orna@cs.huji.ac.il

² LuFG Theor. Informatik, RWTH Aachen, Germany

sattler@cs.rwth-aachen.de

³ Department of Computer Science, Rice University, Houston, TX 77251-1892, U.S.A.

vardi@cs.rice.edu

Abstract. In classical logic, existential and universal quantifiers express that there exists at least one individual satisfying a formula, or that all individuals satisfy a formula. In many logics, these quantifiers have been generalized to express that, for a non-negative integer n , at least n individuals or all but n individuals satisfy a formula. In modal logics, *graded modalities* generalize standard existential and universal modalities in that they express, e.g., that there exist at least n accessible worlds satisfying a certain formula. Graded modalities are useful expressive means in knowledge representation; they are present in a variety of other knowledge representation formalisms closely related to modal logic.

A natural question that arises is how the generalization of the existential and universal modalities affects the decidability problem for the logic and its computational complexity, especially when the numbers in the graded modalities are coded in binary. In this paper we study the *graded μ -calculus*, which extends graded modal logic with fixed-point operators, or, equivalently, extends classical μ -calculus with graded modalities. We prove that the decidability problem for graded μ -calculus is EXPTIME-complete – not harder than the decidability problem for μ -calculus, even when the numbers in the graded modalities are coded in binary.

1 Introduction

In classical logic, existential and universal quantifiers express that there exists *at least one* individual satisfying a formula, or that *all* individuals satisfy a formula. In many logics, these quantifiers have been generalized to express that, for a non-negative integer n , *at least n individuals* or *all but n individuals* satisfy a formula. For example, predicate logic has been extended with so-called *counting quantifiers* $\exists^{\geq n}$ and $\exists^{\leq n}$ [GOR97,GMV99,PST00]. In modal logics, *graded modalities* [Fin72,vdHD95,Tob01] generalize standard existential and universal modalities in that they express, e.g., that there exist at least n accessible worlds satisfying a certain formula. In description logics, *number restrictions* have always played a central role; e.g., they are present in almost all knowledge-representation systems based on description logic [PSMB⁺91,BFH⁺94,Hor98,HM01]. Indeed, in a typical such system, one can describe cars as those vehicles having at least four wheels, and bicycles as those vehicles having exactly two wheels.

A natural question that arises is how the generalization of the existential and universal quantifiers affects the decidability problem for the logic and its computational complexity. The complexity of a variety of description logics with different forms of number restrictions has been investigated; see, e.g. [DLNdN91,HB91,DL94b,BBH96,BS99,Tob00]. It turned

out that, in many cases, one can extend a logic with these forms of counting quantifiers without increasing its computational complexity. On the other hand, in some cases the extension makes the logic much more complex. A prominent example is the guarded fragment of first order logic, which becomes undecidable when extended with a very weak form of counting quantifiers (global functionality conditions on binary relations) [Grä99].

When investigating the complexity of a logic with a form of counting quantifiers, one must decide how the numbers in these quantifiers contribute to the length of a formula, i.e., to the input of a decision procedure. Assuming that these numbers are coded in unary (i.e., $|\exists^{\geq n} x. \varphi(x)| = n + |\varphi(x)|$) might seem odd, but is an assumption often made, for example in description and predicate logics. It reflects the way in which many decision procedures for these logics work: they explicitly generate n individuals for $\exists^{\geq n}$. In contrast, the assumption that the numbers are coded in binary (i.e., $|\exists^{\geq n} x. \varphi(x)| = \lceil \log n \rceil + |\varphi(x)|$) corresponds more closely to our intuition on the length of a formula, but it is not clear whether and how decision procedures can avoid the exponential blow up that a translation from the binary to the unary coding involves.

It is an interesting question whether the complexity of a logic is sensitive to the coding of numbers in counting quantifiers. It seems as if many logics are insensitive to the coding, i.e., both complexities coincide, even though one has to “work harder” for binary coding. For many logics with counting quantifiers, the complexity of the satisfiability problem is known for unary coding only, and is unknown for binary coding. For example, \mathbf{C}^2 (two-variable first-order logic with counting) is known to be NEXPTIME-complete if numbers in counting quantifiers are coded in unary [PST00]. While this coincides with the complexity of first-order two-variable logic without counting [GKV97], the complexity of \mathbf{C}^2 with binary coding is, to the best of our knowledge, unknown so far. Similarly, all the above mentioned complexity results for description and modal logics, with the exception of [Tob00,Tob01], assume unary coding of numbers.

In [Tob00,Tob01], Tobies studies *graded modal logic*, the extension of modal logic with graded modalities. He proves that the decidability problem for this logic is PSPACE-complete — not harder than the decidability problem for classical modal logic [Lad77], even when the numbers in the graded modalities are coded in binary. The binary coding requires additional technical machinery. Indeed, one cannot simply generate (exponentially many) individuals, but use some form of book keeping to keep track and count the individuals required by counting quantifiers.

The μ -calculus [Koz83] extends modal logic with least and greatest fixpoint operators. The extension makes μ -calculus a highly expressive logic, of great theoretical and practical interest (cf. [Eme97]). In this paper, we study the *graded μ -calculus*, i.e., μ -calculus with graded modalities. We show that the decidability problem for graded μ -calculus is EXPTIME-complete — not harder than the decidability problem for classical μ -calculus [FL79], even if the numbers are coded in binary. Our result substantiates the above hypothesis that most logics are insensitive to the coding of numbers, and is interesting for two additional reasons. Firstly, many relevant description, modal, and dynamic logics are fragments of the graded μ -calculus; see, e.g., [Sch94,DL94b,DL94a]. Hence we obtain corresponding EXPTIME upper bounds for these fragments for free. Secondly, other relevant description, modal, and dynamic logics are close relatives of the graded μ -calculus, and we could thus use the techniques developed here to prove EXPTIME upper bounds for these relatives using similar techniques.

Our techniques are based on the automata-theoretic approach [VW86,SE89,Var97]: to develop a decision procedure for a logic with the tree-model property, one first develops an appropriate notion of tree automata and studies their emptiness problem. The satisfiability problem for the logic is then reduced to the automata emptiness problem. We show here that the appropriate notion of automata is that of *graded alternating tree automata*, which generalize standard alternating tree automata by having the ability to count the number of successors that satisfy a certain condition (standard alternating automata can only count up to one). We show that graded μ -calculus has a tree model property and that, given a formula φ , we can construct a graded alternating automaton \mathcal{A}_φ that accepts exactly the (tree abstractions of) models of φ . The size of \mathcal{A}_φ is linear in $|\varphi|$, even if numbers in graded modalities are coded in binary. We then present an EXPTIME decision procedure for the emptiness of graded alternating automaton by an appropriate translation into *graded non-deterministic automata* (with an exponential blow-up in the size of the automaton), and show that emptiness of graded non-deterministic automata can be decided in polynomial time. Like other automata-based decision procedures, the techniques developed here can be re-used: once a suitable class of automata for a certain class of logics is designed (together with the corresponding emptiness test), these automata can be easily re-used for similar logics. In particular, our technique can be easily extended to handle in EXPTIME μ -calculus with *fractional modalities*, where we can express, e.g., that at least half of the accessible worlds satisfy some predicate, as well as all modalities that involve polynomially-checkable conditions on the number of accessible words that satisfy a formula.

2 The Graded μ -Calculus

The *graded μ -calculus* is a propositional modal logic augmented with least and greatest fixpoint operators [Koz83]. Specifically, we consider a μ -calculus where formulas are constructed from atomic propositions with Boolean connectives, the *graded modalities* $\langle n, \alpha \rangle$ (“exist at least n α -successors”) and $[n, \alpha]$ (“all but at most n α -successors”), as well as least (μ) and greatest (ν) fixpoint operators. We assume that μ -calculus formulas are written in positive normal form (negation is applied only to atomic propositions). Formally, given a set AP of atomic propositions, a set Var of propositional variables, and a set Prog of (atomic) programs, the set of formulas of the graded μ -calculus is the smallest set such that the following holds.

- **true**, **false**, p and $\neg p$, for $p \in \text{AP}$, are formulas,
- $x \in \text{Var}$ is a formula, and
- if φ_1 and φ_2 are formulas, α is a program, n is a non-negative integer, and x is a propositional variable, then $\varphi_1 \wedge \varphi_2$, $\varphi_1 \vee \varphi_2$, $\langle n, \alpha \rangle \varphi_1$, $[n, \alpha] \varphi_1$, $\mu y. \varphi_1(y)$, and $\nu y. \varphi_1(y)$ are formulas.

A propositional variable x occurs *free* in a formula if it is not in the scope of a fixpoint operator. Note that x may occur both bound and free in a formula. A *sentence* is formula that contains no free variables. We use λ to denote a fixpoint operator μ or ν . For a λ -formula $\lambda x. \varphi(x)$, we write $\varphi(\lambda x. \varphi(x))$ to denote the formula that is obtained by replacing each free occurrence of x in φ with $\lambda x. \varphi(x)$. We refer to formulas of the form $\langle n, \alpha \rangle \varphi_1$ and $[n, \alpha] \varphi_1$ as *atleast* and *allbut* formulas, respectively.

We say that a formula φ *counts up to* b if the maximal integer occurring in graded modalities in φ is $b - 1$. We refer to b as the *counting bound* of φ . We assume that the

integers in graded modalities are given in binary. The *length* of φ , denoted $|\varphi|$, reflects this fact. Formally, $|\varphi|$ is defined by induction on the structure of φ in a standard way, with $|\langle n, \alpha \rangle \varphi_1| = \lceil \log n \rceil + 1 + |\varphi_1|$, and similarly for $|\lceil n, \alpha \rceil \varphi_1|$.

We define the semantics of graded μ -calculus with respect to a *Kripke structure* $K = \langle \text{AP}, \text{Prog}, W, R, L \rangle$, where AP is the set of atomic propositions, Prog is the set of programs, W is a set of points, $R : \text{Prog} \rightarrow 2^{W \times W}$ assigns to each program a binary relation on W , and $L : \text{AP} \rightarrow 2^W$ assigns to each atomic proposition the set of points in which it holds. For a point w and a program α , let $\text{succ}_R(w, \alpha) = \{u : (w, u) \in R(\alpha)\}$ denote the successors of w in the program α . Given a Kripke structure $K = \langle \text{AP}, \text{Prog}, W, R, L \rangle$ and variables x_1, \dots, x_m , a *valuation* $\mathcal{V} : \{x_1, \dots, x_m\} \rightarrow 2^W$ maps each variable to a subset of W . For a valuation \mathcal{V} , a variable x , and a set of points $W' \subseteq W$, let $\mathcal{V}[x/W']$ denote the valuation that is obtained from \mathcal{V} by assigning W' to x .

A formula φ is interpreted over a Kripke structure $K = \langle \text{AP}, \text{Prog}, W, R, L \rangle$ as a mapping φ^K that associates, with each valuation \mathcal{V} , a subset $\varphi^K(\mathcal{V})$ of W . This mapping is defined inductively as follows.

$$\begin{aligned} \text{true}^K(\mathcal{V}) &= W, \text{false}^K(\mathcal{V}) = \emptyset, \\ \text{for } p \in \text{AP}, \text{ we have } p^K(\mathcal{V}) &= L(p) \text{ and } (\neg p)^K(\mathcal{V}) = W \setminus L(p), \\ (\varphi_1 \wedge \varphi_2)^K(\mathcal{V}) &= (\varphi_1)^K(\mathcal{V}) \cap (\varphi_2)^K(\mathcal{V}), \\ (\varphi_1 \vee \varphi_2)^K(\mathcal{V}) &= (\varphi_1)^K(\mathcal{V}) \cup (\varphi_2)^K(\mathcal{V}), \\ (\langle n, \alpha \rangle \varphi)^K(\mathcal{V}) &= \{u \in W : |\{v \mid (u, v) \in R(\alpha), \text{ and } v \in \varphi^K(\mathcal{V})\}| > n\} \\ (\lceil n, \alpha \rceil \varphi)^K(\mathcal{V}) &= \{u \in W : |\{v \mid (u, v) \in R(\alpha), \text{ and } v \notin \varphi^K(\mathcal{V})\}| \leq n\} \\ (\mu y. \varphi(y))^K(\mathcal{V}) &= \bigcap \{W' \subseteq W \mid \varphi^K(\mathcal{V}[x/W']) \subseteq W'\} \\ (\nu y. \varphi(y))^K(\mathcal{V}) &= \bigcup \{W' \subseteq W \mid \varphi^K(\mathcal{V}[x/W']) \supseteq W'\} \end{aligned}$$

Intuitively, a state w of K satisfies the atleast formula $\langle n, \alpha \rangle \varphi$ iff at least $n + 1$ successors of w in α satisfy φ . Dually, w satisfies the allbut formula $\lceil n, \alpha \rceil \varphi$ iff all but at most n successors of w in α satisfy φ . Note that $\neg \langle n, \alpha \rangle \varphi$ is equivalent to $\lceil n, \alpha \rceil \neg \varphi$. Indeed, $\neg \langle n, \alpha \rangle \varphi$ means that less than $n + 1$ successors of w satisfy φ , that is, at most n successors do not satisfy $\neg \varphi$. The least and greatest fixpoint operators are interpreted as in μ -calculus; thus, for example, $\mu y. p \vee \langle 1, \alpha \rangle y$ is satisfied in a point w if either w satisfies p or w has two different α -successors each of which either satisfies p or has two different α -successors etc., or equivalently, w is a root of a binary tree embedded in the transition relation of α in which each path eventually reaches a point that satisfies p . Note that the interpretation of a sentence is independent of valuations. For a sentence ψ , a Kripke structure $K = \langle \text{AP}, \text{Prog}, W, R, L \rangle$, and $u \in W$, we write $K, w \models \psi$ iff $w \in \psi^K$, and call K a *model* of ψ . A sentence that has a model is called *satisfiable*.

The modalities $\langle n, \alpha \rangle \varphi$ and $\lceil n, \alpha \rceil \varphi$ are natural generalizations of the standard existential and universal next modalities. In particular, $\langle \alpha \rangle \varphi$ and $\lceil \alpha \rceil \varphi$ of modal logic are equivalent to $\langle 0, \alpha \rangle \varphi$ and $\lceil 0, \alpha \rceil \varphi$, respectively, and the *number restrictions* $(\geq nr \varphi)$ and $(\leq nr \varphi)$ for a role r in description logics [HB91] are equivalent to $\langle n - 1, r \rangle \varphi$ and $\lceil n, r \rceil \neg \varphi$, respectively (note that $(\geq 0r \varphi)$ is equivalent to **true**).

For technical convenience, we also restrict our attention to formulas and Kripke structures in which only one program occurs. For the standard μ -calculus, the reduction to a single program uses, for each atomic program α , a new atomic proposition p_α and translates $\langle \alpha \rangle \varphi$ to $\diamond(\varphi \wedge p_\alpha)$, and $\lceil \alpha \rceil \varphi$ to $\Box(\varphi \vee \neg p_\alpha)$. In the presence of graded modalities, we can proceed analogously, translating $\langle n, \alpha \rangle \varphi$ to $\langle n \rangle(\varphi \wedge p_\alpha)$, and $\lceil n, \alpha \rceil \varphi$ to $\lceil n \rceil(\varphi \vee \neg p_\alpha)$, where the graded modalities $\langle n \rangle$ and $\lceil n \rceil$ are evaluated with respect to a single accessibility

ity relation R . We will accordingly refer to Kripke structures $K = \langle AP, W, R, L \rangle$ with $R \subseteq W \times W$, and use $\text{succ}_R(w)$ to denote the successors of w .

Note that we can also add new atomic propositions that would take care of the counting done in graded modalities. Formally, if φ counts up to b , we add propositions c_1, \dots, c_b , conjoin φ with a requirement that exactly one c_i holds in each point, that successors that are labeled by the same c_i agree on their label with respect to all formulae in the closure of φ , and replace an atleast modality $\langle n, \psi \rangle$ by a $\bigvee_{\{j_1, \dots, j_{n+1}\} \subseteq \{1, \dots, b\}} \bigwedge_{1 \leq i \leq n+1} \diamond(\psi \wedge c_{j_i})$, and dually for allbut modalities. The μ -calculus formula we get is satisfiable iff φ is satisfiable, yet the length of each disjunct that replaces a graded modality is exponential in b . Since the bounds in the graded modalities are written in binary, the length of the formula we get is doubly exponential in the length of φ .

A *tree* is a set $T \subseteq \mathbb{N}^*$ such that if $x \cdot c \in T$ where $x \in \mathbb{N}^*$ and $c \in \mathbb{N}$, then also $x \in T$. The elements of T are called *nodes*, and the empty word ε is the *root* of T . For every $x \in T$, the nodes $x \cdot c$ where $c \in \mathbb{N}$ are the *children* of x . The number of children of x is called the *degree* of x , and is denoted $\text{deg}(x)$. The *degree* of a tree is the maximum degree of a node in the tree. A node is a *leaf* if it has no children. A *path* π of a tree T is a set $\pi \subseteq T$ such that $\varepsilon \in \pi$ and for every $x \in \pi$, either x is a leaf or there exists a unique $c \in \mathbb{N}$ such that $x \cdot c \in \pi$. Given an alphabet Σ , a Σ -*labeled tree* is a pair $\langle T, V \rangle$ where T is a tree and $V : T \rightarrow \Sigma$ maps each node of T to a letter in Σ . Note that an infinite word in Σ^ω can be viewed as a Σ -labeled tree in which the degree of all nodes is 1.

In Appendix A, we show that the graded μ -calculus has the tree model property. Thus, if a formula φ is satisfiable, it is also satisfiable in a tree. Moreover, the number of atleast formulas in φ and its counting bound induce a sufficient degree for the tree. Formally, we have the following.

Theorem 1. *Consider a sentence ψ such that ψ has l atleast subsentences, each counting to at most b . If ψ is satisfiable, then ψ is satisfied in a tree whose degree is at most $l \cdot (b + 1)$.*

3 Graded Automata

Automata over infinite trees (tree automata) run over Σ -labeled trees that have no leaves [Tho90]. *Alternating automata* generalize nondeterministic tree automata and were first introduced in [MS87]. For simplicity, we refer first to automata over binary trees (i.e., when $T = \{0, 1\}^*$). Consider a nondeterministic tree automaton $\mathcal{A} = \langle \Sigma, Q, \delta, q_0, \alpha \rangle$, where Σ is the input alphabet, Q is a finite set of states, δ is a transition function, $q_0 \in Q$ is an initial state, and α specifies the acceptance condition (a condition that defines a subset of \mathcal{Q} ; we define the parity acceptance condition below). The transition function $\delta : Q \times \Sigma \rightarrow \mathcal{Q}^2$ maps an automaton state $q \in Q$ and an input letter $\sigma \in \Sigma$ to a set of pairs of states. Each pair suggests a nondeterministic choice for the automaton's next configuration. When the automaton is in a state q as it reads a node x labeled by a letter σ , it proceeds by first choosing a pair $\langle q_1, q_2 \rangle \in \delta(q, \sigma)$, and then splitting into two copies. One copy enters the state q_1 and proceeds to the node $x \cdot 0$ (the left child of x), and the other copy enters the state q_2 and proceeds to the node $x \cdot 1$ (the right child of x).

For a given set Y , let $\mathcal{B}^+(Y)$ be the set of positive Boolean formulas over Y (i.e., Boolean formulas built from elements in Y using \wedge and \vee), where we also allow the formulas **true** and **false** and, as usual, \wedge has precedence over \vee . For a set $X \subseteq Y$ and a formula $\theta \in \mathcal{B}^+(Y)$, we say that X *satisfies* θ iff assigning **true** to elements in X and assigning **false** to elements in $Y \setminus X$ makes θ true.

We can represent the transition relation δ of a nondeterministic automaton on binary trees using $\mathcal{B}^+(\{0, 1\} \times Q)$. For example, $\delta(q, \sigma) = \{\langle q_1, q_2 \rangle, \langle q_3, q_1 \rangle\}$ can be written as $\delta(q, \sigma) = (0, q_1) \wedge (1, q_2) \vee (0, q_3) \wedge (1, q_1)$, meaning that the automaton can choose between two possibilities. In the first, the copy that proceeds to direction 0 enters the state q_1 and the one that proceeds to direction 1 enters the state q_2 . In the second, the copy that proceeds to direction 0 enters the state q_3 and the one that proceeds to direction 1 enters the state q_1 .

In nondeterministic tree automata, each conjunction in δ has exactly one element associated with each direction. In alternating automata over binary trees, $\delta(q, \sigma)$ can be an arbitrary formula from $\mathcal{B}^+(\{0, 1\} \times Q)$. We can have, for instance, a transition $\delta(q, \sigma) = (0, q_1) \wedge (0, q_2) \vee (0, q_2) \wedge (1, q_2) \wedge (1, q_3)$. The above transition illustrates that several copies may go to the same direction and that the automaton is not required to send copies to all the directions.

3.1 Graded Alternating Parity Tree Automata

For $b \geq 0$, let $\langle [b] \rangle = \{\langle 0 \rangle, \langle 1 \rangle, \dots, \langle b \rangle\}$ and $[[b]] = \{[0], [1], \dots, [b]\}$, and let $D_b = \{\varepsilon\} \cup \langle [b] \rangle \cup [[b]]$. A *graded alternating automaton* is an automaton in which the transition function δ maps a state q and a letter σ to a formula in $\mathcal{B}^+(D_b \times Q)$. Intuitively, an atom (ε, q) means that the automaton sends a copy of itself in state q to the current node, an atom $(\langle n \rangle, q)$ means that the automaton sends copies in states q to $n + 1$ different children of the current node, and $([n], q)$ means that the automaton sends copies in state q to all but n children of the current node. When, for instance, the automaton is in state q , reads a node x and $\delta(q, V(x)) = (\langle 3 \rangle, q_1) \wedge (\varepsilon, q_2) \vee ([2], q_3)$, it can either send four copies in state q_1 to four different children of x and send a copy in state q_2 to x , or send copies in state q_3 to $\text{deg}(x) - 2$ children of x . So, while nondeterministic tree automata send exactly one copy to each child, graded automata can send several copies to the same child, they have ε transitions, and extend *symmetric automata* [JW95, Wil99] by specifying the number of children that need to satisfy an existential requirement or are exempt from satisfying a universal one.

Formally, a graded automaton is a tuple $\mathcal{A} = \langle \Sigma, b, Q, \delta, q_0, \alpha \rangle$, where Σ , Q , q_0 , and α are as in alternating automata, b is a *counting bound*, and $\delta : Q \times \Sigma \rightarrow \mathcal{B}^+(D_b \times Q)$ is a transition function. A *run* of \mathcal{A} on an input Σ -labeled tree $\langle T, V \rangle$ is a tree $\langle T_r, r \rangle$ (to be formally defined shortly) in which each node corresponds to a node of T and is labeled by an element of $\mathbb{N}^* \times Q$. A node in T_r , labeled by (x, q) , describes a copy of the automaton that reads the node x of T and visits the state q . Note that many nodes of T_r can correspond to the same node of T ; in contrast, in a run of a nondeterministic automaton on $\langle T, V \rangle$ there is a one-to-one correspondence between the nodes of the run and the nodes of the tree. The labels of a node and its children have to satisfy the transition function. Formally, the run $\langle T_r, r \rangle$ is an $(\mathbb{N}^* \times Q)$ -labeled \mathbb{N} -tree such that $\varepsilon \in T_r$ and $r(\varepsilon) = (\varepsilon, q_0)$, and for all $y \in T_r$ with $r(y) = (x, q)$ and $\delta(q, V(x)) = \theta$, there is a (possibly empty) set $S \subseteq D_b \times Q$, such that S satisfies θ , and for all $(c, s) \in S$, the following holds:

- If $c = \varepsilon$, then there is $j \in \mathbb{N}$ such that $y \cdot j \in T_r$ and $r(y \cdot j) = (x, s)$.
- If $c = \langle n \rangle$, then there are distinct $i_1, \dots, i_{n+1} \in \mathbb{N}$ such that for all $1 \leq j \leq n + 1$, there is $j' \in \mathbb{N}$ such that $y \cdot j' \in T_r$ and $r(y \cdot j') = (x \cdot i_j, s)$.
- If $c = [n]$, then there are distinct $i_1, \dots, i_{\text{deg}(x)-n} \in \mathbb{N}$ such that for all $1 \leq j \leq \text{deg}(x) - n$, there is $j' \in \mathbb{N}$ such that $y \cdot j' \in T_r$ and $r(y \cdot j') = (x \cdot i_j, s)$.

Note that if $\theta = \mathbf{true}$, then y need not have children. This is the reason why T_r may have leaves. Also, since there exists no set S as required for $\theta = \mathbf{false}$, we cannot have a run that takes a transition with $\theta = \mathbf{false}$.

A run $\langle T_r, r \rangle$ is *accepting* if all its infinite paths satisfy the acceptance condition. We consider here the *parity acceptance condition*, where $\alpha = \{F_1, F_2, \dots, F_k\}$ is such that $F_1 \subseteq F_2 \subseteq \dots \subseteq F_k = Q$. The number k of sets in α is called the *index* of the automaton. Given a run $\langle T_r, r \rangle$ and an infinite path $\pi \subseteq T_r$, let $\text{inf}(\pi) \subseteq Q$ be such that $q \in \text{inf}(\pi)$ if and only if there are infinitely many $y \in \pi$ for which $r(y) \in \mathbb{N}^* \times \{q\}$. That is, $\text{inf}(\pi)$ contains exactly all the states that appear infinitely often in π . A path π satisfies a parity acceptance condition $\alpha = \{F_1, F_2, \dots, F_k\}$ iff the minimal index i for which $\text{inf}(\pi) \cap F_i \neq \emptyset$ is even. An automaton accepts a tree if and only if there exists a run that accepts it. We denote by $\mathcal{L}(\mathcal{A})$ the set of all Σ -labeled trees that \mathcal{A} accepts.

Theorem 2. *Given a sentence ψ of the graded μ -calculus¹ that counts up to b , we can construct a graded alternating parity automaton \mathcal{A}_ψ such that*

1. \mathcal{A}_ψ accepts exactly all trees that satisfy ψ .
2. \mathcal{A}_ψ has $|\psi|$ states, index $|\psi|$, and counting bound b .

Proof. The construction generalizes the one for μ -calculus sentences and parity automata [KVVW00]. Given ψ , we define the graded alternating automaton $\mathcal{A}_\psi = \langle 2^{\text{AP}}, b, \text{cl}(\psi), \delta, \psi, \alpha \rangle$, as follows:

For all $\sigma \in 2^{\text{AP}}$, we define:

$$\begin{aligned} \delta(p, \sigma) &= \mathbf{true} \text{ if } p \in \sigma, & \delta(p, \sigma) &= \mathbf{false} \text{ if } p \notin \sigma, \\ \delta(\neg p, \sigma) &= \mathbf{true} \text{ if } p \notin \sigma, & \delta(\neg p, \sigma) &= \mathbf{false} \text{ if } p \in \sigma, \\ \delta(\varphi_1 \wedge \varphi_2, \sigma) &= (\varepsilon, \varphi_1) \wedge (\varepsilon, \varphi_2), & \delta(\varphi_1 \vee \varphi_2, \sigma) &= (\varepsilon, \varphi_1) \vee (\varepsilon, \varphi_2), \\ \delta(\langle n \rangle \varphi, \sigma) &= (\langle n \rangle, \varphi), & \delta([n] \varphi, \sigma) &= ([n], \varphi), \\ \delta(\mu y. f(y), \sigma) &= \delta(f(\mu y. f(y)), \sigma), & \delta(\nu y. f(y), \sigma) &= \delta(f(\nu y. f(y)), \sigma). \end{aligned}$$

For a μ -calculus sentence ψ and a subformula $\varphi = \lambda y. f(y)$ of ψ , we define the *alternation level* of φ in ψ , denoted $\text{al}_\psi(\varphi)$, as follows [BC96].

- If φ is a sentence, then $\text{al}_\psi(\varphi) = 1$.
- Otherwise, let $\xi = \lambda' x. g(x)$ be the innermost μ or ν subformula of ψ that has φ as a strict subformula. Then, if x is free in φ and $\lambda' \neq \lambda$, we have $\text{al}_\psi(\varphi) = \text{al}_\psi(\xi) + 1$. Otherwise, we have $\text{al}_\psi(\varphi) = \text{al}_\psi(\xi)$.

Intuitively, the alternation level of φ in ψ is the number of alternating fixpoint operators we have to “wrap φ with” in order to reach a sub-sentence of ψ .

Let d be the maximal alternation level of subformulas of ψ . Denote by G_i the set of all the ν -formulas in $\text{cl}(\psi)$ of alternation level i . Denote by B_i the set of all μ -formulas in $\text{cl}(\psi)$ of alternation depth less than or equal to i . Now, $\alpha = \{F_0, \dots, F_{2d}\}$, where $F_0 = \emptyset$

¹ A graded μ -calculus sentence is *guarded* if for all $y \in \text{Var}$, all the occurrences of y that are in a scope of a fixpoint modality λ are also in a scope of a graded modality that is itself in the scope of λ . Thus, a μ -calculus sentence is guarded if for all $y \in \text{Var}$, all the occurrences of y are in the scope of a graded modality. For example, the formula $\mu y. (p \vee \langle 0 \rangle y)$ is guarded and the formula $\langle 0 \rangle \mu y. (p \vee y)$ is not guarded. Given a graded μ -calculus formula, we can construct, in linear time, an equivalent guarded formula (see [BB87, KVVW00] for a proof for μ -calculus, which is easily extendible to graded μ -calculus). Accordingly, we assume that all formulas are guarded. This is essential for the correctness of the construction in the proof.

and, for every $1 \leq i \leq d$, we have $F_{2i-1} = F_{2i-2} \cup B_i$ and $F_{2i} = F_{2i-1} \cup G_i$. It is easy to see that $F_1 \subseteq F_2 \subseteq \dots \subseteq F_{2d}$.

3.2 Graded Nondeterministic Parity Tree Automata

For an integer b , a b -bound is pair in $B_b = \{(>, 0), (\leq, 0), (>, 1), (\leq, 1), \dots, (>, b), (\leq, b)\}$. For a set X , a subset $P \subseteq X$, and a tuple $t = \langle x_1, \dots, x_m \rangle \in X^m$, for some $m > 0$, the *weight* of P in t , denoted $weight(P, t)$, is the number of elements in t that are members of P . That is, $weight(P, t) = |\{i : x_i \in P\}|$. For example, $weight(\{1, 2, 3\}, \langle 1, 2, 2, 4, 2 \rangle) = 4$. We say that t satisfies a b -bound $(>, n)$ with respect to P if $weight(P, t) > n$, and t satisfies a b -bound (\leq, n) with respect to P if $weight(P, t) \leq n$.

For a set Y , we use $\mathcal{B}(Y)$ to denote the set of all Boolean formulas over atoms in Y . Each formula $\theta \in \mathcal{B}(Y)$ induces a set $sat(\theta) \subseteq 2^Y$ such that $x \in sat(\theta)$ iff x satisfies θ . For an integer $b \geq 0$, a b -counting constraint for 2^Y is a relation $C \subseteq \mathcal{B}(Y) \times B_b$. For example, if $Y = \{y_1, y_2, y_3\}$, then we can have

$$C = \{\langle y_1 \vee \neg y_2, (\leq, 3) \rangle, \langle y_3, (\leq, 2) \rangle, \langle y_1 \wedge y_3, (>, 1) \rangle\}.$$

A tuple $t = \langle x_1, \dots, x_m \rangle \in (2^Y)^m$ satisfies the b -counting constraint C if for all $\langle \theta, \xi \rangle \in C$, the tuple t satisfies ξ with respect to $sat(\theta)$. Thus, when $\theta \in \mathcal{B}(Y)$ is paired with $(>, n)$, at least $n + 1$ elements of t should satisfy θ , and when θ is paired with (\leq, n) , at most n elements in the tuple satisfy θ . For example, the tuple $t_1 = \langle \emptyset, \{y_1\}, \{y_2\}, \{y_1, y_3\} \rangle$ does not satisfy the constraint C above, as the number of sets in t_1 that satisfy $y_1 \wedge y_3$ is one. On the other hand, the tuple $t_2 = \langle \emptyset, \{y_2\}, \{y_1, y_2, y_3\}, \{y_1, y_3\} \rangle$ satisfies C . Indeed, three sets in t_2 satisfy $y_1 \vee \neg y_2$, two sets satisfy y_3 , and two sets satisfy $y_1 \wedge y_3$.

For a constraint C , the *width* of C is the number of $\theta \in \mathcal{B}(Y)$ for which there is a b -bound ξ such that $\langle \theta, \xi \rangle \in C$. Note that θ may be paired with several b -bounds. Still, it is easy to replace C by an equivalent constraint C' (that is, a tuple t satisfies C iff t satisfies C') in which θ is paired with at most one constraint of the form $(>, n)$ and at most one constraint of the form (\leq, n) . We assume that we work with such *minimized* constraints. Accordingly, for two b -counting constraints C_1 and C_2 , we denote by $C_1 \oplus C_2$ the minimization of $C_1 \cup C_2$. That is, if $\langle \theta, (>, n_1) \rangle \in C_1$ and $\langle \theta, (>, n_2) \rangle \in C_2$, then $C_1 \oplus C_2$ contains only $\langle \theta, (>, \max\{n_1, n_2\}) \rangle$, and dually for constraints of the form (\leq, n) ².

We say that a constraint C is *short* if all the formulas θ that appear in C are of size linear in $|Y|$ and the width of C is at most $|Y|$. We use $\mathcal{C}(Y, b)$ to denote the set of all short b -counting constraints for 2^Y . We assume that the integers in the constraints are coded in binary. Thus, the size of $C \in \mathcal{C}(Y, b)$ is $O(|Y|^2 \lceil \log b \rceil)$.

Lemma 1. *Given a constraint $C \in \mathcal{C}(Y, b)$ and a set $S \subseteq 2^Y$, deciding whether there is a tuple $t \in (2^Y)^*$ such that t satisfies C can be done in space $(1 + \lceil \log(b + 1) \rceil)|Y|$ or time $(2b + 2)^{|Y|}$.*

Proof. Since the width of C is at most $|Y|$, an algorithm that guesses t (element by element) and updates a counter for each θ that participate in C requires space for storing the guess

² To keep the \oplus operator efficient, we do not care for redundancies and contradictions that originate from the relation between the formulas in the constraints. For example, a minimized C may contain both $\langle \theta_1, (>, n) \rangle$ and $\langle \theta_2, (>, n) \rangle$ for θ_1 that implies θ_2 , and it may contain both $\langle \theta, (>, n) \rangle$ and $\langle \theta, (\leq, n) \rangle$

for the current element in 2^Y , and for storing the values of the counters. The algorithm terminates with a positive decision when the values of the counters are such that all the b -bounds in C are satisfied. There are at most $|Y|$ counters, each may count up to at most $b+1$. Thus, the space required is $|Y| + |Y| \lceil \log(b+1) \rceil$. In addition, since the length of each formula θ that participate in C is linear in $|Y|$, its valuation with respect to each element of the tuple can be done in space $\log |Y|$ [Lyn77].

A *graded nondeterministic parity tree automaton* (GNPT, for short) is $\mathcal{A} = \langle \Sigma, b, Q, \delta, q_0, \alpha \rangle$, where Σ and b, q_0 , and α are as in GAPT, and the other components are as follows.

- The state space Q is encoded by a finite set Y of variables; that is, $Q \subseteq 2^Y$.
- The function $\delta : Q \times \Sigma \rightarrow \mathcal{C}(Y, b)$ maps a state and a letter to a b -counting constraint for 2^Y .

Note that, like GAPT, a GNPT is symmetric, in the sense it cannot distinguish between the different children of a node.

A *run* of the graded nondeterministic automaton \mathcal{A} on a Σ -labeled tree $\langle T, V \rangle$ is a Q -labeled tree $\langle T, r \rangle$ such that $r(\varepsilon) = q_0$ and for every $x \in T$, the tuple $\langle r(x \cdot 1), r(x \cdot 2), \dots, r(x \cdot \text{deg}(x)) \rangle$ satisfies $\delta(r(x), V(x))$. The run $\langle T, r \rangle$ is accepting if all its paths satisfy the parity acceptance condition.

We consider two special cases of GNPT.

- In *forall* automata, for each $q \in Q$ and $\sigma \in \Sigma$ there is $s \in Q$ such that $\delta(q, \sigma) = \{ \langle (\neg \theta_s), (\leq, 0) \rangle \}$, where $\theta_s \in \mathcal{B}(Y)$ is such that $\text{sat}(\theta_s) = \{s\}$. Thus, a forall automaton is a notational invariant of a deterministic tree automaton, where the transition function maps q and σ to $\langle s, \dots, s \rangle$.
- In *safety* automata, there is no acceptance condition, and all runs are accepting. Note that this does not mean that safety automata accept all trees, as it may be that on some trees the automaton does not have a run.

Lemma 2. *Given a forall GNPT \mathcal{A}_1 with n_1 states and index k , and a safety GNPT \mathcal{A}_2 with n_2 states and counting bound b , we can define a GNPT \mathcal{A} such that $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}_1) \cap \mathcal{L}(\mathcal{A}_2)$. Moreover, \mathcal{A} has $n_1 n_2$ states, index k , and counting bound b .*

Proof. Let $\mathcal{A}_1 = \langle \Sigma, 0, 2^{Y_1}, \delta_1, q_1^0, \alpha_1 \rangle$ and $\mathcal{A}_2 = \langle \Sigma, b, 2^{Y_2}, \delta_2, q_2^0 \rangle$. Then, $\mathcal{A} = \langle \Sigma, b, 2^{Y_1 \cup Y_2}, \delta, q_1^0 \cup q_2^0, \alpha \rangle$, where for a state $q \in 2^{Y_1 \cup Y_2}$ and $\sigma \in \Sigma$, we define $\delta(q, \sigma) = \delta_1(q \cap Y_1, \sigma) \cup \delta_2(q \cap Y_2, \sigma)$. Finally, $\alpha = \alpha_1 \times 2^{Y_2}$ is obtained from α_1 by replacing each set F_i in α_1 by the set $F_i \times 2^{Y_2}$.

4 The Nonemptiness Problem for GAPT

In this section we solve the nonemptiness problem for GAPT and conclude that the decidability problem for graded μ -calculus can be solved in EXPTIME. We first translate GAPT to GNPT, and then solve the nonemptiness problem for GNPT.

4.1 From GAPT to GNPT

Consider a GAPT $\mathcal{A} = \langle \Sigma, b, Q, \delta, q_0, \alpha \rangle$. Let $D_b = (\{\varepsilon\} \cup \langle [b] \rangle \cup \langle \langle [b] \rangle \rangle)$. Recall that the transition function $\delta : Q \times \Sigma \rightarrow \mathcal{B}^+(D_b \times Q)$ maps a state and a letter to a formula in

$B^+(D_b \times Q)$. A restriction of δ is a partial function $\eta : Q \rightarrow 2^{D_b \times Q}$. For a letter $\sigma \in \Sigma$, we say that a restriction η is *relevant* to σ if for all $q \in Q$ for which $\delta(q, \sigma)$ is satisfiable (i.e., $\delta(q, \sigma)$ is not **false**), the set $\eta(q)$ satisfies $\delta(q, \sigma)$. If $\delta(q, \sigma)$ is not satisfiable, then $\eta(q)$ is undefined. Intuitively, by choosing the atoms that are going to be satisfied, η removes the nondeterminism in δ . Let F be the set of restrictions of δ . A *running strategy* of \mathcal{A} for a Σ -labeled tree $\langle T, V \rangle$ is an F -labeled tree $\langle T, f \rangle$. We say that $\langle T, f \rangle$ is relevant with respect to $\langle T, V \rangle$ if for all $x \in T$, the restriction $f(x)$ is relevant to $V(x)$.

Consider a restriction η relevant to σ . For $q \in Q$, we say that a finite sequence $s = s_0, s_1, \dots, s_{l+1}$ is a *step of η with q and σ* if $s_0 = q$, for all $0 \leq i < l$, we have $(\varepsilon, s_{i+1}) \in \eta(s_i)$, and $(\lambda, s_{l+1}) \in \eta(s_l)$, for $\lambda \in \langle b \rangle \cup [b]$. Thus, s is a step of η with q and σ if by following the restriction η at a node x labeled σ , a run that visits q can continue by first taking l subsequent ε -transitions and visiting s_0, \dots, s_l , and then moving to a child of x in state s_{l+1} . We refer to (λ, s_{l+1}) as the last atom taken in the step. Note that l may be 0. We define the *value* of s , denoted $val(s)$, as the minimal i such that there is $0 < j \leq l + 1$ with $s_j \in F_i$. Note that when s contains only two states, its value is induced by s_1 .

We say that a finite sequence $s = s_0, s_1, \dots, s_l$ is an *ε -lasso of η with q and σ* if $s_0 = q$, for all $0 \leq i \leq l - 1$, we have $(\varepsilon, s_{i+1}) \in \eta(s_i)$, and there is $0 \leq c \leq l$ such that $(\varepsilon, s_c) \in \eta(s_l)$. Thus, s is an ε -lasso of η with q and σ if by following the restriction η at a node x labeled σ , there is $0 \leq c \leq l$ such that a run that visits q can eventually loop forever in s_c, \dots, s_l by taking subsequent ε -transitions. The value of s with a loop starting at c , denoted $val(s, c)$, is the minimal i such that there is $c \leq j \leq l$ with $s_j \in F_i$. We say that s is *rejecting* if there is $0 \leq c \leq l$ such that $val(s)$ is odd.

A *local promise* for the automaton \mathcal{A} is a function $\rho : Q \rightarrow 2^Q$. We extend ρ to sets of states, thus for $P \subseteq Q$, we have $\rho(P) = \bigcup_{q \in P} \rho(q)$. Let G be the set of all local promises. A *promise* of \mathcal{A} for a Σ -labeled tree $\langle T, V \rangle$ is a G -labeled tree $\langle T, g \rangle$. Intuitively, in a run that proceeds according to the promise $\langle T, g \rangle$, if a node $y \cdot j$ has $s \in g(y \cdot j)(q)$ and the run visits its parent y in state q and proceeds by choosing an atom $\langle n \rangle s$ or $[n]s$, for some $0 \leq n \leq b$, then $y \cdot j$ is among the children of y that inherit s .

Consider a Σ -labeled tree $\langle T, V \rangle$, a running strategy $\langle T, f \rangle$ relevant to $\langle T, V \rangle$, and a promise $\langle T, g \rangle$. We say that a $(T \times Q)$ -labeled tree $\langle T_r, r \rangle$ is *consistent* with f and g if $\langle T_r, r \rangle$ suggests a possible run of \mathcal{A} on $\langle T, V \rangle$ such that whenever the run $\langle T_r, r \rangle$ is in state q as it reads a node $x \in T$, the restriction $f(x)(q)$ is defined, the run proceeds according to $f(x)(q)$, and it delivers requirements to each child $x \cdot j$ according to $g(x \cdot j)(q)$. Note that since the counting constraints in $f(x)(q)$ may not be satisfied, $\langle T_r, r \rangle$ may not be a legal run. Formally, $\langle T_r, r \rangle$ is consistent with f and g iff the following holds.

1. $\varepsilon \in T_r$ and $r(\varepsilon) = (\varepsilon, q_0)$.
2. Consider a node $y \in T_r$ with $r(y) = (x, q)$. Then, $f(x)(q)$ is defined, and for all $(c, s) \in f(x)(q)$, the following hold:
 - If $c = \varepsilon$, then there is $j \in \mathbb{N}$ such that $y \cdot j \in T_r$ and $r(y \cdot j) = (x, s)$.
 - If $c = \langle n \rangle$ or $c = [n]$, then for each $j \in \mathbb{N}$ with $s \in g(x \cdot j)(q)$, there is $j' \in \mathbb{N}$ such that $y \cdot j' \in T_r$ and $r(y \cdot j') = (x \cdot j, s)$.

For a node $x \in T$ and a state $q \in Q$, we say that x is *obliged to q* by f, g , and V if x is visited by q in some labeled tree $\langle T_r, r \rangle$ consistent with f and g . Note that while there may be several trees consistent with f and g , they differ only in the order of children of nodes, thus if x is visited by q in some consistent tree, it is visited by all consistent trees.

Let $\Sigma' \subseteq \Sigma \times F \times G$ be such that for all $\langle \sigma, \eta, \rho \rangle \in \Sigma'$, we have that η is relevant to σ . For an infinite sequence $\langle \sigma_0, \eta_0, \rho_0 \rangle, \langle \sigma_1, \eta_1, \rho_1 \rangle, \dots$ of triples in Σ' and a sequence (either finite or infinite) q, q_1, \dots of states, we say that q_0, q_1, \dots is a *trace* induced by $\langle \sigma_0, \eta_0, \rho_0 \rangle, \langle \sigma_1, \eta_1, \rho_1 \rangle, \dots$ if q_0 is the initial state of \mathcal{A} and there is a function $pos : \mathbb{N} \rightarrow \mathbb{N}$ such that $pos(0) = 0$ and for every $i \geq 0$, one of the following holds.

1. $\eta_{pos(i)}(q_i)$ is empty, in which case q_i is the last state in the trace,
2. there is $(\varepsilon, q_{i+1}) \in \eta_{pos(i)}(q_i)$ and $pos(i+1) = pos(i)$, or
3. $\eta_{pos(i)}(q_i)$ contains $(\langle n \rangle, q_{i+1})$ or $([n], q_{i+1})$, $q_{i+1} \in \rho_{pos(i)+1}(q_i)$, and $pos(i+1) = pos(i) + 1$.

Intuitively, q_0, q_1, \dots is a trace induced by $\langle \sigma_0, \eta_0, \rho_0 \rangle, \langle \sigma_1, \eta_1, \rho_1 \rangle, \dots$, if for every path $\pi \subseteq T$ and for every run $\langle T_r, r \rangle$ on a Σ -labeled tree in which π is labeled by $\sigma_0, \sigma_1, \dots$, if $\langle T_r, r \rangle$ is consistent with a running strategy in which π is labeled η_0, η_1, \dots and a promise in which π is labeled ρ_0, ρ_1, \dots , then $\langle T_r, r \rangle$ contains a path that visits the states q_0, q_1, \dots .

Recall that $\Sigma' \subseteq \Sigma \times F \times G$. We refer to a Σ' -labeled tree as $\langle T, (V, f, g) \rangle$, where V, f , and g are the projections of the tree on Σ, F , and G , respectively. We say that a running strategy $\langle T, f \rangle$ and a promise $\langle T, g \rangle$ are *good* for $\langle T, V \rangle$ if all the infinite traces induced by paths in $\langle T, (V, f, g) \rangle$ satisfy the acceptance condition α .

Consider a Σ -labeled tree $\langle T, V \rangle$, a running strategy $\langle T, f \rangle$, and a promise $\langle T, g \rangle$. We say that g *fulfills* f for V if the states promised to be visited by g satisfy the obligations induced by f as it runs on V . Formally, g fulfills f for V if for every node $x \in T$, and state q such that x is obliged to q by f, g , and V , the following hold:

1. For every atom $\langle n \rangle s \in f(x)(q)$, at least $n + 1$ children $x \cdot j$ of x have $s \in g(x \cdot j)(q)$.
2. For every atom $[n] s \in f(x)(q)$, at least $deg(x) - n$ children $x \cdot j$ of x have $s \in g(x \cdot j)(q)$.

Theorem 3. *A GAPT \mathcal{A} accepts $\langle T, V \rangle$ iff there exist a running strategy $\langle T, f \rangle$ and a promise $\langle T, g \rangle$ such that f is relevant for V , f and g are good for $\langle T, V \rangle$, and g fulfills f for V .*

Intuitively, if f and g as above exist, the $(T \times Q)$ -labeled trees that are consistent with f and g suggest legal accepting runs of \mathcal{A} on $\langle T, V \rangle$.

Annotating input trees with restrictions and local promises enables us to transform GAPT to GNPT, with an exponential blow up:

Theorem 4. *Consider a GAPT \mathcal{A} such that \mathcal{A} runs on Σ -labeled trees. There is a GNPT \mathcal{A}' such that \mathcal{A}' runs on Σ' -labeled trees and the following hold:*

1. \mathcal{A}' accepts a tree iff \mathcal{A} accepts its projection on Σ .
2. If \mathcal{A} has n states, index k , and counting bound b , then \mathcal{A}' has $2^{n(2+k \log nk)}$ states, index nk , and b -counting constraints.

Proof. Let $\mathcal{A} = \langle \Sigma, b, Q, \delta, q_0, \alpha \rangle$ with $\alpha = \{F_1, \dots, F_k\}$. The automaton \mathcal{A}' is the intersection of two automata \mathcal{A}'_1 and \mathcal{A}'_2 . The automaton \mathcal{A}'_1 is a forall GNPT and it accepts a tree $\langle T, (V, f, g) \rangle$ iff f and g are good for V . The automaton \mathcal{A}'_2 is a safety GNPT, and it accepts a tree $\langle T, (V, f, g) \rangle$ iff g fulfills f for V . Note that, since Σ' contains only triplets $\langle \sigma, \eta, \rho \rangle$ for which η is relevant to σ , it must be that f is relevant to V . Thus, by Theorem 3, it follows that \mathcal{A}' accepts $\langle T, (V, f, g) \rangle$ iff \mathcal{A} accepts $\langle T, V \rangle$.

In order to define \mathcal{A}_1 , we first define a nondeterministic co-parity word automaton \mathcal{U} over Σ' such that \mathcal{U} accepts a word if some trace it induces is infinite and violates the acceptance condition α . We define $\mathcal{U} = \langle \Sigma', S, M, s_0, F' \rangle$, where

- $S = (Q \times Q \times \{1, \dots, k\}) \cup \{q_{acc}\}$. Intuitively, a state $\langle q, q_{prev}, v \rangle$ indicates that the current state of the trace is q , that it was reached by following a step whose last transition is from the state q_{prev} , and the value of the step is v (note that values are calculated with respect to α). Thus, q corresponds to states q_{i+1} in traces for which $pos(i+1) = pos(i) + 1$. The number v is used for the acceptance condition. In addition, q_{prev} is used for checking the obligation of the current position, given a local promise in the input word.
- For every $\langle q, q_{prev}, v \rangle \in S$ and $\langle \sigma, \eta, \rho \rangle \in \Sigma'$, we distinguish between two cases.
 - If $q \notin \rho(q_{prev})$, then the current position is not obliged to q and $M(\langle q, q_{prev}, v \rangle, \langle \sigma, \eta, \rho \rangle) = \emptyset$.
 - Otherwise, we again distinguish between two cases.
 - * If there is a rejecting ε -lasso of η with q and σ , then $M(\langle q, q_{prev}, v \rangle, \langle \sigma, \eta, \rho \rangle) = \{q_{acc}\}$.
 - * Otherwise, $\langle q', q'_{prev}, v' \rangle \in M(\langle q, q_{prev}, v \rangle, \langle \sigma, \eta, \rho \rangle)$ iff there is a step q, \dots, q'_{prev}, q' of η with q and σ such that the value of the step is v' .

In addition, $M(q_{acc}, \langle \sigma, \eta, \rho \rangle) = \{q_{acc}\}$ for all $\langle \sigma, \eta, \rho \rangle \in \Sigma'$. Intuitively, \mathcal{U} checks whether a possible step of η with q and σ can participate in a rejecting trace. If the current position is not obliged to the current state, no step of η can participate in a trace, so \mathcal{U} gets stuck. Otherwise, if there is a rejecting ε -lasso of η with q and σ , a rejecting trace is found and \mathcal{U} moves to an accepting sink. Otherwise, \mathcal{U} guesses other possible steps of η with q and σ , and moves to a state which remembers the last two states visited in the step (possibly $q'_{prev} = q$), and the value of the step.

- $s'_0 = \langle q_0, q_0, l \rangle$, where l is such that $q_0 \in F_l$. Note that the choice of the second element is arbitrary, as the local promise at the root of the input tree is irrelevant.
- The co-parity condition is $F' = \{F'_1, F'_2, \dots, F'_k\}$, where for $l \geq 2$, we have $F'_l = Q \times Q \times \{l\}$, and $F'_1 = (Q \times Q \times \{1\}) \cup \{q_{acc}\}$. That is, acceptance is determined with respect to the values of the steps taken along the trace. Also, since F' is a co-parity condition, the accepting sink q_{acc} is in F'_1 .

In order to get \mathcal{A}'_1 , we co-determinize \mathcal{U} (note that \mathcal{U} does not have ε -transitions) and expand it to a tree automaton on Σ' . That is, we first construct a deterministic parity word automaton $\tilde{\mathcal{U}}$ that complements \mathcal{U} , and then replace a transition $\tilde{M}(s, \tau) = s'$ in $\tilde{\mathcal{U}}$ by a transition $\tilde{M}_t(s, \tau) = \{\langle -\theta_{s'}, (\leq, 0) \rangle\}$ in \mathcal{A}'_1 , where the states of $\tilde{\mathcal{U}}$ are encoded by some set Y_1 of variables and for every state s' , the formula $\theta_{s'} \in \mathcal{B}(Y_1)$ holds only in the subset of Y_1 that encodes s' . By [Saf89, Tho97], the automaton $\tilde{\mathcal{U}}$ has $(nk)^{nk}$ states and index nk , thus so does \mathcal{A}'_1 . Hence $|Y_1| = nk \log nk$.

It is left to define the safety GNPT \mathcal{A}'_2 . Let $Q_{prev} = \{q_{prev} : q \in Q\}$ be a copy of Q in which each state is tagged with *prev*. The state space of \mathcal{A}'_2 is $Q' = 2^{Q \cup Q_{prev}}$. Intuitively, each state q' of \mathcal{A}'_2 corresponds to a pair $\langle P, P_{prev} \rangle \in Q \times Q$, with $P = q' \cap Q$ and P_{prev} is obtained from $q' \cap Q_{prev}$ by removing the *prev* tags. The element P of q' is a set of “commitments” that the current node should satisfy. The element P_{prev} is used for remembering the state of \mathcal{A} that is visited in the parent node. When \mathcal{A}'_2 is in state $\langle P, P_{prev} \rangle$ and reads the letter $\langle \sigma, \eta, \rho \rangle$, it checks that all the commitments in P are covered by the

local promise $\rho(P_{prev})$ in the input, and it delivers, for each $q \in P$, the requirements on the children as specified in $\eta(q)$.

Consider a state $\langle P, P_{prev} \rangle \in Q'$ and a letter $\langle \sigma, \eta, \rho \rangle \in \Sigma'$. For every $q \in P$, let $C_{\sigma, \eta}^q$ be the b -counting restriction in $\mathcal{C}(Q, b)$ imposed by $\eta(q)$. (If $\eta(q)$ is undefined, we do not care about $C_{\sigma, \eta}^q$, since, as we see shortly, in that case \mathcal{A}'_2 simply gets stuck.) Thus,

$$C_{\sigma, \eta}^q = \{\langle s, (>, n) \rangle : \langle n \rangle s \in \eta(q)\} \cup \{\langle \neg s, (\leq, n) \rangle : [n]s \in \eta(q)\}.$$

Intuitively, $C_{\sigma, \eta}^q$ restricts the tuple of the states that visit the children of the current node, which is visited by $\langle P, P_{prev} \rangle$, so that $\eta(q)$ is satisfied by the first elements of the states. In addition, the second element of the states in the tuple should be the encoding of P tagged with $prev$. This is done by the counting constraint $\{\langle \neg \theta_P^{prev}, (\leq, 0) \rangle\}$, where $\theta_P^{prev} \in \mathcal{B}(Q_{prev})$ is such that the only set that satisfies θ_P^{prev} is the encoding of P tagged with $prev$. Finally, for every $P \in 2^Q$, let $C_{\sigma, \eta}^P = (\bigoplus_{q \in P} C_{\sigma, \eta}^q) \cup \{\langle \neg \theta_P^{prev}, (\leq, 0) \rangle\}$.

Then, $\mathcal{A}'_2 = \langle \Sigma', Q', \delta', \{q_0, q_0\} \rangle$, where for every $\langle P, P_{prev} \rangle \in Q'$ and $\langle \sigma, \eta, \rho \rangle \in \Sigma'$, we have

$$\delta'(\langle P, P_{prev} \rangle, \langle \sigma, \eta, \rho \rangle) = \begin{cases} \emptyset & \text{if } \rho(P_{prev}) \not\subseteq P \text{ or there is } q \in P \text{ for which } \eta(q) \text{ is undefined,} \\ C_{\sigma, \eta}^P & \text{otherwise.} \end{cases}$$

Note that Q' is defined with respect to the $2n$ variables $Q \cup Q_{prev}$. Also, all the formulas θ that are paired to constraints in $C_{\sigma, \eta}^P$ are either s or $\neg s$, for $s \in Q$, or $\neg \theta_P^{prev}$. Hence, the counting constraints in \mathcal{A}'_2 are in $\mathcal{C}(Q \cup Q_{prev}, b)$.

Now, by Lemma 2, we can define the intersection \mathcal{A} of \mathcal{A}'_1 and \mathcal{A}'_2 as a GNPT with $2^{n(2+k \log nk)}$ states, index nk , and b -counting constraints.

4.2 The nonemptiness problem for GNPT

In a nondeterministic parity tree automaton $\mathcal{U} = \langle \Sigma, Q, M, q_0, \alpha \rangle$, the transition function $M : Q \times \Sigma \rightarrow 2^{Q^*}$ maps a state and a letter to a set of possible tuples for the children states. Thus, a run of nondeterministic tree automaton on a tree $\langle T, V \rangle$ is a Q -labeled tree $\langle T, r \rangle$ in which $r(\varepsilon) = q_0$ and for all $x \in T$, the tuple $\langle r(x \cdot 1), r(x \cdot 2), \dots, r(x \cdot \text{deg}(x)) \rangle \in M(r(x), V(x))$. The nonemptiness test for parity tree automata then uses the local test $is_mother : 2^Q \times Q \rightarrow \{\mathbf{true}, \mathbf{false}\}$ that given a set $S \subseteq Q$ and a state q , returns **true** iff there is a tuple $t \in S^*$ and $\sigma \in \Sigma$ such that $t \in M(q, \sigma)$. It is easy to see how the is_mother test is used in a bottom-up nonemptiness algorithm for automata on finite trees, where in order to find the set S of states from which the automaton accepts some tree, one starts with the set S_0 of accepting states then define S_{i+1} as the set of states q such that either q is in S_i or $is_mother(t, q) = \mathbf{true}$. In parity automata, the algorithm is more complicated, as one has to also keep track of the acceptance condition, but the same local test is used. Several nonemptiness algorithms for nondeterministic parity tree automata are known. In particular, the algorithms in [EJS93, KV98] use $O(n^k)$ calls to is_mother , where n is the size of Q and k is the index of the automaton.

Recall that in GNPT, a run $\langle T, r \rangle$ should satisfy $r(\varepsilon) = q_0$ and for all $x \in T$, the tuple $\langle r(x \cdot 1), r(x \cdot 2), \dots, r(x \cdot \text{deg}(x)) \rangle$ satisfies $\delta(r(x), V(x))$, which is a b -counting constraint. Thus, the nonemptiness test is similar, only that the local test $is_mother : 2^Q \times Q \rightarrow \{\mathbf{true}, \mathbf{false}\}$ now returns **true** for a set $S \subseteq Q$ and a state q , iff there is $t \in S^*$ and $\sigma \in \Sigma$ such that t satisfies $\delta(q, \sigma)$. As with nondeterministic automata, the nonemptiness algorithm can do $O(n^k)$ calls to is_mother . Unlike the case for nondeterministic automata,

however, here there is no simple transition function to consult when we perform the local test. In addition, we should take into an account the fact that the GNPT whose emptiness we check have larger alphabets than the GAPT we have started with.

Consider a GAPT $\mathcal{A} = \langle \Sigma, b, Q, \delta, q_0, \alpha \rangle$ with n states, index k , and counting bound b . Let us analyse carefully the complexity of the local *is_mother* test in the GNPT \mathcal{A}' we constructed from \mathcal{A} in Theorem 4. First, \mathcal{A}' has counting constraints in $\mathcal{C}(Y', b)$, for Y' of size $n(2 + k \log nk)$. Hence, by Lemma 1, given S , the check whether there is a tuple $t \in S^*$ such that t satisfies $\delta(q, \sigma')$, for a particular $\sigma' \in \Sigma'$, can be done in time $O(2b + 2)^{n(2+k \log nk)}$. Now, $\Sigma' \subseteq \Sigma \times F \times G$, where F is the set of restrictions for δ and G is the set of all local promises. Let $|\Sigma| = l$. Recall that a restriction relevant to a letter $\sigma \in \Sigma$ maps a state $q \in Q$ to a subset of $D_b \times Q$ that satisfies $\delta(q, \sigma)$. We can restrict our attention to restrictions in which each state is paired with at most one element of $\langle\langle b \rangle\rangle$, one element of $[\langle b \rangle]$, and ε . Thus, $|F|$ is bounded by $(2b + 4)^{n^2}$ and $|G|$ is bounded by 2^{n^2} . It follows that $|\Sigma'| \leq l(2b + 4)^{n^2} 2^{n^2}$, thus *is_mother* can be checked in time $l(b + 2)^{O(n(n+2+k \log nk))}$. Since, as in [EJS93,KV98], the nonemptiness problem can be solved by $O(n^k)$ applications of *is_mother*, we have the following.

Theorem 5. *The nonemptiness problem for \mathcal{A}' can be solved in time $n^k l(b+2)^{O(n(n+2+k \log nk))}$.*

For a graded μ -calculus formula ψ , we get, by Theorem 2, a GAPT \mathcal{A} with n and k bounded by $|\psi|$, and the same counting bound b as ψ . While b and l may be exponential in $|\psi|$, only n and k appear in the exponents in the expression in Theorem 5. This implies the upper bound in the theorem below. The lower bound is due to the fact that the μ -calculus is known to be EXPTIME-hard [FL79].

Corollary 1. *The satisfiability problem for graded μ -calculus is EXPTIME-complete even if the numbers in the graded modalities are coded in binary.*

Note that the space and time bounds in Lemma 1 stay valid for counting constraints that involve richer bounds than $(>, n)$ and (\leq, n) . For example, we can handle bounds of the form $(>, \frac{1}{2})$ or $(\leq, \frac{1}{2})$, bounding the fraction of elements in the tuple that satisfy a predicate (of course, this is applicable only to structures where all points have only finitely many successors). In general, Lemma 1 can handle arbitrary polynomial predicates $\alpha \subseteq \mathbb{N}^2$, where a tuple $t \in (2^Y)^m$ satisfies such a constraint $\langle \theta, \alpha \rangle$ if $\alpha(\text{weight}(\text{sat}(\theta), t), m)$ holds. By defining the corresponding types of alternating automata, we can thus handle μ -calculus formulas with richer types of modalities.

References

- [BB87] B. Banieqbal and H. Barringer. Temporal logic with fixed points. In B. Banieqbal, H. Barringer, and A. Pnueli, editors, *Temporal Logic in Specification*, volume 398 of *Lecture Notes in Computer Science*, pages 62–74. Springer-Verlag, 1987.
- [BBH96] F. Baader, M. Buchheit, and B. Hollunder. Cardinality restrictions on concepts. *Artificial Intelligence*, 88(1–2):195–213, 1996.
- [BC96] G. Bhat and R. Cleaveland. Efficient local model-checking for fragments of the modal μ -calculus. In *Proc. Workshop on Tools and Algorithms for the Construction and Analysis of Systems*, volume 1055 of *Lecture Notes in Computer Science*. Springer-Verlag, 1996.
- [BFH⁺94] F. Baader, E. Franconi, B. Hollunder, B. Nebel, and H.J. Profitlich. An empirical analysis of optimization techniques for terminological representation systems, or: Making KRIS get a move on. *Applied Artificial Intelligence*, 4:109–132, 1994.

- [BS99] F. Baader and U. Sattler. Expressive number restrictions in description logics. *Journal of Logic and Computation*, 9(3):319–350, 1999.
- [DL94a] G. De Giacomo and M. Lenzerini. Boosting the correspondence between description logics and propositional dynamic logics (extended abstract). In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*, 1994.
- [DL94b] G. De Giacomo and M. Lenzerini. Concept language with number restrictions and fixpoints, and its relationship with μ -calculus. In *Proceedings of the Eleventh European Conference on Artificial Intelligence (ECAI-94)*, 1994.
- [DLNdN91] F. Donini, M. Lenzerini, D. Nardi, and W. Nutt. The complexity of concept languages. In *Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning (KR-91)*, Boston, MA, USA, 1991.
- [EJ91] E.A. Emerson and C. Jutla. Tree automata, μ -calculus and determinacy. In *Proc. 32nd IEEE Symp. on Foundations of Computer Science*, pages 368–377, San Juan, October 1991.
- [EJS93] E.A. Emerson, C. Jutla, and A.P. Sistla. On model-checking for fragments of μ -calculus. In *Computer Aided Verification, Proc. 5th International Conference*, volume 697, pages 385–396, Elounda, Crete, June 1993. Lecture Notes in Computer Science, Springer-Verlag.
- [Eme97] E.A. Emerson. Model checking and the μ -calculus. In N. Immerman and Ph.G. Kolaitis, editors, *Descriptive Complexity and Finite Models*, pages 185–214. American Mathematical Society, 1997.
- [Fin72] K. Fine. In so many possible worlds. *Notre Dame Journal of Formal Logics*, 13:516–520, 1972.
- [FL79] M.J. Fischer and R.E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and Systems Sciences*, 18:194–211, 1979.
- [GKV97] E. Grädel, Ph. G. Kolaitis, and M. Y. Vardi. The decision problem for 2-variable first-order logic. *Bulletin of Symbolic Logic*, 3:53–69, 1997.
- [GMV99] H. Ganzinger, C. Meyer, and M. Veanes. The two-variable guarded fragment with transitive relations. In *Proceedings of the 14th Annual IEEE Symposium on Logic in Computer Science (LICS-99)*, pages 24–34. IEEE Computer Society Press, 1999.
- [GOR97] E. Grädel, M. Otto, and E. Rosen. Two-variable logic with counting is decidable. In *Proceedings of the Twelfth Annual IEEE Symposium on Logic in Computer Science (LICS-97)*, 1997. Available via <http://speedy.informatik.rwth-aachen.de/www/papers.html>.
- [Grä99] E. Grädel. On the restraining power of guards. *Journal of Symbolic Logic*, 64:1719–1742, 1999.
- [HB91] B. Hollunder and F. Baader. Qualifying number restrictions in concept languages. In *Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning (KR-91)*, pages 335–346, Boston, MA, USA, 1991.
- [HM01] V. Haarslev and R. Möller. RACER system description. In *Proceedings of the International Joint Conference on Automated Reasoning (IJCAR-01)*, volume 2083 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 2001.
- [Hor98] I. Horrocks. Using an Expressive Description Logic: FaCT or Fiction? In *Proceedings of the Sixth International Conference on the Principles of Knowledge Representation and Reasoning (KR-98)*, 1998.
- [JW95] D. Janin and I. Walukiewicz. Automata for the modal μ -calculus and related results. In *Proc. 20th International Symp. on Mathematical Foundations of Computer Science*, Lecture Notes in Computer Science, pages 552–562. Springer-Verlag, 1995.
- [Koz83] D. Kozen. Results on the propositional μ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
- [KV98] O. Kupferman and M.Y. Vardi. Weak alternating automata and tree automata emptiness. In *Proc. 30th ACM Symp. on Theory of Computing*, pages 224–233, Dallas, 1998.

- [KVV00] O. Kupferman, M.Y. Vardi, and P. Wolper. An automata-theoretic approach to branching-time model checking. *Journal of the ACM*, 47(2):312–360, March 2000.
- [Lad77] R. E. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM Journal of Control and Optimization*, 6(3):467–480, 1977.
- [Lyn77] N. Lynch. Log space recognition and translation of parenthesis languages. *Journal ACM*, 24:583–590, 1977.
- [MS87] D.E. Muller and P.E. Schupp. Alternating automata on infinite trees. *Theoretical Computer Science*, 54:267–276, 1987.
- [PSMB⁺91] P. Patel-Schneider, D. McGuinness, R. Brachman, L. Resnick, and A. Borgida. The CLASSIC knowledge representation system: Guiding principles and implementation rationale. *SIGART Bulletin*, 2(3):108–113, 1991.
- [PST00] L. Pacholski, W. Szwast, and L. Tendera. Complexity results for first-order two-variable logic with counting. *SIAM Journal of Computing*, 29(4):1083–1117, 2000.
- [Saf89] S. Safra. *Complexity of automata on infinite objects*. PhD thesis, Weizmann Institute of Science, Rehovot, Israel, 1989.
- [Sch94] K. Schild. Terminological cycles and the propositional μ -calculus. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Proceedings of the Fourth International Conference on the Principles of Knowledge Representation and Reasoning (KR-94)*, pages 509–520, Bonn, 1994. Morgan Kaufmann, Los Altos.
- [SE89] R.S. Streett and E.A. Emerson. An automata theoretic decision procedure for the propositional μ -calculus. *Information and Computation*, 81(3):249–264, 1989.
- [Tho90] W. Thomas. Automata on infinite objects. *Handbook of Theoretical Computer Science*, pages 165–191, 1990.
- [Tho97] W. Thomas. Languages, automata, and logic. *Handbook of Formal Language Theory*, III:389–455, 1997.
- [Tob00] S. Tobies. The complexity of reasoning with cardinality restrictions and nominals in expressive description logics. *Journal of Artificial Intelligence Research*, 12:199–217, May 2000.
- [Tob01] S. Tobies. PSPACE reasoning for graded modal logics. *Journal of Logic and Computation*, 11(1):85–106, 2001.
- [Var97] M.Y. Vardi. What makes modal logic so robustly decidable? In N. Immerman and Ph.G. Kolaitis, editors, *Descriptive Complexity and Finite Models*, pages 149–183. American Mathematical Society, 1997.
- [vdHD95] W. van der Hoek and M. De Rijke. Counting objects. *Journal of Logic and Computation*, 5(3):325–345, 1995.
- [VW86] M.Y. Vardi and P. Wolper. Automata-theoretic techniques for modal logics of programs. *Journal of Computer and System Science*, 32(2):182–221, April 1986.
- [Wil99] T. Wilke. CTL⁺ is exponentially more succinct than CTL. In C. Pandu Ragan, V. Raman, and R. Ramanujam, editors, *Proc. 19th conference on Foundations of Software Technology and Theoretical Computer Science*, volume 1738 of *Lecture Notes in Computer Science*, pages 110–121. Springer-Verlag, 1999.

A Tree Model Property

In this section we show that the graded μ -calculus has the tree model property. Thus, if a formula φ is satisfiable, it is also satisfiable in a tree. Moreover, the number of at least formulas in φ and its counting bound induce a sufficient branching degree for the tree.

We first need some definitions. The *closure* of a sentence ψ , denoted $\text{cl}(\psi)$, is the smallest set of sentences that satisfies the following:

- $\psi \in \text{cl}(\psi)$,

- if $\varphi_1 \wedge \varphi_2 \in \text{cl}(\psi)$ or $\varphi_1 \vee \varphi_2 \in \text{cl}(\psi)$, then $\{\varphi_1, \varphi_2\} \subseteq \text{cl}(\psi)$,
- if $\langle n \rangle \varphi \in \text{cl}(\psi)$ or $[n] \varphi \in \text{cl}(\psi)$, then $\varphi \in \text{cl}(\psi)$, and
- if $\lambda x. \varphi(x) \in \text{cl}(\psi)$, then $\varphi(\lambda x. \varphi(x)) \in \text{cl}(\psi)$.

An *atom* of ψ is a maximal consistent set A of formulas in $\text{cl}(\psi)$. Formally, A satisfies the following properties:

- if $p \in \text{AP}$ occurs in ψ , then $p \in A$ iff $\neg p \notin A$,
- if $\varphi_1 \wedge \varphi_2 \in \text{cl}(\psi)$, then $\varphi_1 \wedge \varphi_2 \in A$ iff $\{\varphi_1, \varphi_2\} \subseteq A$,
- if $\varphi_1 \vee \varphi_2 \in \text{cl}(\psi)$, then $\varphi_1 \vee \varphi_2 \in A$ iff $\{\varphi_1, \varphi_2\} \cap A \neq \emptyset$, and
- if $\lambda x. \varphi(x) \in \text{cl}(\psi)$, then $\lambda x. \varphi(x) \in A$ iff $\varphi(\lambda x. \varphi(x)) \in A$.

The set of atoms of ψ is denoted $\text{at}(\psi)$.

A *pre-model* $\langle K, \pi \rangle$ for a sentence ψ consists of a Kripke structure $K = \langle \text{AP}, W, R, L \rangle$ and a mapping $\pi : W \rightarrow \text{at}(\psi)$ that satisfies the following properties:

- there is $u_0 \in W$ with $\psi \in \pi(u_0)$,
- for $p \in \text{AP}$, if $p \in \pi(u)$ then $u \in L(p)$, and if $\neg p \in \pi(u)$ then $u \notin L(p)$,
- if $\langle n \rangle \varphi \in \pi(u)$, then there is $V \subseteq \text{succ}_R(u)$ such that $|V| > n$ and $\varphi \in \pi(v)$ for all $v \in V$, and
- if $[n] \varphi \in \pi(u)$, then there is $V \subseteq \text{succ}_R(u)$ such that $|V| \leq n$ and $\varphi \in \pi(v)$ for all $v \in \text{succ}_R(u) \setminus V$.

A *choice function* for a pre-model $\langle K, \pi \rangle$ of ψ is a partial function $\text{ch} : W \times \text{cl}(\psi) \rightarrow \text{cl}(\psi) \cup 2^W$ such that for all $u \in W$, the following hold.

- If $\varphi_1 \vee \varphi_2 \in \pi(u)$, then $\text{ch}(u, \varphi_1 \vee \varphi_2) \in \{\varphi_1, \varphi_2\} \cap \pi(u)$.
- If $\langle n \rangle \varphi \in \pi(u)$, then $\text{ch}(u, \langle n \rangle \varphi) = V \subseteq \text{succ}_R(u)$ such that $|V| > n$ and $\varphi \in \pi(v)$ for all $v \in V$.
- if $[n] \varphi \in \pi(u)$, then $\text{ch}(u, [n] \varphi) = V \subseteq \text{succ}_R(u)$ such that $|V| \leq n$ and $\varphi \in \pi(v)$ for all $v \in \text{succ}_R(u) \setminus V$.

An *adorned pre-model* $\langle K, \pi, \text{ch} \rangle$ consists of a pre-model $\langle K, \pi \rangle$ together with a choice function ch . For an adorned pre-model $\langle \langle \text{AP}, W, R, L \rangle, \pi, \text{ch} \rangle$ of ψ , the *derivation relation* $\rightsquigarrow \subseteq (W \times \text{cl}(\psi)) \times (W \times \text{cl}(\psi))$ is defined as follows:

- if $\varphi_1 \vee \varphi_2 \in \pi(u)$, then $(u, \varphi_1 \vee \varphi_2) \rightsquigarrow (u, \text{ch}(\varphi_1 \vee \varphi_2))$
- if $\varphi_1 \wedge \varphi_2 \in \pi(u)$, then $(u, \varphi_1 \wedge \varphi_2) \rightsquigarrow (u, \varphi_1)$ and $(u, \varphi_1 \wedge \varphi_2) \rightsquigarrow (u, \varphi_2)$.
- if $\langle n \rangle \varphi \in \pi(u)$, then $(u, \langle n \rangle \varphi) \rightsquigarrow (v, \varphi)$ for each $v \in \text{ch}(u, \langle n \rangle \varphi)$,
- if $[n] \varphi \in \pi(u)$, then $(u, [n] \varphi) \rightsquigarrow (v, \varphi)$ for each with $v \in \text{succ}_R(u) \setminus \text{ch}([n] \varphi)$,
- if $\lambda x. \varphi(x) \in \pi(u)$, then $(u, \lambda x. \varphi(x)) \rightsquigarrow (u, \varphi(\lambda x. \varphi(x)))$

Intuitively, the derivation relation \rightsquigarrow is such that in order to derive the satisfaction of φ in u as suggested by ch , one needs to satisfy φ' in u' , for all φ' and u' with $(u, \varphi) \rightsquigarrow (u', \varphi')$. For a least-fixed point sentence, one should avoid a cyclic dependency in the derivation chain. Formally, we say that a least-fixed point sentence $\mu x. \varphi(x)$ is *regenerated* from point u to point v in an adorned pre-model $\langle K, \pi, \text{ch} \rangle$ if there is a sequence $(\rho_1, u_1), \dots, (\rho_k, u_k) \in (\text{cl}(\psi) \times W)^*$ with $k \geq 2$ such that $\rho_1 = \rho_k = \mu x. \varphi(x)$, $u = u_1$, $v = u_k$, the formula $\mu x. \varphi(x)$ is a sub-sentence of each ρ_i in the sequence, and for all $1 \leq i < k$, we have $(\rho_i, u_i) \rightsquigarrow (\rho_{i+1}, u_{i+1})$. We then say that $\langle K, \pi, \text{ch} \rangle$ is *well-founded* if there is no least fixed point sentence $\mu x. \varphi(x) \in \text{cl}(\psi)$ and an infinite sequence \emptyset, u_1, \dots such that, for each $i \geq 0$, $\mu x. \varphi(x)$ is regenerated from u_i to u_{i+1} .

Lemma 3. *A sentence ψ has a model K iff ψ has a well-founded adorned pre-model $\langle K, \pi, \text{ch} \rangle$.*

Proof (sketch): Assume first that ψ has a well-founded adorned pre-model $\langle K, \pi, \text{ch} \rangle$. The proof that K is then a model of ψ is very similar to the analogous claim for μ -calculus formulas [SE89]. For the other direction, we should show that given a model K of ψ , we can tag each state of K with an atom and a choice function and get a well-founded adorned pre-model. This direction is harder, as it is not clear that the same choice can be made in different visits to the same state of K . In [EJ91] (see also [Tho97]), a *memoryless winning strategy* is defined for parity games. The strategy being memoryless means that the same choice is made by the strategy in different visits to the same position of the game. Since μ -calculus formulas correspond to parity games, this implies the correctness of the lemma for a μ -calculus formula ψ . As we now show, graded μ -calculus formulas also correspond to parity games. A position of the game is a pair $\langle U, \varphi \rangle$ where $U \subseteq W$ is a set of points and φ is a formula. Essentially, the game is played between player 1, who tries to show that φ is satisfied in all the points in U , and player 2, who tries to show the opposite. Assume that the current position of the game is $\langle U, \varphi \rangle$. When the set U is not a singleton, player 2 chooses $u \in U$ and the game moves to position $\langle u, \varphi \rangle$. When $U = \{u\}$ is a singleton, the game proceeds as follows. If $\varphi = \varphi_1 \wedge \varphi_2$, then player 2 chooses φ_i and the game moves to position $\langle u, \varphi_i \rangle$. On the other hand, when $\varphi = \varphi_1 \vee \varphi_2$, player 1 chooses the φ_i to move to. Formulas of the form $\langle n \rangle \varphi'$ and $[n] \varphi'$ also involve a disjunction, and player 1 proceeds by picking the set of successors that should satisfy φ' , thus the new position is $\langle U', \varphi' \rangle$, where the size of $U' \subseteq \text{succ}_R(u)$ is $n + 1$ or $\text{deg}(u) - n$, respectively. Recall that then, player 2 chooses the point in U' to continue with. Finally, when φ is of the form $\lambda x. \varphi'(x)$, the new position is $\langle u, \varphi'(\lambda x. \varphi'(x)) \rangle$. Note that in parity games that correspond to μ -calculus formulas (equivalently, in the game described above with $n = 0$), the set U is always a singleton. The parity winning condition is similar to the one for μ -calculus. Now, as in [SE89], the a memoryless strategy in the game induces the required choice function and implies the existence of a well-founded adorned pre-model. \square

Theorem 6. *Consider a sentence ψ such that ψ has l atleast subsentences, each counting to at most b . If ψ is satisfiable, then ψ is satisfied in a tree with branching degree bounded by $l \cdot (b + 1)$.*

Proof. Assume that ψ be satisfiable. By Lemma 3, ψ has a well-founded adorned pre-model $\langle K, \pi, \text{ch} \rangle$. Let $K = \langle \text{AP}, W, R, L \rangle$ and let $w_0 \in W$ be such that $w_0 \in \psi^K$. Also, let $\langle n_1 \rangle \varphi_1, \dots, \langle n_l \rangle \varphi_l$ be all the atleast formulas in $\text{cl}(\psi)$. Recall that $b = \max\{n_i \mid 1 \leq i \leq l\}$. We embody $\langle K, \pi, \text{ch} \rangle$ in a tree with branching degree bounded by $l(b + 1)$. For that, we define a partial mapping $\tau : \{1, \dots, l(b + 1)\}^* \rightarrow W$ inductively, together with an adorned pre-model $\langle K', \pi', \text{ch}' \rangle$, where $K' = \langle \text{AP}, W', R', L' \rangle$ is such that W' is the subset of $\{1, \dots, l(b + 1)\}^*$ for which τ is defined. The functions L' and π' are directly induced by L and π . Formally, for a point $x \in W'$, we define $\pi'(x) = \pi(\tau(x))$, and for an atomic proposition $p \in \text{AP}$, we define $x \in L'(p)$ iff $\tau(x) \in L(p)$. The choice function ch' for disjunctions is also directly induced by ch . Formally, $\text{ch}'(x, \varphi_1 \vee \varphi_2) = \text{ch}(\tau(x), \varphi_1 \vee \varphi_2)$.

It is left to define R' and the choice function ch' for atleast and allbut subsentences. This is done inductively, together with the definition of τ . The induction proceeds on the length of the nodes $x \in \{1, \dots, l(b + 1)\}^*$. For the induction base, we set $\tau(\varepsilon) = w_0$. Let

$i \geq 0$ be such that $\tau(x)$ is already defined for nodes in $x \in \{1, \dots, l(b+1)\}^i$. Then, for all $x \in \{1, \dots, l(b+1)\}^i$ for which $\tau(x)$ is defined and for all $1 \leq j \leq l$, we do the following.

- If $\langle n_j \rangle \varphi_j \in \pi'(x)$, then, since $\langle K, \pi, \text{ch} \rangle$ is a pre-model and $\pi'(x) = \pi(\tau(x))$, we have that $\text{ch}(\tau(x), \langle n_j \rangle \varphi_j) = \{v_1, \dots, v_{n_j+1}\} \subseteq \text{succ}_R(\tau(x))$ is such that for all $1 \leq r \leq n_j + 1$, we have $\varphi_j \in \pi(v_r)$. For all $1 \leq r \leq n_j + 1$, we do the following.
 - (i) If there is some j' such that $\tau(x \cdot j') = v_r$, we add $x \cdot j'$ to the set $\text{ch}'(x, \langle n_j \rangle \varphi_j)$.
 - (ii) Otherwise, we set $\tau(x \cdot (j-1)l + r) = v_r$, add $\langle x, x((j-1)l + r) \rangle$ to R' , and add $x \cdot ((j-1)l + r)$ to $\text{ch}'(x, \langle n_j \rangle \varphi_j)$.
- for each x with $\tau(x) \in W$ and $[n]\varphi \in \pi'(x)$, we set $\text{ch}'(x, [n]\varphi) = \{x \cdot j \mid \tau(x \cdot j) \in \text{ch}(\tau(x), [n]\varphi)\}$.

Since $\langle K, \pi, \text{ch} \rangle$ is an adorned pre-model, so is $\langle K', \pi', \text{ch}' \rangle$. Moreover, if a sentence $\mu x.\varphi(x)$ is regenerated from x to y in $\langle K', \pi', \text{ch}' \rangle$, then $\mu x.\varphi(x)$ is also regenerated from $\tau(x)$ to $\tau(y)$ in $\langle K, \pi, \text{ch} \rangle$. Thus, since $\langle K, \pi, \text{ch} \rangle$ is well-founded, so is $\langle K', \pi', \text{ch}' \rangle$. Finally, since all the transitions in R' are from a node in the tree $\{1, \dots, l \cdot (b+1)\}^*$ to its children in the tree, then K' is the tree we seek, and ψ is satisfied in its root ε .