On closure ordinals for the modal μ-calculus

Bahareh Afshari[1] and Graham E. Leigh[2]

1   **Department of Computer Science, University of Oxford, Parks Road, Oxford OX1 3QD, UK**
    `bahareh.afshari@cs.ox.ac.uk`
2   **Faculty of Philosophy, University of Oxford, Woodstock Road, Oxford OX2 6GG, UK**
    `graham.leigh@philosophy.ox.ac.uk`

─── **Abstract** ───

The closure ordinal of a formula of modal μ-calculus $\mu X\varphi$ is the least ordinal κ, if it exists, such that the denotation of the formula and the κ-th iteration of the monotone operator induced by φ coincide across all transition systems (finite and infinite). It is known that for every $\alpha < \omega^2$ there is a formula φ of modal logic such that $\mu X\varphi$ has closure ordinal α [3]. We prove that the closure ordinals arising from the alternation-free fragment of modal μ-calculus (the syntactic class capturing $\Sigma_2 \cap \Pi_2$) are bounded by ω^2. In this logic satisfaction can be characterised in terms of the existence of tableaux, trees generated by systematically breaking down formulæ into their constituents according to the semantics of the calculus. To obtain optimal upper bounds we utilise the connection between closure ordinals of formulæ and embedded order-types of the corresponding tableaux.

1   Introduction

Modal μ-calculus is often referred to as the "mother of all temporal logics". Indeed the majority of temporal logics, including LTL (Linear Time Logic), CTL (Computational Tree Logic) and their various extensions, can be easily interpreted and analysed in μ-calculus making the study of this logic of high interest in the research community. The defining feature of the modal μ-calculus is the expression of fixpoints. In this calculus the syntax of modal logic is extended by least and greatest fixpoint quantifiers (μ and ν) that bind propositional variables. The formulæ $\mu X\varphi$ and $\nu X\varphi$ are interpreted respectively as the least and greatest fixpoints of the monotone operator induced by φ. In analogy to the hierarchies defined in second order logic, one can alternate the fixpoint quantifiers to define a hierarchy of formulæ. Although we have a relatively good understanding of least and greatest fixpoints, when nested their meaning and behaviour is easily lost. As a result many fundamental properties of this calculus have remained unanswered even after decades of attention from logicians and computer scientists.

An interesting open problem for μ-calculus is that of closure ordinals, the number of iterations required for a fixpoint to close across all structures. Given an arbitrary formula, its closure ordinal may not exist, such as in the case of $\mu X\square X$. On the other hand mere syntactic analysis suggests that the fixpoint iterations in this context cannot exhaust the power of ordinals beyond certain levels. Hence one may ask the following question.

For which ordinals α is there a formula of modal μ-calculus with closure ordinal α?

In the case of finite ordinals the formulæ $\mu X.\,(\Diamond X \wedge \Box^n \bot) \vee \Box \bot$, which express that all paths in a model of the formula have length at most n, are guaranteed to close across all structures after n iterations. By expressing the existence of arbitrarily long finite paths, through the formula $\mu X.\,\Diamond X \vee \Box \bot$ for example, transfinite closure ordinals are obtained. In fact it is known that for every $\alpha < \omega^2$ there is a formula φ of modal logic such that $\mu X \varphi$ has closure ordinal α [3].

In this paper we establish optimal upper bounds on closure ordinals, showing that no formula of the alternation-free fragment can have a closure ordinal equal or greater than ω^2, even if iterations of all quantifiers occurring in the formula are taken into account. We begin with a syntactic analysis on a fragment of the Σ_1-formulæ in section 2. This study, despite applying only to operators induced by particular formulæ of modal logic, provides the motivation for the general solution. The main result of the paper is given in section 3 and consists of a semantic analysis of the problem by means of tableaux constructions. We present a strong characterisation of closure ordinals in terms of order-types of tableaux for formulæ without genuine dependencies between their alternating fixpoint quantifiers. This correspondence will prove sufficient to bound closure ordinals of these formulæ by their logical complexity.

1.1   Syntax and semantics of modal μ-formulæ

Let VAR be an infinite set of propositional variables and PROP an infinite set of propositional constants. The set of μ-formulæ is defined inductively as follows.

$$\varphi := p \mid \bar{p} \mid X \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \Box \varphi \mid \Diamond \varphi \mid \mu X \, \varphi \mid \nu X \, \varphi$$

where $p \in \text{PROP}$ and $X \in \text{VAR}$. Also define $\bot := p \wedge \bar{p}$ and $\top := p \vee \bar{p}$ for some propositional constant p. A variable X in φ is called a μ-variable (respectively, ν-variable) if the quantifier μX (resp. νX) occurs in φ. We assume that all quantifiers occur uniquely. This can be achieved through implicit α-conversion.

A *transition system* is a tuple $T = (S, \rightarrow, \lambda)$ where $(S, \rightarrow)$ is a directed graph and $\lambda \colon S \to \mathcal{P}(\text{PROP})$ is an assignment of propositional constants to states. Given a transition system $T = (S, \rightarrow, \lambda)$ and a valuation $\mathcal{V} \colon \text{VAR} \to \mathcal{P}(S)$ of free variables, the set of states satisfying a formula φ, denoted by $\|\varphi\|_{\mathcal{V}}^T$, is defined inductively as follows.

$\|p\|_{\mathcal{V}}^T = \{x \in S : p \in \lambda(x)\}$
$\|\bar{p}\|_{\mathcal{V}}^T = \{x \in S : p \notin \lambda(x)\}$
$\|X\|_{\mathcal{V}}^T = \mathcal{V}(X)$
$\|\varphi \wedge \psi\|_{\mathcal{V}}^T = \|\varphi\|_{\mathcal{V}}^T \cap \|\psi\|_{\mathcal{V}}^T$
$\|\varphi \vee \psi\|_{\mathcal{V}}^T = \|\varphi\|_{\mathcal{V}}^T \cup \|\psi\|_{\mathcal{V}}^T$
$\|\Box\varphi\|_{\mathcal{V}}^T = \{x \in S : \forall y (x \rightarrow y \Rightarrow y \in \|\varphi\|_{\mathcal{V}}^T)\}$
$\|\Diamond\varphi\|_{\mathcal{V}}^T = \{x \in S : \exists y (x \rightarrow y \wedge y \in \|\varphi\|_{\mathcal{V}}^T)\}$
$\|\mu X \varphi(X)\|_{\mathcal{V}}^T = \bigcap \{U \subseteq S : \|\varphi\|_{\mathcal{V}[X \mapsto U]}^T \subseteq U\}$
$\|\nu X \varphi(X)\|_{\mathcal{V}}^T = \bigcup \{U \subseteq S : U \subseteq \|\varphi\|_{\mathcal{V}[X \mapsto U]}^T\}$

In the above $\mathcal{V}[X \mapsto U]$ is the valuation that maps X into U and agrees with $\mathcal{V}$ on all other variables. Note that a formula φ gives rise to a function $f_\varphi \colon \mathcal{P}(S) \to \mathcal{P}(S)$ given by $U \mapsto \{x \in S : x \in \|\varphi(X)\|_{\mathcal{V}[X \mapsto U]}^T\}$. As f_φ is a monotone function on the powerset lattice

$\langle \mathcal{P}(S), \subseteq \rangle$, by the Knaster-Tarski Theorem its least (and greatest) fixpoint exists, and is equal to the least prefixed point (resp. greatest postfixed point) of f_φ, the set $\|\mu X \varphi\|_\mathcal{V}^T$ (resp. $\|\nu X \varphi\|_\mathcal{V}^T$).

1.2   Alternation-free fragment

The alternation of fixpoint quantifiers is the major source of potency, and a fundamental measure of logical strength in the study of fragments of μ-calculus. The number of genuine alternations between least and greatest fixpoint quantifiers is called the *depth* of the formula. Bradfield [1] showed that there are modal fixpoint properties which require arbitrarily large depth, and hence the modal μ-calculus alternation hierarchy is strict. Formally, the Niwiński hierarchy is defined as follows. A formula φ is in the classes Π_0 and Σ_0 if it contains no fixpoint quantifiers, i.e. it is a formula of modal logic. The class Σ_{n+1} (Π_{n+1}) is the closure of $\Sigma_n \cup \Pi_n$ under the following rules.

- If $\varphi, \psi \in \Sigma_{n+1}$ (Π_{n+1}), then $\varphi \wedge \psi, \varphi \vee \psi, \Box \varphi, \Diamond \varphi \in \Sigma_{n+1}$ (Π_{n+1}).
- If $\varphi \in \Sigma_{n+1}$ (Π_{n+1}), then $\mu X \varphi \in \Sigma_{n+1}$ ($\nu X \varphi \in \Pi_{n+1}$).
- If $\varphi, \psi \in \Sigma_{n+1}$ (Π_{n+1}), then $\varphi(\psi) \in \Sigma_{n+1}$ (Π_{n+1}), provided the free variables of ψ do not become bound by quantifiers in φ.

In comparison the *alternation-free fragment* of the modal μ-calculus is the class of formulæ with no real dependencies between alternating fixpoint quantifiers. This fragment is the closure of $\Sigma_1 \cup \Pi_1$ under Boolean and modal operators and substitutions that preserve the alternation depth. Despite the restrictions imposed, this class of properties still forms a remarkably expressive fragment encompassing the majority of logics used in the verification of systems. It is known that this class coincides with the collection of all formulæ semantically equivalent to both a Σ_2-formula and a Π_2-formula [5]. Moreover, this fragment is the limit of the weak index hierarchy as introduced in [6]; thus, the languages defined by alternation-free formulæ are also referred to as *weakly definable languages.*

1.3   Trees

A *tree* is a pair $t = (V, \rightarrow)$ with a distinguished node ρ_t such that $(V, \rightarrow)$ is a connected directed graph, there are no transitions into ρ_t and for every $v \in V \setminus \{\rho_t\}$ there is exactly one $v_0 \in V$ such that $v_0 \rightarrow v$. The node ρ_t is referred to as the *root* of the tree and any node without outgoing transitions is called a *leaf*. For a tree t and a node v in t, we write $t{\restriction}_v$ to denote the sub-tree rooted at v. If there is no cause for confusion we identify a tree with its domain. Tree $t_0 = (V_0, \rightarrow_0)$ is a *pruning* of $t = (V, \rightarrow)$ if $V_0 \subseteq V$, $\rightarrow_0 = \rightarrow \cap V_0^2$ and if $u \rightarrow v \notin V_0$ then $\{w \in V_0 : u \rightarrow_0 w\} = \emptyset$.

A *path* through a tree $t = (V, \rightarrow)$ is an enumerable set $\mathbb{P} \subseteq V$ such that $\rho_t \in \mathbb{P}$, if $v_0 \rightarrow v \in \mathbb{P}$ then $v_0 \in \mathbb{P}$, and for every $v \in \mathbb{P}$ either v is a leaf or there exists exactly one $u \in V$ such that $v \rightarrow u$ and $u \in \mathbb{P}$. For a path $\mathbb{P}$ given by a sequence $\rho_t = v_0 \rightarrow v_1 \rightarrow v_2 \rightarrow \ldots \rightarrow v_n \rightarrow \ldots$, we write $\mathbb{P}(n)$ to denote v_n. For nodes $u, v \in t$ we write $u <_t v$ (resp. $u \leq_t v$) if for some path $\mathbb{P}$ through t and $i < j$ (resp. $i \leq j$), $\mathbb{P}(i) = u$ and $\mathbb{P}(j) = v$.

A *tree transition system* (TTS) is a transition system $T = (S, \rightarrow, \lambda)$ for which $(S, \rightarrow)$ is a tree. We say a TTS T *satisfies* φ, written $T \models \varphi$, if $\rho_T \in \|\varphi\|_\mathcal{V}^T$. In this case T is a *model* of φ and φ is *satisfiable*. Note that modal μ-calculus has the *tree model property*, namely every satisfiable formula has a model which is a TTS (see e.g. [2]).

1.4 Closure ordinals

The definition of semantics for μ-formulæ can be generalised to also take into account approximations to fixpoint variables. For each formula φ, set of bound variables $\mathcal{X}$ occurring in φ and ordinal α, we define a set $\|\varphi^\alpha\|_{\mathcal{V}}^T$ by induction on α. Let $T = (S, \to, \lambda)$ be a transition system and $\mathcal{V}$ a valuation on T. For every α, define

$$\|p^\alpha\|_{\mathcal{V}}^T = \|p\|_{\mathcal{V}}^T$$
$$\|\bar{p}^\alpha\|_{\mathcal{V}}^T = \|\bar{p}\|_{\mathcal{V}}^T$$
$$\|Z^\alpha\|_{\mathcal{V}}^T = \mathcal{V}(Z)$$
$$\|(\varphi \wedge \psi)^\alpha\|_{\mathcal{V}}^T = \|\varphi^\alpha\|_{\mathcal{V}}^T \cap \|\psi^\alpha\|_{\mathcal{V}}^T$$
$$\|(\varphi \vee \psi)^\alpha\|_{\mathcal{V}}^T = \|\varphi^\alpha\|_{\mathcal{V}}^T \cup \|\psi^\alpha\|_{\mathcal{V}}^T$$
$$\|(\Box\varphi)^\alpha\|_{\mathcal{V}}^T = \{x \in S : \forall y(x \to y \Rightarrow y \in \|\varphi^\alpha\|_{\mathcal{V}}^T)\}$$
$$\|(\Diamond\varphi)^\alpha\|_{\mathcal{V}}^T = \{x \in S : \exists y(x \to y \wedge y \in \|\varphi^\alpha\|_{\mathcal{V}}^T)\}$$
$$\|(\mu X\varphi)^\alpha\|_{\mathcal{V}}^T = \begin{cases} \bigcup_{\gamma<\alpha} \|\varphi[\mu X\varphi/X]^\gamma\|_{\mathcal{V}}^T, & \text{if } X \in \mathcal{X}, \\ \|\mu X\varphi^\alpha\|_{\mathcal{V}}^T, & \text{otherwise.} \end{cases}$$
$$\|(\nu X\varphi)^\alpha\|_{\mathcal{V}}^T = \begin{cases} \bigcap_{\gamma<\alpha} \|\varphi[\nu X\varphi/X]^\gamma\|_{\mathcal{V}}^T, & \text{if } X \in \mathcal{X}, \\ \|\nu X\varphi^\alpha\|_{\mathcal{V}}^T, & \text{otherwise.} \end{cases}$$

For every formula φ there exists an ordinal κ such that $\|\varphi\|_{\mathcal{V}}^T = \|\varphi^\kappa\|_{\mathcal{V}}^T = \|\varphi^{\kappa+1}\|_{\mathcal{V}}^T$. The least such κ is called the *closure ordinal of φ with respect to T and $\mathcal{X}$* and is denoted $CO_{T,\mathcal{X}}(\varphi)$. Note that a formula may have different closure ordinals depending on the transition system on which it is evaluated as well as the particular collection of variables analysed. For example the formula $\mu X\Box X$ is satisfied by all well-founded trees; its closure ordinal with respect to $\{X\}$ in each case is the order-type of the tree.

▶ **Definition 1.1** (Closure Ordinal). The *closure ordinal* of a closed formula φ with respect to a non-empty set $\mathcal{X}$ of variables, denoted by $CO_{\mathcal{X}}(\varphi)$, is the ordinal $\sup_T CO_{T,\mathcal{X}}(\varphi)$, if this ordinal exists.

2 Syntactic analysis

Let $\overline{\text{PROP}} := \{\bar{p} : p \in \text{PROP}\}$ and $P_1, P_1', P_2, P_2', \ldots, P_n, P_n'$ be finite subsets of $\text{PROP} \cup \overline{\text{PROP}}$. Each such set, when referred to as a formula, denotes the conjunction of its elements. We say a formula of modal logic is *primary* if it is of the form

$$(P_1 \wedge \Box P_1' \wedge \nabla_1 X) \vee (P_2 \wedge \Box P_2' \wedge \nabla_2 X) \vee \ldots \vee (P_n \wedge \Box P_n' \wedge \nabla_n X) \vee \Box\bot \tag{1}$$

where $\nabla_i \in \{\Diamond, \Box\}$ for each i. Czarnecki's analysis in [3] establishes that every ordinal below ω^2 is the closure ordinal of the least fixpoint of some primary formula. In this section we establish a strong converse: if the primary formula given in (1) has closure ordinal α, then $\alpha < \omega.(n+1)$. For the following let ψ denote the formula in (1) and $\varphi = \mu X\psi$.

▶ **Lemma 2.1.** Fix a transition system T and a valuation $\mathcal{V}$. Suppose κ is a limit ordinal. If $x \in \|\varphi^{\kappa+1}\|_{\mathcal{V}}^T \setminus \|\varphi^\kappa\|_{\mathcal{V}}^T$, then there is no $j \leq n$ such that $x \in \|P_j \wedge \Box P_j' \wedge \nabla_j \varphi^\kappa\|_{\mathcal{V}}^T$ and $\nabla_j = \Diamond$.

Proof. Suppose $T = (S, \to, \lambda)$ and let $\|\varphi^\alpha\|$ abbreviate $\|\varphi^\alpha\|_{\mathcal{V}}^T$. Suppose $x \in \|\varphi^{\kappa+1}\| \setminus \|\varphi^\kappa\|$. By way of contradiction suppose also $x \in \|P_j \wedge \Box P_j' \wedge \nabla_j \varphi^\kappa\|$ and $\nabla_j = \Diamond$ for some $j \leq n$. If $\{y \in S : x \to y\} = \emptyset$ then $x \in \|\varphi^1\| \subseteq \|\varphi^\kappa\|$ which cannot be, so let $x \to y$ be such that $y \in \|\varphi^\kappa\|$. Thus there exists $\gamma < \kappa$ such that $y \in \|\varphi^\gamma\|$, and hence $x \in \|\varphi^{\gamma+1}\| \subseteq \|\varphi^\kappa\|$ yielding a contradiction. ◀
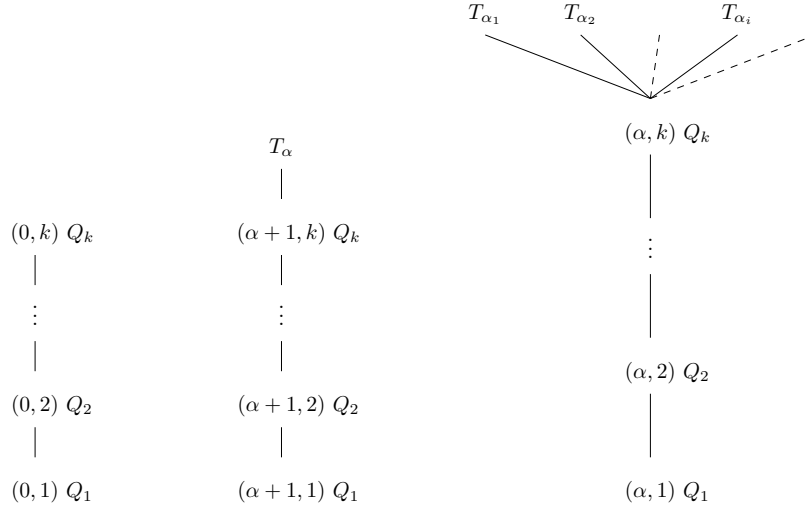
■ **Figure 1** T_0, $T_{\alpha+1}$ and T_α (in the case $\alpha = \sup_i \alpha_i$) in the proof of lemma 2.3.

▶ **Corollary 2.2.** If $\nabla_i = \Diamond$ for every $i \leq n$ then the closure ordinal of $\mu X \psi$ exists and is no greater than ω.

▶ **Lemma 2.3.** Suppose there exist consistent sets of propositions Q_1, Q_2,…,Q_{k+1} and numbers $i_1, i_2, \ldots, i_k < n$ such that $P_{i_j} \wedge \Box P'_{i_j} \wedge \nabla_{i_j} X$ is a subformula of ψ with $P_{i_j} \subseteq Q_j$ and $P'_{i_j} \subseteq Q_{j+1}$ for each $j \leq k$. Furthermore, suppose $\nabla_{i_k} = \Box$ and there is no $j \leq n$ such that $P_j \subseteq Q_k$, $P'_j \subseteq Q_{k+1}$ and $\nabla_j = \Diamond$. If $Q_{k+1} = Q_1$, then $\mu X \psi$ does not have a closure ordinal.

Proof. Let $\lambda \colon \mathrm{ON} \times \{i : i \leq k+1\} \to \mathcal{P}(\mathrm{PROP})$, where ON is the class of all ordinals, be defined by $p \in \lambda((\alpha, j))$ if and only if $p \in Q_j$. Furthermore, let $T_0^\alpha = (S_0^\alpha, \to_0^\alpha, \lambda)$ be the TTS where

$$S_0^\alpha = \{(\alpha, j) : 0 < j \leq k\},$$
$$\to_0^\alpha = \{((\alpha, j), (\alpha, j+1)) : 0 < j < k\}.$$

For each countable ordinal α we define a tree T_α as follows. Let $T_0 = T_0^0$ and $T_{\alpha+1} = (S_{\alpha+1}, \to_{\alpha+1}, \lambda)$ where $S_{\alpha+1} = S_0^{\alpha+1} \cup S_\alpha$ and $\to_{\alpha+1} = \to_0^{\alpha+1} \cup \to_\alpha \cup \{((\alpha+1, k), (\alpha, 1))\}$. If α is a limit ordinal, then $S_\alpha = S_0^\alpha \cup \bigcup_{\beta < \alpha} S_\beta$ and $\to_\alpha = \to_0^\alpha \cup \bigcup_{\beta < \alpha} \to_\beta \cup \{(\alpha, k), (\beta, 1)) : \beta < \alpha\}$.

Let f be the function $\kappa \mapsto k.\kappa$. We will show that for each $\kappa \leq \alpha$ and $0 \leq j < k$,

$$(\kappa, k - j) \in \|\varphi^{f(\kappa)+j+1}\|_{\mathcal{V}}^{T_\alpha} \setminus \|\varphi^{f(\kappa)+j}\|_{\mathcal{V}}^{T_\alpha} \tag{2}$$

whereby it will be clear that the formula φ does not possess a closure ordinal. The argument proceeds by transfinite induction on $\kappa \leq \alpha$ with an auxiliary induction on $j < k$. If $j \neq 0$ then (2) follows from the fact that $(\kappa, k - (j-1))$ is the unique successor of $(\kappa, k - j)$ and the definition of λ. Thus suppose $j = 0$, whence three sub-cases manifest:

- $\kappa = 0$. Then $f(\kappa) = 0$ and (κ, k) is a leaf of T_α, so (2) trivially holds.
- $\kappa = \kappa' + 1$. By the definition of T_α, (κ, k) has a unique successor, namely $(\kappa', 1)$, whence (2) follows from the induction hypothesis
- κ limit. The successors of $(\kappa, k - j)$ in this case are the nodes $(\gamma, 1)$ for $\gamma < \kappa$. By the induction hypothesis we know $(\gamma, 1) \in \|\varphi^{f(\gamma+1)}\|_{\mathcal{V}}^{T_\alpha} \setminus \|\varphi^{f(\gamma)+k-1}\|_{\mathcal{V}}^{T_\alpha}$ for each $\gamma < \kappa$. Notice

that $\|\varphi^{f(\kappa)}\|_{\mathcal{V}}^{T_\alpha} = \bigcup_{\gamma < \kappa} \|\varphi^{f(\gamma)}\|_{\mathcal{V}}^{T_\alpha}$. Since $P_{i_k} \subseteq Q_k$, $P'_{i_k} \subseteq Q_1$ and $\nabla_{i_k} = \square$, it follows that $(\kappa, k) \in \|\varphi^{f(\kappa)+1}\|_{\mathcal{V}}^{T_\alpha}$. If, however, $(\kappa, k) \in \|\varphi^{f(\kappa)}\|_{\mathcal{V}}^{T_\alpha}$, then $(\kappa, k) \in \|\varphi^{f(\gamma)+k}\|_{\mathcal{V}}^{T_\alpha}$ for some $\gamma < \kappa$. But then $(\gamma, 1) \in \|\varphi^{f(\gamma)+k-1}\|_{\mathcal{V}}^{T_\alpha}$ by the assumption on φ. ◄

▶ **Proposition 2.4.** The closure ordinal of φ is strictly less than ω^2.

Proof. Let α be the closure ordinal of φ and suppose $\alpha \geq \omega^2$. Fix $N \geq 2^{|\varphi|+1}$ where $|\varphi|$ denotes the number of symbols occurring in φ. Let T be a TTS such that for every $i \leq N$, $\|\varphi^{\omega \cdot i}\|_{\mathcal{V}}^T$ is a proper subset of $\|\varphi^{\omega \cdot i+1}\|_{\mathcal{V}}^T$. Then there exists a path $\mathbb{P}$ through T, $m_N < m_{N-1} < \cdots < m_0 < \omega$ and a function $f : \omega \times \omega \to \omega$ such that for every $i \leq N$ and $j < m_i - m_{i+1}$, $f(i,j) \leq n$ and

$$\mathbb{P}(m_i - j) \in \|P_{f(i,j)} \wedge \square P'_{f(i,j)} \wedge \nabla_{f(i,j)} \varphi^{\omega \cdot i+j}\|_{\mathcal{V}}^T \setminus \|\varphi^{\omega \cdot i+j}\|_{\mathcal{V}}^T.$$

Define for each $j < \omega$, $Q_j = \lambda_T(\mathbb{P}(j)) \cup \{\bar{p} : p \notin \lambda_T(\mathbb{P}(j))\}$. For some $i_0 < i_1 \leq N$ it must be the case that

$$Q_{m_{i_0}} \cap \text{PROP}_\varphi = Q_{m_{i_1}} \cap \text{PROP}_\varphi$$

where $\text{PROP}_\varphi = \bigcup_{i \leq n} (P_i \cup P'_i)$. The sequence $Q_{m_{i_1}}, \ldots, Q_{m_{i_0}}$ therefore fulfils the hypothesis of lemma 2.3 whence, contrary to our assumption, φ does not have a closure ordinal. ◄

The above analysis can also be applied to formulæ of the form

$$(\psi_1 \wedge \nabla_1 X) \vee (\psi_2 \wedge \nabla_2 X) \vee \ldots \vee (\psi_n \wedge \nabla_n X) \vee \square\bot \tag{3}$$

where $\psi_1, \ldots, \psi_n$ are closed formulæ of modal logic. Replacing literals with arbitrary modal formulæ in each disjunct alters the "proposition paths" that can occur. Therefore, in order to find a repetition as in the proof of proposition 2.4, one will need to look at larger segments of a suitable model. As such a proof would be technically cumbersome, in the next section we will employ a semantic analysis which will include (3) and extend the bounds to formulæ of the alternation-free fragment of μ-calculus.

3 Semantic analysis

For the remainder of the paper, formulæ are assumed to be closed and guarded unless otherwise stated. A formula φ is *guarded* if in every subformula $\sigma Z.\psi$ of φ, every occurrence of the bound variable Z in ψ appears within the scope of a modal operator. The restriction to the guarded fragment is not significant as every formula is equivalent to one in guarded form (see e.g. [7]). Moreover, by following the approach of [4] it is possible to carry out the analysis below for unguarded formulæ.[1]

Upper-case Greek letters such as Γ and Δ denote *sequents*, finite sets of formulæ. $\square\Gamma$ abbreviates the set $\{\square\varphi : \varphi \in \Gamma\}$ and $\Diamond\Gamma$ is defined analogously. We write Γ, φ for $\Gamma \cup \{\varphi\}$, and Γ, Δ to denote $\Gamma \cup \Delta$. The *Fischer-Ladner closure* of a formula φ, denoted by $\text{FL}(\varphi)$, is the smallest set such that

- $\varphi \in \text{FL}(\varphi)$,
- if $\psi_0 \circ \psi_1 \in \text{FL}(\varphi)$ where $\circ \in \{\vee, \wedge\}$ then $\psi_0, \psi_1 \in \text{FL}(\varphi)$,
- if $\nabla\psi \in \text{FL}(\varphi)$ where $\nabla \in \{\Diamond, \square\}$ then $\psi \in \text{FL}(\varphi)$,
- if $\sigma X\psi \in \text{FL}(\varphi)$ where $\sigma \in \{\mu, \nu\}$ then $\psi[\sigma X\psi/X] \in \text{FL}(\varphi)$.

Note that $|\text{FL}(\varphi)| \leq |\varphi|$ where $|\varphi|$ denotes the number of symbols occurring in φ. For a sequent Γ we set $\text{FL}(\Gamma) = \bigcup_{\gamma \in \Gamma} \text{FL}(\gamma)$.

---

[1] We would like to thank the anonymous referee for drawing our attention to [4].

3.1   Tableaux

▶ **Definition 3.1.** Given a TTS T and a sequent Γ, a *pre-tableau* for (T, Γ) is a tree $t = (V, \to)$ together with functions $\tau_t : t \to T$ and $\lambda_t : t \to \mathcal{P}(\mathrm{FL}(\Gamma))$ such that the following conditions are satisfied.

- $\tau_t(\rho_t) = \rho_T$ and $\lambda_t(\rho_t) = \Gamma$.
- If $v \in t$ is a leaf then $\lambda_t(v) = \Box\Xi, \Theta$ where $\Theta \subseteq \mathrm{PROP} \cup \overline{\mathrm{PROP}}$, and either $\Xi = \emptyset$ or $\tau_t(v)$ is a leaf of T.
- If $\tau_t(u) = \tau_t(v)$ then either $u \leq_t v$ or $v \leq_t u$.
- For every $v \in t$, $\lambda_t(v) \cap \mathrm{PROP} \subseteq \lambda_T(\tau_t(v)) \subseteq \{p \in \mathrm{PROP} : \bar{p} \notin \lambda_t(v)\}$.
- For every $v_0 \to v_1 \in t$ with $\tau_t(v_i) = x_i$ and $\lambda_t(v_i) = \Gamma_i$, one of the following conditions hold.

  ($\wedge$) $x_0 = x_1$ and there are formulæ φ_0, φ_1 such that $\varphi_0 \wedge \varphi_1 \in \Gamma_0$ and $\Gamma_1 = (\Gamma_0 \setminus \{\varphi_0 \wedge \varphi_1\}) \cup \{\varphi_0, \varphi_1\}$. The formula $\varphi_0 \wedge \varphi_1$ is called *active at v_0* and both φ_0 and φ_1 *residual at v_1*.

  ($\vee$) $x_0 = x_1$ and there are formulæ φ_0, φ_1 such that $\varphi_0 \vee \varphi_1 \in \Gamma_0$ and $\Gamma_1 = (\Gamma_0 \setminus \{\varphi_0 \vee \varphi_1\}) \cup \{\varphi_i\}$. The formula $\varphi_0 \vee \varphi_1$ is called *active at v_0* and φ_i *residual at v_1*.

  (σX) $x_0 = x_1$ and there is a formula φ and $\sigma \in \{\mu, \nu\}$ such that $\sigma X \varphi \in \Gamma_0$ and $\Gamma_1 = (\Gamma_0 \setminus \{\sigma X \varphi\}) \cup \{\varphi[\sigma X \varphi / X]\}$. The formula $\sigma X \varphi$ is called *active at v_0* and $\varphi(\sigma X \varphi)$ *residual at v_1*.

  (mod) $x_0 \to_T x_1$ and $\Gamma_0 = \Box\Xi, \Diamond\Delta, \Theta$ with $\Theta \subseteq \mathrm{PROP} \cup \overline{\mathrm{PROP}}$ and $\Xi \subseteq \Gamma_1 \subseteq \Xi \cup \Delta$. All formulæ in Γ_0 are considered *active at v_0* and all formulæ in Γ_1 *residual at v_1*.

  In the cases ($\wedge$), ($\vee$) and (σX) above, $|\{u : v_0 \to u\}| = 1$, while in the case of (mod), $\bigcup_{v_0 \to u} \lambda_t(u) = \Xi \cup \Delta$ and $\{\tau_t(u) : v_0 \to u\} = \{y : x_0 \to_T y\}$.

▶ **Remark 3.2.** Exactly one of the four conditions ($\wedge$), ($\vee$), (σX) and (mod) can apply to a non-leaf node of a pre-tableau; henceforth we will refer to them as *tableaux rules*. Note that in a pre-tableau branching only occurs at a (mod)-rule and may be infinite.

Suppose t is a pre-tableau for (T, Γ) and $\Psi = \{(\psi_i, v_i) : i \in I\} \subseteq \mathrm{FL}(\Gamma) \times t$ where I is an initial segment of natural numbers. Ψ is called a *trace* from (ψ, v) if $(\psi_0, v_0) = (\psi, v)$ and there exists a path $\mathbb{P}$ in t and natural number n such that for every $i \in I$,

- $v_i = \mathbb{P}(n + i)$,
- $\psi_i \in \lambda_t(v_i)$,
- if v_i is a leaf or $\psi_i \in \mathrm{PROP} \cup \overline{\mathrm{PROP}}$ is active at v_i then $i + 1 \notin I$,
- if $i + 1 \in I$ and ψ_i is active at v_i then ψ_{i+1} is an immediate subformula of ψ_i that is residual at v_{i+1},
- if $i + 1 \in I$ and ψ_i is not active at v_i then $\psi_{i+1} = \psi_i$.

In each infinite trace (i.e. if I is infinite) there exists a variable that appears infinitely often and subsumes all other infinitely occurring variables. If this unique variable is a μ-variable then the trace is called a *μ-trace*; otherwise it is a *ν-trace*.

▶ **Definition 3.3.** A pre-tableau for (T, Γ) is a *tableau* if every infinite trace is a ν-trace.

The following theorem which provides a characterisation of satisfaction in terms of the existence of tableaux is folklore; see for example [7].

▶ **Theorem 3.4.** $T \models \bigwedge \Gamma$ if and only if there is a tableau for (T, Γ).

3.2 Order-types of tableaux

Fix a TTS T and a sequent Γ. To each tableau for (T, Γ) and set of μ-variables $\mathcal{X}$ one can assign an order-type with respect to $\mathcal{X}$ in a natural way. The *order-type of ψ at a node v*, denoted by $\alpha_{\psi, v, \mathcal{X}}$, is defined recursively as follows. If there exists a trace $\Psi = \{(\psi_i, v_i) : i \in I\}$ from (ψ, v) such that for infinitely many $i \in I$, ψ_i has the form $\mu X \psi'$ for some $X \in \mathcal{X}$, or there are no traces $\Psi = \{(\psi_i, v_i) : i \in I\}$ from (ψ, v) for which ψ_i has the form $\mu X \psi'$ for some $i \in I$ and $X \in \mathcal{X}$, then $\alpha_{\psi, v, \mathcal{X}} = 0$. Otherwise,

- if $\psi = \mu X \psi'$ is active at v and $X \in \mathcal{X}$ then $\alpha_{\psi, v, \mathcal{X}} = \alpha_{\psi', u, \mathcal{X}} + 1$ where u is the unique successor of v in the tableau,
- if ψ is not of the form $\mu X \psi'$ for some $X \in \mathcal{X}$ or not active at v then $\alpha_{\psi, v, \mathcal{X}}$ is the supremum of $\alpha_{\psi_1, v_1, \mathcal{X}}$ for which there exists a trace $\Psi = \{(\psi_i, v_i) : i \in I\}$ from (ψ, v).

▶ **Definition 3.5.** The *order-type with respect to $\mathcal{X}$* of a tableau t for (T, Γ) is the ordinal $\sup\{\alpha_{\varphi, \rho_t, \mathcal{X}} : \varphi \in \Gamma\}$. A tableau is an *$\alpha$-tableau with respect to $\mathcal{X}$* if its order-type with respect to $\mathcal{X}$ is no greater than α.

To establish the connection between the closure ordinal of a formula and order-types of the corresponding tableaux we show that if φ is alternation-free and $\mathcal{X}$ a set of μ-variables,

$$x \in \|\varphi^\alpha\|_\mathcal{V}^T \text{ iff there exists an } \alpha\text{-tableau for } (T\!\restriction_x, \varphi) \text{ with respect to } \mathcal{X}.$$

We will prove the result for $\mathcal{X} = \{X\}$; the above statement is a direct generalisation of the next lemma.

▶ **Lemma 3.6.** Suppose $\psi(Y)$ is a formula with at most Y free and X a variable not occurring in ψ. Let $\mathcal{X} = \{X\}$ and T be a TTS. Then $x \in \|\psi(\mu X \varphi)^\alpha\|_\mathcal{V}^T$ if and only if there exists an α-tableau for $(T\!\restriction_x, \psi(\mu X \varphi))$ with respect to $\mathcal{X}$.

Proof. By transfinite induction on α. For the base case suppose $\alpha = 0$. We want to show

$$x \in \|\psi(Z)\|_{\mathcal{V}[Z \mapsto \emptyset]}^T \text{ iff there exists a } 0\text{-tableau } (T\!\restriction_x, \psi(\mu X \varphi)).$$

Notice $x \in \|\psi(Z)\|_{\mathcal{V}[Z \mapsto \emptyset]}^T$ if and only if there is a tableau for $(T\!\restriction_x, \psi(\bot))$. Consider a tableau for $(T\!\restriction_x, \psi(\bot))$. Since $\bot$ cannot appear in the label of any node, this tableau can be used to create a tableau for $(T\!\restriction_x, \psi(\mu X \varphi))$ in a trivial way: replace $\bot$ by $\mu X \varphi$ at relevant positions. The order-type of the emerging tableau is 0 as $\mu X \varphi$ can never appear in any trace. Conversely, since a tableau of order-type 0 means the (μX)-rule is never applied, replacing occurrences of $\mu X \varphi$ by $\bot$ in a tableau for $(T\!\restriction_x, \psi(\mu X \varphi))$ yields a tableau for $\psi(\bot)$.

For the successor case we want to show

$$x \in \|\psi(\mu X \varphi)^{\alpha+1}\|_\mathcal{V}^T \text{ iff there exists an } (\alpha+1)\text{-tableau } (T\!\restriction_x, \psi(\mu X \varphi)).$$

Note that $x \in \|\psi(\mu X \varphi)^{\alpha+1}\|_\mathcal{V}^T$ if and only if $x \in \|(\psi \circ \varphi)(\mu X \varphi)^\alpha\|_\mathcal{V}^T$, if and only if there exists an α-tableau for $(T\!\restriction_x, (\psi \circ \varphi)(\mu X \varphi))$ by the induction hypothesis. Hence it suffices to show how to construct an $(\alpha+1)$-tableau for $(T\!\restriction_x, \psi(\mu X \varphi))$ from an α-tableau for $(T\!\restriction_x, \psi \circ \varphi(\mu X \varphi))$ and vice versa. Given an α-tableau t for $(T\!\restriction_x, \psi \circ \varphi(\mu X \varphi))$, along every path look for the first node v with $\lambda_t(v) = \Gamma, \varphi(\mu X \varphi)$ for some Γ, and replace all occurrences of $\varphi(\mu X \varphi)$ by $\mu X \varphi$ in nodes $u \leq_t v$. The sequent at v has therefore become $\Gamma, \mu X \varphi$. Between v and its successors, insert a new node labelled by $\Gamma, \varphi(\mu X \varphi)$. The added transition is a valid (μX)-rule so the resulting tableau is readily seen to be a tableau for $(T\!\restriction_x, \psi(\mu X \varphi))$. Moreover, all traces from $(\mu X \varphi, v)$ have order-type at most $\alpha + 1$ and indeed, the tableau

for $\big(T{\restriction}_x, \psi(\mu X\varphi)\big)$ has order-type $\alpha + 1$. Similarly, by replacing occurrences of $\mu X\varphi$ by $\varphi(\mu X\varphi)$ at the relevant nodes in an $(\alpha{+}1)$-tableau for $\big(T{\restriction}_x, \psi(\mu X\varphi)\big)$ and removing the first application of a (μX)-rule on every trace one obtains an α-tableau for $\big(T{\restriction}_x, \psi \circ \varphi(\mu X\varphi)\big)$.

For the limit case suppose $x \in \|\psi(\mu X\varphi)^\alpha\|_\mathcal{V}^T$. Let q be a fresh proposition and T^q a new TTS obtained by adjusting the labelling so that q holds at all nodes belonging to $\|(\mu X\varphi)^\alpha\|_\mathcal{V}^T$ i.e.

$$\lambda_{T^q}(x) = \left\{ \begin{array}{ll} \lambda_T(x) \cup \{q\}, & \text{if } x \in \|(\mu X\varphi)^\alpha\|_\mathcal{V}^T, \\ \lambda_T(x), & \text{otherwise.} \end{array} \right.$$

Since $\|\psi(\mu X\varphi)^\alpha\|_\mathcal{V}^T = \|\psi(q)\|_\mathcal{V}^{T^q}$ and $\psi(q)$ is closed, there is a tableau t for $(T^q{\restriction}_x, \psi(q))$ of order-type 0. It is possible that there are nodes of this tableau at which q is active. The key to obtaining a tableau for $\big(T{\restriction}_x, \psi(\mu X\varphi)\big)$ lies in replacing the occurrences of q at these nodes by tableaux for $\mu X\varphi$ of the relevant order-type. Suppose $\lambda_t(v) = \Box\Gamma, \Diamond\Delta, \Theta, q$, $\tau_t(v) = y$, q is active at v and for no $u <_t v$ is q active at u. Let $\beta < \alpha$ be such that $y \in \|(\mu X\varphi)^\beta\|_\mathcal{V}^T$. By the main induction hypothesis there is a β-tableau for $(T{\restriction}_y, \mu X\varphi)$. We can combine this tableau with the sub-tableau $t{\restriction}_v$ to obtain a β-tableau t_v for $(T^q{\restriction}_y, \Box\Gamma, \Diamond\Delta, \Theta, \mu X\varphi)$. Now we replace $t{\restriction}_v$ by t_v in t, substitute each occurrence of q by $\mu X\varphi$ in the trace from the root to (q,v) and repeat the procedure. In the limit a tableau for $(T{\restriction}_x, \psi(\mu X\varphi))$ is obtained. Moreover, the order-type of this tableau can be no greater than α.

The converse direction is equally straight forward. $\blacktriangleleft$

▶ **Corollary 3.7.** Suppose φ is a closed formula and $\mathcal{X}$ a set of μ-variables occurring in φ. For an arbitrary TTS T, set α_T to be 0 if $T \not\models \varphi$, and otherwise the infimum of the order-types of all possible tableaux for (T, φ) with respect to $\mathcal{X}$. Then $CO_\mathcal{X}(\varphi) = \sup\{\alpha_T : T \text{ a TTS}\}$.

With corollary 3.7 in mind, in order to rule out certain ordinals being closure ordinals we require a notion of minimality of order-types for tableaux.

▶ **Definition 3.8.** A tableau t for (T, Γ) is *minimal* if there are no tableau for (T, Γ) with smaller order-type, and *absolutely minimal* if for every node $v \in t$, $t{\restriction}_v$ is a minimal tableau for $(T{\restriction}_{\tau_t(v)}, \lambda_t(v))$.

▶ **Remark 3.9.** If $T \models \varphi$ then a minimal tableau t for (T, φ) exists. Moreover, as $T{\restriction}_{\tau_t(v)} \models \bigwedge \lambda_t(v)$ for each $v \in t$, the existence of an absolutely minimal tableau for (T, φ) is also guaranteed.

As a refinement of lemma 3.6 for limit ordinals we have the following.

▶ **Proposition 3.10.** Suppose φ is a formula with closure ordinal $\omega.\alpha > 0$ with respect to a set $\mathcal{X}$ of μ-variables. Then there exists a TTS T and a minimal tableau for $(T, \Box\varphi)$ with order-type $\omega.\alpha$ with respect to $\mathcal{X}$.

Proof. By corollary 3.7, for every $\beta < \omega.\alpha$ there exists a TTS T_β such that every tableau for (T_β, φ) has order-type greater than β. Let T be the TTS obtained by extending the disjoint union of $\{T_\beta : \beta < \omega.\alpha\}$ by a fresh node ρ_T whose immediate successors are $\{\rho_{T_\beta} : \beta < \omega.\alpha\}$. As $T \models \Box\varphi$, there exists a tableau for $(T, \Box\varphi)$. Moreover, every minimal tableau for $(T, \Box\varphi)$ has order-type $\omega.\alpha$ with respect to $\mathcal{X}$. $\blacktriangleleft$

3.3   Closure ordinals for the alternation-free fragment

In this section we determine upper bounds on the closure ordinals of alternation-free formulæ. The analysis breaks into two parts. First we prove that if an alternation-free formula φ has closure ordinal strictly less than ω^2 with respect to its external μ-variables, then this ordinal

is bounded by $\omega.2^{2^{|\varphi|+2}}$. Although primary formulæ can yield ordinals arbitrary close to ω^2 (from below), in the second part we show that the closure ordinal of any alternation-free formula is strictly less than ω^2.

We need only consider order-types for tableaux with respect to particular classes of μ-variables. Given a formula φ, a set of variables $\mathcal{X}$ of φ is called *principal* if whenever $X \in \mathcal{X}$ appears within the scope of a quantifier σY in φ, also $Y \in \mathcal{X}$. Let $\mathcal{X}_\varphi$ denote the largest principal set containing only μ-variables of φ.

An *ordinal assignment* on a tree t is a function $o \colon t \to \mathrm{ON}$ such that if x, y are nodes in t and $x \leq_t y$ then $o(y) \leq o(x)$. A tableau t for (T, Γ) induces a natural ordinal assignment on itself, denoted o_t, setting $o_t(u) = \sup\{\alpha_{\psi,u,\mathcal{X}_\Gamma} : \psi \in \lambda_t(u)\}$ for every $u \in t$, where $\mathcal{X}_\Gamma = \bigcup_{\varphi \in \Gamma} \mathcal{X}_\varphi$. Furthermore, the same tableau induces an ordinal assignment on T, also denoted o_t, by defining $o_t(x) = \sup\{o_t(u) : u \in t \wedge \tau_t(u) = x\}$ for each $x \in T$. The *order-type of a tableau* t, denoted $o(t)$, is the ordinal $o_t(\rho_t)$. A tableau is an α-*tableau* if its order-type is no greater than α.

▶ **Lemma 3.11.** If $T \models \varphi$ is a TTS with an infinite path $x_1 <_T x_2 <_T \cdots$ then there exists k such that for every $\Gamma \subseteq \mathrm{FL}(\varphi)$, every absolutely minimal tableau t for (T, Γ) and every $l > k$, $o_t(x_l) = 0$.

Proof. Suppose the contrary, namely for every i there exists $\Gamma_i \subseteq \mathrm{FL}(\varphi)$ and absolutely minimal tableau t_i for (T, Γ_i) such that $o_{t_i}(x_i) > 0$. For each m and i, let $\Delta_i^m \subseteq \mathcal{P}(\mathrm{FL}(\varphi))$ be the collection of sequents that are associated with x_m by t_i,

$$\Delta_i^m = \{\Delta : \exists u \in t_i(\tau_{t_i}(u) = x_m \wedge \lambda_{t_i}(u) = \Delta)\}.$$

For each m, there exists an infinite set $I \subseteq \omega$ with $\Delta_i^m = \Delta_j^m$ for every $i, j \in I$. Thus it is possible to define a sequence $(S_m)_{n \in \omega}$ such that for each m,

1. S_m is an infinite set,
2. $S_{m+1} \subseteq S_m$,
3. for every $i, j \in S_m$, $\Delta_i^m = \Delta_j^m$.

As for each i the tableau t_i is absolutely minimal, we have in fact

$$\forall i, j \in S_m \, o_{t_i}(x_m) = o_{t_j}(x_m)$$

for every m. Let $f \colon \omega \to S_0$ be a strictly increasing function such that $f(m) \in S_m$ for every m and set $\alpha_m = o_{t_{f(m)}}(x_m)$. Then the sequence $(\alpha_m)_{m \in \omega}$ is a weakly decreasing sequence of ordinals as

$$\begin{aligned}
\alpha_{m+1} &= o_{t_{f(m+1)}}(x_{m+1}) \\
&\leq o_{t_{f(m+1)}}(x_m), \quad \text{since } x_m <_T x_{m+1}, \\
&= o_{t_{f(m)}}(x_m), \quad \text{since } S_{m+1} \subseteq S_m, \\
&= \alpha_m.
\end{aligned}$$

As $f(m) \geq m$, we also have that $\alpha_m = o_{t_{f(m)}}(x_m) \geq o_{t_{f(m)}}(x_{f(m)}) > 0$. Thus, the sequence $(\alpha_{m.|\varphi|})_{m \in \omega}$ forms an infinite, strictly decreasing sequence of ordinals.   ◀

Given T, Γ and a non-empty collection S of tableaux for (T, Γ), we define the S-*pruning of* T to be the TTS T' that alters T by setting, for each propositional constant q not appearing in Γ, $q \notin \lambda_{T'}(x)$ iff for some $s \in S$ and all $y <_T x$, $o_s(y) > 0$. If S is the collection of all absolutely minimal tableaux for (T, Γ), we write $\Gamma \star T$ for the S-pruning of T.

▶ Lemma 3.12 (Well-foundedness lemma). If T is a TTS, Γ is a finite set of formulæ all satisfied by T and q is a propositional constant not occurring in Γ then $\{x \in \Gamma \star T : q \notin \lambda_{\Gamma \star T}(x)\}$ forms a well-founded initial sub-tree of T.

Proof. Immediate consequence of lemma 3.11.   ◄

The next three lemmata relate tableaux on $\Gamma \star T$ and T. Let T be a TTS, Γ a sequent and q a propositional constant not occurring in Γ.

▶ Lemma 3.13. If $y \in T$ and $o_s(y) \leq \alpha$ for every absolutely minimal tableau s for (T, Γ) then the set $\{x \in \Gamma \star T : q \notin \lambda_{\Gamma \star T}(x) \wedge y \leq_T x\}$ forms a well-founded tree of order-type no greater than $|\Gamma|.(1 + \alpha)$.

Proof. By transfinite induction on α. Notice that if $\tau_s(u) = y$ and $o_s(u) > 0$ then every trace in $s|_u$ must pass through a (μX)-rule for which $\mu X \varphi$ is active, within the first $|\Gamma|$ occurrences of a (mod)-rule.   ◄

▶ Lemma 3.14. If $\{x \in \Gamma \star T : q \notin \lambda_{\Gamma \star T}(x) \wedge y \leq_T x\}$ forms a non-empty (well-founded) tree of order-type $\omega.\alpha$ then

1. for every $\Delta \subseteq \Gamma$ and every absolutely minimal tableau t for (T, Δ), $o_t(y) \leq \omega.\alpha$,
2. there exists an absolutely minimal tableau s for (T, Γ) such that $o_s(y) = \omega.\alpha$.

Proof. 1 can be proved via transfinite induction, noting that since Γ is a set of guarded formulæ, between any two applications of the (σY)-rule on the same trace, the (mod)-rule must have been applied.

We prove 2. Suppose, in search of a contradiction, that for every absolutely minimal tableau s for (T, Γ), $o_s(y) < \omega.\alpha$. Consider the ordinal

$$\delta = \sup\{o_s(y) : s \text{ is an absolutely minimal tableau for } (T, \Gamma)\}.$$

By lemma 3.13 it must be the case that $\delta = \omega.\alpha$. But then for every $\beta < \alpha$ there exists an absolutely minimal tableau s for (T, Γ) such that $\beta < o_s(y) < \delta$; contradiction.   ◄

For a formula $\varphi \in \Gamma$, let φ_q denote the formula resulting from replacing in φ each $X \in \mathcal{X}_\varphi$ by $\bar{q} \wedge X$, and set $\Gamma_q = \{\varphi_q : \varphi \in \Gamma\}$.

▶ Lemma 3.15. There exists an α-tableau for (T, Γ) iff there is an α-tableau for $(\Gamma \star T, \Gamma_q)$.

Proof. Suppose t is an α-tableau for (T, Γ). Then there exists an absolutely minimal tableau t' for (T, Γ) with $o(t') \leq \alpha$. An $o(t')$-tableau for $(\Gamma \star T, \Gamma_q)$ can be readily constructed from t'. For the converse, let t be a tableau for $(\Gamma \star T, \Gamma_q)$. By the definition of Γ_q, it follows that if $o_t(y) > 0$ then $y \in \{x \in \Gamma \star T : q \notin \lambda_{\Gamma \star T}(x)\}$ whence t can be modified to yield a tableau for (T, Γ) with the same order-type.   ◄

Lemma 3.15 together with lemma 3.14 provide immediate upper bounds on the order-types of sub-tableaux for $(\Gamma \star T, \Gamma_q)$. We can now expand on these properties to obtain a more fine-grained version of lemma 3.14.

If B is a collection of nodes in a tableau s, ee write $v \leq_s B$ if for some $u \in B$, $v \leq_s u$. Let s be an arbitrary tableau, s_0 a pruning of s and suppose $A \subseteq s_0$ is the collection of leaves of s_0 that are inner nodes of s. A *filter over* (s, s_0) is a set $B \subseteq A$ such that for every $v \leq_s B$ if $\{u : v \to_s u \text{ and } u \leq_s A\}$ is infinite, so is $\{u : v \to_s u \text{ and } u \leq_s B\}$. An *ordinal for the filter B* is any α such that for every $v \leq_s B$, if $\{u : v \leq_s u \in A\}$ is infinite then for every $\beta < \alpha$ there is $w \in A$ such that $v \leq_s w$ and $\beta \leq o_s(w)$. It follows that for any tableau s and pruning s_0:
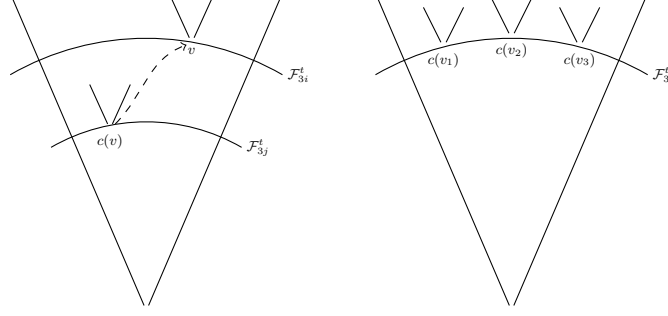
Figure 2 Tableaux t and $\hat{t}$ in the proof of lemma 3.18.

▶ **Lemma 3.16.** If $o(s) < \alpha + o(s_0)$ then there is no filter over s with ordinal $\alpha + \omega$.

▶ **Lemma 3.17.** If every ordinal for every filter over s is bounded by α, then $o(s) \leq \alpha + o(s_0)$.

Proof. Both lemmata are proved by transfinite induction on $o(s)$. For the second lemma, notice that for $u \leq_s B$, if $o_{s_0}(u) = \omega.\beta$ and for every $v > u$, $o_{s_0}(v) < \omega.\beta$, then for $o_s(u) > \alpha + o_{s_0}(u)$ to be the case we must have $o_s(v) > \alpha + o_{s_0}(u)$ for some $v \geq_{s_0} u$. ◀

We are now ready to prove the core lemma.

▶ **Lemma 3.18.** Let $N = 2^{2^{|\varphi|+2}}$. If there is a minimal tableau for (T, φ) of order-type $\alpha \in [\omega.N, \omega^2)$ then there exists a TTS $\hat{T}$ and a minimal tableau for $(\hat{T}, \varphi)$ with order-type strictly greater than α.

Proof. Suppose $\alpha = \omega.m_1 + m_2$ and q is a constant not appearing in φ. Let $T' = \varphi \star T$. For each $i \leq m_1$ define

$$\mathcal{F}_i = \{y \in T' : \{x \in T : q \notin \lambda_{T'}(x) \land y \leq_T x\} \text{ is a tree of order-type } \omega.i\}.$$

Since there is a minimal tableau for (T, φ) of order-type $\alpha \geq \omega.N$, the set $\mathcal{F}_i$ is non-empty for every $i \leq N$. Moreover, by lemma 3.15 there exists a tableau for (T', φ_q) with order-type precisely α. Denote this tableau by t and set $\mathcal{F}_i^t = \{v \in t : \tau_t(v) \in \mathcal{F}_i\}$. Let

$$\Delta_i = \{\Delta : \text{there exists } v \in \mathcal{F}_i^t \text{ and an } \omega.i\text{-tableau for } (T'\!\upharpoonright_{\tau_t(v)}, \Delta_q)\}.$$

Notice Δ_i is non-empty for each $0 < i \leq N$. Moreover, as $\Delta_i \subseteq \mathcal{P}(\mathrm{FL}(\varphi))$ and $m_1 \geq N$, there exists $0 < i < j \leq m_1$ such that $\Delta_{3i} = \Delta_{3j}$ and $\Delta_{3i-1} = \Delta_{3j-1}$. To each $v \in \mathcal{F}_{3i}^t$ is therefore associated a node $c(v) \in \mathcal{F}_{3j}^t$ such that for every $\Delta \subseteq \mathrm{FL}(\varphi)$,

1. there is a tableau for $(T'\!\upharpoonright_{\tau_t(v)}, \Delta_q)$ if and only if there is a tableau for $(T'\!\upharpoonright_{\tau_t(c(v))}, \Delta_q)$,
2. there exists an $\omega.(3i-1)$-tableau for $(T'\!\upharpoonright_{\tau_t(v)}, \Delta_q)$ if and only if there exists an $\omega.(3j-1)$-tableau for $(T'\!\upharpoonright_{\tau_t(c(v))}, \Delta_q)$.

Let $\hat{t}$ be the tableau obtained from t by replacing each node $v \in \mathcal{F}_{3i}^t$ by $t\!\upharpoonright_{c(v)}$. $\hat{t}$ is a tableau for $(\hat{T}, \varphi_q)$ where $\hat{T}$ is obtained from T' by replacing the sub-tree at each $\tau_t(v) \in \mathcal{F}_{3i}$ by $T'\!\upharpoonright_{\tau_t(c(v))}$. Denote by A the set of nodes of $\hat{t}$ corresponding to this change.

Let $\hat{s}$ be an absolutely minimal tableau for $(\hat{T}, \varphi)$ and $\hat{A} = \{u \in \hat{s} : \exists v \in A \ \tau_{\hat{s}}(u) = \tau_{\hat{t}}(v)\}$. It suffices to prove that $o(\hat{s}) > \alpha = \omega.m_1 + m_2$. Since $(\Delta_{3i}, \Delta_{3i-1}) = (\Delta_{3j}, \Delta_{3j-1})$, lemma 3.14 implies that for every $u \in \hat{A}$ there is a tableau, say t_u, for $(T'\!\upharpoonright_{\tau_{\hat{s}}(u)}, \lambda_{\hat{s}}(u))$ with $o(t_u) \leq \omega.3i$, and $o(t_u) \leq \omega.(3i-1)$ if $o_{\hat{s}}(u) \leq \omega.(3j-1)$. From $\hat{s}$ we define a new tableau

s for (T', φ_q) replacing the sub-tableau $\hat{s}\restriction_u$ by t_u for each $u \in \hat{A}$. We remark that s and $\hat{s}$ have a common initial part, namely the pruning $s_0 = s \cap \{v : v \leq_s \mathcal{F}_{3i}\}$.

Assume $o(\hat{s}) \leq \alpha$. Every ordinal for a filter over (s, s_0) is no greater than $\omega.3i$ by lemma 3.14, so by lemma 3.17, $o(s_0) \geq \omega.(m_1 - 3i) + m_2$. Notice also that $o(s_0) < \omega.(m_1 - 3i) + \omega$. But then $o(\hat{s}) \leq \omega.m_1 + m_2 < \omega.(3i + 1) + o(s_0)$ and lemma 3.16 implies that every ordinal for a filter over $\hat{s}$ is strictly below $\omega.(3i + 2)$. Since $3i + 2 \leq 3j - 1$, in forming s a sub-tableau of order-type $< \omega.(3i + 2)$ at A is replaced by a tableau of order-type $\omega.(3i - 1)$. Therefore every filter over (s, s_0) has ordinal $\leq \omega.(3i - 1)$, whence

$$o(s) \leq \omega.(3i - 1) + o(s_0)$$
$$< \omega.(3i - 1) + \omega.(m_1 - 3i) + \omega \leq \alpha$$

Thus by lemma 3.15 there exists a tableau for (T, φ) with order-type $\beta < \alpha$, yielding a contradiction. ◀

▶ **Corollary 3.19.** Let φ be a closed formula of alternation-free μ-calculus. If φ has closure ordinal $\alpha < \omega^2$ with respect to $\mathcal{X}_\varphi$, in fact $\alpha < \omega.N$ where $N = 2^{2^{|\varphi|+2}}$.

Proof. Suppose $CO_{\mathcal{X}_\varphi}(\varphi) = \alpha \in [\omega.N, \omega^2)$. Proposition 3.10 implies the existence of a TTS T and an absolutely minimal tableau t for $(T, \Box\varphi)$ with order-type α. By lemma 3.18 there exists a TTS $\hat{T} \models \Box\varphi$ and a minimal tableau $\hat{s}$ for $(\hat{T}, \Box\varphi)$ with order-type greater than α, whence lemma 3.6 implies $CO_{\mathcal{X}_\varphi}(\varphi) \geq CO_{\hat{T}, \mathcal{X}_\varphi}(\varphi) > \alpha$. ◀

It remains to rule out closure ordinals of ω^2 or greater. To achieve this a more general version of the argument in the preceding proof is required.

▶ **Lemma 3.20.** If t is a minimal tableau for (T, φ) and $o(t) \geq \omega^2$, then there exists a TTS $\hat{T}$ and a minimal tableau for $(\hat{T}, \varphi)$ with order-type strictly greater than $o(t)$.

Proof. Suppose t is a minimal tableau for (T, φ) and $\omega^2 \leq \omega.\alpha_t \leq o(t) < \omega.(\alpha_t + 1)$. Set $T_0 = \varphi \star T$. Let q not appear in φ and for each $k < \omega$ let the set $\mathcal{F}_k$ be defined analogously to the previous lemma as the collection of nodes in $\varphi \star T$ such that the sub-tree $\{x \in T_0 : q \notin \lambda_{T_0}(y) \wedge y \leq_T x\}$ has order-type $\omega.k$. Now $\mathcal{F}_k$ is non-empty for every $k < \omega$, so there exist infinitely many indices, $0 < i < j(1) < j(2) < \ldots$ such that $j(n + 1) \geq j(n) + 2$ and $(\Delta_i, \Delta_{i-1}) = (\Delta_{j(n)}, \Delta_{j(n)-1})$ for every n. Let $c_m \colon \mathcal{F}_i \to \mathcal{F}_{j(m)}$ be the function such that for each $x \in \mathcal{F}_i$, $\Delta \subseteq \mathrm{FL}(\varphi)$ and every $m < \omega$,

1. there is a tableau for $(T_0\restriction_x, \Delta_q)$ if and only if there is a tableau for $(T_0\restriction_{c_m(x)}, \Delta_q)$,
2. there is a tableau for $(T_0\restriction_x, \Delta_q)$ with order-type $\omega.(i - 1)$ if and only if there is a tableau for $(T_0\restriction_{c_m(x)}, \Delta_q)$ with order-type $\omega.(j(m) - 1)$.

Beginning with c_m, one can define iterated versions, c_m^α for each α: for $i \in \mathcal{F}_k$ with $k \geq i$, $c_m^0(x) = T_0\restriction_x$ and $c_m^1(x)$ is defined to be the result of replacing in $T_0\restriction_x$ each node $y \in \mathcal{F}_i$ by the tree $c_m(y)$; $c_m^{\alpha+1}(x)$ is the tree $c_m^1(x)$ in which each node $y \in \mathcal{F}_i$ is replaced by $c_m^\alpha(y)$; for a limit ordinal α, $c_m^\alpha(x)$ is the tree $c_m^1(x)$ in which, given a bijection $g_0 \colon \mathcal{F}_i \to \omega$,

- if $\alpha = \omega.\gamma + \omega$ then each node $y \in \mathcal{F}_i$ is replaced by the tree $c_{g_0(y)}^{\omega.\gamma+g_0(y)}(y)$,
- if $\alpha = \omega.\alpha_0$, α_0 is a limit ordinal and $g_1 \colon \mathcal{F}_i \to \alpha_0$ is a bijection, then each node $y \in \mathcal{F}_i$ is replaced by the tree $c_{g_0(y)}^{\omega.g_1(y)}(y)$.

The following two lemmata are obtained by generalising the argument in lemma 3.18 making essential use of lemmata 3.16 and 3.17.

▶ **Sub-lemma 1.** There exists a tableau for $(c_m^\alpha(x), \Delta_q)$ if and only if there exists a tableau for $(T_0\restriction_{c_m(x)}, \Delta_q)$.

▶ **Sub-lemma 2.** If $x \in \mathcal{F}_i$ and there exists a tableau for $(c_m^\alpha(x), \Delta_q)$ with order-type $< \omega.\alpha$ then there exists an $\omega.(i-1)$-tableau for $(T_0 \upharpoonright_x, \Delta_q)$.

The construction of the trees $c_m^\alpha(x)$ and the two previous sub-lemmata suffice to prove the main lemma. By lemma 3.15, t naturally induces an absolutely minimal tableau for (T_0, φ_q) of the same order-type. Let $T_0^{\alpha_t}$ be the tree obtained by replacing each sub-tree $T_0 \upharpoonright_y$ for $y \in \mathcal{F}_i$ by $c_i^{\alpha_t + \omega}(y)$. It is easy to see that $T_0^{\alpha_t} \models \varphi \wedge \varphi_q$.

Let $\hat{s}$ be an arbitrary absolutely minimal tableau for $(T_0^{\alpha_t}, \varphi_q)$ and s the collapse of $\hat{s}$ to a tableau for (T_0, φ_q): on each path replace the first $v \in \hat{s}$ such that $T_0^{\alpha_t} \upharpoonright_{\tau_{\hat{s}}(v)} = c_i^{\alpha_t + \omega}(y)$ for some $y \in \mathcal{F}_i$ by the tableau for $(T_0 \upharpoonright_y, \lambda_{\hat{s}}(v))$ given by sub-lemma 2, if $o_{\hat{s}}(v) < \omega.\alpha_t + \omega$, and by sub-lemma 1 otherwise. Let S_0 denote the collection of absolutely minimal tableaux for T_0, and set S_0' to be the collection of tableaux for (T_0, φ_q) that arise as the collapse, in the manner described above, of an absolutely minimal tableau for $(T_0^{\alpha_t}, \varphi_q)$. If there is a minimal tableau for $(T_0^{\alpha_t}, \varphi_q)$ with order-type strictly greater than $o(t)$ then we are done. Otherwise, for every $r' \in S_0'$ there exists $r \in S_0$ such that for all x, if $o_{r'}(x) = \omega.i$ then $o_r(x) > \omega.i$. Now set T_1 to be the S_0'-pruning of T_0. T_1 has the same domain as T_0 and hence T. Moreover, if $\{z \in T_1 : q \notin \lambda_{T_1}(z) \wedge x \leq_T z\}$ has order-type $\omega.i$ then there exists $x <_T y$ such that the tree $\{z \in T_0 : q \notin \lambda_{T_0}(z) \wedge y \leq_T z\}$ has order-type $\omega.i$. Let the set S_1 comprise all absolutely minimal tableaux for (T_1, φ_q). Any $r \in S_1$ is also a tableau for (T_0, φ_q) and hence also for (T, φ). Thus consider tableaux for $(T_1^{\alpha_t}, \varphi_q)$ and set S_1' to be the collection of tableaux that are obtained from the collapse of absolutely minimal tableaux for $(T_1^{\alpha_t}, \varphi_q)$. Define S_2 to be the set of absolutely minimal tableaux for the S_1'-pruning of T_1. Similarly define S_3, S_4, etc. Every tableau in S_{n+1} "moves" the $\omega.i$-frontier of T closer to the root. Thus, either for some n there exists a minimal tableau for $(T_n^{\alpha_t}, \varphi_q)$ with order-type strictly greater than $o(t)$, or for every n there exists $x \in T$ and tableau $s_j \in S_j$ for every $j \leq n$ such that $o_{s_j}(x) < o_{s_{j+1}}(x)$. As the latter will yield a contradiction, we are done. ◀

As a consequence of lemmata 3.18 and 3.20 the closure ordinals of μ-formulæ will be sufficiently bounded.

▶ **Theorem 3.21.** Let $\mathcal{X}$ be a principal set of μ-variables for a closed and alternation-free formula φ. Then the closure ordinal of φ with respect to $\mathcal{X}$, if it exists, is strictly less than $\omega.2^{2^{|\varphi|+2}}$.

▶ **Corollary 3.22.** Suppose φ is a closed formula in the alternation-free fragment of the μ-calculus and $\mathcal{X}$ is a principal set of ν-variables only. Then $CO_{\mathcal{X}}(\varphi) < \omega.2^{2^{|\varphi|+2}}$ if the former ordinal exists.

Proof. Let $\bar{\varphi}$ denote the dual of φ and let $\mathcal{X}$ be a set of ν-variables principal in φ. That $CO_{\mathcal{X}}(\varphi) = CO_{\mathcal{X}}(\bar{\varphi})$ follows from the dual semantics of the μ-calculus, whence theorem 3.21 implies $CO_{\mathcal{X}}(\varphi) < \omega.2^{2^{|\varphi|+2}}$. ◀

▶ **Theorem 3.23.** Let φ be a closed alternation-free formula in guarded form and let $\mathcal{X}$ be the set of variables occurring in φ. If $CO_{\mathcal{X}}(\varphi)$ exists then $CO_{\mathcal{X}}(\varphi) < \omega^2$.

Proof sketch. Suppose $\varphi \in \Sigma_{n+1}$ in the weak hierarchy has closure ordinal κ with respect to the set of all variables in φ. By theorem 3.21 all μ-variables that do not appear under the scope of a ν-variable close off at some ordinal $\alpha < \omega^2$. Moreover, the structure of φ will induce, for each closed weak Π_n sub-formula ψ, a particular class of transition systems, say $\mathcal{T}$, such that ψ has closure ordinal κ with respect to trees in $\mathcal{T}$. In the case $n = 1$, by relativising the previous arguments to the class $\mathcal{T}$, one may deduce ψ has closure ordinal,

say α_ψ, strictly less than ω^2 with respect to $\mathcal{T}$. As the closure ordinal of φ is no greater than the sum of α and ordinals α_ψ, $CO_\mathcal{X}(\varphi) < \omega^2$. ◄

A profound consequence of lemma 3.20 and corollary 3.22 and one that also applies to theorem 3.23, is that there is no essential dependency between closure ordinals and alternation depth for the alternation-free fragment: the choice of N in these results depends only on the logical complexity of φ and the dependency on the alternation depth of φ is essentially trivial, necessitating a smaller increase in bounds than for the connectives and quantifiers. Whether this remains the case for formulæ outside the alternation-free fragment is unclear.

—— **References** ——

1   J.C. Bradfield. The modal mu-calculus alternation hierarchy is strict. *Theoretical Computer Science* 195(2)(1998), 133–153.

2   J.C. Bradfield and C. Stirling. Modal mu-calculi. In: P. Blackburn, J. Van Benthem and F. Wolter eds. *Handbook of Modal Logic*, Studies in Logic and Practical Reasoning 3, Elsevier, 721–756, 2006.

3   M. Czarnecki. How fast can the fixpoints in modal μ-calculus be reached? In: L. Santocanale, ed. *Fixed Points in Computer Science 2010, Brno, August 2010*, 35–39, 2010.

4   O. Friedmann and M. Lange. The modal μ-calculus caught off guard. In: K. Brünnler and G. Metcalfe eds. *Proceedings of the 20th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, Bern, July 2011*, LNCS 6793, 149–163, 2011.

5   O. Kupferman and M.Y. Vardi. $\Pi_2 \cap \Sigma_2 \equiv$ AFMC. In: J.C.M. Baeten, J.K. Lenstra, J. Parrow and G.J. Woeginger, eds. *Proceedings of the 30th International Colloquium on Automata, Languages and Programming, Eindhoven, July 2003*, LNCS 2719, 697–713, 2003.

6   D.E. Muller, A. Saoudi and P.E. Schupp. Alternating automata, the weak monadic theory of the tree, and its complexity. In: L. Kott, ed. *Proceedings of the 13th International Colloquium on Automata, Languages and Programming, Rennes, July 1986*, LNCS 226, 275–283, 1986.

7   D. Niwiński and I. Walukiewicz. Games for the μ-calculus. *Theoretical Computer Science* 163(1–2)(1996), 99–116.