Closure Ordinals of the Two-Way Modal μ-Calculus

Gian Carlo Milanese[1](✉) and Yde Venema[2]

[1] University of Milano-Bicocca, Milan, Italy
giancarlo.milanese@gmail.com
[2] Institute for Logic, Language and Computation, Universiteit van Amsterdam, Amsterdam, The Netherlands
y.venema@uva.nl
https://staff.fnwi.uva.nl/y.venema/

Abstract. The closure ordinal of a μ-calculus formula $\varphi(x)$ is the least ordinal α, if it exists, such that, in any model, the least fixed point of $\varphi(x)$ can be computed in at most α many steps, by iteration of the meaning function associated with $\varphi(x)$, starting from the empty set. In this paper we focus on closure ordinals of the two-way modal μ-calculus. Our main technical contribution is the construction of a two-way formula φ_n with closure ordinal ω^n for an arbitrary $n \in \omega$. Building on this construction, as our main result we define a two-way formula φ_α with closure ordinal α for an arbitrary $\alpha < \omega^\omega$.

Keywords: Modal logic · Fixed points · Closure ordinals · Two-way μ-calculus

1 Introduction

The modal μ-calculus μML, introduced by Kozen [11] in the form known today, is an extension of basic modal logic with explicit least- and greatest fixed point operators. The addition of these operators significantly increases the expressive power of the formalism, enabling it to deal with various forms of *recursion*, as required by applications in for instance the area of program verification. In fact, the modal μ-calculus was shown to be expressively complete with respect to the bisimulation-invariant fragment of monadic second-order logic [10], and it embeds many other logics such as PDL, CTL, and CTL*. Despite this expressive power, the modal μ-calculus has remarkably fine computational properties, such as a quasi-polynomial model checking problem [3] and a satisfiability problem that can be solved in exponential time [6].

In addition, the system admits a nice logical meta-theory: it has the finite model property, uniform interpolation, and a decent model theory [5,8,12]. The set of all valid μ-calculus formulas admits an elegant axiomatisation, which was already introduced by Kozen in his original paper [11], and proved to be complete

R. Iemhoff et al. (Eds.): WoLLIC 2019, LNCS 11541, pp. 498–515, 2019.
https://doi.org/10.1007/978-3-662-59533-6_30

some years later by Walukiewicz [15]. Recently, a cut-free proof system was introduced by Afshari and Leigh [2].

Over the years, the modal μ-calculus has developed into the 'canonical' or 'universal' modal fixed point logic. This status motivates the full development of the meta-logical theory of the logic $\mu\mathtt{ML}$, and of its variants such as the *two-way μ-calculus*, which features both forward- and backward modalities, to be interpreted by the corresponding directions of the model's accessibility relation.

A relatively recent line of research on the modal μ-calculus concerns its *closure ordinals*. For an introduction to this notion, consider a formula $\varphi(x)$ (with only positive occurrences of the variable x) and a Kripke model $\mathbb{S} = (S, R, V)$. We may define a monotone function $\varphi_x^{\mathbb{S}} : \wp(S) \to \wp(S)$, which intuitively expresses how in $\mathbb{S}$ the meaning of φ depends on the valuation of x. The formula $\mu x.\varphi$ is then interpreted in $\mathbb{S}$ as the least fixed point of this map $\varphi_x^{\mathbb{S}}$ – that is, as the least subset $L \subseteq S$ such that $\varphi_x^{\mathbb{S}}(L) = L$ – and the point is that this least fixed point can be 'computed' by performing an iterative process involving the function $\varphi_x^{\mathbb{S}}$. Starting from the empty set, we define the following ordinal-indexed sequence $(\varphi_{\mathbb{S}}^{\alpha})_{\alpha\in\mathrm{On}}$ of subsets of S:

$$\varphi_{\mathbb{S}}^{0} := \emptyset, \quad \varphi_{\mathbb{S}}^{\beta+1} := \varphi_x^{\mathbb{S}}(\varphi_{\mathbb{S}}^{\beta}), \quad \varphi_{\mathbb{S}}^{\lambda} := \bigcup_{\beta<\lambda} \varphi_x^{\mathbb{S}}(\varphi_{\mathbb{S}}^{\beta}),$$

where λ denotes an arbitrary limit ordinal. By monotonicity of the function $\varphi_x^{\mathbb{S}}$, the sequence $(\varphi_{\mathbb{S}}^{\alpha})_{\alpha\in\mathrm{On}}$ converges: there must be a least ordinal α such that $\varphi_{\mathbb{S}}^{\alpha} = \varphi_{\mathbb{S}}^{\alpha+1}$. The element $\varphi_{\mathbb{S}}^{\alpha}$ of the sequence then coincides with the least fixed point of $\varphi_x^{\mathbb{S}}$ so that we say that the function $\varphi_x^{\mathbb{S}}$ converges to its least fixed point in α many steps.

The *closure ordinal* of a formula $\varphi(x)$ is the least ordinal α such that the function $\varphi_x^{\mathbb{S}}$ converges to its least fixed point in at most α many steps across every model $\mathbb{S}$, if such an ordinal exists. In other words, we are interested in the least number of steps that a least fixed point formula needs to converge to its meaning in every model. Not every formula will have a closure ordinal; for instance, let $\mathbb{S}_\alpha = (S_\alpha, >)$ be the structure where S_α is the set of all ordinals smaller than α and the accessibility relation is the converse order relation on these ordinals. It is not hard to see that on this model, the formula $\Box x$ needs exactly α steps to converge to its least fixed point. Clearly then, this formula does not have a closure ordinal across all models.

Intuitively, the closure ordinal of a formula is some measure of its complexity. For instance, a (basic) modal logic formula $\varphi(x)$ has a finite closure ordinal if and only if $\mu x.\varphi$ is definable in (basic) modal logic [14]. Another interesting example is obtained if we involve the first infinite ordinal ω: call a formula $\varphi(x)$ *constructive in* x it has a closure ordinal $\alpha \leq \omega$. The name 'constructive' is taken loosely here, motivated by the observation that since $\varphi_{\mathbb{S}}^{\omega} = \bigcup_{n<\omega} \varphi_{\mathbb{S}}^{n}$ in every model, a formula $\varphi(x)$ is constructive iff for each model $\mathbb{S}$ and for each point s in $\mathbb{S}$ we only need *finitely* many iterations of the map $\varphi_x^{\mathbb{S}}$ in order to find out whether s satisfies the formula $\mu x.\varphi$ or not.

Generally, there are many interesting questions to ask about closure ordinals, and at this moment only few of these have been answered. In fact, it seems that we can summarize our knowledge in one paragraph. Otto [14] proved that it is decidable whether a modal μ-calculus formula can equivalently be expressed in (basic) modal logic. As a corollary, we can also decide whether a formula of modal logic has a finite closure ordinal. Czarnecki [4] showed how to construct a formula φ_α with closure ordinal α for an arbitrary $\alpha < \omega^2$. An interesting result by Afshari and Leigh [1] confirms the intuition that closure ordinals are an indication of the complexity of a formula: they proved that the closure ordinals reached by formulas in the alternation-free fragment of the μ-calculus are all smaller than the ordinal ω^2. Gouveia and Santocanale [9] presented a formula with closure ordinal ω_1 and proved that closure ordinals are closed under ordinal sum.

In this paper we contribute to the theory of closure ordinals by taking a look at the two-way modal μ-calculus. After recalling the syntax and semantics of the logic and providing some definitions concerning closure ordinals in Sect. 2, in the following section we show how to define a formula φ_n with closure ordinal ω^n for every $n \in \omega$ (some of the technical proofs of this section are delayed to the appendix of the paper). In Sect. 4 we build on this result by proving that every ordinal smaller than ω^ω is a closure ordinal in the two-way setting. One way to achieve this is via transferring a result by Gouveia and Santocanale [9] – stating that the class of closure ordinals is closed under taking ordinal sum – to the two-way setting. We also define, given a representation $\alpha = \omega^n \cdot k_1 + \omega^{n-1} \cdot k_2 + \ldots + \omega \cdot k_n + k_{n+1}$ of an arbitrary ordinal $\alpha < \omega^\omega$, an explicit formula φ_α with closure ordinal α. We finish the paper with mentioning some questions for further research.

Source. The results in this paper are taken from the MSc thesis [13], which was written by the first author under the supervision of the second.

2 Preliminaries

Definition 1. *The language* μ`TML` *of the two-way modal* μ-calculus *is given by the following grammar:*

$$\varphi ::= \bot \mid p \mid \neg\varphi \mid \varphi \vee \varphi \mid F\varphi \mid P\varphi \mid \mu x.\varphi$$

where $p, x \in$ `PROP` *and the formation of the formula* $\mu x.\varphi$ *is subject to the constraint that the variable* x *is positive in* φ*, that is, every occurrence of* x *in* φ *is under the scope of an even number of negations.*

We can define $\top$, $\wedge$ and the box operators by letting $G\varphi := \neg F\neg\varphi$ and $H\varphi := \neg P\neg\varphi$, as well as the greatest fixed point operator as $\nu x.\varphi := \neg\mu x.\neg\varphi(\neg x)$. The intended interpretation of a formula $F\varphi$ is 'φ is true at some (one-step) *future* state', while that of $P\varphi$ is 'φ is true at some (one-step) *past* state'.

Formulas of this language will be interpreted in two-way models. These can be defined as Kripke models featuring a pair of accessibility relations that are each other's converse, where we recall that the *converse* of a relation R is the relation $R^{-1} := \{(s,t) \mid (t,s) \in R\}$. It will be more convenient to simply identify two-way models with standard Kripke models with one single relation, and make sure that the diamonds F and P access this relation in its two different directions.

Definition 2. *A Kripke model is a triple $\mathbb{S} = (S, R, V)$ where S, the domain or underlying set, is a set of points or states, R is a binary relation on S, and V is a valuation on S, that is, a function $V : \mathtt{PROP} \to \wp(S)$.*

Given a model $\mathbb{S} = (S, R, V)$, a propositional variable x and a subset $X \subseteq S$, we define $V[x \mapsto X]$ as the valuation given by $V[x \mapsto X](p) = X$ if $p = x$, and $V[x \mapsto X](p) = V(p)$ otherwise. We denote the model $(S, R, V[x \mapsto X])$ by $\mathbb{S}[x \mapsto X]$.

Given a subset $S' \subseteq S$, the submodel of $\mathbb{S}$ induced by S' *is the model $\mathbb{S}' = (S', R', V')$, where $R' = R \cap (S' \times S')$ and $V'(p) = V(p) \cap S'$ for all $p \in \mathtt{PROP}$.*

We now inductively define the meaning of a formula φ in a model $\mathbb{S}$ as the set of states where this formula is true, or satisfied. At the same time we define the function $\varphi_x^{\mathbb{S}}$, which intuitively expresses how in $\mathbb{S}$ the meaning of the formula φ varies depending on the meaning of the variable x.

Definition 3. *Given a $\mu\mathtt{TML}$-formula φ and a model $\mathbb{S} = (S, R, V)$, we define the meaning $[\![\varphi]\!]^{\mathbb{S}}$ of φ in $\mathbb{S}$, together with the function $\varphi_x^{\mathbb{S}} : \wp(S) \to \wp(S)$ mapping a subset $X \subseteq S$ to $[\![\varphi]\!]^{\mathbb{S}[x \mapsto X]}$, by the following simultaneous induction:*

$$
\begin{array}{llll}
[\![\bot]\!]^{\mathbb{S}} & = \emptyset, & [\![p]\!]^{\mathbb{S}} & = V(p), \\
[\![\varphi \vee \psi]\!]^{\mathbb{S}} & = [\![\varphi]\!]^{\mathbb{S}} \cup [\![\psi]\!]^{\mathbb{S}}, & [\![\neg\varphi]\!]^{\mathbb{S}} & = S \backslash [\![\varphi]\!]^{\mathbb{S}}, \\
[\![F\varphi]\!]^{\mathbb{S}} & = \{s \in S \mid R[s] \cap [\![\varphi]\!]^{\mathbb{S}} \neq \emptyset\}, & [\![P\varphi]\!]^{\mathbb{S}} & = \{s \in S \mid R^{-1}[s] \cap [\![\varphi]\!]^{\mathbb{S}} \neq \emptyset\}, \\
[\![\mu x.\varphi]\!]^{\mathbb{S}} & = \bigcap\{U \subseteq S \mid \varphi_x^{\mathbb{S}}(U) \subseteq U\}, & &
\end{array}
$$

where $R[s] := \{t \in S \mid (s,t) \in R\}$ and similarly for R^{-1}. For an element $s \in S$ we write $\mathbb{S}, s \Vdash \varphi$ if $s \in [\![\varphi]\!]^{\mathbb{S}}$.

Let $\varphi \in \mu\mathtt{TML}$ be a formula in which the variable x occurs only positively and let $\mathbb{S}$ be a model. By induction on φ one can prove that $\varphi_x^{\mathbb{S}} : \wp(S) \to \wp(S)$ is a monotone operation. Consequently, by the Knaster-Tarski theorem we obtain that $[\![\mu x.\varphi]\!]^{\mathbb{S}}$ is the least fixed point of $\varphi_x^{\mathbb{S}}$, denoted by $\mathrm{LFP}.\varphi_x^{\mathbb{S}}$. As we saw in the introduction, the meaning of $\mu x.\varphi$ in a model $\mathbb{S}$ can also be computed by performing an iteration of the function $\varphi_x^{\mathbb{S}}$ starting from the empty set, resulting in the ordinal-indexed sequence $(\varphi_{\mathbb{S}}^{\alpha})_{\alpha \in \mathrm{On}}$. When the model $\mathbb{S}$ is clear from context we will write φ^{α} instead of $\varphi_{\mathbb{S}}^{\alpha}$; we shall also exclusively take x as the fixed point variable of the formulas that we are looking at, so that we need not mention this explicitly in the sequel.

In this paper we are interested in the number of times we need to iterate the function $\varphi_{\mathbb{S}}^{\alpha}$ before we reach its least fixed point.

Definition 4. *Let $\varphi(x)$ be a formula which is positive in x. Then for a Kripke model $\mathbb{S}$, we let $\gamma_x(\varphi, \mathbb{S})$ denote the closure ordinal of φ in $\mathbb{S}$ with respect to x, that is, the least ordinal α such that $\varphi^{\alpha}_{\mathbb{S}} = \varphi^{\alpha+1}_{\mathbb{S}}$.*

The closure ordinal of φ with respect to the variable x is the least ordinal α such that $\gamma_x(\varphi, \mathbb{S}) \leq \alpha$ for every model $\mathbb{S}$, if it exists. If α is the closure ordinal of some (two-way) formula, we say that α is a (two-way) closure ordinal.

When proving results about closure ordinals an equivalent characterisation, given in Proposition 1, is often useful.

Proposition 1. *An ordinal α is the closure ordinal of $\varphi(x)$ if and only if (1) $\gamma_x(\varphi, \mathbb{S}) \leq \alpha$ for every model $\mathbb{S}$ and (2) $\gamma_x(\varphi, \mathbb{S}) = \alpha$ for some model $\mathbb{S}$.*

Proof. The only nontrivial observation in the proof concerns the case, in the direction from left to right, where the closure ordinal α of φ is a limit ordinal. In order to prove (2), let B be the set of ordinals $\beta < \alpha$ for which there is a model $\mathbb{S}_\beta$ with $\gamma_x(\varphi, \mathbb{S}_\beta) = \beta$. This set must be cofinal in α, and it is then easy to show that if we take $\mathbb{S}$ to be the disjoint union of the collection $\{\mathbb{S}_\beta \mid \beta \in B\}$, we find $\gamma_x(\varphi, \mathbb{S}) = \alpha$ as required.

Example 1. The closure ordinal of $\varphi := (G\bot \vee Fx)$ is ω. It is not hard to prove that $\gamma_x(\varphi, \mathbb{S}) \leq \omega$ for every model $\mathbb{S}$, and that φ converges to its least fixed point in exactly ω steps in the model $\mathbb{S}$ depicted in Fig. 1. Indeed, one can show that $\varphi^n = \{m \in \omega \mid m < n\}$ for all $n \in \omega$: the iteration of φ in $\mathbb{S}$ traverses the ordinal ω by adding each finite ordinal to the iteration, one by one. After ω many steps in the iteration we observe that $\varphi^\omega = \bigcup_{n<\omega} \varphi^n = \{0, 1, 2, \ldots\} = \omega$ and $\varphi^{\omega+1} = \varphi^{\mathbb{S}}_x(\varphi^\omega) = \varphi^\omega$, so that the iteration converges in exactly ω steps.

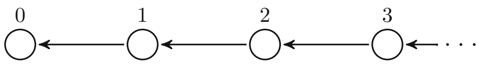

Fig. 1. Model where $G\bot \vee Fx$ converges in ω many steps

3 Two-Way Formulas: Closure Ordinal ω^n

In this section we define a two-way formula φ_n with closure ordinal ω^n for an arbitrary $n \in \omega$. We first need to define *colours*, which are essentially conjunctions of literals as specified in the next definition.

Definition 5. *Fix a subset $\{q_i \mid i \in \omega\}$ of propositional variables. For every $0 < n < \omega$ we define the colour c_n as the conjunction of literals $c_n := \bigwedge_{0<i<n} \neg q_i \wedge q_n$.*

For example, $c_1 = q_1$ and $c_2 = \neg q_1 \wedge q_2$. Clearly, $c_i \wedge c_j \equiv \bot$ for every $i \neq j$.

We now define, for all $0 < n < \omega$, a two-way formula φ_n.

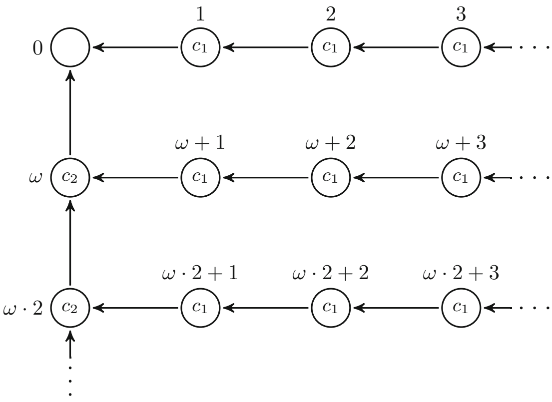

Fig. 2. Model corresponding to ω^2

Definition 6. *By induction on $i \in \omega$ we define the formulas π_i^∞ as follows:*

$$\begin{aligned}\pi_0^\infty &:= \top,\\ \pi_{i+1}^\infty &:= \nu y_{i+1}.(P(y_{i+1} \wedge c_{i+1}) \wedge \pi_i^\infty).\end{aligned}$$

For all $n \in \omega$ let φ_n be the formula

$$\varphi_n := G\bot \vee (c_1 \wedge Fx) \vee \bigvee_{i=2}^{n} (c_i \wedge \pi_{i-1}^\infty \wedge F(\nu y.P(y \wedge x \wedge c_{i-1}))).$$

Example 2. Consider the formula

$$\varphi_2 = G\bot \vee (c_1 \wedge Fx) \vee (c_2 \wedge \pi_1^\infty \wedge F(\nu y.P(y \wedge x \wedge c_1)))$$

and the model $\mathbb{S}$ depicted in Fig. 2, consisting of ω many copies of ω, thus intuitively corresponding to the ordinal ω^2.

The formula φ_2 crucially involves the formula $\nu y.P(y \wedge x \wedge c_1)$, which expresses the existence of an infinite R^{-1}-path of points where x and c_1 are always true starting from the R^{-1}-next state, and which allows the iteration to move from a copy of ω to the next, as we shall now see. The iteration of φ_2 in this model starts similarly as the one in Example 1, by including the state 0 through the disjunct $G\bot$ and then adding, one by one, each state labelled with a finite ordinal through the disjunct $(c_1 \wedge Fx)$. After ω many steps in the iteration we have $\varphi^\omega = \{0, 1, 2, \dots\} = \omega$, so that every state labelled with a finite ordinal is inside the iteration. Now it holds that $\mathbb{S}[x \mapsto \varphi^\omega], \omega \Vdash (c_2 \wedge \pi_1^\infty \wedge F(\nu y.P(y \wedge x \wedge c_1)))$, so that $\varphi^{\omega+1} = [\![\varphi]\!]^{\mathbb{S}[x \mapsto \varphi^\omega]} = \omega \cup \{\omega\}$: the state ω is added to the iteration. The iteration continues through the disjunct $(c_1 \wedge Fx)$, with $\varphi^{\omega+n} = \omega \cup \{\omega, \dots, \omega + (n-1)\}$, arriving at $\varphi^{\omega\cdot 2} = \omega \cup \{\omega, \omega+1, \omega+2, \dots\}$, at which point the state $\omega \cdot 2$ will satisfy $(c_2 \wedge \pi_1^\infty \wedge F(\nu y.P(y \wedge x \wedge c_1)))$, and so on. The iteration will progress in a similar way, traversing all the copies of ω and converging in exactly ω^2 steps.

The following example concerns the formulas of shape π_i^∞ that appear as subformulas of φ_n. These formulas will make sure that the models of φ_n have a particular grid-like shape: we will need that whenever a state s in a model makes φ_n true and has colour c_i, then this state is the starting point of an infinite R^{-1}-path of points where c_{i-1} is always true, and from every point on this path an infinite R^{-1}-path starts of points where c_{i-2} is always true, and so on.

Example 3. Consider for instance

$$\pi_3^\infty = \nu y_3.(P(y_3 \wedge c_3) \wedge \nu y_2.(P(y_2 \wedge c_2) \wedge \nu y_1.(P(y_1 \wedge c_1) \wedge \top))).$$

This formula expresses the existence of an infinite R^{-1}-path $t_0 t_1 t_2 \ldots$ such that *(i)* c_3 is true at every t_i with $i > 0$; *(ii)* every t_i makes $\nu y_2.(P(y_2 \wedge c_2) \wedge \nu y_1.P(y_1 \wedge c_1))$ true, so from each t_i there is an infinite R^{-1}-path $u_0 u_1 u_2 \ldots$ where $u_0 = t_i$ and c_2 is true at every u_j with $j > 0$; *(iii)* every u_j makes $\nu y_1.P(y_1 \wedge c_1)$ true, so from each u_j there exists a R^{-1}-path $v_0 v_1 \ldots$, with $v_0 = u_j$, such that c_1 is true at v_k for every $k > 0$. For example, the point 0 in the model of Fig. 3 makes π_3^∞ true (as does every state of the form $\omega^2 \cdot n$ for $n \in \omega$).

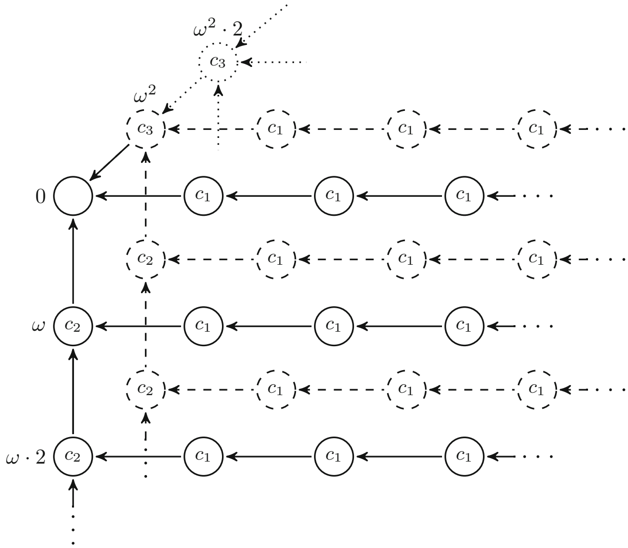

Fig. 3. Model corresponding to ω^3

Example 4. As a further example, consider the formula

$$\varphi_3 = \varphi_2 \vee (c_3 \wedge \pi_2^\infty \wedge F(\nu y.P(y \wedge x \wedge c_2)))$$

and the model pictured in Fig. 3, which consists of ω many copies of the model from Fig. 2, each attached to a state of the infinite path $0R^{-1}\omega^2 R^{-1}\omega^2 \cdot 2 \ldots$.

The iteration of φ_3 in this model starts similarly as the one in Example 2, but after ω^2 many steps, when the first copy of ω^2 is inside the approximating set φ^{ω^2}, the state ω^2 will satisfy the disjunct $(c_3 \wedge \pi_2^\infty \wedge F(\nu y.P(y \wedge x \wedge c_2)))$ of φ_3, so that the iteration will move to the second copy of ω^2 and continue in an analogous way, with convergence in exactly ω^3 steps.

The last example also suggests a recipe for constructing a model where the formula φ_n converges in exactly ω^n steps. For $n = 4$, we could take an infinite R^{-1}-chain of c_4-states, where to each such state is attached a copy of ω^3 (that is, a copy of the model of Fig. 3): the disjunct $(c_4 \wedge \pi_3^\infty \wedge F(\nu y.P(y \wedge x \wedge c_3)))$ of the formula φ_4 would allow the iteration to move between the copies of ω^3, exactly as the disjunct $(c_3 \wedge \pi_2^\infty \wedge F(\nu y.P(y \wedge x \wedge c_2)))$ of φ_3 allowed the iteration to move from a copy of ω^2 to the next. For $n = 5$ we could consider an infinite R^{-1}-chain of c_5-states, where to each such state is attached a copy of the model we have just described, and so on.

Lemma 1. *Let $0 < n < \omega$ be a finite ordinal. Then there is a model $\mathbb{S}$ where $\gamma_x(\varphi_n, \mathbb{S}) = \omega^n$.*

Up to this point we have only focused on one of the two conditions that the ordinal ω^n must satisfy in order to qualify as the closure ordinal of φ_n, namely the one concerning convergence in exactly ω^n steps in some model. It is less intuitive to see why ω^n should be an upper bound for the number of steps in the iteration of φ_n in an arbitrary model. Indeed, the previous models present a particular grid-like structure, which allows the iteration to progress in a very controlled way: if a state is added to the iteration at some step α, then the chain of c_1-states attached to it is included in the iteration in at most ω more steps, the chain of c_2-states attached to it is included in the iteration in at most ω^2 more steps, and the chain of c_3-states attached to it is included in the iteration in at most ω^3 more steps (in case these chains exist). This is formulated in a more general way in the next lemma, which states that if we have an infinite R^{-1}-path of c_i-states presenting the desired grid-like structure (that is, each satisfying π_{i-1}^∞) and the first state of this path belongs to the approximating set φ_n^α, then all the states of the path will be inside the iteration after at most ω^i more steps. Put differently, if a state t_0 in a model satisfies π_i^∞ and is in the approximating set φ_n^α, then all the states forming the R^{-1}-path that witnesses the truth of π_i^∞ at t_0 will belong to $\varphi_n^{\alpha+\omega^i}$.

Lemma 2. *Let $\mathbb{S} = (S, R, V)$ be a model and let $n \in \omega$. For $1 \leq i \leq n$, let $t_0 t_1 t_2 \ldots$ be an infinite R^{-1}-path such that*

$$\mathbb{S}, t_0 \Vdash \pi_{i-1}^\infty \text{ and, for all } j > 0, \mathbb{S}, t_j \Vdash c_i \wedge \pi_{i-1}^\infty.$$

Then, for any ordinal α: if $t_0 \in \varphi_n^\alpha$ then $t_j \in \varphi_n^{\alpha+\omega^{i-1}\cdot j+1}$ for all $j \in \omega$.

In order to make sure that something similar also happens in an arbitrary model, the presence in φ_n of the subformulas π_{i+1}^∞'s from Definition 6 is necessary: by the definition of φ_n, if a state s in a model $\mathbb{S}$ satisfies $(\varphi_n \wedge c_i \wedge F\top)$, then it must also satisfy π_{i-1}^{∞}, so that the model $\mathbb{S}$ will present the desired grid-like structure. This fact and Lemma 2 are essential for proving that indeed φ_n converges to its least fixed point in at most ω^n steps in every model.

Lemma 3. *For an arbitrary model $\mathbb{S}$ and $0 < n < \omega$: $\gamma_x(\varphi_n, \mathbb{S}) \leq \omega^n$.*

By Proposition 1, and the Lemmas 1 and 3, the following is immediate.

Theorem 1. *For all $0 < n < \omega$, the two-way closure ordinal of $\varphi_n(x)$ is ω^n.*

The proofs of all the statements of this section can be found in the appendix.

4 Two-Way Formulas: Closure Ordinals Below ω^ω

This section is devoted to the main result of our paper, stating that every ordinal below ω^ω is a closure ordinal in the two-way setting. In the next subsection we transfer a result by Gouveia and Santocanale [9] to the two-way setting. That is, we show that for two-way formulas $\varphi_0(x)$ and $\varphi_1(x)$ with closure ordinals α_0 and α_1, respectively, we can define a two-way formula $\psi(x)$ with closure ordinal $\alpha_0 + \alpha_1$. From this observation and Theorem 1, our main result follows, since every ordinal α below ω^ω can be written as a finite sum of ordinals of the form ω^n. In the following subsection we improve on this result by defining, for an arbitrary ordinal $\alpha < \omega^\omega$, an explicit two-way formula φ_α with closure ordinal α.

4.1 Two-Way Formulas: Sum of Ordinals

In the introduction we already mentioned that Gouveia and Santocanale showed the class of closure ordinals to be closed under taking ordinal sums [9]. We will now see that their result also holds in the two-way setting.

Theorem 2. *There is an effective construction transforming a pair of two-way formulas $\varphi_0(x)$ and $\varphi_1(x)$ into a formula ψ such that, if $\varphi_0(x)$ and $\varphi_1(x)$ have closure ordinals α_0 and α_1, respectively, then $\psi(x)$ has closure ordinal $\alpha_0 + \alpha_1$.*

Our proof follows the approach from [9], but we provide some proof details here in order to keep our presentation self-contained, and because we can make some simplifications in the two-way setting. We confine ourselves to a proof sketch, focusing on intuitions rather than on technicalities. One concept we will need is that of a *strong* closure ordinal.

Definition 7. *An ordinal α is a* strong *closure ordinal for a (two-way) μ-calculus formula $\varphi(x)$ if $\gamma(\varphi, \mathbb{S}) \leq \alpha$ for all models $\mathbb{S}$, while there is a model $\mathbb{S} = (S, R, V)$ such that*

$$S = [\![\mu x.\varphi]\!]^{\mathbb{S}} = \varphi_{\mathbb{S}}^{\alpha} \text{ and } \varphi_{\mathbb{S}}^{\alpha} \neq \varphi_{\mathbb{S}}^{\beta} \text{ for every } \beta < \alpha.$$

Proposition 2. *If α is the two-way closure ordinal of some formula $\varphi(x)$, then it is a strong closure ordinal for the formula $\widehat{\varphi}(x) := (\mu x.\varphi) \rightarrow \varphi(x \wedge \mu x.\varphi)$.*

Proof. As in [9] the key observation is that in any model $\mathbb{S} = (S, R, V)$ we have

$$\widehat{\varphi}^{\gamma}_{\mathbb{S}} = (S \setminus \text{LFP}.\varphi^{\mathbb{S}}_{x}) \cup \varphi^{\gamma}_{\mathbb{S}},$$

for any $\gamma \geq 1$—this claim can be proved by a straightforward transfinite induction. Consequently, for $\gamma = \alpha$ we obtain $\widehat{\varphi}^{\alpha}_{\mathbb{S}} = (S \setminus \text{LFP}.\varphi^{\mathbb{S}}_{x}) \cup \varphi^{\alpha}_{\mathbb{S}} = (S \setminus \text{LFP}.\varphi^{\mathbb{S}}_{x}) \cup \text{LFP}.\varphi^{\mathbb{S}}_{x} = S$.

We now turn to the proof of Theorem 2. Throughout this subsection we let $\varphi_0(x)$ and $\varphi_1(x)$ be two-way formulas with closure ordinals α_0 and α_1, respectively. Our aim is to define a two-way formula $\psi(x)$ with closure ordinal $\alpha_0 + \alpha_1$. Because of Proposition 2 we may without loss of generality assume that α_0 is a *strong* closure ordinal for φ_0.

The idea underlying the definition of $\psi(x)$ is that in any model $\mathbb{S}$, in order to calculate the least fixed point of $\psi(x)$, one may first focus on φ_0 and then move on to φ_1. More precisely, with each model $\mathbb{S} = (S, R, V)$ we will associate two submodels $\mathbb{S}_0$ and $\mathbb{S}_1$ such that

$$\gamma(\psi, \mathbb{S}) \leq \gamma(\varphi_0, \mathbb{S}_0) + \gamma(\varphi_1, \mathbb{S}_1). \tag{1}$$

This implies that ψ has a closure ordinal β indeed, and that $\beta \leq \alpha_0 + \alpha_1$. To prove that $\beta \geq \alpha_0 + \alpha_1$ we will employ a special model $\mathbb{S}$ such that, for each i, $\mathbb{S}_i$ is a model witnessing that α_i is a strong closure ordinal for φ_i.

For the details of the construction of the submodels $\mathbb{S}_0$ and $\mathbb{S}_1$, note that the formula ψ will use one fresh variable p (so that in particular, p occurs neither in φ_0 nor in φ_1), and write $\texttt{PROP}_p = \texttt{PROP} \cup \{p\}$. Now, given a $\texttt{PROP}_p$-model $\mathbb{S} = (S, R, V)$, we define $S_0 = S \setminus V(p)$ and $S_1 = V(p)$, and for $i = 0, 1$ let $\mathbb{S}_i$ be the submodel of $\mathbb{S}$ induced by the set S_i (and with V_i restricted to the set $\texttt{PROP}$).

Syntactically, we need the following definition.

Definition 8. *Let $p \notin \texttt{PROP}$ be a fresh variable and set $\texttt{p}_0 := \neg p$ and $\texttt{p}_1 := p$. For $i \in \{0, 1\}$ we define the* restriction *of φ to $\texttt{p}_i$ as follows:*

$\texttt{tr}_i(y)$	$:= \texttt{p}_i \wedge y$	$\texttt{tr}_i(\psi_0 \wedge \psi_1)$	$:= \texttt{tr}_i(\psi_0) \wedge \texttt{tr}_i(\psi_1)$
$\texttt{tr}_i(\neg y)$	$:= \texttt{p}_i \wedge \neg y$	$\texttt{tr}_i(\psi_0 \vee \psi_1)$	$:= \texttt{tr}_i(\psi_0) \vee \texttt{tr}_i(\psi_1)$
$\texttt{tr}_i(\bot)$	$:= \bot$	$\texttt{tr}_i(F\psi)$	$:= \texttt{p}_i \wedge F(\texttt{p}_i \wedge \texttt{tr}_i(\psi))$
$\texttt{tr}_i(\top)$	$:= \texttt{p}_i$	$\texttt{tr}_i(G\psi)$	$:= \texttt{p}_i \wedge G(\texttt{p}_i \to \texttt{tr}_i(\psi))$
$\texttt{tr}_i(\mu z.\psi)$	$:= \mu z.\texttt{tr}_i(\psi)$	$\texttt{tr}_i(P\psi)$	$:= \texttt{p}_i \wedge P(\texttt{p}_i \wedge \texttt{tr}_i(\psi))$
$\texttt{tr}_i(\nu z.\psi)$	$:= \nu z.\texttt{tr}_i(\psi)$	$\texttt{tr}_i(H\psi)$	$:= \texttt{p}_i \wedge H(\texttt{p}_i \to \texttt{tr}_i(\psi))$

We need the following properties of these restriction formulas.

Proposition 3. *Let $\varphi(x)$ be a formula in the two-way μ-calculus and let $\mathbb{S} = (S, R, V)$ be an arbitrary model. Then for $i = 0, 1$ we have*

1. $[\![\texttt{tr}_i(\varphi)]\!]^{\mathbb{S}} = [\![\varphi]\!]^{\mathbb{S}_i}$
2. *with x free in φ, $(\texttt{tr}_i(\varphi))^{\alpha}_{\mathbb{S}} = \varphi^{\alpha}_{\mathbb{S}_i}$, for every ordinal α.*

We are now ready for the definition of the formula $\psi(x)$. Consider the following formulas (which are somewhat simpler than the corresponding one-way formulas of [9]):

$\psi_0 \quad := \neg p \wedge \mathtt{tr}_0(\varphi_0)(x)$
$\psi_1 \quad := p \wedge \mathtt{tr}_1(\varphi_1)(x) \wedge G(\neg p \to x)$
$\psi(x) := \psi_0(x) \vee \psi_1(x)$.

To compute the least fixed point of the formula $\psi(x)$ on an arbitrary model $\mathbb{S}$, first consider its disjunct $\psi_0(x) = \neg p \wedge \mathtt{tr}_0(\varphi_0)(x)$. By Proposition 3 we may think of the computation of its least fixed point as taking place in the $\neg p$-part S_0 of $\mathbb{S}$, parallel to that of $\mu x.\varphi_0$ in $\mathbb{S}_0$, and so this computation finishes after $\gamma(\varphi_0, \mathbb{S}_0)$ steps. Similarly, the iterative process approximating the least fixed point of the formula $\psi_1' := p \wedge \mathtt{tr}_0(\varphi_1)(x)$ can be fully executed in the p-part S_1 of $\mathbb{S}$, and this computation would finish after $\gamma(\varphi_1, \mathbb{S}_1)$ steps. The formula $\psi_1(x)$, however, has an additional conjunct, viz., the formula $G(\neg p \to x)$; this ensures that a point in S_1 will only be included in an approximating set $\psi^{\alpha+1}$ if each of its successors in S_0 has been included in the set ψ^{α}. As a consequence, the computation of the S_1-part of the least fixed point of $\psi(x)$ need not be (fully) operational before the computation of the S_0-part is completed. Nevertheless, once the latter computation has terminated indeed, the conjunct $G(\neg p \to x)$ evaluates to true in every state in S_1, and so from that moment on at most $\gamma(\varphi_1, \mathbb{S}_1)$ steps are needed to finish the computation of $[\![\mu x.\psi]\!]^{\mathbb{S}}$. This finishes a proof sketch of the statement (1).

It remains to provide a model $\mathbb{S}$ where the closure ordinal of $\psi(x)$ is actually identical to $\alpha_0 + \alpha_1$. For this purpose, consider two models $\mathbb{S}_0$ and $\mathbb{S}_1$ such that $\gamma(\varphi_i, \mathbb{S}_i) = \alpha_i$ for $i = 0, 1$. Additionally, we require that $[\![\mu x.\varphi_0]\!]^{\mathbb{S}_0} = S_0$—such a model exists by our assumption that α_0 is a *strong* closure ordinal for φ_0. Now take the disjoint union of $\mathbb{S}_0$ and $\mathbb{S}_1$, add an arrow from every state of S_1 to every state of S_0, and set $V(p) := S_1$. Call the resulting model $\mathbb{S}$; it is easy to see that this definition does not cause notational confusion, since the models $\mathbb{S}_0$ and $\mathbb{S}_1$ are identical to the submodels of $\mathbb{S}$ induced by the sets $S_0 = [\![\neg p]\!]^{\mathbb{S}}$ and $S_1 = [\![p]\!]^{\mathbb{S}}$, respectively. The crux of this construction is that in the model $\mathbb{S}$, because every state s in S_1 has the *full* set S_0 among its successors, and we need *exactly* α_0 steps to get all S_0-points in the least fixed point of ψ, we can only start adding S_1-states to the least fixed point of $\psi(x)$ after we have added *all* S_0-states, that is, at stage $\alpha_0 + 1$. It is then not hard to see that another α_1 steps are needed to include all S_1-states, so that all in all we need exactly $\alpha_0 + \alpha_1$ steps for $\psi(x)$ to converge. This shows that the closure ordinal of the formula $\psi(x)$ is $\alpha_0 + \alpha_1$ indeed.

4.2 An Explicit Formula for Every Ordinal Below ω^ω

In this section we shall provide, for every ordinal $\alpha < \omega^\omega$, an explicit two-way formula φ_α with closure ordinal α.

In the case α is *finite*, it is not hard to see that n is the closure ordinal of $\varphi_n := (Gx \wedge G^n\bot)$ for every $n \in \omega$, so that in the sequel we confine attention to the infinite case. Recall that every ordinal α with $\omega \leq \alpha < \omega^\omega$ can be written in a unique normal form

$$\alpha = \omega^n \cdot k_1 + \omega^{n-1} \cdot k_2 + \ldots + \omega \cdot k_n + k_{n+1} \tag{2}$$

for some finite ordinals $n, k_1, \ldots, k_{n+1}$ with $n, k_1 > 0$. We may then use the Theorems 1 and 2 to construct, for every such ordinal α, an explicit two-way formula taking α as its closure ordinal.

Alternatively, in Definition 11 below we provide a different two-way formula φ_α with closure ordinal α; this definition is directly based on the normal form (2). In order to achieve this, we need to define a second set of colours.

Definition 9. *Fix a subset $\{p_i \mid i \in \omega\}$ of propositional variables that is disjoint from the set $\{q_i \mid i \in \omega\}$ from Definition 5. For every $0 < n < \omega$ we define the colour f_n as the conjunction of literals $f_n := \bigwedge_{0<i<n} \neg p_i \wedge p_n$.*

Definition 10. *For every $i, k \in \omega$ we define a formula $\pi^\infty_{i,k}$ inductively on i as follows:*

$$\begin{aligned} \pi^\infty_{0,k} &:= f_k \\ \pi^\infty_{i+1,k} &:= \nu y_{i+1}.(P(y_{i+1} \wedge c_{i+1} \wedge f_k \wedge G f_k) \wedge \pi^\infty_i). \end{aligned}$$

We finally state the definition of the formula φ_α.

Definition 11. *For $n, k \in \omega$ define the formulas*

$$\begin{aligned} \varphi_{(n,k)} &:= (Fx \wedge c_1 \wedge f_k \wedge G f_k) \vee \\ &\quad \bigvee_{i=2}^{n} (c_i \wedge f_k \wedge G f_k \wedge \pi^\infty_{i-1,k} \wedge F(\nu y. f_k \wedge P(y \wedge x \wedge G f_k \wedge c_{i-1}))), \\ \chi_k &:= (Gx \wedge f_{k+1} \wedge G f_k). \end{aligned}$$

Now let, for $n > 0$, $\alpha = \omega^n \cdot k_1 + \omega^{n-1} \cdot k_2 + \ldots + \omega \cdot k_n + k_{n+1}$. For all $0 \leq m \leq n$ define $k(\overrightarrow{m}) := \sum_{i=0}^{m} k_i$, where we let $k_0 := 0$. The formula φ_α is defined by letting

$$\begin{aligned} \psi &:= \bigvee_{i=0}^{k_{n+1}-1} (Gx \wedge \bigwedge_{j=0}^{i} G^j f_{k(\overrightarrow{n})+1} \wedge G^{i+1} f_{k(\overrightarrow{n})}), \\ \varphi_\alpha &:= G\bot \vee \bigvee_{k=1}^{k(\overrightarrow{n})-1} \chi_k \vee \bigvee_{m=0}^{n-1} \left(\bigvee_{k=k(\overrightarrow{m})+1}^{k(\overrightarrow{m+1})} \varphi_{(n-m,k)} \right) \vee \psi. \end{aligned}$$

Example 5. Consider the formulas

$$\begin{aligned} \varphi_{(2,i)} &:= (Fx \wedge c_1 \wedge f_i \wedge G f_i) \vee \\ &\quad\quad (c_2 \wedge f_i \wedge G f_i \wedge \pi^\infty_{1,i} \wedge F(\nu y. f_i \wedge P(y \wedge x \wedge G f_i \wedge c_1))), \\ \varphi_{\omega^2 \cdot 2} &:= G\bot \vee (Gx \wedge f_2 \wedge G f_1) \vee \varphi_{(2,1)} \vee \varphi_{(2,2)} \end{aligned}$$

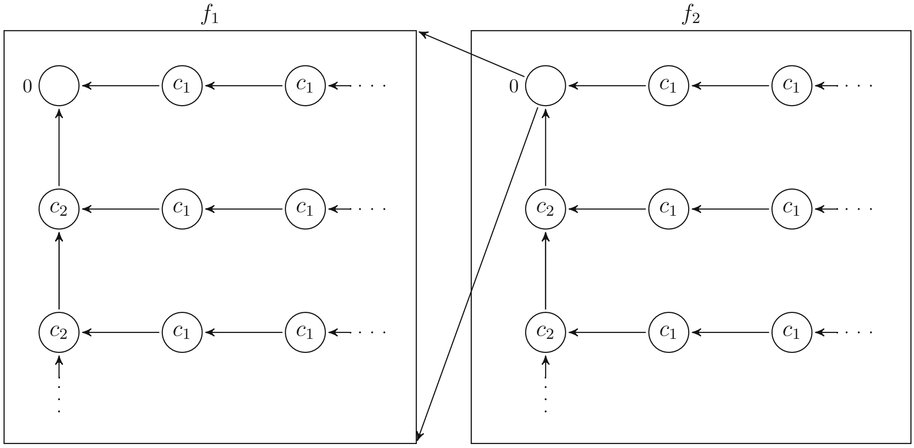

Fig. 4. Model corresponding to $\omega^2 \cdot 2$

and a model $\mathbb{S}$ consisting of two submodels $\mathbb{S}_1$ and $\mathbb{S}_2$, both copies of the model of Fig. 2, such that from the point corresponding to 0 in $\mathbb{S}_2$ there is an arrow to every state of $\mathbb{S}_1$, and moreover $S_1 = [\![f_1]\!]^{\mathbb{S}_1}$ and $S_2 = [\![f_2]\!]^{\mathbb{S}_2}$, as shown in Fig. 4. The colours f_1 and f_2 work similarly as the *fuses* used by Czarnecki in [4]: these force the iteration of $\varphi_{\omega^2 \cdot 2}$ to first traverse the f_1-copy $\mathbb{S}_1$ of ω^2 through the disjuncts $G\bot$ and $\varphi_{(2,1)}$, then move to the state 0 of the f_2-copy $\mathbb{S}_2$ of ω^2 through the disjunct $(Gx \wedge f_2 \wedge Gf_1)$, and finally traverse all $\mathbb{S}_2$ through the disjunct $\varphi_{(2,2)}$.

Theorem 3 below states that the two-way formula φ_α has closure ordinal α indeed. Due to space limitations, we have to omit the rather tedious proof; the interested reader can find the details in [13].

Theorem 3. *For every ordinal α with $\omega \leq \alpha < \omega^\omega$, the two-way closure ordinal of φ_α is α.*

5 Further Research

The research regarding closure ordinals of the μ-calculus has barely scratched the surface, and many questions remain open. We point out some possible future research lines.

Generally, we would like to understand better which ordinals feature as closure ordinals, and which ones don't. In particular, is there a two-way formula with closure ordinal at least ω^ω? Is there a standard (i.e., 'one-way') formula with a countable closure ordinal α at least ω^2? The approach taken here does not seem to work in the one-way setting—we refer to [13] for the details.

Another research direction involves decidability results. Given a formula $\varphi(x)$, is it decidable whether it has a closure ordinal, and can this be read off from its syntactic shape? Given an ordinal α, is it decidable whether a formula has closure ordinal α?

A more specific question concerns the *number of proposition letters* that is needed to characterize closure ordinals. In our approach we need an infinite set of atomic propositions to capture *all* ordinals below ω^ω. It is an interesting question to see whether this can be done with a *finite* set as well. We conjecture that this is indeed the case, by replacing the colors and fuses of Definition 11 by suitably chosen (basic) two-way formulas.[1]

Gouveia and Santocanale proved that closure ordinals are closed under ordinal sum [9] and we have transferred this result to the two-way setting. Is the class of closure ordinals closed under other ordinal operations as well, such as multiplication? Conversely, one may ask whether the formulas $\varphi \vee \psi, \varphi \wedge \psi, \varphi[\psi/x], \ldots$ have a closure ordinal whenever $\varphi(x)$ and $\psi(x)$ do.[2]

Finally, we mentioned the property of constructivity in the introduction. An interesting research direction involves the relationship between this property and that of continuity, where a formula $\varphi(x)$ is said to be *(Scott) continuous* in the variable x if, for an arbitrary model $\mathbb{S}$: $\mathbb{S}, s \Vdash \varphi$ iff $\mathbb{S}[x \mapsto V(p) \cap F], s \Vdash \varphi$, for some finite subset $F \subseteq S$. In particular, the second author [7,8] has formulated the question whether for every formula $\varphi(x)$ that is constructive in x one may find some formula $\psi(x)$ that is continuous in x, and equivalent to $\varphi(x)$ 'modulo an application of the least fixed point operator' (i.e., such that $\mu x.\varphi \equiv \mu x.\psi$). Some evidence supporting a positive answer can be found in [8,13].

A Proof of the Main Result in Section 3

The statement of Theorem 1 from Sect. 3 is a direct consequence of the following lemmas.

Lemma 1. *Let $0 < n < \omega$ be a finite ordinal. Then there is a model $\mathbb{S}$ where $\gamma_x(\varphi_n, \mathbb{S}) = \omega^n$.*

Proof. For the rest of the proof we adopt the following notation: since every ordinal $\alpha < \omega^n$ can be written as $\omega^{n-1} \cdot k_1 + \ldots + \omega \cdot k_{n-1} + k_n$, we also denote α as $(k_1, \ldots, k_n)$. From now on, if we write $\alpha = (k_1, \ldots, k_n)$ we mean that $\alpha = \omega^{n-1} \cdot k_1 + \ldots + \omega \cdot k_{n-1} + k_n$. Also, if a tuple $(k_1, \ldots, k_n)$ is of the form $(k_1, \ldots, k_i, 0, \ldots, 0)$, we mean that $k_j = 0$ for $i + 1 \leq j \leq n$.
Fix $n > 0$ and let $\varphi := \varphi_n$ as an abbreviation. We define $\mathbb{S} = (S, R, V)$ to be the model where:

- $S := \omega^n = \{(k_1, \ldots, k_n) \mid k_j \in \omega\}$;
- $R := \bigcup_{1 \leq i \leq n} \{((k_1, \ldots, k_i + 1, 0, \ldots, 0), (k_1, \ldots, k_i, 0, \ldots, 0)) \mid k_j \in \omega\}$;
- for $1 \leq i \leq n$, $V(q_i) := \{(k_1, \ldots, k_{n-i+1} + 1, 0, \ldots, 0) \mid k_j \in \omega\}$.

[1] This suggestion was raised by one of the referees.

[2] One of the reviewers pointed out that the formulas $p \wedge \Box(\neg p \wedge x)$ and $\neg p \wedge \Box(p \wedge x)$ both have closure ordinals, but their disjunction, behaving similarly to the formula $\Box x$, does not.

Note that $R[(0,\dots,0)] = \emptyset$ and that $(0,\dots,0)$ falsifies q_i for every $1 \leq i \leq n$.

Before proving the key claim we make an observation about notation. Note that an ordinal $\beta < \omega^n$ can both be seen as an *element* $\beta \in S = \omega^n$ of the model and as a *subset* $\beta = \{\gamma \mid \gamma < \beta\} \subseteq S = \omega^n$. To avoid confusion, until the end of the proof we write β when we consider it as an element of the domain, and S_β when we consider it as a subset of the domain ($S_\beta = \beta$ holds in any case).

Claim. For every $\alpha < \omega^n$, $\varphi^\alpha = S_\alpha$.

Proof of Claim. The proof goes by induction on α. The case for $\alpha = 0$ is immediate. If α is a limit ordinal, then $\varphi^\alpha = \bigcup_{\beta<\alpha} \varphi^\beta =_{IH} \bigcup_{\beta<\alpha} S_\beta = S_\alpha$.

Now suppose that $\alpha = \beta + 1$. We want to show that $\varphi^{\beta+1} = S_{\beta+1}$. We have that $\varphi^{\beta+1} = \varphi_x^{\mathbb{S}}(\varphi^\beta) =_{IH} \varphi_x^{\mathbb{S}}(S_\beta)$: we show

$$\varphi_x^{\mathbb{S}}(S_\beta) = S_{\beta+1}. \tag{3}$$

For the $\supseteq$ inclusion of (3) it suffices to show that $\mathbb{S}[x \mapsto S_\beta], \beta \Vdash \varphi$, since $S_{\beta+1} = S_\beta \cup \{\beta\}$ and $S_\beta = \varphi^\beta \subseteq \varphi^{\beta+1} = \varphi_x^{\mathbb{S}}(\varphi^\beta)$. If $\beta = 0 = (0,\dots,0)$ we are done. If $\beta = (k_1,\dots,k_n+1)$, then $\beta \in V(q_1)$ and $(k_1,\dots,k_n) \in S_\beta \cap R[\beta]$, so $\mathbb{S}[x \mapsto S_\beta], \beta \Vdash c_1 \wedge Fx$ and $\beta \in \varphi_x^{\mathbb{S}}(S_\beta)$.

Otherwise let $\beta = (k_1,\dots,k_i+1,0,\dots,0)$ for some $1 \leq i < n$, so that $\beta \in V(q_{n-i+1})$. Note that

$$\begin{aligned}&(k_1,\dots,k_i,k,0,\dots,0) \in S_\beta \text{ for all } k \in \omega,\\ &(k_1,\dots,k_i,0,0,\dots,0) \in R[\beta] \text{ and}\\ &(k_1,\dots,k_i,k,0,\dots,0) \in R[(k_1,\dots,k_i,k+1,0,\dots,0)] \cap V(q_{n-i}) \text{ for all } k > 0.\end{aligned}$$

By construction of the model $\beta \Vdash \pi^\infty_{n-i}$ also holds: then $\mathbb{S}[x \mapsto S_\beta], \beta \Vdash c_{n-i+1} \wedge \pi^\infty_{n-i} \wedge F(\nu y.P(x \wedge y \wedge c_{n-i}))$, so $\beta \in \varphi_x^{\mathbb{S}}(S_\beta)$.

Now we move to the $\subseteq$ inclusion of (3). Let $\gamma \in \varphi_x^{\mathbb{S}}(S_\beta)$. We want to show that $\gamma \in S_{\beta+1}$. Since $\mathbb{S}[x \mapsto S_\beta], \gamma \Vdash \varphi$ holds, we proceed by case distinction as to which disjunct of φ is satisfied by γ. If $\gamma \Vdash G\bot$ then $\gamma = 0 \in S_{\beta+1}$. If $\gamma \Vdash c_1 \wedge Fx$, then $\gamma \in V(q_1)$, so that $\gamma = (k_1,\dots,k_n+1)$ and $\gamma' = (k_1,\dots,k_n) \in R[\gamma] \cap S_\beta$: as $\gamma' \in S_\beta$, then $\gamma = \gamma' + 1 \in S_{\beta+1}$.

Now suppose $\gamma \Vdash c_i \wedge \pi^\infty_{i-1} \wedge F(\nu y.P(y \wedge x \wedge c_{i-1}))$ for some $2 \leq i \leq n$. Then $\gamma \in V(q_i)$, so $\gamma = (k_1,\dots,k_{n-i+1}+1,0,\dots,0)$. For $j \in \omega$ let

$$\delta_j := (k_1,\dots,k_{n-i+1},j,0,\dots,0).$$

By construction $\delta_0 \in R[\gamma]$ and $\delta_j \in R[\delta_{j+1}]$ for all $j \geq 0$. Since $\mathbb{S}[x \mapsto S_\beta], \gamma \Vdash c_i \wedge \pi^\infty_{i-1} \wedge F(\nu y.P(y \wedge x \wedge c_{i-1}))$ then $\delta_j \in S_\beta$ for all $j > 0$. Hence

$$\beta > (k_1,\dots,k_{n-i+1},j,0,\dots,0) \text{ for all } j > 0,$$

implying $\beta \geq (k_1,\dots,k_{n-i+1}+1,0,\dots,0) = \gamma$, so $\gamma \in S_{\beta+1}$. $\lhd$

Now that we have the claim, it follows that there is a $\gamma \in \varphi^{\omega^n} \setminus \varphi^\beta$ for each $\beta < \omega^n$.

Proposition 4. *For all $m, n \in \omega$, if $m \geq n$, then $\pi_m^\infty \models \pi_n^\infty$. Moreover, if $\mathbb{S}$ is a model, $\mathbb{S}, s \Vdash \pi_m^\infty$ for some state s, and $t_0 t_1 \ldots$ is an R^{-1}-path witnessing the truth of π_m^∞ at s, then $t_j \Vdash \pi_m^\infty$ for all $j \in \omega$.*

Lemma 2. *Let $\mathbb{S} = (S, R, V)$ be a model and let $n \in \omega$. For $1 \leq i \leq n$, let $t_0 t_1 t_2 \ldots$ be an infinite R^{-1}-path such that*

$$\mathbb{S}, t_0 \Vdash \pi_{i-1}^\infty \text{ and, for all } j > 0, \mathbb{S}, t_j \Vdash c_i \wedge \pi_{i-1}^\infty.$$

Then, for any ordinal α: if $t_0 \in \varphi_n^\alpha$ then $t_j \in \varphi_n^{\alpha+\omega^{i-1}\cdot j+1}$ for all $j \in \omega$.

Proof. We prove the statement by induction on $1 \leq i \leq n$.

As the base case take $i = 1$, so that by assumption we have an infinite R^{-1}-path $t_0 t_1 t_2 \ldots$ such that $\mathbb{S}, t_j \Vdash c_1$ for all $j > 0$. Let $t_0 \in \varphi_n^\alpha$. We want to show that, for all $j \in \omega$, $t_j \in \varphi_n^{\alpha+j+1}$: we prove this by induction on $j \in \omega$. If $j = 0$, then $t_0 \in \varphi_n^\alpha \subseteq \varphi_n^{\alpha+1}$. Next, inductively assume that $t_j \in \varphi_n^{\alpha+j+1}$: then, since $t_j \in R[t_{j+1}]$, it follows that $\mathbb{S}[x \mapsto \varphi_n^{\alpha+j+1}], t_{j+1} \Vdash (c_1 \wedge Fx)$, so $t_{j+1} \in \varphi_n^{\alpha+(j+1)+1}$.

For the inductive step assume that the statement holds for i. We prove it for $i+1$, where $i < n$. Suppose then that $t_0 t_1 t_2 \ldots$ is an infinite R^{-1}-path such that $t_0 \Vdash \pi_i^\infty$ and for all $j > 0$, $t_j \Vdash c_{i+1} \wedge \pi_i^\infty$. Let $t_0 \in \varphi_n^\alpha$. We want to show that

$$\text{for every } j \in \omega, t_j \in \varphi_n^{\alpha+\omega^i\cdot j+1}.$$

The proof of this last statement goes by induction on $j \in \omega$. The base case with $j = 0$ follows immediately, as by assumption $t_0 \in \varphi_n^\alpha$.

Now suppose that $t_j \in \varphi_n^{\alpha+\omega^i\cdot j+1}$: we show that $t_{j+1} \in \varphi_n^{\alpha+\omega^i\cdot(j+1)+1}$. By assumption $t_j \in R[t_{j+1}]$ and $t_j \Vdash \pi_i^\infty$, which in particular means that there is an infinite R^{-1}-path $u_0 u_1 \ldots$ (with $u_0 = t_j$) such that, for all $k > 0$, $u_k \Vdash c_i$. But then this path satisfies the conditions of the inductive hypothesis: by Proposition 4, since $u_0 \Vdash \pi_i^\infty$, then $u_0 \Vdash \pi_{i-1}^\infty$, and for every $k > 0$, $u_k \Vdash c_i \wedge \pi_{i-1}^\infty$. Then, by inductive hypothesis, since $u_0 = t_j \in \varphi_n^{\alpha+\omega^i\cdot j+1}$ it follows that, for every $k \in \omega$, $u_k \in \varphi_n^{\alpha+\omega^i\cdot j+1+\omega^{i-1}\cdot k+1}$. Since for all $k \in \omega$ it holds that

$$\begin{aligned}
\omega^i \cdot j + 1 + \omega^{i-1} \cdot k + 1 &< \omega^i \cdot j + 1 + \omega^i && (\text{as } \omega^{i-1} \cdot k + 1 < \omega^i \text{ for } i > 0) \\
&= \omega^i \cdot j + \omega^i && (1 + \omega^i = \omega^i \text{ for } i > 0) \\
&= \omega^i \cdot (j+1)
\end{aligned}$$

then also

$$\alpha + \omega^i \cdot j + 1 + \omega^{i-1} \cdot k + 1 < \alpha + \omega^i \cdot (j+1).$$

It follows that $u_k \in \varphi_n^{\alpha+\omega^i\cdot(j+1)}$ for all $k \in \omega$, so that

$$\mathbb{S}[x \mapsto \varphi_n^{\alpha+\omega^i\cdot(j+1)}], t_{j+1} \Vdash c_{i+1} \wedge \pi_i^\infty \wedge F(\nu y.P(x \wedge y \wedge c_i)).$$

We conclude that $t_{j+1} \in \varphi_n^{\alpha+\omega^i\cdot(j+1)+1}$ as desired.

Lemma 3. *For an arbitrary model* $\mathbb{S}$ *and* $0 < n < \omega$*:* $\gamma_x(\varphi_n, \mathbb{S}) \leq \omega^n$*.*

Proof. It is sufficient to prove that $\varphi_n^{\omega^n+1} \subseteq \varphi_n^{\omega^n}$ for every model $\mathbb{S}$. Let $s \in \varphi_n^{\omega^n+1}$, that is, $\mathbb{S}[x \mapsto \varphi_n^{\omega^n}], s \Vdash \varphi_n$. We proceed by case distinction as to which disjunct of φ_n is satisfied by s to prove that $s \in \varphi_n^{\omega^n}$. If $s \Vdash G\bot$ then $s \in (\varphi_n)_x^{\mathbb{S}}(\emptyset) \subseteq \varphi_n^{\omega^n}$, while if $s \Vdash c_1 \wedge Fx$, then there is a $t \in R[s]$ such that $t \in \varphi_n^{\alpha}$ for some $\alpha < \omega^n$, so that $s \in \varphi_n^{\alpha+1} \subseteq \varphi_n^{\omega^n}$.

Now suppose $s \Vdash c_i \wedge \pi_{i-1}^{\infty} \wedge F(\nu y.P(y \wedge x \wedge c_{i-1}))$ for some $2 \leq i \leq n$. Then in particular there is a point $t \in R[s]$ and a R^{-1}-path $t_0 t_1 \dots$ such that: *(i)* $t \in R[t_0]$, *(ii)* for all $j \in \omega$, $t_j \in \varphi_n^{\omega^n}$ and $t_j \Vdash c_{i-1}$. In particular, $t_0 \in \varphi_n^{\alpha}$ for some $\alpha < \omega^n$. Observe that $\varphi_n \wedge c_{i-1} \wedge F\top \models \pi_{i-2}^{\infty}$: this implies that $t_j \Vdash \pi_{i-2}^{\infty}$ for all $j \in \omega$, since $t_j \in \varphi_n^{\omega^n}$, $t_j \Vdash c_{i-1}$ and $R[t_j] \neq \emptyset$. This means that we can apply Lemma 2 and it follows that $t_j \in \varphi_n^{\alpha+\omega^{i-2}\cdot j+1} \subseteq \varphi_n^{\alpha+\omega^{i-1}}$ for all $j \in \omega$. Hence $\mathbb{S}[x \mapsto \varphi^{\alpha+\omega^{i-1}}], s \Vdash \varphi_n$ and $s \in \varphi_n^{\alpha+\omega^{i-1}+1} \subseteq \varphi_n^{\omega^n}$ (since $i \leq n$ and $\alpha < \omega^n$ imply $\alpha + \omega^{i-1} + 1 < \omega^n$).

References

1. Afshari, B., Leigh, G.E.: On closure ordinals for the modal μ-calculus. In: Computer Science Logic 2013 (CSL 2013). Leibniz International Proceedings in Informatics (LIPIcs), vol. 23, pp. 30–44. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2013). https://doi.org/10.4230/LIPIcs.CSL.2013.30. http://drops.dagstuhl.de/opus/volltexte/2013/4188
2. Afshari, B., Leigh, G.: Cut-free completeness for modal μ-calculus. In: Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2017), pp. 1–12 (2017). https://doi.org/10.1109/LICS.2017.8005088
3. Calude, C., Jain, S., Khoussainov, B., Li, W., Stephan, F.: Deciding parity games in quasipolynomial time. In: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, pp. 252–263 (2017). https://doi.org/10.1145/3055399.3055409
4. Czarnecki, M.: How fast can the fixpoints in modal μ-calculus be reached? In: Fixed Points in Computer Science 2010 (FICS 2010), pp. 35–39, August 2010. https://hal.archives-ouvertes.fr/hal-00512377/document#page=36
5. D'Agostino, G., Hollenberg, M.: Logical questions concerning the μ-calculus. J. Symb. Log. **65**, 310–332 (2000)
6. Emerson, E.A., Jutla, C.S.: The complexity of tree automata and logics of programs (extended abstract). In: Proceedings of the 29th Annual Symposium on Foundations of Computer Science, pp. 328–337. IEEE Computer Society Press (1988). https://doi.org/10.1109/SFCS.1988.21949
7. Fontaine, G.: Continuous fragment of the μ-calculus. In: Kaminski, M., Martini, S. (eds.) CSL 2008. LNCS, vol. 5213, pp. 139–153. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-87531-4_12
8. Fontaine, G., Venema, Y.: Some model theory for the modal μ-calculus: syntactic characterisations of semantic properties. Log. Methods Comput. Sci. **14**(1) (2018)

9. Gouveia, M.J., Santocanale, L.: $\aleph_1$ and the modal μ-calculus. In: 26th EACSL Annual Conference on Computer Science Logic (CSL 2017). Leibniz International Proceedings in Informatics (LIPIcs), vol. 82, pp. 38:1–38:16. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2017). https://doi.org/10.4230/LIPIcs.CSL.2017.38. http://drops.dagstuhl.de/opus/volltexte/2017/7692. An updated version can be found at https://arxiv.org/abs/1704.03772v2
10. Janin, D., Walukiewicz, I.: On the expressive completeness of the propositional μ-calculus with respect to monadic second order logic. In: Montanari, U., Sassone, V. (eds.) CONCUR 1996. LNCS, vol. 1119, pp. 263–277. Springer, Heidelberg (1996). https://doi.org/10.1007/3-540-61604-7_60
11. Kozen, D.: Results on the propositional μ-calculus. Theor. Comput. Sci. **27**, 333–354 (1983)
12. Kozen, D.: A finite model theorem for the propositional μ-calculus. Stud. Log. **47**, 233–241 (1988)
13. Milanese, G.: An exploration of closure ordinals in the modal μ-calculus. Master's thesis, Institute for Logic, Language and Computation, University of Amsterdam (2018)
14. Otto, M.: Eliminating recursion in the μ-calculus. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 531–540. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-49116-3_50
15. Walukiewicz, I.: Completeness of Kozen's axiomatisation of the propositional μ-calculus. Inf. Comput. **157**(1), 142–182 (2000). http://www.sciencedirect.com/science/article/pii/S0890540199928365