An exploration of closure ordinals in the modal μ-calculus

MSc Thesis *(Afstudeerscriptie)*

written by

Gian Carlo Milanese
(born February 5th, 1994 in Alessandria, Italy)

under the supervision of **Prof. Dr. Yde Venema**, and submitted to the Board of Examiners in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:
October 31, 2018

Members of the Thesis Committee:
Dr. Alexandru Baltag
Dr. Nick Bezhanishvili
Dr. Paul Dekker *(chair)*
Dr. Gaëlle Fontaine
Prof. Dr. Yde Venema *(supervisor)*

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

In this thesis we explore closure ordinals of the modal μ-calculus. The closure ordinal of a μ-calculus formula $\varphi(x)$ is the least ordinal α, if it exists, such that the iteration of the meaning function $\varphi_x^{\mathbb{S}}$ starting from the emptyset converges to its least fixed point in at most α many steps on every model $\mathbb{S}$. Our goal is to write an accessible introduction to this field by recalling some fundamental recent results and contributing with a few of our own. We provide a syntactic characterisation of the fragment of the μ-calculus corresponding to the semantic property of normality in a finite set of variables (which coincides with the property of having closure ordinal 0), we provide a syntactic characterisation of the fragment of the μ-calculus corresponding to the property of continuity on finitely branching models, we construct a formula φ_n with closure ordinal ω^n on bidirectional models for an arbitrary $n \in \omega$, we construct a formula φ_α with closure ordinal α on bidirectional models for an arbitrary $\alpha < \omega^\omega$, and we prove that the set of closure ordinals on bidirectional models is closed under ordinal sum (as an adaptation to this setting of the original result by Gouveia and Santocanale).

Acknowledgements

First of all, I want to express my gratitude to Yde Venema for his active support and enthusiasm for the project: the results of this thesis would not have been possible without his guidance. Working with Yde has been inspiring to say the least. I have genuinely enjoyed listening to his insights during our meetings, and he has always been able to motivate me and push me in the right direction whenever I felt stuck or lost.

I wish to thank all the members of the thesis committee for reading my thesis and asking many interesting questions during the defense. Each of these questions could constitute a direction for further research that I really hope I will be able to investigate.

Furthermore, I would like to thank my academic mentor, Benno van den Berg, for his advice in the two years of the master, and all the ILLC staff for making the Master of Logic such a special programme.

Very special thanks to Alinda for helping me proofread every draft of this thesis, and for always being by my side.

Finally, I want to thank all of my family for their presence and for supporting me and my choices.

Contents

Introduction

The central theme of this thesis is the study of closure ordinals of formulas of the modal μ-calculus.

Introduced by Kozen in the form that is studied today [15], the modal μ-calculus is an extension of modal logic with a least fixed point operator μ and a greatest fixed point operator ν. The addition of these operators enriches modal logic with inductive definitions, which results in a significant increase in expressive power. Janin and Walukiewicz proved that a formula of monadic second-order logic is invariant under bisimulation if and only if it is equivalent to a formula of the μ-calculus [14], extending the celebrated theorem by Johan van Benthem which characterises modal logic as the bisimulation invariant fragment of first-order logic [18]. Moreover, the μ-calculus is an extension of several temporal logics like PDL (propositional dynamic logic) and CTL (computation tree logic).

Besides its inherent mathematical and logical interest, a certain balance between expressiveness and computational feasibility makes the modal μ-calculus a central logic in certain areas of computer science, as it can be used as a specification language. Indeed, this logic enjoys the finite model property [16] and problems such as model checking (deciding whether a given formula is true in a given finite model) and satisfiability (deciding whether a given formula is true in some model) for formulas of this logic are decidable with relatively low computational complexity [8].

Another engaging feature of the μ-calculus is that its semantics can be specified both by algebraic means and in an intuitive game-theoretic format based on parity games. The latter is often essential when trying to decipher the meaning of a formula of the modal μ-calculus: the alternation of fixed point operators and the presence of bound variables can sometimes result in formulas that are difficult to read. We also note that an axiomatization for the μ-calculus has been formulated by Kozen [15] and proven to be complete, partially by Kozen and fully by Walukiewicz [20].

The close connection with automata theory [13] is also important to mention: the equivalence between formulas of the μ-calculus and modal automata allows the application of automata theoretic methods for proving results about this logic. One example of a fundamental result proved with an automata theoretic approach is the already mentioned theorem by Janin and Walukiewicz [14].

A classic example of a μ-calculus formula is $\mu x.p \vee \langle d \rangle x$ (for the following example in this paragraph we refer to [19]). This formula is equivalent to the formula $\langle d^* \rangle p$ from PDL, expressing the existence of a finite R_d-path leading to a state that satisfies p. It is well known that, for every transition system (or

Kripke model) $\mathbb{S}$ and state s:

$$\mathbb{S}, s \Vdash \langle d^* \rangle p \leftrightarrow (p \vee \langle d \rangle \langle d^* \rangle p).$$

In a sense, the formula $\langle d^* \rangle p$ is a *fixed point* of the expression

$$x \leftrightarrow p \vee \langle d \rangle x.$$

The modal μ-calculus allows to refer explicitly to the least fixed point and to the greatest fixed point of this equation, by writing $\mu x.p \vee \langle d \rangle x$ and $\nu x.p \vee \langle d \rangle x$, respectively. Other examples of μ-calculus formulas are $\nu x.p \wedge \Diamond\Diamond x$, expressing the existence of a path where p is true at every even position, and $\nu y.\mu x.(p \wedge \Diamond y) \vee \Diamond x$, expressing the existence of a path where p is true infinitely often.

A recent line of research on the modal μ-calculus has focused on the closure ordinals of its formulas. Given a formula φ in this language, a Kripke model $\mathbb{S} = (S, R, V)$ and a variable x that occurs positively in φ, we can define a 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 meaning of $\mu x.\varphi$ in $\mathbb{S}$ can then be computed by performing an iteration of the function $\varphi_x^{\mathbb{S}}$ starting from the emptyset, resulting in an ordinal-indexed sequence $(\varphi_\mu^\alpha)_{\alpha \in \mathrm{On}}$ of subsets of S defined by letting

$$\begin{array}{lcl} \varphi_\mu^0 & := & \emptyset, \\ \varphi_\mu^{\beta+1} & := & \varphi_x^{\mathbb{S}}(\varphi_\mu^\beta), \\ \varphi_\mu^\lambda & := & \bigcup_{\beta<\lambda} \varphi_x^{\mathbb{S}}(\varphi_\mu^\beta), \end{array}$$

where λ denotes an arbitrary limit ordinal.

The positivity assumption on x implies that the function $\varphi_x^{\mathbb{S}}$ is monotone and that the sequence $(\varphi_\mu^\alpha)_{\alpha \in \mathrm{On}}$ converges: there must be an ordinal α such that $\varphi_\mu^\alpha = \varphi_\mu^{\alpha+1}$. The element φ_μ^α of the sequence coincides with the least fixed point of $\varphi_x^{\mathbb{S}}$ – that is, the least subset $L \subseteq S$ such that $\varphi_x^{\mathbb{S}}(L) = L$ – 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.

Our investigation of closure ordinals in the modal μ-calculus stems from a number of recent results in this field.

1999 Martin Otto [17] proved that it is decidable whether a modal μ-calculus formula can equivalently be expressed in (basic) modal logic. As a corollary, whether a formula of modal logic has a finite closure ordinal is also decidable.

2008 Gaëlle Fontaine [9] presented a syntactic characterization of the continuous fragment of the modal μ-calculus. The property of continuity is often studied in relation to that of constructivity (a formula is constructive if its closure ordinal is at most ω), since every continuous formula is also constructive. The connection between these two properties constitutes a very intriguing research direction.

2010 Marek Czarnecki [6] showed how to construct a formula φ_α with closure ordinal α for an arbitrary $\alpha < \omega^2$.

2013 Bahareh Afshari and Graham E. Leigh [1] proved that if a formula in the alternation-free fragment of the modal μ-calculus has a closure ordinal, this is strictly less than ω^2.

2017 Maria João Gouveia and Luigi Santocanale [12] presented a syntactic characterization of the $\aleph_1$-continuous fragment of the modal μ-calculus, defined a formula with closure ordinal ω_1 and proved that closure ordinals of the μ-calculus are closed under ordinal sum.

Our goal is to offer an accessible introduction to this interesting area of research by recalling some fundamental recent results and contributing with a few of our own. The thesis is structured as follows.

Chapter 1

We present the modal μ-calculus after recalling some basic facts about fixed points of functions on complete lattices. While this presentation should be accessible to readers that are not already familiar with the μ-calculus, we refer to [19] for a more detailed introduction to the subject.

Chapter 2

We begin our exploration of closure ordinals of the modal μ-calculus. We start by defining a fragment $\mu\mathtt{ML}_X^N$ of the modal μ-calculus that syntactically characterises the property of normality in a finite set of variables, which coincides with the property of having 0 as closure ordinal.

Contribution (Theorem 2.2.15) *Every formula in $\mu\mathtt{ML}_X^N$ is normal in X. Moreover, there is an effective translation which, given a $\mu\mathtt{ML}$-formula φ, computes an equivalent formula $\varphi^d \in \mu\mathtt{DML}$ such that*

$$\varphi \textit{ is normal in } X \textit{ iff } \varphi^d \in \mu\mathtt{ML}_X^N.$$

We then move to formulas with finite closure ordinals and to formulas that need at most ω many iterations to converge to their least fixed point across all models. We discuss the connection between the properties of continuity and constructivity. The property of constructivity is particularly interesting: in an arbitrary model an individual state is included in the iteration of the fixed point of a constructive formula in at most finitely many steps. As a minor contribution, we conclude the chapter by formulating the fragment $\mu\mathtt{ML}_x^D$ of the μ-calculus that characterises the property of continuity on finitely branching models, which coincides with the fragment of the μ-calculus that characterises the finite depth property [10, 11].

Chapter 3

We show how to construct formulas with closure ordinals larger than ω. We discuss the intuitions behind Czarnecki's construction of a formula with closure ordinal α for every $\alpha < \omega^2$ [6] and the difficulties involved with finding a formula with closure ordinal ω^2. We show however that the latter and greater countable closure ordinals can be obtained in the setting of bidirectional models.

Contribution (Theorem 3.3.9) *For all $0 < n < \omega$ there is a formula φ_n with closure ordinal ω^n on bidirectional models.*

Contribution (Theorem 3.4.10) *For every $\omega \leq \alpha < \omega^\omega$ there is a formula φ_α with closure ordinal α on bidirectional models.*

As a minor contribution, we also provide an adaptation to the setting of bidirectional models of the result by Gouveia and Santocanale that the set of closure ordinals is closed under ordinal sum. We conclude the chapter by discussing the first uncountable ordinal ω_1 as a closure ordinal.

Part of what makes closure ordinals of formulas of the modal μ-calculus interesting and fun is the challenge involved in controlling the number of iterations that a formula needs in order to converge to its least fixed point – coming up with disjuncts that allow the iteration to progress on the one hand, and disjuncts that keep the iteration from blowing up on the other – and the visual aspect tied with imagining the iteration of a formula traversing infinite ordinal numbers (in Chapter 3 we provide numerous figures to guide the intuition). We hope that our thesis will encourage some reader to pick up the many open questions of this field.

Chapter 1

The modal μ-calculus

1.1 Basic theory of fixed points

This short section is meant to recall some basic definitions and results about lattices. We present some well-known characterisations of the least and the greatest fixed points of a monotone function on a complete lattice that we will need in order to define the algebraic semantics of the modal μ-calculus and the notion of closure ordinal of a formula. A game-theoretic characterisation of these fixed points when the complete lattice is a powerset algebra is also mentioned. We start with a few definitions.

Definition 1.1.1 A *partially ordered set* $\mathbb{P} = (P, \leq)$, or *poset*, is a set P together with a binary relation $\leq$ on P satisfying, for all $p, q, r \in P$:

- $p \leq p$;
- if $p \leq q$ and $q \leq p$, then $p = q$;
- if $p \leq q$ and $q \leq r$, then $p \leq r$.

In other words, $\leq$ is *reflexive*, *antisymmetric* and *transitive*.

If $X \subseteq P$, an element $p \in P$ is an *upper bound* of X if $x \leq p$ for every $x \in X$, and a *lower bound* of X if $p \leq x$ for every $x \in X$. An element p is called the *least upper bound* of X if it is an upper bound of X, and whenever p' is an upper bound of X, then $p \leq p'$. The least upper bound of a set is also called its *supremum*, or *join*. Dually, p is called the *greatest lower bound* of X if it is a lower bound of X and $p' \leq p$ for every other lower bound p' of X. The greatest lower bound is also called *infimum* or *meet*.

Definition 1.1.2 A *lattice* is a poset where every two elements p and q have a least upper bound and a greatest lower bound, denoted respectively by $p \vee q$ and $p \wedge q$.

A *complete lattice* is a poset where a least upper bound and a greatest lower bound exist for *every* subset X of the poset: these are denoted respectively by $\bigvee X$ and $\bigwedge X$. Note that $p \wedge q = \bigwedge\{p, q\}$ and $p \vee q = \bigvee\{p, q\}$. In particular in a complete lattice there exists a least element $\bot := \bigvee \emptyset$ and a greatest element $\top := \bigwedge \emptyset$. Complete lattices will usually be denoted by $\mathbb{C} = (C, \bigvee, \bigwedge)$.

Example 1.1.3 For any set S, $(\wp(S), \subseteq)$ is a complete lattice with $\bigvee X := \bigcup X$ and $\bigwedge X := \bigcap X$ for all $X \subseteq S$, and where the least and the greatest elements are, respectively, $\emptyset$ and S.

Definition 1.1.4 For two partial orders $\mathbb{P} = (P, \leq)$ and $\mathbb{P}' = (P', \leq')$, a function $f : P \to P'$ is *monotone* if $p \leq q$ implies $f(p) \leq f(q)$ for all $p, q \in P$.

Definition 1.1.5 Let $\mathbb{P} = (P, \leq)$ be a partial order and $f : P \to P$ a function. An element $p \in P$ is called

- a *prefixpoint* of f if $f(p) \leq p$,
- a *postfixpoint* of f if $p \leq f(p)$,
- a *fixed point* (or *fixpoint*) of f if $f(p) = p$.

The sets of prefixpoints, postfixpoints and fixed points of f are denoted by $\mathrm{PRE}(f)$, $\mathrm{POS}(f)$ and $\mathrm{FIX}(f)$, respectively.

In case they exist, the least element of $\mathrm{FIX}(f)$ is called *least fixed point* of f and denoted by $\mathrm{LFP}.f$, while its greatest element is called *greatest fixed point* of f and denoted by $\mathrm{GFP}.f$.

It is important to note that, while the least and the greatest fixed points of a function do not always exist, when the function is a monotone map on a complete lattice the next theorem guarantees that they can be identified with, respectively, the meet of the set of all prefixpoints and the join of the set of all postfixpoints of f.

Theorem 1.1.6 (Knaster-Tarski) *If $f : C \to C$ is a monotone function on a complete lattice $\mathbb{C} = (C, \bigvee, \bigwedge)$, then f has both a least fixed point and a greatest fixed point, which are given by*

$$\begin{aligned} \mathrm{LFP}.f &= \bigwedge \mathrm{PRE}(f), \\ \mathrm{GFP}.f &= \bigvee \mathrm{POS}(f). \end{aligned}$$

Alternatively, $\mathrm{LFP}.f$ and $\mathrm{GFP}.f$ can be approximated by two sequences that start from the bottom and the top elements of a complete lattice.

Definition 1.1.7 Let $f : C \to C$ be a function on a complete lattice $\mathbb{C} = (C, \bigvee, \bigwedge)$. We define two sequences $(f_\mu^\alpha)_{\alpha \in \mathrm{On}}$ and $(f_\nu^\alpha)_{\alpha \in \mathrm{On}}$ of elements of $\mathbb{C}$ by ordinal induction:

$$\begin{array}{rclrcl} f_\mu^0 &:=& \bot, & f_\nu^0 &:=& \top, \\ f_\mu^{\alpha+1} &:=& f(f_\mu^\alpha), & f_\nu^{\alpha+1} &:=& f(f_\nu^\alpha), \\ f_\mu^\lambda &:=& \bigvee\limits_{\alpha<\lambda} f_\mu^\alpha, & f_\nu^\lambda &:=& \bigwedge\limits_{\alpha<\lambda} f_\mu^\alpha, \end{array}$$

where λ denotes an arbitrary limit ordinal.

As the following two propositions state, when f is a monotone function on a complete lattice there must be an ordinal α such that f_μ^α coincides with $\mathrm{LFP}.f$ Since ordinal approximations of least fixed points will play an important role in the rest of the thesis we provide a proof of this fact.

Proposition 1.1.8 If $f : C \to C$ is a monotone function on a complete lattice $\mathbb{C} = (C, \bigvee, \bigwedge)$, then $f_\mu^\alpha \leq f_\mu^\beta$ whenever $\alpha < \beta$.

Proof. First of all, note that $f_\mu^\alpha \leq f_\mu^{\alpha+1}$ for every ordinal α. By induction on α, if $\alpha = \beta + 1$, by inductive hypothesis and monotonicity we immediately obtain $f_\mu^{\beta+1} = f(f_\mu^\beta) \leq f(f_\mu^{\beta+1}) = f_\mu^{\beta+2}$. If α is a limit, by definition $f_\mu^\alpha = \bigvee_{\beta<\alpha} f_\mu^\beta$, so that $f_\mu^\beta \leq f_\mu^\alpha$ for all $\beta < \alpha$. Moreover, by inductive hypothesis $f_\mu^\beta \leq f_\mu^{\beta+1} = f(f_\mu^\beta)$ for all $\beta < \alpha$, so that by monotonicity of f we have $f_\mu^\beta \leq f(f_\mu^\beta) \leq f(f_\mu^\alpha) = f_\mu^{\alpha+1}$ for every $\beta < \alpha$, implying that $f_\mu^{\alpha+1}$ is an upper bound of $\{f_\mu^\beta \mid \beta < \alpha\}$: as $f_\mu^\alpha = \bigvee_{\beta<\alpha} f_\mu^\beta$ is the least upper bound of the same set, we conclude $f_\mu^\alpha \leq f_\mu^{\alpha+1}$.

Now we prove the statement of the proposition by induction on β. If $\beta = \gamma+1$ then either $\alpha = \gamma$ and the statement follows, or $\alpha < \gamma$ and $f_\mu^\alpha \leq f_\mu^\gamma \leq f_\mu^{\gamma+1}$ by inductive hypothesis. If β is a limit then $\alpha < \gamma$ for some $\gamma < \beta$, and $f_\mu^\alpha \leq f_\mu^\gamma \leq \bigvee_{\gamma<\beta} f_\mu^\gamma = f_\mu^\beta$, with the first inequality given by the inductive hypothesis. □

Corollary 1.1.9 Let $f : C \to C$ be a monotone function on a complete lattice $\mathbb{C} = (C, \bigvee, \bigwedge)$. Then there is some α such that LFP.$f = f_\mu^\alpha$.

Proof. An easy ordinal induction, similar to those in the previous proof, shows that $f_\mu^\alpha \leq$ LFP.f for every ordinal α. For the other direction, the fact that the lattice $\mathbb{C}$ has a size $|C|$ implies that there must be an ordinal α of size at most $|C|$ such that $f_\mu^\alpha = f_\mu^{\alpha+1} = f(f_\mu^\alpha)$, giving LFP.$f \leq f_\mu^\alpha$. □

When $f_\mu^\alpha = f_\mu^{\alpha+1}$ for some ordinal α, we say that f *converges to its least fixed point* in (at most) α steps, while if α is the least ordinal such that $f_\mu^\alpha = f_\mu^{\alpha+1}$, then we say that f *converges to its least fixed point in exactly* α steps. Analogously, one can prove that $f_\nu^\alpha \geq f_\nu^\beta$ whenever $\alpha < \beta$ and that there must be some α such that GFP.$f = f_\nu^\alpha$ when f is a monotone function on a complete lattice.

When the complete lattice under consideration is actually the powerset algebra $(\wp(S), \subseteq)$, for some set S, a game-theoretic characterisation of the least and the greatest fixed points of a monotone function $f : \wp(S) \to \wp(S)$ is also possible. We refer to [19] for more precise definitions about board games: for our purposes it is enough to know that the games we are going to define are played by two players, denoted by $\exists$ (Éloise) and $\forall$ (Abélard), who, depending on the position of the game, can make moves according to Table 1.1. A move consists in choosing the next position, which can be either an element or a subset of the set S, while matches are finite or infinite sequences of positions. The games are played until either player cannot make a move, in which case we say that the player got stuck, but it is also possible for a match of the game to continue forever. The next definition fixes the notation and states the winning conditions for the two games.

Definition 1.1.10 Let S be some set and $f : \wp(S) \to \wp(S)$ be a monotone function. The following table specifies the positions and admissible moves for the two players $\exists$ and $\forall$ of the *unfolding games* $\mathcal{U}^\mu(f)$ and $\mathcal{U}^\nu(f)$.

Position	Player	Set of admissible moves
$s \in S$	$\exists$	$\{A \in \wp(S) \mid s \in f(A)\}$
$A \in \wp(S)$	$\forall$	A

Table 1.1: Unfolding game

Finite matches are lost by the player who got stuck. The only difference between these two games consists in the winning conditions for infinite matches: infinite matches of $\mathcal{U}^\nu(f)$ are won by $\exists$, while those of $\mathcal{U}^\mu(f)$ are won by $\forall$. A *strategy* for a player $\Pi \in \{\exists, \forall\}$ is a method that dictates which move Π should play depending on the position of the game. A strategy is *winning* for Π from a certain position if every match of the game started at the given position is won by Π whenever Π adopts this strategy. A position is *winning* for Π if Π has a winning strategy for the game initialised in that position. The set of winning positions for a player $\Pi \in \{\exists, \forall\}$ in $\mathcal{U}^\eta(f)$, with $\eta \in \{\mu, \nu\}$ is denoted by $\mathrm{Win}_\Pi(\mathcal{U}^\eta(f))$.

Theorem 1.1.11 *Let S be some set and $f : \wp(S) \to \wp(S)$ be a monotone operation. Then:*

1. $\mathrm{LFP}.f = \mathrm{Win}_\exists(\mathcal{U}^\mu(f)) \cap S$,
2. $\mathrm{GFP}.f = \mathrm{Win}_\exists(\mathcal{U}^\nu(f)) \cap S$.

A detailed proof can be found in [19].

1.2 The modal μ-calculus

The focus of this section is the modal μ-calculus. We define its language and show how to interpret it on Kripke models both via an algebraic semantics and a game-theoretic semantics. What differentiates the μ-calculus from modal logic is the presence of *fixed point operators* μx and νx, which significantly enhance its expressive power.

The algebraic semantics determines the meaning of a formula $\mu x.\varphi$ or $\nu x.\varphi$ in a Kripke model $\mathbb{S}$ as the *least fixed point* or the *greatest fixed point*, respectively, of the function $\varphi_x^{\mathbb{S}} : \wp(S) \to \wp(S)$, which intuitively expresses the dependence inside $\mathbb{S}$ of the meaning of φ on the valuation of the variable x. While this definition is very clear, often it is not very easy to work with, for instance when trying to determine the meaning of a concrete formula. An alternative is to look at the ordinal approximations of the least and the greatest fixed points of $\varphi_x^{\mathbb{S}}$, a tool which we will often employ in the study of closure ordinals of μ-calculus formulas.

The game-theoretic semantics, on the other hand, provides a very intuitive tool when dealing with formulas of the form $\eta x.\varphi$, where η is either μ or ν. In the context of a game a position corresponding to the bound variable x will, in some sense, move the game back to the formula φ: we say that the variable x is *unfolded*. Depending on whether x is bound by a ν operator or a μ operator, the truth or falsehood of $\eta x.\varphi$ in a model is linked to matches of the game where x is unfolded infinitely often, similarly to how in the unfolding game of the previous section infinite matches of $\mathit{Unf}^\nu(f)$ are won by $\exists$, while those of $\mathit{Unf}^\mu(f)$ are won by $\forall$.

All of this will hopefully become clearer when more precise definitions are stated: we now start with the language of the modal μ-calculus.

1.2.1 Language of the μ-calculus

Convention 1.2.1 Throughout the text we fix an infinite set `PROP` of propositional variables and a finite set `D` of atomic actions.

Definition 1.2.2 The language $\mu\mathtt{ML}_\mathtt{D}$ of the *polymodal μ-calculus* is given by the following grammar:

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

where $p, x \in \mathtt{PROP}$, $d \in \mathtt{D}$ 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.

A formula of the μ-calculus is in *negation normal form* if it belongs to the language defined by the following grammar:

$$\varphi ::= \bot \mid \top \mid p \mid \neg p \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Diamond_d \varphi \mid \Box_d \varphi \mid \mu x.\varphi \mid \nu x.\varphi$$

where $p, x \in \mathtt{PROP}$, $d \in \mathtt{D}$ and the formation of the formulas $\mu x.\varphi$ and $\nu x.\varphi$ is subject to the constraint that the variable x is positive in φ.

When the set $\mathtt{D}$ is a singleton we will denote the modal operators by $\Diamond$ and $\Box$, and speak of the language $\mu\mathtt{ML}$ of the modal μ-calculus.

Working with the language of the μ-calculus in negation normal form is especially convenient for the game theoretic definition of its semantics, and we will usually assume that a formula belongs to this language. In the smaller setting the missing connectives can be defined by letting $\nu x.\varphi := \neg\mu x.\neg(\varphi[\neg x/x])$ and the standard abbreviations for $\top$, $\wedge$ and $\Box_d$; we also treat $\to$ as a defined connective. The language $\mathtt{ML}$ of (basic) modal logic is defined in an analogous way, without the rules for the fixed point operators. We continue with some syntactic definitions.

Definition 1.2.3 The set $\mathit{Sfor}(\varphi)$ of *subformulas* of φ is defined in the usual familiar way, with the clause for a fixed point formula being $\mathit{Sfor}(\eta x.\varphi) := \{\eta x.\varphi\} \cup \mathit{Sfor}(\varphi)$, $\eta \in \{\mu, \nu\}$. The *size* $|\varphi|$ of a formula φ is the size of $\mathit{Sfor}(\varphi)$. We write $\psi \trianglelefteq \varphi$ when ψ is a subformula of φ.

Definition 1.2.4 An occurrence of a variable x in a formula φ is *bound* if x is under the scope of an operator μx or νx and *free* otherwise; $FV(\varphi)$ and $BV(\varphi)$ denote the sets of *free* and *bound variables* of φ.

Definition 1.2.5 Let φ and ψ be $\mu\mathtt{ML}_\mathtt{D}$-formulas, and z a variable that is free in φ. We let $\varphi[\psi/z]$ denote the formula obtained from φ by substituting the formula ψ for z in φ, under the assumption that no free variable in ψ will become bound in the process. If z is clear from context we will also write $\varphi(\psi)$.

Definition 1.2.6 A formula $\varphi \in \mu\mathtt{ML}_\mathtt{D}$ is *clean* if no two distinct (occurrences of) fixed point operators in φ bind the same variable, and no variable has both free and bound occurrences in φ. If x is a bound variable of the clean formula φ, let $\varphi_x = \eta_x x.\delta_x$ denote the unique subformula of φ where x is bound by the operator $\eta_x x$, $\eta_x \in \{\mu, \nu\}$. If $\eta_x = \mu$ we say that x is a μ-variable, and a ν-variable otherwise.

Observe that through a suitable renaming of bound variables we can find, for every formula φ, an equivalent clean formula, so that we will often assume a formula to be clean without loss of generality.

Definition 1.2.7 Let φ be a clean formula. The *dependency order* $\leq_\varphi$ on the bound variables of φ is defined by letting $x \leq_\varphi y$ if $\delta_x \trianglelefteq \delta_y$.

1.2.2 Algebraic semantics

Formulas of the modal μ-calculus will be interpreted in Kripke models.

Definition 1.2.8 A *Kripke model* of type $\mathtt{D}$ is a triple $\mathbb{S} = (S, (R_d)_{d\in\mathtt{D}}, V)$ where S, the *domain* or *underlying set*, is a set of *points* or *states*, $R_d \subseteq S \times S$ is a binary relation, called *accessibility relation*, for each $d \in \mathtt{D}$, and V is a valuation on S, that is, a function $V : \mathtt{PROP} \to \wp(S)$. A *pointed model* is a model $\mathbb{S}$ together with a designated point $s \in S$, and is denoted by $(\mathbb{S}, s)$. A model $\mathbb{S}$ is *image-finite*, or *finitely branching*, if $R_d[s] := \{t \in S \mid (s,t) \in R_d\}$ is finite for every element s of $\mathbb{S}$ and $d \in \mathtt{D}$.

Given a model $\mathbb{S} = (S, (R_d)_{d\in\mathtt{D}}, V)$, a propositional variable x and a subset $X \subseteq S$, we define $V[x \mapsto X]$ as the valuation given by

$$V[x \mapsto X](p) := \begin{cases} V(p) & \text{if } p \neq x \\ X & \text{otherwise} \end{cases}$$

and we denote the model $(S, (R_d)_{d\in\mathtt{D}}, V[x \mapsto X])$ by $\mathbb{S}[x \mapsto X]$.

Definition 1.2.9 Let $\mathbb{S} = (S, (R_d)_{d\in\mathtt{D}}, V)$ be a Kripke model of type $\mathtt{D}$. A subset $S' \subseteq S$ is *(upward) closed* if $s \in S'$ and $t \in R_d[s]$ imply $t \in S'$ for all $d \in \mathtt{D}$. A subset $S' \subseteq S$ is *downward closed* if $s \in S'$ and $s \in R_d[t]$ imply $t \in S'$ for all $d \in \mathtt{D}$.

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

Tree models, as defined below, often play an important role when proving results about the μ-calculus.

Definition 1.2.10 A pointed model $(\mathbb{S}, s)$ is a *tree model* with root s if $S = \bigcup_{d\in\mathtt{D}} R_d^*[s]$ (with R_d^* denoting the reflexive-transitive closure of R_d) and every state $t \neq s$ has a unique predecessor. A *sibling* of a node t in a tree model is a node $t' \neq t$ with the same predecessor of t.

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 1.2.11 Given a $\mu\mathtt{ML}_\mathtt{D}$-formula φ and a model $\mathbb{S} = (S, (R_d)_{d\in\mathtt{D}}, V)$, 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:

$$\begin{array}{lcl}
[\![\bot]\!]^{\mathbb{S}} & = & \emptyset \\
[\![\top]\!]^{\mathbb{S}} & = & S \\
[\![p]\!]^{\mathbb{S}} & = & V(p) \\
[\![\neg p]\!]^{\mathbb{S}} & = & S \backslash V(p) \\
[\![\varphi \vee \psi]\!]^{\mathbb{S}} & = & [\![\varphi]\!]^{\mathbb{S}} \cup [\![\psi]\!]^{\mathbb{S}} \\
[\![\varphi \wedge \psi]\!]^{\mathbb{S}} & = & [\![\varphi]\!]^{\mathbb{S}} \cap [\![\psi]\!]^{\mathbb{S}} \\
[\![\Diamond_d \varphi]\!]^{\mathbb{S}} & = & \{s \in S \mid R_d[s] \cap [\![\varphi]\!]^{\mathbb{S}} \neq \emptyset\} \\
[\![\Box_d \varphi]\!]^{\mathbb{S}} & = & \{s \in S \mid R_d[s] \subseteq [\![\varphi]\!]^{\mathbb{S}}\} \\
[\![\mu x.\varphi]\!]^{\mathbb{S}} & = & \bigcap \mathrm{PRE}(\varphi_x^{\mathbb{S}}) \\
[\![\nu x.\varphi]\!]^{\mathbb{S}} & = & \bigcup \mathrm{POS}(\varphi_x^{\mathbb{S}})
\end{array}$$

For an element $s \in S$ we write $\mathbb{S}, s \Vdash \varphi$ if $s \in [\![\varphi]\!]^{\mathbb{S}}$.

Remark Let $\varphi \in \mu\mathtt{ML}_{\mathtt{D}}$ be a formula in which the variable x occurs only positively and $\mathbb{S}$ 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 $[\![\mu x.\varphi]\!]^{\mathbb{S}} = \text{LFP}.\varphi_x^{\mathbb{S}}$ and $[\![\nu x.\varphi]\!]^{\mathbb{S}} = \text{GFP}.\varphi_x^{\mathbb{S}}$.

Definition 1.2.12 Let φ and ψ be $\mu\mathtt{ML}_{\mathtt{D}}$-formulas. We say that ψ is a *local consequence* of φ (notation: $\varphi \models \psi$) if $\mathbb{S}, s \Vdash \varphi$ implies $\mathbb{S}, s \Vdash \psi$ for every pointed model $(\mathbb{S}, s)$. We say that φ and ψ are *equivalent* (notation: $\varphi \equiv \psi$) if $\varphi \models \psi$ and $\psi \models \varphi$.

It is often difficult to decipher the meaning of a fixed point formula just by looking at it, especially if alternation of fixed point operators is involved. Fortunately, the modal μ-calculus also admits an equivalent and much more intuitive game-theoretic semantics, which will be presented in the next section. For simple formulas one could also look at the ordinal fixed point approximation of Definition 1.1.7.

Example 1.2.13 Consider the formula $\varphi(x) := \Diamond x \vee \Box\bot$ and the model $\mathbb{S}$ with domain ω, empty valuation for each proposition letter and $R := \{(n+1, n) \mid n \in \omega\}$ as accessibility relation. This is depicted in the following picture[1].

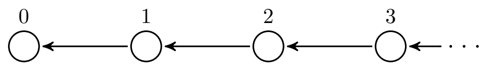

By the previous remark and Corollary 1.1.9 we know that $[\![\mu x.\varphi]\!]^{\mathbb{S}} = \text{LFP}.\varphi_x^{\mathbb{S}} = (\varphi_x^{\mathbb{S}})_\mu^\alpha$ for some α. Let us look at the stages of the ordinal approximation more closely. With $(\varphi_x^{\mathbb{S}})_\mu^0 = \emptyset$ by definition, observe that

$$(\varphi_x^{\mathbb{S}})_\mu^1 = \varphi_x^{\mathbb{S}}(\emptyset) = [\![\Diamond x \vee \Box\bot]\!]^{\mathbb{S}[x\mapsto\emptyset]} = [\![\Box\bot]\!]^{\mathbb{S}[x\mapsto\emptyset]} = \{0\},$$
$$(\varphi_x^{\mathbb{S}})_\mu^2 = \varphi_x^{\mathbb{S}}((\varphi_x^{\mathbb{S}})_\mu^1) = \varphi_x^{\mathbb{S}}(\{0\}) = [\![\Diamond x \vee \Box\bot]\!]^{\mathbb{S}[x\mapsto\{0\}]} = \{0, 1\},$$

and so on. In other words, each finite iteration of the least fixed point approximation adds one state to the meaning of $\mu x.\varphi$ in $\mathbb{S}$ (a simple induction shows that $(\varphi_x^{\mathbb{S}})_\mu^n = \{m \in \omega \mid m < n\}$ for all $n \in \omega$), so that $\text{LFP}.\varphi_x^{\mathbb{S}} = \bigcup_{n<\omega} (\varphi_x^{\mathbb{S}})_\mu^n = (\varphi_x^{\mathbb{S}})_\mu^\omega$.

Example 1.2.14 Now consider the formula $\varphi(x) := \Box x$ and the model $\mathbb{S}$ where the domain is now $\omega \cup \{\omega\}$, the valuation is empty for each proposition letter and $R := \{(n + 1, n) \mid n \in \omega\} \cup \{(\omega, n) \mid n \in \omega\}$, as depicted in the following picture.

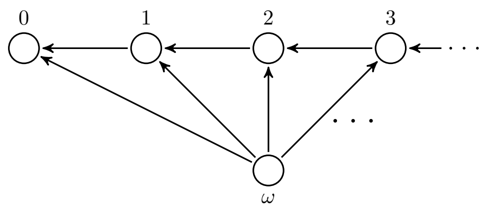

[1]Clicking on the pictures in the digital version of this thesis (with compatible software) will start an animation showing the progress of the iteration of the formula. A pause in the animation indicates that an infinite number of steps has passed. When a state becomes coloured in the animation it means that it is included in the iteration, where each *colour* corresponds to a different *phase* of the iteration, which depends on the *disjunct* that allows the iteration to progress.

We can observe that $(\varphi_x^{\mathbb{S}})^1_\mu = \varphi_x^{\mathbb{S}}(\emptyset) = [\![\Box x]\!]^{\mathbb{S}[x\mapsto\emptyset]} = \{0\}$, $(\varphi_x^{\mathbb{S}})^2_\mu = \varphi_x^{\mathbb{S}}(\{0\}) = \{0,1\}$ and, in general, one can prove that $(\varphi_x^{\mathbb{S}})^n_\mu = \{m \in \omega \mid m < n\}$ for every $n \in \omega$, which implies $(\varphi_x^{\mathbb{S}})^\omega_\mu = \omega$. We can go further and compute $(\varphi_x^{\mathbb{S}})^{\omega+1}_\mu = [\![\Box x]\!]^{\mathbb{S}[x\mapsto\omega]} = \omega\cup\{\omega\}$: $\Box x$ becomes true at point ω only after ω-many steps in the computation of LFP.$\varphi_x^{\mathbb{S}}$. Clearly $(\varphi_x^{\mathbb{S}})^{\omega+1}_\mu = (\varphi_x^{\mathbb{S}})^{\omega+2}_\mu$ and LFP.$\varphi_x^{\mathbb{S}} = (\varphi_x^{\mathbb{S}})^{\omega+1}_\mu$.

Note that it is not always the case that the least fixed point of a formula $\varphi(x)$ in a model corresponds to its full underlying set. Consider for instance the formula $\varphi(x) := \Diamond x \vee \Box\bot$ of Example 1.2.13 and the model $\mathbb{S}$ where the domain is again ω, the valuation is empty for every proposition letter, but now the accessibility relation is $R := \{(0,0),(2,0),(2,1)\} \cup \{(n+1,n) \mid 2 \leq n < \omega\}$: in this case we will have LFP.$\varphi_x^{\mathbb{S}} = \omega\backslash\{0\}$.

1.2.3 Game-theoretic semantics

Satisfiability of a formula φ in a model $\mathbb{S} = (S, (R_d)_{d\in\mathtt{D}}, V)$ can also be established by the existence of winning strategies in a two-player board game, called the *evaluation game* and denoted by $\mathcal{E}(\varphi,\mathbb{S})$. Positions in this game are pairs $(\psi, s) \in \mathit{Sfor}(\varphi) \times S$. From a position of this shape, one of the two players $\exists$ and $\forall$ can make a move according to Table 1.2. Intuitively they have the following goals: $\exists$'s goal is to show that the formula ψ is true at s, while $\forall$'s goal is to show that it is false. Matches are finite or infinite sequences of positions that are consistent with the rules of the game, with finite matches being lost by the player who got stuck, that is, who could not make a move.

Before stating the more precise definitions and the winning conditions for infinite matches, we give more intuitions on how to interpret Table 1.2. Consider $\mathcal{E}(\varphi,\mathbb{S})$. A position of the form $(\psi_1 \vee \psi_2, s)$ belongs to $\exists$, who can move to either position (ψ_1, s) or (ψ_2, s): since her goal is to show that $\psi_1 \vee \psi_2$ is true at s, she should move to the disjunct that witnesses the truth of the formula. Similarly and dually, from a position of the form $(\psi_1 \wedge \psi_2, s)$, $\forall$ can move to either (ψ_1, s) or (ψ_2, s), with the intent to show which of the two conjuncts makes the whole formula false at s. From position $(\Diamond_d\psi, s)$ $\exists$ should move to a position (ψ, t), with t a R_d-successor of s that satisfies ψ, in order to show that $\Diamond_d\psi$ is true at s, while from position $(\Box_d\psi, s)$ $\forall$'s strategy is to find a R_d-successor t of s where ψ is false and move to (ψ, t) to show that $\Box_d\psi$ cannot be true at s: in both cases, if $R_d[t] = \emptyset$ one of the players will get stuck and lose. If the game is at a position of the form (p, s) with $p \in FV(\varphi)$, depending on whether or not $s \in V(p)$ it will be $\forall$'s or $\exists$'s turn to move: in any case, the set of admissible moves is empty, so that either player will be stuck, with for instance $\forall$ being stuck and losing the game if $s \in V(p)$, that is, if p is true at s. When the current position of the match is (x, s), with x a bound variable of φ, the game automatically moves to position (δ_x, s), δ_x being the unique subformula of φ where x is bound by $\eta_x x$, and the match will then continue from there: in this case we say that the variable x is *unfolded.* The game also performs an automatic move from a position of the form $(\eta_x x.\delta_x, s)$ to position (δ_x, s). It can happen that one or more bound variables are unfolded infinitely many times during a play of the evaluation game, leading to an infinite match. The winner of such a match will depend on whether the variable with the highest ranking (with respect to the dependency order $\leq_\varphi$ from Definition 1.2.7) is a μ-variable or a ν-variable: $\forall$ will be the winner in the first case, and $\exists$ in the latter (note that this is analogous to the winning conditions of the unfolding game, which is a key observation in the proof of the equivalence between the algebraic and the game-theoretic semantics). We now move to the definitions and give a few examples of matches.

Definition 1.2.15 Given a clean $\mu\mathtt{ML_D}$-formula φ and a model $\mathbb{S}$ we define the *evaluation game* $\mathcal{E}(\varphi, \mathbb{S})$ as a board game with players $\exists$ and $\forall$ moving a token around positions of the form $(\psi, s) \in Sfor(\varphi) \times S$. The rules determining the admissible moves for a certain player at a given position are given in Table 1.2. $\mathcal{E}(\varphi, \mathbb{S})@(\varphi, s)$ denotes the instantiation of this game where the starting position is fixed as (φ, s).

Position	Player	Set of admissible moves
$(\bot, s)$	$\exists$	$\emptyset$
$(\top, s)$	$\forall$	$\emptyset$
(p, s), with $p \in FV(\varphi)$ and $s \notin V(p)$	$\exists$	$\emptyset$
(p, s), with $p \in FV(\varphi)$ and $s \in V(p)$	$\forall$	$\emptyset$
$(\neg p, s)$, with $p \in FV(\varphi)$ and $s \notin V(p)$	$\forall$	$\emptyset$
$(\neg p, s)$, with $p \in FV(\varphi)$ and $s \in V(p)$	$\exists$	$\emptyset$
$(\psi_1 \wedge \psi_2, s)$	$\forall$	$\{(\psi_1, s), (\psi_2, s)\}$
$(\psi_1 \vee \psi_2, s)$	$\exists$	$\{(\psi_1, s), (\psi_2, s)\}$
$(\Diamond_d \psi, s)$	$\exists$	$\{(\psi, t) \mid t \in R_d[s]\}$
$(\Box_d \psi, s)$	$\forall$	$\{(\psi, t) \mid t \in R_d[s]\}$
$(\eta_x x.\delta_x, s)$	$-$	$\{(\delta_x, s)\}$
(x, s), with $p \in BV(\varphi)$	$-$	$\{(\delta_x, s)\}$

Table 1.2: Evaluation game

Definition 1.2.16 Let φ be a clean $\mu\mathtt{ML_D}$ formula and $\mathbb{S}$ a model. A *match* of the game $\mathcal{E}(\varphi, \mathbb{S})$ is a finite or infinite sequence of positions

$$\Sigma = (\varphi_i, s_i)_{i<\kappa}, \quad \kappa \leq \omega$$

which is consistent with the rules of the evaluation game, that is, Σ is a path through the game graph given in Table 1.2. A *full match* is an infinite match, or a finite match in which one of the player got stuck at the last position. Full matches will be referred to simply as *matches*, while a match that is not full will be called *partial*. Given an infinite match Σ, we let $Unf^{\infty}(\Sigma) \subseteq BV(\varphi)$ denote the set of bound variables of φ that are unfolded infinitely often during Σ.

We recall that a *strategy* for a player $\Pi \in \{\exists, \forall\}$ in an initialised game is a method that dictates which move Π should play depending on the position of the game. A strategy is *winning* for Π if every match of the game started at the given position is won by Π whenever Π adopts this strategy. A position is *winning* for Π if Π has a winning strategy for the game initialized in that position.

Remark Let φ be a clean $\mu\mathtt{ML_D}$-formula and $\mathbb{S}$ a model. For any infinite match Σ of $\mathcal{E}(\varphi, \mathbb{S})$, the set $Unf^{\infty}(\Sigma)$ is finite and directed with respect to the dependency order of Definition 1.2.7 (that is, for any $x, y \in Unf^{\infty}(\Sigma)$ there is $z \in Unf^{\infty}(\Sigma)$ such that $x \leq_{\varphi} z$ and $y \leq_{\varphi} z$). From this it follows that $Unf^{\infty}(\Sigma)$ has a maximum.

Definition 1.2.17 Let φ be a clean $\mu\mathtt{ML}_\mathtt{D}$ formula. The winning conditions of the game $\mathcal{E}(\varphi, \mathbb{S})$ are as follows:

- finite matches are lost by the payer who got stuck,
- an infinite match Σ is won by $\exists$ if $\max(\mathit{Unf}^\infty(\Sigma))$ is a ν-variable, and by $\forall$ otherwise.

The set of winning positions for $\exists$ in $\mathcal{E}(\varphi, \mathbb{S})$ is denoted by $\mathrm{Win}_\exists(\mathcal{E}(\varphi, \mathbb{S}))$.

Example 1.2.18 As a trivial example, consider the formula $\eta x.x$. Given a pointed model $(\mathbb{S}, s)$, a match of $\mathcal{E}(\eta x.x, \mathbb{S})@(\eta x.x, s)$ consists of an infinite sequence of automatic moves: $(\eta x.x, s)(x, s)(x, s)\ldots$, with the winner being $\exists$ if $\eta = \nu$ and $\forall$ otherwise. We observe then that $\mu x.x$ is equivalent to $\bot$, while $\nu x.x$ is equivalent to $\top$.

Example 1.2.19 Consider the formula $\varphi(x) := \Diamond x \vee \Box\bot$, an arbitrary pointed model $(\mathbb{S}, s)$ and the game $\mathcal{E}(\mu x.\varphi, \mathbb{S})@(\mu x.\varphi, s)$. First of all, note that from a position of the form $(\Box\bot, t)$ it would be $\forall' s$ turn to move: if $R[t] = \emptyset$ he would be stuck and lose, otherwise he would have to move to $(\bot, u)$, for some $u \in R[t]$, where $\exists$ would be stuck by the rules of the game. In other words, $(\Box\bot, t)$ is a winning position for $\exists$ for any $t \in S$ with $R[t] = \emptyset$. On the other hand, a position of the form $(\Diamond x, t)$ belongs to $\exists$, who has to move to (x, u) for some $u \in R[t]$ if it exists, otherwise she is stuck: in the first case the game then automatically moves to $(\Diamond x \vee \Box\bot, u)$. We observe then that from a position $(\Diamond x \vee \Box\bot, t)$ $\exists$ should move to $(\Diamond x, t)$ and then to (x, u) if $R[t] \neq \emptyset$, and to $(\Box\bot, t)$ otherwise. In the first case the game continues from $(\Diamond x \vee \Box\bot, u)$ in a similar manner. If $\exists$ wants to win she should move to $(\Box\bot, v)$ as soon as a blind state v occurs in the game: if this never happens the μ-variable x will be unfolded infinitely often and she will lose the game. Similar considerations show that $\exists$ has a winning strategy in $\mathcal{E}(\mu x.\varphi, \mathbb{S})@(\mu x.\varphi, s)$ iff there is a finite path starting from s.

We now state the fundamental result regarding the equivalence of the two semantics for the modal μ-calculus.

Theorem 1.2.20 (Adequacy) *Let φ be a clean $\mu\mathtt{ML}_\mathtt{D}$-formula. Then for all pointed models $(\mathbb{S}, s)$:*

$$s \in [\![\varphi]\!]^{\mathbb{S}} \text{ if and only if } (\varphi, s) \in \mathit{Win}_\exists(\mathcal{E}(\varphi, \mathbb{S})).$$

We refer the reader to [19] for a detailed proof of this result.

1.2.4 Bisimulation

We conclude this chapter with some notes on bisimulations and bisimulation invariance.

Definition 1.2.21 Let $\mathbb{S} = (S, (R_d)_{d\in\mathtt{D}}, V)$ and $\mathbb{S}' = (S', (R'_d)_{d\in\mathtt{D}}, V')$ be two models of the same type $\mathtt{D}$. A *bisimulation* of type $\mathtt{D}$ between $\mathbb{S}$ and $\mathbb{S}'$ is a relation $Z \subseteq S \times S'$, satisfying, for every $(s, s') \in Z$:

- for all $p \in \mathtt{PROP}$, $s \in V(p)$ iff $s' \in V(p)$;
- for every $d \in \mathtt{D}$, for every $t \in R_d[s]$ there is a $t' \in R'_d[s']$ such that $(t, t') \in Z$;

- for every $d \in \mathtt{D}$, for every $t' \in R'_d[s']$ there is a $t \in R_d[s]$ such that $(t, t') \in Z$.

If there is a bisimulation Z between two models $\mathbb{S}$ and $\mathbb{S}'$ and $(s, s') \in Z$, we say that the states s and s' are bisimilar and write $\mathbb{S}, s \underline{\leftrightarrow}_{\mathtt{D}} \mathbb{S}', s'$, or just $\mathbb{S}, s \underline{\leftrightarrow} \mathbb{S}', s'$ when $\mathtt{D}$ is clear.

Like modal logic, the modal μ-calculus enjoys the property that the truth of its formulas is invariant under bisimulation.

Theorem 1.2.22 (Bisimulation Invariance) *Let $\mathbb{S}$ and $\mathbb{S}'$ be two models of the same type $\mathtt{D}$ such that $\mathbb{S}, s \underline{\leftrightarrow} \mathbb{S}', s'$ for some $s \in S$ and $s' \in S'$. Then, for every $\varphi \in \mu\mathtt{ML}_{\mathtt{D}}$:*

$$\mathbb{S}, s \Vdash \varphi \textit{ iff } \mathbb{S}', s' \Vdash \varphi.$$

An immediate consequence of this result is that whenever a $\mu\mathtt{ML}_{\mathtt{D}}$ formula is satisfiable in a pointed model $(\mathbb{S}, s)$, it is satisfiable at the root of a tree model that can be obtain by *unravelling* the original model $\mathbb{S}$ from the state s. The next definition generalises the unravelling construction that we assume the reader to be familiar with.

Definition 1.2.23 Let κ be a countable cardinal with $1 \leq \kappa \leq \omega$, and $(\mathbb{S}, s)$ be a pointed model of type $\mathtt{D}$. A *κ-path through* $\mathbb{S}$ is a finite sequence of the form

$$s_0 d_1 k_1 s_1 \cdots s_{n-1} d_n k_n s_n \quad (n \geq 0),$$

where $s_i \in S, d_i \in \mathtt{D}$ and $0 < k_i < \kappa$ for each i, and such that $s_{i+1} \in R_{d_{i+1}}[s_i]$ for each $i < n$. $\mathit{Paths}^{\kappa}(\mathbb{S})$ denotes the set of all such paths, and $\mathit{Paths}^{\kappa}_s(\mathbb{S})$ denotes the set of those starting at s. Given a κ-path ρ, we let $\mathit{last}(\rho) \in S$ denote its last element.

The tree model $\mathbb{E}_{\kappa}(\mathbb{S}, s) = (\mathit{Paths}^{\kappa}_s(\mathbb{S}), (R^{\kappa}_d)_{d \in \mathtt{D}}, V^{\kappa})$ is the *κ-expansion* of $\mathbb{S}$ *around* s, where:

$$\begin{array}{lll} V^{\kappa}(p) & := & \{\rho \in \mathit{Paths}^{\kappa}_s(\mathbb{S}) \mid \mathit{last}(\rho) \in V(p)\}; \\ R^{\kappa}_d & := & \{(\rho, \rho{\cdot}dkt) \in \mathit{Paths}^{\kappa}_s(\mathbb{S}) \times \mathit{Paths}^{\kappa}_s(\mathbb{S}) \mid (\mathit{last}(\rho), t) \in R_d, k < \kappa\}. \end{array}$$

Proposition 1.2.24 For any cardinal κ with $1 \leq \kappa \leq \omega$ the graph of the function *last* is a bisimulation between $(\mathbb{E}_{\kappa}(\mathbb{S}, s), s)$ and $(\mathbb{S}, s)$.

Remark Observe that the κ-expansion of a model is κ-expanded, which means that, for every state s and distinct $d, d' \in \mathtt{D}$, $R_d[s] \cap R_{d'}[s] = \emptyset$, and every $t \in R_d[s]$ has at least $\kappa - 1$ many bisimilar siblings $t' \in R_d[s]$. Note also that the unravelling of a model can be identified with its 1-expansion.

From the previous proposition and bisimulation invariance the following theorem follows.

Theorem 1.2.25 (Tree Model Property) *Let $\varphi \in \mu\mathtt{ML}_{\mathtt{D}}$: if φ is satisfiable, then it is satisfiable at the root of a tree model.*

Actually, something even better is true: whenever φ is satisfiable, it is satisfiable at the root of a finitely branching tree model. Again, we refer the reader to [19] for the proof of the next theorem.

Theorem 1.2.26 (Bounded Tree Model property) *Let $\varphi \in \mu\mathtt{ML}_{\mathtt{D}}$: if φ is satisfiable, then it is satisfiable at the root of a finitely branching tree model, in which every state has at most $|\varphi|$-many successors.*

Chapter 2

Closure ordinals up to ω

In this chapter we begin our exploration of closure ordinals of formulas of the modal μ-calculus. Essentially, given a formula φ, we are interested in how many iterations of $\varphi_x^{\mathbb{S}}$, starting from $\emptyset$, are needed in order for this function to converge to its least fixed point across all models. More precisely, we are looking for the least α such that $(\varphi_x^{\mathbb{S}})_\mu^\alpha = (\varphi_x^{\mathbb{S}})_\mu^{\alpha+1}$ for every model $\mathbb{S}$.

As we will see, not every formula has a closure ordinal. On the other hand, there are interesting classes of formulas for which a closure ordinal always exists. For instance, consider formulas that are *continuous in x*. In relation to the modal μ-calculus the property of continuity in a variable x has been studied extensively [9, 10, 11] and it is a standard result that every continuous formula converges to its least fixed point in at most ω iterations, a property that is known as *constructivity*. We will talk more about the connection between continuity and constructivity in Section 2.4, where we also present the syntactic characterisation of the property of continuity over finitely branching models.

Another interesting class of formulas is that of $\aleph_1$-continuous formulas, introduced in [12] together with its syntactic characterisation. $\aleph_1$-continuous formulas converge to their least fixed point in at most ω_1 many steps, ω_1 being the first uncountable ordinal. Closure ordinals greater than ω will be discussed in Chapter 3.

The properties of continuity and $\aleph_1$-continuity for μ-calculus formulas have been shown to be *decidable* in, respectively, [9, 10, 11] and [12]. However, whether it is decidable if a μ-calculus formula always reaches its least fixed point in at most ω or ω_1 steps is an open question. Otto [17] showed that the *boundedness* problem for a modal logic formula is decidable, that is, given a formula φ of modal logic it is decidable whether there exists a $n \in \omega$ such that $(\varphi_x^{\mathbb{S}})_\mu^n = (\varphi_x^{\mathbb{S}})_\mu^{n+1}$ for every model $\mathbb{S}$. We will discuss bounded formulas in Section 2.3.

After stating the definition of closure ordinal of a μ-calculus formula and providing a few examples in the next section, in Section 2.2 we focus on a class of formulas that all have 0 as their (very trivial) closure ordinal and we provide a syntactic characterisation.

2.1 Definition and examples

Convention 2.1.1 While the following definitions and results would also apply to the polymodal language of the modal μ-calculus, for notational convenience in the rest of this chapter we restrict to the monomodal setting.

Definition 2.1.2 The *closure ordinal* of a μ-calculus formula φ with respect to the variable x is the least ordinal α such that $(\varphi_x^{\mathbb{S}})_\mu^\alpha = (\varphi_x^{\mathbb{S}})_\mu^{\alpha+1}$ for every model $\mathbb{S}$, if it exists. If α is the closure ordinal of φ with respect to x we write $\mathtt{cl}_x(\varphi) = \alpha$; if α is the closure ordinal of some formula, we say that α *is a closure ordinal*.

Convention 2.1.3 From now on we always assume that when a variable x is free in a formula φ, then x is positive in φ. We will also usually leave the reference to the variable x more implicit and say, for instance, that α is the closure ordinal of $\varphi(x)$. Moreover, since the focus of the thesis is on the ordinal approximation of *least* fixed points, and since the free variable x in a formula $\varphi(x)$ will usually be clear from context, for notational convenience we will often write $\varphi_{\mathbb{S}}^\alpha$ instead of $(\varphi_x^{\mathbb{S}})_\mu^\alpha$ (sometimes φ^α if $\mathbb{S}$ is also clear).

When proving results about closure ordinals an equivalent characterisation, given in Proposition 2.1.5, is often useful. To prove its equivalence we will need the following fact, which is easily proven by transfinite induction and using properties of disjoint unions.

Proposition 2.1.4 Let I be an index set, $\varphi(x)$ be a formula and $\mathbb{S}_i = (S_i, R_i, V_i)$ be a model for every $i \in I$. If $\mathbb{S} = (S, R, V)$ is the disjoint union of all the models $\mathbb{S}_i$'s, then, for every $i \in I$ and ordinal α: $\varphi_{\mathbb{S}_i}^\alpha = \varphi_{\mathbb{S}}^\alpha \cap S_i$.

Proposition 2.1.5 An ordinal α is the closure ordinal of $\varphi(x)$ if and only if the following two conditions are satisfied:

(1) for every model $\mathbb{S}$ the least fixed point of $\varphi_x^{\mathbb{S}}$ is always reached in *at most* α steps, that is, $\mathrm{LFP}.\varphi_x^{\mathbb{S}} = \varphi_{\mathbb{S}}^\alpha = \varphi_{\mathbb{S}}^{\alpha+1}$, and

(2) there exists a model $\mathbb{S}$ where the least fixed point of $\varphi_x^{\mathbb{S}}$ is reached in *exactly* α steps, that is, $\mathrm{LFP}.\varphi_x^{\mathbb{S}} = \varphi_{\mathbb{S}}^\alpha \neq \varphi_{\mathbb{S}}^\beta$ for all $\beta < \alpha$.

Proof. We begin with assuming items (1) and (2). From the first item it immediately follows that $\varphi_{\mathbb{S}}^\alpha = \varphi_{\mathbb{S}}^{\alpha+1}$ for every model $\mathbb{S}$. To prove that α is also the least ordinal satisfying this equality, suppose that there is an ordinal γ such that $\varphi_{\mathbb{S}}^\gamma = \varphi_{\mathbb{S}}^{\gamma+1}$ for every model $\mathbb{S}$. We want $\alpha \leq \gamma$. By item (2) let $\mathbb{S}$ be a model where $\varphi_{\mathbb{S}}^\beta \subsetneq \varphi_{\mathbb{S}}^\alpha$ for every $\beta < \alpha$: then $\varphi_{\mathbb{S}}^\gamma = \varphi_{\mathbb{S}}^{\gamma+1}$ implies that it cannot be the case that $\gamma < \alpha$, and we conclude that $\alpha \leq \gamma$.

Now suppose that $\mathtt{cl}_x(\varphi) = \alpha$. We want to prove that items (1) and (2) hold. We focus on the second item since the first one follows immediately. Assume towards a contradiction that for every model $\mathbb{S}$ there exists an ordinal $\beta < \alpha$ such that $\varphi_{\mathbb{S}}^\beta = \varphi_{\mathbb{S}}^{\beta+1}$: in particular for every $\mathbb{S}$ there exists a *least* ordinal $\beta < \alpha$ satisfying this equality. Define the set

$$B := \{\beta < \alpha \mid \text{for some } \mathbb{S}, \beta \text{ is the least ordinal such that } \varphi_{\mathbb{S}}^\beta = \varphi_{\mathbb{S}}^{\beta+1}\}$$

and let $\bar{\beta} := \sup(B)$. Since α is an upper bound of B it follows that $\bar{\beta} \leq \alpha$, and since $\varphi_{\mathbb{S}}^{\bar{\beta}} = \varphi_{\mathbb{S}}^{\bar{\beta}+1}$ for every model $\mathbb{S}$, the assumption that $\alpha = \mathtt{cl}_x(\varphi)$ implies that $\alpha \leq \bar{\beta}$ and finally $\bar{\beta} = \alpha$. We proceed with a case distinction.

If $\alpha = \gamma + 1$ is a successor ordinal, then γ is an upper bound of B and we would have $\alpha = \bar{\beta} \leq \gamma$, which is absurd.

Otherwise, suppose α is a limit. If B is not a cofinal subset of α then there is a $\gamma < \alpha$ which is an upper bound of B, implying $\alpha = \bar{\beta} \leq \gamma$, so assume instead that B is a cofinal subset of α, that is, for all $\gamma < \alpha$ there exists a $\beta \in B$ such that $\gamma \leq \beta$. Since α is a limit and strictly greater than every element of B, this can be strengthened to

$$\text{for all } \gamma < \alpha \text{ there exists a } \beta \in B \text{ such that } \gamma < \beta. \tag{2.1}$$

For each $\gamma < \alpha$ denote by β_γ an element of B that witnesses (2.1), and by $\mathbb{S}_\gamma = (S_\gamma, R_\gamma, V_\gamma)$ a model where β_γ is the least ordinal such that $\varphi_{\mathbb{S}_\gamma}^{\beta_\gamma} = \varphi_{\mathbb{S}_\gamma}^{\beta_\gamma+1}$. Now take the model $\mathbb{S}$ as the disjoint union of all $\mathbb{S}_\gamma$'s: we claim that $\varphi_{\mathbb{S}}^{\gamma} \subsetneq \varphi_{\mathbb{S}}^{\alpha}$ for every $\gamma < \alpha$. Indeed, let $\gamma < \alpha$ be arbitrary: by (2.1) let $\beta_\gamma \in B$ with $\gamma < \beta_\gamma$ be the ordinal such that $\varphi_{\mathbb{S}_\gamma}^{\delta} \subsetneq \varphi_{\mathbb{S}_\gamma}^{\beta_\gamma}$ for every $\delta < \beta_\gamma$. In particular let $s \in \varphi_{\mathbb{S}_\gamma}^{\beta_\gamma}$ such that $s \notin \varphi_{\mathbb{S}_\gamma}^{\gamma}$: by Proposition 2.1.4 then $s \in \varphi_{\mathbb{S}}^{\beta_\gamma} \cap S_\gamma \subseteq \varphi_{\mathbb{S}}^{\alpha}$ and $s \notin \varphi_{\mathbb{S}}^{\gamma} \cap S_\gamma$, finally giving $s \in \varphi_{\mathbb{S}}^{\alpha}$ but $s \notin \varphi_{\mathbb{S}}^{\gamma}$.

In conclusion, against our initial assumption, $\mathbb{S}$ is a model where there is no $\beta < \alpha$ such that $\varphi_{\mathbb{S}}^{\beta} = \varphi_{\mathbb{S}}^{\beta+1}$, which is the desired contradiction. □

From now on, when we say that α is the closure ordinal of $\varphi(x)$ we will interchangeably use both definitions, and we will often prove that $\mathtt{cl}_x(\varphi) = \alpha$ by showing that $\varphi_x^{\mathbb{S}}$ converges to its least fixed point in *at most* α steps on every model $\mathbb{S}$ and by constructing a model $\mathbb{S}$ where convergence happens in *exactly* α steps. We now consider some examples.

Example 2.1.6 In Example 1.2.13 we presented a model where the formula $\varphi := \Diamond x \vee \Box\bot$ converges to its least fixed point in exactly ω steps. To see that φ converges in at most ω many steps on every model, and so that $\mathtt{cl}_x(\varphi) = \omega$, let $\mathbb{S}$ be arbitrary. We want to show $\varphi^\omega = \varphi^{\omega+1}$. We only need to prove that $\varphi^{\omega+1} \subseteq \varphi^\omega$, so let $s \in \varphi^{\omega+1}$, which means that $\mathbb{S}[x \mapsto \varphi^\omega], s \Vdash \Diamond x \vee \Box\bot$. If $\mathbb{S}[x \mapsto \varphi^\omega], s \Vdash \Box\bot$ then $s \in \varphi_x^{\mathbb{S}}(\emptyset) = \varphi^1 \subseteq \varphi^\omega$, otherwise if $\mathbb{S}[x \mapsto \varphi^\omega], s \Vdash \Diamond x$ then there is a $t \in R[s]$ such that $t \in \varphi^\omega$, implying $t \in \varphi^n$ for some $n \in \omega$ and $\mathbb{S}[x \mapsto \varphi^n], s \Vdash \Diamond x$. We can conclude that $s \in \varphi_x^{\mathbb{S}}(\varphi^n) = \varphi^{n+1} \subseteq \varphi^\omega$.

Example 2.1.7 We have mentioned that not every formula has a closure ordinal: a very simple example is the formula $\varphi := \Box x$. Indeed, for every ordinal α we can construct a model where $\Box x$ converges in exactly $\alpha + 1$ steps. Explicitly, define $\mathbb{S}_\alpha$ to be the model with domain $\alpha \cup \{\alpha\}$, accessibility relation $R := \{(\beta + 1, \beta) \mid \beta < \alpha\} \cup \{(\lambda, \beta) \mid \beta < \lambda \leq \alpha, \lambda \text{ a limit ordinal}\}$ and empty valuation for every propositional variable: here it holds that $\varphi_{\mathbb{S}_\alpha}^{\beta} = \beta$ for every $\beta \leq \alpha + 1$, so that $\varphi_x^{\mathbb{S}_\alpha}$ converges to its least-fixed point in exactly $\alpha + 1$ steps. In Example 1.2.14 we showed this for $\alpha = \omega$.

2.2 Normal formulas

In this section we focus on a class of formulas that immediately converge to their least fixed points. These formulas satisfy $\varphi_x^{\mathbb{S}}(\emptyset) = \emptyset$ on every model $\mathbb{S}$, so that they have 0 as their closure ordinal. In other words, these are formulas that are always false in a model whenever the valuation of the variable x is empty. This property is known as *normality* in the variable x and is usually mentioned in relation to other semantic properties, for instance when differentiating between *full* and *complete additivity* [5, 11].

Definition 2.2.1 A function $f : \wp(S)^n \to \wp(S)$ is called *normal in the ith coordinate* if

$$f(S_1, \ldots, S_{i-1}, \emptyset, S_{i+1}, \ldots, S_n) = \emptyset$$

for all $S_1, \ldots, S_n \subseteq S$ and *normal (in the product)* if

$$f(\emptyset, \ldots, \emptyset) = \emptyset.$$

In order to state the definition of what it means for a formula to be normal in a finite set of propositional variables we need to slightly generalise the definition of the function $\varphi_x^{\mathbb{S}}$.

Definition 2.2.2 Let $\varphi \in \mu\mathtt{ML}$ be a formula and $X = \{x_1, \ldots, x_n\} \subseteq \mathtt{PROP}$ a finite set of propositional variables. For every Kripke model $\mathbb{S} = (S, R, V)$, define the function $\varphi_X^{\mathbb{S}} : \wp(S) \to \wp(S)$ by letting $\varphi_X^{\mathbb{S}}(S') := [\![\varphi]\!]^{\mathbb{S}[x_1 \mapsto S', \ldots, x_n \mapsto S']}$ for all $S' \subseteq S$. If $X = \{x\}$ is a singleton we write $\varphi_x^{\mathbb{S}}$ instead of $\varphi_{\{x\}}^{\mathbb{S}}$.

Convention 2.2.3 For a finite set $X = \{x_1, \ldots, x_n\}$ of propositional variables, a model $\mathbb{S}$ and $S' \subseteq S$, we write $\mathbb{S}[X \mapsto S']$ instead of $\mathbb{S}[x_1 \mapsto S', \ldots, x_n \mapsto S']$.

Definition 2.2.4 A formula $\varphi \in \mu\mathtt{ML}$ is *normal in* $X = \{x_1, \ldots, x_n\} \subseteq \mathtt{PROP}$ if, for every model $\mathbb{S}$, $\varphi_X^{\mathbb{S}}$ is normal (in the product). This is equivalent to requiring that $\mathbb{S}[X \mapsto \emptyset], s \not\Vdash \varphi$ for every pointed model $(\mathbb{S}, s)$.

We are going to provide a syntactic characterisation of this property, in the sense that we are going to prove, for every $\mu\mathtt{ML}$-formula φ, that φ is normal in X if and only if it is equivalent to some formula in the following restricted fragment of the language of the modal μ-calculus.

Definition 2.2.5 Given a finite subset $X \subseteq \mathtt{PROP}$, we define the fragment $\mu\mathtt{ML}_X^N$ by the following grammar:

$$\varphi ::= \bot \mid p \mid \varphi \vee \varphi \mid \varphi \wedge \psi \mid \Diamond\varphi \mid \mu x.\varphi' \mid \nu x.\varphi$$

where $p \in X$, $\psi \in \mu\mathtt{ML}$ is an arbitrary formula, and $\varphi' \in \mu\mathtt{ML}_{X \cup \{x\}}^N$.

The first ingredient needed for our characterisation is that every formula in this fragment is indeed normal in X.

Proposition 2.2.6 For every finite $X \subseteq \mathtt{PROP}$: if $\chi \in \mu\mathtt{ML}_X^N$, then χ is normal in X.

Proof. The proof goes by induction on $\chi \in \mu\mathtt{ML}^N_X$: we only consider the cases for the fixed point operators. First suppose $\chi = \mu x.\varphi$, where $\varphi \in \mu\mathtt{ML}^N_{X\cup\{x\}}$: we want to show that $(\mu x.\varphi)^{\mathbb{S}}_X(\emptyset) = \emptyset$. Observe that by inductive hypothesis $\varphi^{\mathbb{S}}_{X\cup\{x\}}(\emptyset) = \emptyset$, or equivalently $\varphi^{\mathbb{S}[X\mapsto\emptyset]}_x(\emptyset) = \emptyset$, so LFP.$\varphi^{\mathbb{S}[X\mapsto\emptyset]}_x = \emptyset$.

Now let $\chi = \nu x.\varphi$, where $\varphi \in \mu\mathtt{ML}^N_X$. Assume $G :=$ GFP.$\varphi^{\mathbb{S}[X\mapsto\emptyset]}_x \neq \emptyset$ towards a contradiction and let $s \in G$. It follows that $s \in G = \varphi^{\mathbb{S}[X\mapsto\emptyset]}_x(G)$ and so that $\mathbb{S}[X \mapsto \emptyset, x \mapsto G], s \Vdash \varphi$. But by inductive hypothesis $\varphi^{\mathbb{S}[x\mapsto G]}_X(\emptyset) = \emptyset$, which gives the desired contradiction. □

In order to also prove that every formula that is normal in X is equivalent to some formula in the fragment $\mu\mathtt{ML}^N_X$ we are going to take advantage of the fact that the modal μ-calculus admits a *disjunctive normal form.*

Convention 2.2.7 For the rest of the section it will be convenient to distinguish an infinite proper subset $\mathtt{VAR} \subseteq \mathtt{PROP}$ of *fixed point variables*: these will not be allowed to appear in a conjunction of literals, but will be the only variables that may be bound by a fixed point operator in a disjunctive formula. We are then dealing with three different sets of variables at the moment: the set $\mathtt{PROP}$ of all propositional variables, a subset $\mathtt{VAR} \subseteq \mathtt{PROP}$ of fixed point variables and a finite subset $X \subseteq \mathtt{PROP}$ for which we are characterising normality.

Definition 2.2.8 Let $\mathtt{VAR} \subseteq \mathtt{PROP}$. The set $\mathtt{CL}$ of literal conjunctions is defined by the following grammar:

$$\pi ::= \bot \mid \top \mid p \mid \neg p \mid \pi \wedge \pi$$

where $p \in \mathtt{PROP}\backslash\mathtt{VAR}$.

The set $\mu\mathtt{DML}$ of *disjunctive* μ-calculus formulas is given by the following grammar:

$$\varphi ::= \bot \mid \top \mid x \mid \varphi \vee \varphi \mid \pi \bullet \nabla\Phi \mid \mu x.\varphi \mid \nu x.\varphi$$

where $x \in \mathtt{VAR}$ and $\pi \in \mathtt{CL}$.

This language is interpreted on Kripke models, where the meaning of $\pi \bullet \nabla\Phi$ is defined as follows. Let $\mathbb{S}$ be an arbitrary model:

$$\begin{array}{lcl} [\![\pi \bullet \nabla\Phi]\!]^{\mathbb{S}} & = & [\![\pi]\!]^{\mathbb{S}} \cap [\![\nabla\Phi]\!]^{\mathbb{S}} \\ [\![\nabla\Phi]\!]^{\mathbb{S}} & = & \{s \in S \mid R[s] \subseteq \bigcup\{[\![\varphi]\!]^{\mathbb{S}} \mid \varphi \in \Phi\}\} \cap \\ & & \{s \in S \mid \text{for all } \varphi \in \Phi, [\![\varphi]\!]^{\mathbb{S}} \cap R[s] \neq \emptyset\} \end{array}$$

In other words, $\mathbb{S}, s \Vdash \nabla\Phi$ if every successor of t satisfies some formula in Φ, and every formula in Φ is true at some successor of s. Note that the operator $\bullet$ semantically behaves as a conjunction, and that that $\nabla\Phi$ can be expressed using the modal operators $\Diamond$ and $\Box$:

$$\nabla\Phi \equiv \Box \bigvee \Phi \wedge \bigwedge \Diamond\Phi.$$

where $\Diamond\Phi := \{\Diamond\varphi \mid \varphi \in \Phi\}$. Conversely, it is also true that $\Diamond\varphi \equiv \nabla\{\varphi, \top\}$ and $\Box\varphi \equiv \nabla\emptyset \vee \nabla\{\varphi\}$.

For the game-theoretic characterisation of the meaning of a disjunctive μ-calculus formula in a model $\mathbb{S}$ we need the definition of a ∇-marking. Given a set of formulas Φ and a point s in a model, a ∇-marking is a map $m : R[s] \to \wp(\Phi)$ such that

- for all $\varphi \in \Phi$ there exists a $t \in R[s]$ such that $\varphi \in m(t)$,
- for all $t \in R[s]$ there exists a $\varphi \in \Phi$ such that $\varphi \in m(t)$.

Table 2.1 provides the set of admissible moves for a player from positions of the form $(\pi \bullet \nabla\Phi, s)$, $(\nabla\Phi, s)$ and $m : R[s] \to \wp(\Phi)$.

Position	Player	Set of admissible moves
$(\pi \bullet \nabla\Phi, s)$	$\forall$	$\{(\pi, s), (\nabla\Phi, s)\}$
$(\nabla\Phi, s)$	$\exists$	$\{m : R[s] \to \wp(\Phi) \mid m \text{ is a } \nabla\text{-marking}\}$
$m : R[s] \to \wp(\Phi)$	$\forall$	$\{(\psi, t) \mid \psi \in m(t)\}$

Table 2.1: Additional rules for the evaluation game

Theorem 2.2.9 (Janin & Walukiewicz, [13]) *There is an effective algorithm that rewrites a modal fixed point formula $\varphi \in \mu\mathtt{ML}$ into an equivalent disjunctive formula $\varphi^d \in \mu\mathtt{DML}$.*

An interesting property of disjunctive formulas that we will soon need is that in order to know whether a disjunctive formula $\nu x.\varphi$ is satisfiable, it is enough to know if $\varphi(\top)$ is satisfiable [13, 19]. We sketch a proof of this fact.

Proposition 2.2.10 If $\nu x.\varphi$ is disjunctive, then $\nu x.\varphi$ is satisfiable if and only if $\varphi(\top)$ is satisfiable.

Proof. We prove the right to left direction. By Theorem 1.2.22 and Proposition 1.2.24 let $(\mathbb{S}, r)$ be an ω-expanded tree such that $\mathbb{S}, r \Vdash \varphi(\top)$. In order to show that $\nu x.\varphi(x)$ is satisfiable we construct a pointed model $(\mathbb{S}', r)$ and define a winning strategy for $\exists$ in $\mathcal{E}(\nu x.\varphi(x), \mathbb{S})@(\nu x.\varphi(x), r)$. By assumption we know that $\exists$ has a winning strategy f in $\mathcal{E}(\varphi(\top), \mathbb{S})@(\varphi(\top), r)$.

<u>Claim.</u> $\exists$ has a winning strategy f in $\mathcal{E}(\varphi(\top), \mathbb{S})@(\varphi(\top), r)$ such that, for all f-guided partial matches $\sigma = (\varphi(\top), r) \cdots (\nabla\Phi, s)$ and for all $t \in R[s]$, there exists a unique $\psi \in \Phi$ such that $\sigma \cdot (\psi, t)$ is an f-guided continuation of σ.

<u>Proof of Claim.</u> To see that the claim holds note that, since $\mathbb{S}$ is ω-expanded, at position $(\nabla\Phi, s)$ $\exists$ can pick a marking $m : R[s] \to \wp(\Phi)$ such that for all $t \in R[s]$, $|m(t)| = 1$. $\lhd$

We can then assume without loss of generality that $\exists$'s winning strategy f in $\mathcal{E}(\varphi(\top), \mathbb{S})$ from position $(\varphi(\top), r)$ satisfies the condition of the claim.

Our goal is to construct a model $\mathbb{S}'$ that satisfies $\nu x.\varphi(x)$. The idea is that we consider all nodes t of $\mathbb{S}$ such that $(\top, t)$ – with $\top$ a substitution instance of x in φ – is an f-reachable position, prune $\mathbb{S}$ from t and attach a copy of $\mathbb{S}$ to t (identifying t with the root of $\mathbb{S}$). From this latter node $\exists$ can use a strategy that is analogous to the strategy f in order to satisfy $\varphi(\top)$, so that we can look at positions t' in this new model such that $(\top, t')$ is f-reachable from $(\varphi(\top), t)$ and repeat the above process. We do this infinitely many times and obtain a model $\mathbb{S}'$ that we claim satisfies $\nu x.\varphi(x)$. The danger is that by pruning the tree after the state t we might cut some states that are essential in some match of $\mathcal{E}(\varphi(\top)), \mathbb{S})$, in the sense that there might be some f-guided partial match $\sigma' = (\varphi(\top), r) \cdots (\psi, t)$ that requires the successors of t in order to be completed and won by $\exists$. The next claim assures that this cannot happen.

Claim. Let $\sigma = (\varphi(\top), r) \cdots (\top, t)$ – with $\top$ a substitution instance of x in φ – be an f-guided match of $\mathcal{E}(\varphi(\top), \mathbb{S})$ from position $(\varphi(\top), r)$. If $\sigma' = (\varphi(\top), r) \cdots (\psi, t)$ is an f-guided partial match of $\mathcal{E}(\varphi(\top), \mathbb{S})$ from position $(\varphi(\top), r)$ where ψ is not a propositional letter, then σ' is a prefix of σ.

Proof of Claim. Suppose this is not the case, so that σ and σ' are such that there is a position (χ, s) with

$$\sigma = (\varphi(\top), r) \cdots (\chi, s)(\chi_0, s') \cdots (\top, t),$$
$$\sigma' = (\varphi(\top), r) \cdots (\chi, s)(\chi_1, s'') \cdots (\psi, t),$$

and $\chi_0 \neq \chi_1$. We argue by a case distinction on the shape of χ that this cannot hold. Obviously χ cannot be a literal, $\top$ or $\bot$, as both matches would end at position (χ, s). If $\chi = \varphi \vee \varphi'$, then (χ, s) is a position for $\exists$ and it must be the case that $\chi_0 = \chi_1$ and $s = s' = s''$ against assumption. Now suppose $\chi = \nabla\Phi$. By our previous claim, for each $u \in R[s]$ there is a unique $\psi \in \Phi$ such that (ψ, u) is f-reachable: since $\mathbb{S}$ is a tree, there is a unique path from s to t, and so there is a unique successor u of s such that $(u, t) \in R^*$, implying that it must be the case that $s' = s''$ and $\chi_0 = \chi_1$. If χ is of the form $\eta y.\chi'$, then $\chi_0 = \chi_1 = \chi'$ and $s = s' = s''$ and there is no branching in this case too; similarly if χ is a variable y. $\lhd$

This guarantees that the construction of $\mathbb{S}'$ as described can be performed. It is also easy to see why $\mathbb{S}', r \Vdash \nu x.\varphi(x)$. $\square$

We state the following corollary for future reference.

Corollary 2.2.11 Let $X \subseteq$ `PROP` be a finite set of propositional variables and $\nu x.\varphi$ be a disjunctive formula. If $\varphi(\top)$ is satisfiable in a model with an empty valuation for every $p \in X$, then $\nu x.\varphi$ is also satisfiable in a model $\mathbb{S}$ where $V(p) = \emptyset$ for every $p \in X$.

For every finite $X \subseteq$ `PROP` we now define a subset $\mu\mathtt{DML}^N_X$ of μ`DML`: it will turn out that every disjunctive formula that is normal in X belongs to this subset. Recall that in Convention 2.2.7 we have distinguished a subset `VAR` $\subseteq$ `PROP` of variables that may be bound by a fixed point operator, but that cannot occur in a conjunction of literals in a disjunctive formula.

Definition 2.2.12 Let $X \subseteq$ `PROP` be a finite set of propositional variables. We define inductively the set $\mathtt{CL}^N_X \subseteq$ `CL` to be the smallest set such that

$\bot \in \mathtt{CL}^N_X$		always,
$p \in \mathtt{CL}^N_X$	if	$p \in X$,
$\pi \wedge \pi' \in \mathtt{CL}^N_X$	if	$\pi \in \mathtt{CL}^N_X$, or $\pi' \in \mathtt{CL}^N_X$, or
		π and π' are inconsistent,

where $p \notin$ `VAR` and we say that $\pi, \pi' \in$ `CL` *are inconsistent* if for some $q \in$ `PROP`: $q \in \pi$ and $\neg q \in \pi'$, or $q \in \pi'$ and $\neg q \in \pi$ (recall that by the definition of `CL` in Definition 2.2.8 such a q cannot be an element of `VAR`).

Define inductively the set $\mu\mathtt{DML}^N_X \subseteq \mu$`DML` to be the smallest set such that

$\bot \in \mu\mathtt{DML}^N_X$		always,
$x \in \mu\mathtt{DML}^N_X$	if	$x \in X$ and $x \in \mathtt{VAR}$,
$\varphi \vee \varphi' \in \mu\mathtt{DML}^N_X$	if	$\varphi \in \mu\mathtt{DML}^N_X$ and $\varphi' \in \mu\mathtt{DML}^N_X$,
$\mu x.\varphi \in \mu\mathtt{DML}^N_X$	if	$\varphi \in \mu\mathtt{DML}^N_{X\cup\{x\}}$,
$\nu x.\varphi \in \mu\mathtt{DML}^N_X$	if	$\varphi \in \mu\mathtt{DML}^N_X$,
$\pi \bullet \nabla\Phi \in \mu\mathtt{DML}^N_X$	if	$\pi \in \mathtt{CL}^N_X$ or there is some $\varphi \in \Phi$ such that $\varphi \in \mu\mathtt{DML}^N_X$.

Lemma 2.2.13 For all $\varphi \in \mu\mathtt{DML}$ and all finite $X \subseteq \mathtt{PROP}$:

$$\varphi \text{ is normal in } X \text{ if and only if } \varphi \in \mu\mathtt{DML}^N_X.$$

Proof. We focus on left to right direction, of which we prove the contrapositive: for all $\varphi \in \mu\mathtt{DML}$ and all finite $X \subseteq \mathtt{PROP}$, if $\varphi \notin \mu\mathtt{DML}^N_X$, then φ is not normal in X. We proceed by induction on $\varphi \in \mu\mathtt{DML}$, considering only the less trivial cases.

Suppose $\varphi = \pi \bullet \nabla\Phi \notin \mu\mathtt{DML}^N_X$: by definition of $\mu\mathtt{DML}^N_X$ it holds that $\pi \notin \mathtt{CL}^N_X$ and for all $\psi \in \Phi$, $\psi \notin \mu\mathtt{DML}^N_X$.

We first prove that, for all $\pi \in \mathtt{CL}$, if $\pi \notin \mathtt{CL}^N_X$, then π is not normal in X. The cases where π is a literal or $\bot$ are easily dealt with. Now suppose $\pi = \pi' \wedge \pi'' \notin \mathtt{CL}^N_X$: then $\pi' \notin \mathtt{CL}^N_X$, $\pi'' \notin \mathtt{CL}^N_X$ and π' and π'' are not inconsistent. This implies that all $p \in X$ and $\bot$ are not conjuncts of π and that π is satisfiable: being a conjunction of literals, π is satisfiable in a model $(\{s\}, \emptyset, V)$ consisting of a single irreflexive point, where for every propositional variable q: $V(q) = \{s\}$ if $q \in \pi$ and $V(q) = \emptyset$ otherwise. This proves that π is not normal in X.

Going back to $\pi \bullet \nabla\Phi$, let $\Phi = \{\varphi_1, \ldots, \varphi_n\}$. By assumption and induction hypothesis, for all $1 \leq i \leq n$, φ_i is not normal in X, hence there is a pointed model $(\mathbb{S}_i, s_i)$ such that $\mathbb{S}_i[X \mapsto \emptyset], s_i \Vdash \varphi_i$. Without loss of generality we assume that the underlying sets $S_1, \ldots, S_n$ of these models are disjoint. Consider the model

$$\mathbb{S} = (\{s\} \cup \bigcup_{1\leq i\leq n} S_i,\ \{(s, s_i) \mid 1 \leq i \leq n\} \cup \bigcup_{1\leq i\leq n} R_i,\ V)$$

where $s \notin \bigcup_i S_i$ and such that $V(p) = \emptyset$ for all $p \in X$ and, for all $q \notin X$: $V(q) = \{s\} \cup \bigcup_i V_i(q)$ if $q \in \pi$ and $V(q) = \bigcup_i V_i(q)$ otherwise. In other worlds, $\mathbb{S}$ is the disjoint union of all the models $(\mathbb{S}_i, s_i)$'s with one more point s added as a "root" under it. We claim that $\mathbb{S}, s \Vdash \pi \bullet \nabla\Phi$. It should be clear that $\mathbb{S}, s \Vdash \pi$. Note that $\mathbb{S}, s_i \Vdash \varphi_i$ iff $\mathbb{S}_i[X \mapsto \emptyset], s_i \Vdash \varphi_i$. Let $t \in R[s]$ be arbitrary: $R[s] = \{s_1, \ldots, s_n\}$, so $t = s_i$ for some $1 \leq i \leq n$ and $\mathbb{S}, s_i \Vdash \varphi_i$, hence $\mathbb{S}, t \Vdash \psi$ for some $\psi \in \Phi$. Now let $\varphi_i \in \Phi$ be arbitrary: in this case we have that s_i is such that $\mathbb{S}, s_i \Vdash \varphi_i$. It follows that $\mathbb{S}, s \Vdash \pi \bullet \nabla\Phi$ and, since $V(p) = \emptyset$ for all $p \in X$, we conclude that $\pi \bullet \nabla\Phi$ is not normal in X.

Now suppose φ is of the form $\mu x.\psi \notin \mu\mathtt{DML}^N_X$, which means that $\psi \notin \mu\mathtt{DML}^N_{X\cup\{x\}}$. By induction hypothesis there exists a pointed model $(\mathbb{S}, s)$ such that $\mathbb{S}[X \cup \{x\} \mapsto \emptyset], s \Vdash \psi$, which implies that $\mathbb{S}[X \mapsto \emptyset], s \Vdash \psi[\bot/x]$. But then in the evaluation game $\mathcal{E}(\mu x.\psi, \mathbb{S}[X \mapsto \emptyset])@(\mu x.\psi, s)$ $\exists$ can use (almost) the same winning strategy f she has in $\mathcal{E}(\psi[\bot/x], \mathbb{S}[X \mapsto \emptyset])@(\psi[\bot/x], s)$: since a position of the form $(\bot, t)$ is never reached in the latter game if $\exists$ uses the strategy f, a corresponding position (x, t) is not f-reachable in the first game, hence f is a winning strategy in $\mathcal{E}(\mu x.\psi, \mathbb{S}[X \mapsto \emptyset])@(\mu x.\psi, s)$. We conclude that $\mu x.\psi$ is not normal in X, since $\mathbb{S}[X \mapsto \emptyset], s \Vdash \mu x.\psi$.

Finally, suppose φ is of the form $\nu x.\psi \notin \mu\mathtt{DML}^N_X$, meaning that $\psi \notin \mu\mathtt{DML}^N_X$. By induction hypothesis there is a pointed model $(\mathbb{S}, s)$ such that $\mathbb{S}[X \mapsto \emptyset], s \Vdash \psi$. Then, since x is positive in ψ, by monotonicity of $\psi^{\mathbb{S}}_x$ it holds that $\mathbb{S}[X \mapsto \emptyset], s \Vdash \psi[\top/x]$. By Corollary 2.2.11 $\nu x.\psi$ is satisfiable in a model with empty valuation for every $p \in X$, so we conclude that $\nu x.\psi$ is not normal in X. $\square$

The following proposition can be proved by induction on $\varphi \in \mu\mathtt{DML}$, considering $\nabla\Phi$ as an abbreviation for $\Box \bigvee \Phi \wedge \bigwedge \Diamond\Phi$.

Proposition 2.2.14 For every $\varphi \in \mu\mathtt{DML}$: if $\varphi \in \mu\mathtt{DML}^N_X$, then $\varphi \in \mu\mathtt{ML}^N_X$.

Finally, we combine the previous results to obtain the desired characterisation of the normal fragment of the modal μ-calculus.

Theorem 2.2.15 *Every formula in $\mu\mathtt{ML}^N_X$ is normal in X. Moreover, there is an effective translation which, given a $\mu\mathtt{ML}$-formula φ, computes an equivalent formula $\varphi^d \in \mu\mathtt{DML}$ such that*

$$\varphi \text{ is normal in } X \text{ iff } \varphi^d \in \mu\mathtt{ML}^N_X.$$

Proof. The first part of the statement follows from Proposition 2.2.6. Now let $\varphi \in \mu\mathtt{ML}$ be normal in X. By Theorem 2.2.9, $\varphi \equiv \varphi^d$ for some $\varphi^d \in \mu\mathtt{DML}$ that can be effectively obtained from φ. Since $\varphi \equiv \varphi^d$, φ^d is normal in X and so by Lemma 2.2.13 $\varphi^d \in \mu\mathtt{DML}^N_X$. By Proposition 2.2.14 then $\varphi^d \in \mu\mathtt{ML}^N_X$ and the statement of the proposition follows. $\square$

We conclude this section by mentioning that it is decidable whether a μ-calculus formula is normal.

Proposition 2.2.16 The problem whether a given $\mu\mathtt{ML}$-formula φ is normal in a finite set of variables $X \subseteq \mathtt{PROP}$ is decidable.

Proof. By a theorem of Emerson and Jutla it is decidable whether a given $\mu\mathtt{ML}$-formula is satisfiable [8]. As a consequence of this fact, it is decidable whether two formulas φ and ψ are equivalent. To obtain the statement of the proposition we observe that a formula $\varphi(x_1, \ldots, x_n)$ is normal in $X = \{x_1, \ldots, x_n\}$ if and only if $\varphi(\bot, \ldots, \bot) \equiv \bot$. $\square$

2.3 Bounded formulas

We now start to take into consideration formulas with more interesting closure ordinals. In this section we focus on formulas that need finitely many steps in order to converge to their least fixed point across all models, so that they have a finite closure ordinal.

Definition 2.3.1 A formula $\varphi \in \mu\mathtt{ML}$ is *bounded in x* if for some $n \in \omega$ the least fixed point of $\varphi^{\mathbb{S}}_x$ is always reached in n steps on every model $\mathbb{S}$. More precisely, φ is bounded in x if there is an $n \in \omega$ such that $\varphi^n_{\mathbb{S}} = \varphi^{n+1}_{\mathbb{S}}$ for every model $\mathbb{S}$.

Remark Observe that if a formula $\varphi(x)$ is such that for all models $\mathbb{S}$ there is an $n \in \omega$ such that $\varphi^n_{\mathbb{S}} = \varphi^{n+1}_{\mathbb{S}}$, then there is an $n \in \omega$ such that $\varphi^n_{\mathbb{S}} = \varphi^{n+1}_{\mathbb{S}}$ for all $\mathbb{S}$. To see why, by contraposition for each $n \in \omega$ let $\mathbb{S}_n$ be a model where $\varphi^n_{\mathbb{S}_n} \neq \varphi^{n+1}_{\mathbb{S}_n}$: then the disjoint union $\mathbb{S}$ of all these models would be such that $\varphi^n_{\mathbb{S}} \neq \varphi^{n+1}_{\mathbb{S}}$ for all $n \in \omega$.

Example 2.3.2 Fix $n \in \omega$ and let $\varphi := \Box x \wedge \Box^n \bot$, where the formula $\Box^n \psi$ is inductively defined by $\Box^0 \psi := \psi$ and $\Box^{m+1}\psi := \Box(\Box^m \psi)$. We claim that $\mathtt{cl}_x(\varphi) = n$ and hence that φ is bounded. To see this, first take an arbitrary model $\mathbb{S}$: one can show that, for every $0 \leq m \leq n$, if $\mathbb{S}[x \mapsto \varphi^n], s \Vdash \Box x \wedge \Box^m \bot$, then $s \in \varphi^m_{\mathbb{S}} \subseteq \varphi^n_{\mathbb{S}}$, so that $\varphi^{n+1}_{\mathbb{S}} = \varphi^n_{\mathbb{S}}$. A model where convergence happens in exactly n steps has domain $n = \{0, \ldots, n-1\}$ and accessibility relation $R := \{(m+1, m) \mid m < n-1\}$.

We continue with some syntactic considerations. Denote by $\mu\mathtt{ML}_\infty$ the extension of the modal μ-calculus where infinite disjunctions and conjunctions are allowed.

Definition 2.3.3 Let $\varphi(x) \in \mu\mathtt{ML}$ be a formula. By ordinal induction we define the formula $\widehat{\varphi}_\alpha \in \mu\mathtt{ML}_\infty$ as follows:

$$\begin{array}{lll} \widehat{\varphi}_0 & := & \bot, \\ \widehat{\varphi}_{\alpha+1} & := & \varphi[\widehat{\varphi}_\alpha / x], \\ \widehat{\varphi}_\lambda & := & \bigvee\limits_{\alpha<\lambda} \widehat{\varphi}_\alpha, \end{array}$$

where λ is an arbitrary limit ordinal.

Observe that $\widehat{\varphi}_n$ is a $\mu\mathtt{ML}$-formula for every $n \in \omega$, but in general $\widehat{\varphi}_\alpha$ is not, since it involves infinite disjunctions. Before the next proposition, we note that for a model $\mathbb{S}$, an ordinal γ and a set $\{\varphi_\alpha \mid \alpha < \gamma\}$ of formulas, the meaning of $\bigvee_{\alpha<\gamma} \varphi_\alpha$ in $\mathbb{S}$ is defined by letting $[\![\bigvee_{\alpha<\gamma} \varphi_\alpha]\!]^{\mathbb{S}} := \bigcup_{\alpha<\gamma} [\![\varphi_\alpha]\!]^{\mathbb{S}}$.

Proposition 2.3.4 Let $\varphi(x) \in \mu\mathtt{ML}$ be a formula and α an ordinal. Then $\varphi^\alpha_{\mathbb{S}} = \varphi^{\alpha+1}_{\mathbb{S}}$ on every model $\mathbb{S}$ if and only if $\mu x.\varphi \equiv \widehat{\varphi}_\alpha$.

Proof. This is an immediate consequence of $\varphi^\alpha_{\mathbb{S}} = [\![\widehat{\varphi}_\alpha]\!]^{\mathbb{S}}$ for every ordinal α and model $\mathbb{S}$, which is provable by induction on α. □

Corollary 2.3.5 A formula $\varphi \in \mu\mathtt{ML}$ is bounded in x if and only if there is an $n \in \omega$ such that $\mu x.\varphi \equiv \widehat{\varphi}_n$.

Now, recall that $\mathtt{ML}$ denotes the language of (basic) modal logic: we say that a formula φ is $\mathtt{ML}$-*definable* if there is a $\psi \in \mathtt{ML}$ such that $\varphi \equiv \psi$. The statement of the next proposition is found in [17].

Proposition 2.3.6 A formula $\varphi \in \mathtt{ML}$ is bounded in x iff $\mu x.\varphi$ is $\mathtt{ML}$-definable.

Finally, consider the next result by Otto [17].

Theorem 2.3.7 (Otto) *The following problem is decidable: given a formula in the modal μ-calculus, decide whether this formula can equivalently be expressed in plain modal logic.*

By Proposition 2.3.6 it immediately follows that whether a formula φ of modal logic is bounded in some variable x is also decidable [17].

2.4 Constructive and continuous formulas

We now move to formulas that need (at most) ω steps in the approximation of their least fixed point in order to converge, so that their closure ordinal is at most ω.

Definition 2.4.1 A formula φ is *constructive in* x if the least fixed point of $\varphi_x^{\mathbb{S}}$ is always reached in ω steps on every model $\mathbb{S}$. More precisely, φ is constructive in x if $\varphi_{\mathbb{S}}^{\omega} = \varphi_{\mathbb{S}}^{\omega+1}$ for every model $\mathbb{S}$.

Example 2.4.2 Clearly, if a formula is bounded in x it is also constructive in x. A more interesting example of a formula that is constructive (but not bounded) in x is the formula $\Diamond x \vee \Box\bot$ from Example 2.1.6.

An important property that is often mentioned in relation to constructivity is that of *continuity*. Essentially, a formula φ is continuous in a variable x if, whenever φ is satisfied in a pointed model $(\mathbb{S}, s)$, a finite subset of $V(p)$ is enough in order for φ to be satisfied in the same model. For the next definition, we write $\mathbb{S}[x{\upharpoonright}F]$ as an abbreviation for $\mathbb{S}[x \mapsto V(p) \cap F]$.

Definition 2.4.3 A formula $\varphi \in \mu\mathtt{ML}$ is *continuous* in x if

$$\mathbb{S}, s \Vdash \varphi \text{ iff } \mathbb{S}[x{\upharpoonright}F], s \Vdash \varphi, \text{ for some finite subset } F \subseteq S$$

for every pointed model $(\mathbb{S}, s)$.

Example 2.4.4 An example of a formula that is continuous in x is again $\Diamond x \vee \Box\bot$, or also simply $\Diamond x$: whenever $\Diamond x$ is satisfied in a pointed model $(\mathbb{S}, s)$ it is enough to restrict the valuation of x to a single successor of s in order for $\Diamond x$ to be true at s in $\mathbb{S}$. Another example is the formula $\mu z.\Diamond z \vee x$, which expresses the existence of a point where x is true at a finite distance from the current state (this can be argued along the same lines of Example 1.2.19). Finally, an example of a continuous formula that involves a conjunction is $\mu z.(\Box\bot \vee \Diamond z) \wedge x$, expressing the existence of a finite path of points that satisfy x.

Example 2.4.5 A non-example is the formula $\Box x$. Consider a pointed model $(\mathbb{S}, s)$ where $\mathbb{S}, s \Vdash \Box x$ and $R[s]$ is infinite: clearly there is no finite subset $F \subseteq S$ such that $\mathbb{S}[x{\upharpoonright}F], s \Vdash \Box x$.

Remark Observe that we do not need to specify that the formula φ is monotone in x (or that x occurs positively in φ) in Definition 2.4.3, because continuity in x implies monotonicity in x.

We mention that by a result of Fontaine [9] there is a nice syntactic fragment $\mu\mathtt{ML}_x^C$ of the modal μ-calculus characterising the property of continuity in x, given in the next definition. Note how the presence of the $\Box$ operator is heavily restricted, and similarly the presence of the greatest fixed point operator: for instance, a formula like $\nu z.\Diamond z \wedge x$, which expresses the existence of an *infinite path* where x is always true, is clearly not continuous in x.

Definition 2.4.6 Given a finite set $X \subseteq \mathtt{PROP}$, define the fragment $\mu\mathtt{ML}_X^C$ by the following grammar:

$$\varphi ::= p \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Diamond\varphi \mid \mu z.\varphi'$$

where $p \in X$, ψ is a X-free formula and $\varphi' \in \mu\mathtt{ML}^C_{X\cup\{z\}}$. In case X is a singleton, say, $X = \{x\}$ we will write $\mu\mathtt{ML}^C_x$ rather than $\mu\mathtt{ML}^C_{\{x\}}$.

Moreover, there are effective translations mapping a formula $\varphi \in \mu\mathtt{ML}$ to a formula $\varphi^C \in \mu\mathtt{ML}^C_x$ in such a way that φ is continuous in x if and only if $\varphi \equiv \varphi^C$: a direct translation is given in [10], translations involving automata can be found in [9] and [11].

Theorem 2.4.7 (Fontaine & Venema) *Every formula in $\mu\mathtt{ML}^C_x$ is continuous in x. Moreover, there is an effective translation which, given a $\mu\mathtt{ML}$-formula φ, computes a formula $\varphi^C \in \mu\mathtt{ML}^C_x$ such that*

$$\varphi \text{ is continuous in } x \text{ iff } \varphi \equiv \varphi^C,$$

and it is decidable whether a given formula φ is continuous in x.

The property of continuity in x derives its name from its connection with the topological notion of *Scott continuity*. In our context it is enough to note that a function $f : C \to C'$ between two complete lattices $\mathbb{C}$ and $\mathbb{C}'$ is *Scott continuous* if $f(\bigvee D) = \bigvee' f[D]$ whenever $D \subseteq C$ is directed, where a subset $D \subseteq C$ of a complete lattice is said to be *directed* if for every $d_1, d_2 \in D$ there is a $d \in D$ such that $d_1 \leq d$ and $d_2 \leq d$. The reader can find a full proof of the next proposition in [9].

Proposition 2.4.8 A formula $\varphi \in \mu\mathtt{ML}$ is continuous in x if and only if the map $\varphi^{\mathbb{S}}_x : \wp(S) \to \wp(S)$ is Scott continuous.

Using this equivalent definition it is easy to prove that, if a formula φ is continuous in x, then it is constructive in x, so that $\mathtt{cl}_x(\varphi) \leq \omega$.

Proposition 2.4.9 If $\varphi \in \mu\mathtt{ML}$ is continuous in x, then it is constructive in x.

Proof. Let $\mathbb{S}$ be an arbitrary model. We want to prove that $\varphi^{\omega}_{\mathbb{S}} = \varphi^{\omega+1}_{\mathbb{S}}$. First note that, as a consequence of Proposition 1.1.8, the set $\{\varphi^n_{\mathbb{S}} \mid n < \omega\}$ is directed. Then

$$\varphi^{\omega+1}_{\mathbb{S}} = \varphi^{\mathbb{S}}_x(\varphi^{\omega}_{\mathbb{S}}) = \varphi^{\mathbb{S}}_x(\bigcup_{n<\omega} \varphi^n_{\mathbb{S}}) = \bigcup_{n<\omega} \varphi^{\mathbb{S}}_x(\varphi^n_{\mathbb{S}}) = \bigcup_{n<\omega} \varphi^{n+1}_{\mathbb{S}} \subseteq \varphi^{\omega}_{\mathbb{S}},$$

where the equality $\varphi^{\mathbb{S}}_x(\bigcup_{n<\omega} \varphi^n_{\mathbb{S}}) = \bigcup_{n<\omega} \varphi^{\mathbb{S}}_x(\varphi^n_{\mathbb{S}})$ is given by the assumption that $\varphi^{\mathbb{S}}_x$ is Scott continuous. □

The converse of the last proposition is notoriously not true.

Example 2.4.10 For every $n \in \omega$, the formula $\Box x \wedge \Box^n \bot$ (as was shown in Example 2.3.2) is bounded in x, implying that it is constructive in x: however, it is clearly not continuous in x. Another example is the formula $\nu z.\Diamond z \wedge x$, which is actually normal in x.

Now, say that two formulas $\varphi, \psi \in \mu\mathtt{ML}$ are *μx-equivalent* (notation: $\varphi \equiv_{\mu x} \psi$) if $\mu x.\varphi \equiv \mu x.\psi$. Note that, while the formulas in the previous example are not continuous in x, they are μx-equivalent to some formula that is continuous in x: indeed $\Box x \wedge \Box^n \bot \equiv_{\mu x} \Box^n \bot$ and $\nu z.\Diamond z \wedge x \equiv_{\mu x} \bot$. On the basis of this observation, an interesting open question regarding the link between continuity and constructivity has been formulated [11].

Question (Venema) *Can we find, for any formula $\varphi \in \mu\mathtt{ML}$ which is constructive in x, a μx-equivalent formula ψ that is continuous in x?*

While unfortunately we do not provide an answer to this question here, we present another example of a formula that is constructive in x, but not continuous in x. We believe that this example is interesting for the following reason: it involves a formula that is constructive in x, *but neither continuous in x nor bounded in x* (thus refuting the tempting idea that every constructive formula is either continuous or bounded) but, at the same time, *it is μx-equivalent to a formula that is continuous in x*, thus making a further argument in favour of a positive answer to Venema's question[1].

Example 2.4.11 Consider the formula

$$\varphi := (\Box\bot \wedge \neg p) \vee (\Box x \wedge \Box\neg p \wedge p) \vee (\Diamond x \wedge \Box p \wedge p).$$

which clearly is not continuous in x. We now show that $\mathtt{cl}(\varphi(x)) = \omega$, which implies that φ is constructive (but not bounded) in x.

We begin by proving that $\varphi^{\omega}_{\mathbb{S}} = \varphi^{\omega+1}_{\mathbb{S}}$ in any model $\mathbb{S}$. Let $\mathbb{S}$ be arbitrary and let $s \in \varphi^{\omega+1}_{\mathbb{S}}$, that is, $\mathbb{S}[x \mapsto \varphi^{\omega}_{\mathbb{S}}], s \Vdash \varphi$. We proceed by case distinction as to which disjunct of φ is satisfied by s to prove $s \in \varphi^{\omega}_{\mathbb{S}}$. If $s \Vdash \Box\bot \wedge \neg p$, then $s \in \varphi^{\mathbb{S}}_x(\emptyset) \subseteq \varphi^1_{\mathbb{S}} \subseteq \varphi^{\omega}_{\mathbb{S}}$. Otherwise, if $s \Vdash \Box x \wedge \Box\neg p \wedge p$, then $R[s] \subseteq \varphi^{\omega}_{\mathbb{S}} \cap (S \backslash V(p))$, hence every $t \in R[s]$ is such that $\mathbb{S}[x \mapsto \varphi^n_{\mathbb{S}}], t \Vdash \varphi \wedge \neg p$ for some $n \in \omega$, implying $\mathbb{S}[x \mapsto \varphi^n_{\mathbb{S}}], t \Vdash \Box\bot \wedge \neg p$: then $R[s] \subseteq [\![\Box\bot \wedge \neg p]\!]^{\mathbb{S}[x \mapsto \emptyset]} \subseteq \varphi^1_{\mathbb{S}}$ and so $s \in \varphi^2_{\mathbb{S}} \subseteq \varphi^{\omega}_{\mathbb{S}}$. Finally, if $s \Vdash \Diamond x \wedge \Box p \wedge p$, then $t \in \varphi^{\omega}_{\mathbb{S}}$ for some $t \in R[s]$, implying $t \in \varphi^n_{\mathbb{S}}$ for some $n \in \omega$ and $s \in \varphi^{n+1}_{\mathbb{S}} \subseteq \varphi^{\omega}_{\mathbb{S}}$. This finishes the case distinction, and we conclude that $s \in \varphi^{\omega}_{\mathbb{S}}$.

For a model where φ converges to its least fixed point in exactly ω steps, let $\mathbb{S} = (S, R, V)$ be the model where $S = \omega$, $R = \{(n+1, n) \mid n \in \omega\}$ and $V(p) = \omega \backslash \{0\}$. Here $\varphi^n_{\mathbb{S}} = \{m \in \omega \mid m < n\}$ holds, so that $\varphi^n_{\mathbb{S}} \subsetneq \varphi^{\omega}_{\mathbb{S}}$ for every $n < \omega$.

Now consider the formula

$$\psi := \Box\bot \vee (\Diamond x \wedge \Box\neg p \wedge p \wedge \Box\Box\bot) \vee (\Diamond x \wedge \Box p \wedge p),$$

which is continuous in x: we claim that $\varphi \equiv_{\mu x} \psi$. We need to show that $\mathrm{LFP}.\varphi^{\mathbb{S}}_x = LFP.\psi^{\mathbb{S}}_x$ on every model $\mathbb{S}$. Let $\mathbb{S}$ be arbitrary. Observe that, since both formula are constructive in x, it is enough to show that $\varphi^{\omega}_{\mathbb{S}} = \psi^{\omega}_{\mathbb{S}}$.

We start by proving that, for all $n \in \omega$, $\varphi^n \subseteq \psi^n$. For $n = 0$ this is obvious, so inductively assume that $\varphi^n \subseteq \psi^n$: we want $\varphi^{n+1} \subseteq \psi^{n+1}$. Let $s \in \varphi^{n+1}$, meaning that $\mathbb{S}[x \mapsto \varphi^n], s \Vdash \varphi$: we proceed by case distinction as to which disjunct of φ is satisfied by s to prove $s \in \psi^{n+1}$. The only interesting case is when $\mathbb{S}[x \mapsto \varphi^n], s \Vdash \Box x \wedge \Box\neg p \wedge p$. If $R[s] = \emptyset$ then $s \Vdash \Box\bot$, so $s \in \psi^1 \subseteq \psi^{n+1}$. Otherwise, observe that $R[s] \subseteq \varphi^n \cap (S \backslash V(p))$, so for every $t \in R[s]$ there is an $m < n$ such that $\mathbb{S}[x \mapsto \varphi^m], t \Vdash \varphi \wedge \neg p$. Then it must be the case that $t \Vdash \Box\bot \wedge \neg p$ for every $t \in R[s]$: in particular $R[t] = \emptyset$ for every such t. Since

[1] After the final draft of this chapter was concluded, we found out that Czarnecki also provided an example of a formula that is constructive in x, but neither bounded in x nor continuous in x: the formula $\Diamond\Diamond x \vee (\Box x \wedge \Box\Box\bot)$ [7]. While it is not mentioned in his presentation, this formula is μx-equivalent to $\Diamond\Diamond x \vee (\Box\bot \wedge \Box\Box\bot) \vee (\Diamond x \wedge \Box\Box\bot)$, which is continuous in x, thus providing yet another argument in favour of a positive answer to Venema's question.

$R[s] \neq \emptyset$, then $s \Vdash \Diamond x \wedge \Box\neg p \wedge p \wedge \Box\Box\bot$, so $\mathbb{S}[x \mapsto \varphi^n], s \Vdash \psi$. By induction hypothesis and monotonicity: $s \in \psi_x^{\mathbb{S}}(\varphi_n) \subseteq \psi_x^{\mathbb{S}}(\psi^n) = \psi^{n+1}$.

Now we prove that for all $n \in \omega$, $\psi^n \subseteq \varphi^n$. Similarly as above, we prove that $\psi^{n+1} \subseteq \varphi^{n+1}$ under the inductive hypothesis that $\psi^n \subseteq \varphi^n$. Let $s \in \psi^{n+1}$, equivalently $\mathbb{S}[x \mapsto \psi^n], s \Vdash \psi$: we proceed by case distinction as to which disjunct of ψ is satisfied by s to prove $s \in \varphi^{n+1}$. We consider two cases here, as the third is immediate. First suppose $s \Vdash \Box\bot$: if $s \in V(p)$, then $\mathbb{S}[x \mapsto \emptyset], s \Vdash \Box x \wedge \Box\neg p \wedge p$, and so $s \in \varphi^1 \subseteq \varphi^{n+1}$; otherwise, if $s \notin V(p)$, then $\mathbb{S}[x \mapsto \emptyset], s \Vdash \Box\bot \wedge \neg p$, so again $s \in \varphi^1 \subseteq \varphi^{n+1}$. For the second case suppose $\mathbb{S}[x \mapsto \psi^n], s \Vdash \Diamond x \wedge \Box\neg p \wedge p \wedge \Box\Box\bot$ and note that in particular this implies that $n > 0$. Since $s \Vdash \Box\neg p \wedge \Box\Box\bot$, then $R[t] = \emptyset$ and $t \notin V(p)$ for all $t \in R[s]$, meaning that $\mathbb{S}[x \mapsto \emptyset], t \Vdash \Box\bot \wedge \neg p$ for all such t. In other words $R[s] \subseteq \varphi^1$, which gives $\mathbb{S}[x \mapsto \varphi^1], s \Vdash \Box x \wedge \Box\neg p \wedge p$, so $s \in \varphi^2 \subseteq \varphi^{n+1}$.

We have obtained that $\varphi^\omega_{\mathbb{S}} = \psi^\omega_{\mathbb{S}}$ for every $\mathbb{S}$, and it follows that φ and ψ are indeed μx-equivalent.

We conclude this chapter by making some observations about closure ordinals and continuity in the setting of finitely branching models. Our first remark is that in this setting more formulas have a closure ordinal, as the next example shows.

Example 2.4.12 As was shown in Example 2.1.7, the formula $\varphi := \Box x$ does not have a closure ordinal on the class of all models. This is not the case in the setting of finitely branching models, as ω is the closure ordinal of $\Box x$ on finitely branching models. Indeed, if $\mathbb{S}$ is an arbitrary finitely branching model and $s \in \varphi^{\omega+1}_{\mathbb{S}} = [\![\Box x]\!]^{\mathbb{S}[x \mapsto \varphi^\omega_{\mathbb{S}}]}$, then $R[t] \subseteq \varphi^\omega_{\mathbb{S}}$, but $R[t]$ being finite implies that $R[t] \subseteq \varphi^n_{\mathbb{S}}$ for some n. Convergence in exactly ω steps happens, for instance, in the model from Example 1.2.13.

More can be said about the $\Box$ operator on finitely branching models. For instance, it is easy to show that whenever φ is continuous in x on finitely branching models, then $\Box\varphi$ is also continuous in x on finitely branching models. This immediately suggests the following definition and theorem.

Definition 2.4.13 Given a finite set $X \subseteq \texttt{PROP}$, define the fragment $\mu\texttt{ML}^D_X$ by the following grammar:

$$\varphi ::= p \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Diamond\varphi \mid \Box\varphi \mid \mu z.\varphi'$$

where $p \in X$, ψ is a X-free formula and $\varphi' \in \mu\texttt{ML}^D_{X\cup\{z\}}$. In case X is a singleton, say, $X = \{x\}$ we will write $\mu\texttt{ML}^D_x$ rather than $\mu\texttt{ML}^D_{\{x\}}$.

Theorem 2.4.14 *Every formula in $\mu\texttt{ML}^D_x$ is continuous in x on finitely branching models. Moreover, there is an effective translation which, given a $\mu\texttt{ML}$-formula φ, computes a formula $\varphi^D \in \mu\texttt{ML}^D_x$ such that*

φ is continuous in x on finitely branching models iff $\varphi \equiv \varphi^D$,

and it is decidable whether a given formula φ is continuous in x on finitely branching models.

The reason for the name $\mu\texttt{ML}^D_X$ for this fragment of the μ-calculus is that it coincides with the fragment corresponding to the *finite depth property* (or finite path property) of [10, 11].

Definition 2.4.15 A formula $\varphi \in \mu\mathtt{ML}$ has the *finite depth property* for $x \in \mathtt{PROP}$ if φ is monotone in x, and, for every tree model $(\mathbb{S}, s)$,

$$\mathbb{S}, s \Vdash \varphi \text{ iff } \mathbb{S}[x{\upharpoonright}U], s \Vdash \varphi, \text{ for some noetherian subtree } U \subseteq S,$$

where we call U a *noetherian subtree* of $\mathbb{S}$ if it is downward closed and it contains no infinite paths.

Indeed, the proof of Theorem 2.4.14 is easily obtained from the proof of the corresponding theorem in [11] (in particular, via a straightforward variation of Proposition 6.9 in [11], using the fact that the modal μ-calculus enjoys the bounded tree model property). The next corollary immediately follows.

Corollary 2.4.16 A formula $\varphi \in \mu\mathtt{ML}$ has the finite depth property for x if and only if it is continuous in x on finitely branching models.

Chapter 3

Closure ordinals above ω

In this chapter we show how to construct formulas that need a very large number of steps in order to converge to their least fixed point.

In the first section we start by recalling Czarnecki's definition of a formula φ_α with closure ordinal α for every $\omega \leq \alpha < \omega^2$ [6]. In this construction an important role is played by the subformula $\Box x$: while its presence is often problematic when closure ordinals are involved, in this case its occurrence in φ_α is essential to control the number of iterations. We conclude the section by mentioning the result by Gouveia and Santocanale [12] that closure ordinals are closed under ordinal sum, from which it also follows that every ordinal strictly less than ω^2 is a closure ordinal.

In Section 3.2 we discuss the ordinal ω^2. It is a result by Afshari and Leigh [1] that formulas in the alternation-free fragment of the modal μ-calculus cannot have a closure ordinal equal or greater than ω^2. While this suggests a way to find formulas having at least ω^2 as their closure ordinal, a very intuitive candidate turns out to converge to its least fixed point in many more steps than desired.

In order to reach ω^2 and greater countable closure ordinals we will move to the setting of bidirectional models: for every $n \in \omega$, a formula φ_n with closure ordinal ω^n on bidirectional models is constructed in Section 3.3, while in Section 3.4 and Section 3.5 we adopt methodologies similar to Czarnecki's and Gouveia and Santocanale's to show that in this class of models every ordinal strictly less than ω^ω is the closure ordinal of some formula.

Finally, in Section 3.6 we discuss ω_1 as a closure ordinal. In particular, we mention Gouveia and Santocanale's $\aleph_1$-continuous fragment of the modal μ-calculus [12] and we show that the candidate for a formula with closure ordinal ω^2 actually has ω_1 as its closure ordinal. This fact will constitute the basis for an interesting open question about closure ordinals of the modal μ-calculus.

3.1 Closure ordinals below ω^2

We start by showing how to control the number of iterations that a formula needs in order to converge to its least fixed point across all models, up to (not including) ω^2 many steps. Note that every ordinal strictly below ω^2 is of the form $\omega \cdot n + m$, where $n, m \in \omega$: we discuss some ideas and intuitions on how to achieve these closure ordinals before stating Czarnecki's result [6].

Suppose that we want to define a formula that converges in $\omega \cdot 2$ steps. We will keep the discussion informal and try to simultaneously come up with a formula φ and a model with domain $\omega \cdot 2$ where φ converges in exactly $\omega \cdot 2$ many steps: our goal is to make sense of the reason why the formula φ defined in the end has indeed closure ordinal $\omega \cdot 2$. Intuitively we can think that we want some formula φ whose least fixed point iteration builds up two copies of ω, step by step, so that $\varphi^\alpha = \{\beta \mid \beta < \alpha\} = \alpha$ for every $\alpha < \omega \cdot 2$. The formula φ will consist of a finite number of disjuncts, with each disjunct witnessing a different phase of the iteration. We can build the first copy of ω starting from the number 0, by making it a blind state in a model, and let $\Box\bot$ be the first disjunct of φ, so that 0 will immediately be added to the iteration. We can continue building the first copy of ω by adding finite ordinals to the model: we let each positive number n see its predecessor $n - 1$, and let $\Diamond x$ be the second disjunct of φ, so that every finite ordinal is added to the iteration of φ one by one. So far we have the formula $\Box\bot \vee \Diamond x$ from Example 2.1.6, which we know has closure ordinal ω, and the following model[1] (as a reminder on how to compute the iterations of a formula in a model we suggest to look back at Examples 1.2.13 and 1.2.14).

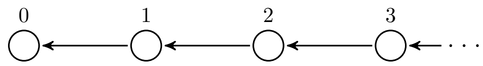

Suppose now that we are at step ω in the iteration of φ, so that every finite ordinal has been added to the iteration, and we want to start building the second copy of ω starting from the ordinal ω itself. Since we want the state ω to be added to the iteration after every finite ordinal has already been added, the formula $\Box\bot \vee \Diamond x$ is not suitable: if we make the state ω a dead end it will be added to the iteration at the first step through the disjunct $\Box\bot$, and otherwise it will be included via the disjunct $\Diamond x$ as soon as one of its successors is added to the iteration at some finite step. At this point it is tempting to consider the formula $\Box\bot \vee \Box x$ instead, together with the following model.

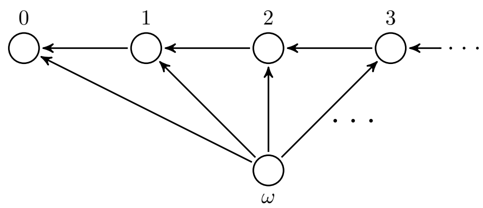

Here again we have that each finite ordinal in the model is added to the iteration in finitely many steps, building up the first copy of ω one after the other, and only after ω many steps can the state ω be added to the iteration, since ω needs to satisfy the disjunct $\Box x$. Clearly however, the problem now is that the formula $\Box\bot \vee \Box x$ does not have a closure ordinal, as it involves the subformula $\Box x$ in an unrestricted way. The solution, as given by Czarnecki [6], is to assign different *colours* to the two copies of ω. Let, for instance, $r := p$ and $b := \neg p$

[1] We recall that clicking on the pictures in the digital version of this thesis (with compatible software) will start an animation showing the progress of the iteration of the formula. A pause in the animation indicates that an infinite number of steps has passed. When a state becomes coloured in the animation it means that it is included in the iteration, where each *colour* corresponds to a different *phase* of the iteration, which depends on the *disjunct* that allows the iteration to progress.

be two literals corresponding to the colours red and blue, respectively, and we let the first copy of ω be red, while the second copy blue, meaning that all (and only) the finite ordinals will make p true in the model we are building. Now, after adding the number 0 to the iteration by using the disjunct $\Box\bot$, we can consider the formula $(\Diamond x \wedge r \wedge \Box r)$ as the second disjunct of φ, so that every finite red ordinal is added one by one, like before. Since ω is not red, it will not be added during this phase. In order for ω to be added to the iteration after every finite ordinal we consider the third disjunct of φ to be $(\Box x \wedge b \wedge \Box r)$: here the occurrence of $\Box x$ is safely restricted by $b \wedge \Box r$. We then define the last disjunct of φ to be $(\Diamond x \wedge b \wedge \Box b)$, so that the blue ordinals $\omega+1, \omega+2, \dots$ will become part of the iteration after the state ω. Note that this last formula in particular is always false at state ω, which falsifies $\Box b$, so that there is no danger that this state will be added to the iteration through $\Diamond x$. The situation is depicted in the following figure.

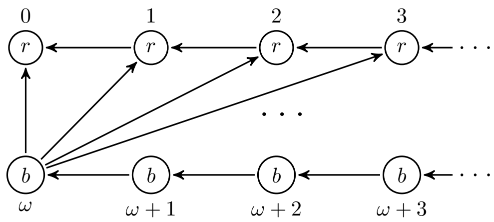

We have thus arrived at the formula

$$\varphi := \Box\bot \vee (\Diamond x \wedge r \wedge \Box r) \vee (\Box x \wedge b \wedge \Box r) \vee (\Diamond x \wedge b \wedge \Box b).$$

The model we have constructed is, in a sense, a minimal model where φ converges in exactly $\omega \cdot 2$ steps, but what guarantees that this formula will converge in *at most* $\omega \cdot 2$ steps on *every model*? Intuitively, suppose that we were to add one more state $\omega \cdot 2$ to the model in the previous picture, and we want it to be included in the iteration of φ in strictly more than $\omega \cdot 2$ steps. One crucial observation to see why this is not possible is that we cannot exploit the subformula $\Box x$ to force this new point to be added to the iteration after every current point in the model has already been included, because in φ the occurrence of $\Box x$ is restricted by $b \wedge \Box r$: if, for instance, we make the state $\omega \cdot 2$ see everything in the model pictured above, it will falsify $b \wedge \Box r$, so that it will not be included in the iteration through $\Box x$ after $\omega \cdot 2$ iterations.

To fix the intuition, as a further example suppose that we now want to construct a formula with closure ordinal $\omega \cdot 3$. In this case we can consider three colours $r := p_1 \wedge p_2$ (red), $b := \neg p_1 \wedge p_2$ (blue) and $g := \neg p_1 \wedge \neg p_2$ (green) to assign to three copies of ω and define

$$\varphi := \Box\bot \vee (\Diamond x \wedge r \wedge \Box r) \vee (\Box x \wedge b \wedge \Box r) \vee (\Diamond x \wedge b \wedge \Box b) \vee (\Box x \wedge g \wedge \Box b) \vee (\Diamond x \wedge g \wedge \Box g).$$

The three copies of ω are built by the iteration as follows: *(i)* the disjunct $\Box\bot$ takes care of the first red state 0; *(ii)* the disjunct $(\Diamond x \wedge r \wedge \Box r)$ builds up the first red copy of ω; *(iii)* the disjunct $(\Box x \wedge b \wedge \Box r)$ allows to move from the first copy of ω to the first point ω of the second blue copy of ω; *(iv)* the disjunct $(\Diamond x \wedge b \wedge \Box b)$ builds up, step by step, the second copy of ω; *(v)* similarly, the disjuncts $(\Box x \wedge g \wedge \Box b)$ and $(\Diamond x \wedge g \wedge \Box g)$ allow to move from the second to the third green copy of ω, and to build it step by step, respectively. The model is pictured below.

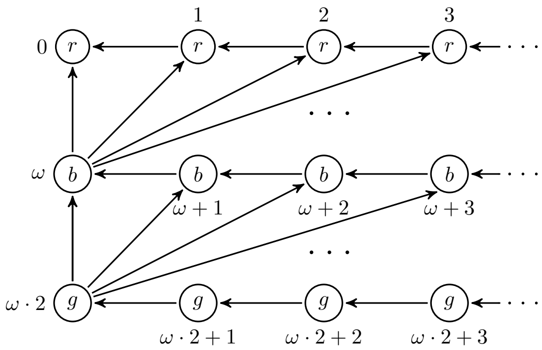

If we want to add a finite number of iterations after one or more copies of ω, for instance to achieve the closure ordinal $\omega + 3$, we can again consider colours $r := p$ and $b := \neg p$ and the formula $\Box\bot \vee (\Diamond x \wedge r \wedge \Box r) \vee (\Box x \wedge b \wedge \Box r) \vee (\Box x \wedge b \wedge \Box b \wedge \Box\Box r) \vee (\Box x \wedge b \wedge \Box b \wedge \Box\Box b \wedge \Box\Box\Box r)$. The disjuncts of shape $(\Box x \wedge \bigwedge_{j=0}^{i} \Box^j b \wedge \Box^{i+1} r)$ express the finite distance between the state where they are satisfied and the red copy of ω. As a model where this formula converges to its least fixed point in exactly $\omega + 3$ steps we can take the following.

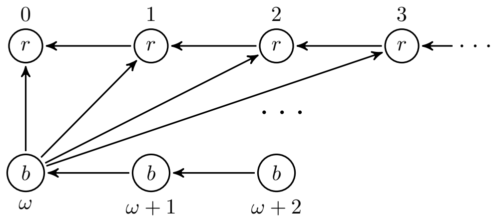

Figure 3.1: Adding finitely many iterations

We now move to the general definition of a formula $\varphi_\alpha(x)$ with closure ordinal α for every $\omega \leq \alpha < \omega^2$, as it is stated in [6]. We will follow Czarnecki's terminology and talk of *fuses* instead of *colours* (this will also be convenient for later sections, where we will use colours instead, or both fuses and colours).

Definition 3.1.1 For every $n \in \omega$ we define the *fuse* f_n as the conjunction of literals $f_n := \bigwedge_{0<i\leq n} \neg p_i \wedge p_{n+1}$.

For example, $f_0 = p_1$ and $f_2 = \neg p_1 \wedge \neg p_2 \wedge p_3$. Clearly, $f_i \wedge f_j \equiv \bot$ for all $i \neq j$.

Definition 3.1.2 Let $0 < n < \omega$ and $m \in \omega$. For an ordinal $\alpha = \omega \cdot n + m$, define the formula φ_α as follows:

$$\begin{aligned}
\psi_{\omega\cdot n} &:= \bigvee_{i=0}^{n-1} (\Diamond x \wedge f_i \wedge \Box f_i) \vee \bigvee_{i=0}^{n-2} (\Box x \wedge f_{i+1} \wedge \Box f_i), \\
\psi_{\omega\cdot n+m} &:= \psi_{\omega\cdot n} \vee \bigvee_{i=0}^{m-1} (\Box x \wedge \bigwedge_{j=0}^{i} \Box^j f_n \wedge \Box^{i+1} f_{n-1}), \\
\varphi_{\omega\cdot n+m} &:= \psi_{\omega\cdot n+m} \vee \Box\bot.
\end{aligned}$$

A detailed proof of the next theorem can be found in [6]

Theorem 3.1.3 (Czarnecki) *For every $\omega \leq \alpha < \omega^2$, the closure ordinal of $\varphi_\alpha(x)$ is α.*

To conclude this section we mention that Gouveia and Santocanale [12] have found a way to construct a formula ψ with closure ordinal $\alpha+\beta$ whenever α and β are closure ordinals. Since every ordinal below ω^2 has the shape $\omega \times n + m$, and $0, 1$ and ω are closure ordinals, their result also implies that every ordinal strictly below ω^2 is the closure ordinal of some formula. In Section 3.5 we will show that a similar construction can be carried out in the setting of bidirectional models: together with the result of Section 3.3, it will follow that every ordinal strictly below ω^ω is a closure ordinal on bidirectional models.

3.2 The case of ω^2

In the previous section we presented Czarnecki's definition of a formula $\varphi_\alpha(x)$ with closure ordinal α for every $\alpha < \omega^2$. We now discuss the possibility of finding a formula with closure ordinal ω^2 and our (failed) attempt at defining it: to this end we start by making a few observations.

The first observation involves a very interesting result due to Afshari and Leigh regarding the ordinal ω^2 in connection to the modal μ-calculus [1]. In order to state it we first need to define the modal μ-calculus alternation hierarchy.

Definition 3.2.1 A formula $\varphi \in \mu\mathtt{ML}$ belongs to the classes Σ_0 and Π_0 if 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 $\varphi(x), \psi \in \Sigma_{n+1}$ (Π_{n+1}), then $\varphi(\psi) \in \Sigma_{n+1}$ (Π_{n+1}), provided the free variables of ψ do no become bound by fixed point operators in φ.

The *alternation-free fragment* of the μ-calculus is the closure of $\Sigma_1 \cup \Pi_1$ under Boolean connectives, modal operators and substitutions that preserve the alternation depth. In other words, a formula belongs to the alternation-free fragment of the μ-calculus if there is no alternation, or nesting, of fixed point operators μ and ν.

Theorem 3.2.2 (Afshari & Leigh) *Let $\varphi(x)$ be an alternation-free formula: if $\mathtt{cl}_x(\varphi)$ exists then $\mathtt{cl}_x(\varphi) < \omega^2$.*

The original theorem in [1] actually requires φ to be guarded, that is, in every subformula of φ of the form $\eta y.\delta$, every occurrence of the variable y in δ must occur under the scope of a modal operator. We mention that every formula φ of the modal μ-calculus is equivalent to a guarded one, and it can be checked that the latter can be taken as alternation-free if so is the original formula, so that the guardedness assumption is without loss of generality: a proof of this fact can be found in [15, 19]. The original statement of this theorem is also more general, since in [1] the closure ordinal of a formula is defined with respect to a finite set of variables, rather than a single one.

To better understand Theorem 3.2.2 we must note that a slightly different notation is adopted in [1]: essentially, while in this thesis we say that 'α is the closure ordinal of $\varphi(x)$', if we adopted the conventions of [1] we would say that 'α is the closure ordinal of $\mu x.\varphi$' instead. It follows that in our setting we can interpret this result as stating that *if $\mu x.\varphi$ is an alternation-free formula of the modal μ-calculus, then the closure ordinal of $\varphi(x)$ is strictly less than ω^2, if it exists.* The result by Afshari and Leigh proves very useful in our context, as it suggests that if we are trying to construct formulas with closure ordinals at least ω^2 we should look for formulas that involve the greatest fixed point operator ν in a non-trivial manner.

The second observation concerns the role of colours and of the $\Box$ operator in the subformula $\Box x$ of Czarnecki's formulas: informally, these provide a way to move the iteration from a copy of ω to the next *finitely many times* (as only finitely many colours can be used in a formula). Since our goal at the moment is to find a formula with closure ordinal ω^2, we should try to generalise this process and find a formula that allows moving from one copy of ω to the next copy ω many times, building up through its iteration the ordinal ω^2 (which can be visualised as ω many copies of ω, one next to the other).

Finally, we assemble these facts with the goal of designing a formula that makes it possible to move the fixed point iteration from a copy of ω to the next, without the need of infinitely many colours. Consider the model in the following picture, depicting one copy of ω where p is true at every positive finite ordinal (as we will see, we only need p as a colour), and the formula $\varphi := \Box\bot \vee (\Diamond x \wedge p)$. Suppose that we are at step ω in the iteration of φ, so that every state in the model is included in the current valuation of x.

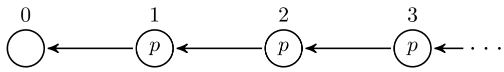

We want to add a new state ω to this model and add a disjunct to the formula φ that makes it possible for the state ω to be added to the iteration of φ only *after* the step ω. In other words, we want a formula that becomes true at state ω after ω many steps of the iteration. We note that at step ω in the computation of the least fixed point of φ infinitely many states of the model have been added to the iteration of φ, and that these states build an infinite $R^{\smile}$-path (where $R^{\smile}$ denotes the converse of the relation R), so that we could try to add a disjunct that expresses the *existence of an infinite path* of points where x is always true. The easiest way to achieve this is to move to the bimodal language of the modal μ-calculus, so that we will consider two relations R_a and R_b, together with modalities $\Diamond_a$ and $\Diamond_b$. We can then consider the following model.

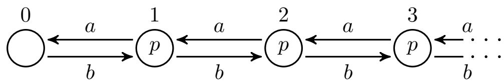

The formula expressing the existence of an infinite R_b-path of points where x and p are always true starting from the R_b-next state – which crucially involves the greatest fixed point operator – is $\nu y.\Diamond_b(p \wedge x \wedge y)$: observe that this formula is satisfied by the state 0 after ω many steps of the iteration. We can now add the state ω to the model in the following way.

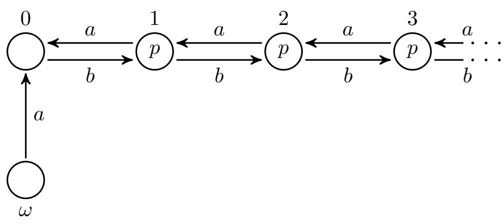

Consider then the formula

$$\varphi := \Box_a\bot \vee (\Diamond_a x \wedge p) \vee (\neg p \wedge \Diamond_a(\nu y.\Diamond_b(y \wedge x \wedge p)))$$

and note that the subformula $(\neg p \wedge \Diamond_a(\nu y.\Diamond_b(y \wedge x \wedge p)))$ is true at state ω after ω many iterations of φ. We can now build the complete model corresponding to the ordinal ω^2 as in Figure 3.2. After the state ω has been added to the iteration

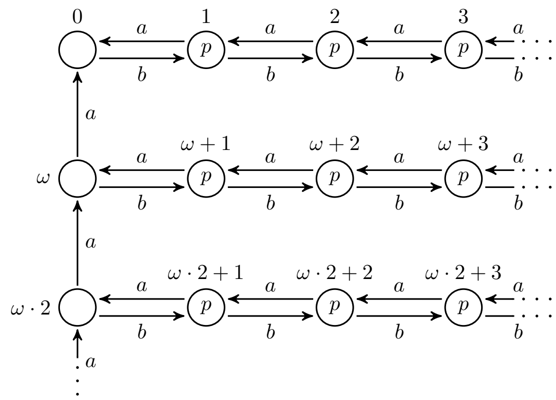

Figure 3.2: Model corresponding to the ordinal ω^2

through the disjunct $(\neg p \wedge \Diamond_a(\nu y.\Diamond_b(y \wedge x \wedge p)))$, the states $\omega+1, \omega+2, \ldots$ will be included, one by one, through the disjunct $(\Diamond_a x \wedge p)$. Once every state of the form $\omega + n$ has been added to the iteration of φ, the state $\omega \cdot 2$ will make $(\neg p \wedge \Diamond_a(\nu y.\Diamond_b(y \wedge x \wedge p)))$ true, and this process starts again. We thus observe that every limit ordinal of the form $\omega \cdot n$ is added to the iteration of φ in this model at step $\omega \cdot n + 1$: for all $\alpha < \omega^2$ we have that $\varphi^\alpha = \{\beta \mid \beta < \alpha\}$ and the least fixed point of φ in this model is φ^{ω^2}.

Unfortunately, the model we have built has the very strong property that R_b is (almost) the converse of R_a: in fact, as we will see in Section 3.6, it is possible to define a model where φ converges to its least fixed point in exactly ω_1 steps, where ω_1 is the first uncountable ordinal. As a model where φ converges in strictly more than ω^2 steps consider the one depicted in Figure 3.3: observe that here the infinite R_b-path $t_0t_1t_2\ldots$ witnesses the truth of $(\neg p \wedge \Diamond_a(\nu y.\Diamond_b(y \wedge x \wedge p)))$ at state ω^2 only after ω^2 many steps in the iteration of φ. On the other hand, this result suggests that we might be able to obtain ω^2 as a closure ordinal in the setting of bidirectional models, that is, models with two accessibility relations, one being the converse of the other. Indeed, in the next section we find a formula φ_n with closure ordinal ω^n on bidirectional models for all $n \in \omega$.

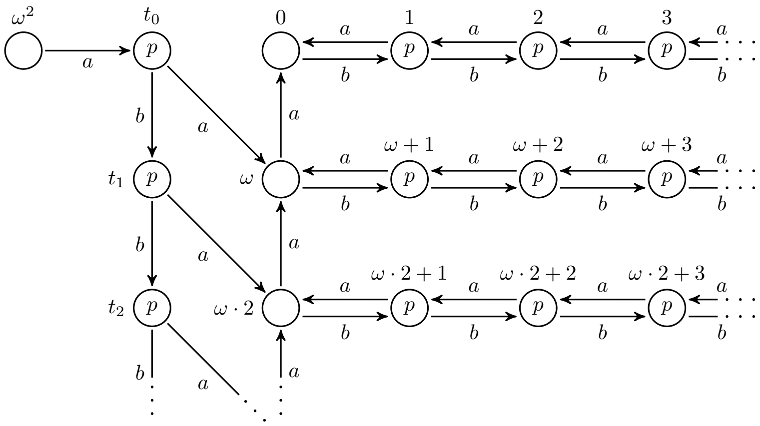

Figure 3.3: Counterexample to φ having closure ordinal ω^2

3.3 Bidirectional models: the case of ω^n

In this section we restrict our focus to the class of bidirectional models: these are models with two relations R and R' such that $R' = R^{\smallsmile} := \{(s,t) \mid (t,s) \in R\}$ is the converse of R. Since one relation is completely determined by the other it is sufficient to only specify one of them, so that we will denote a bidirectional model by $\mathbb{S} = (S, R, V)$. We will adopt the *basic temporal language* [3] (with fixed point operators) as the language of formulas to be interpreted on bidirectional models: this consists of a diamond modality F for the relation R and a diamond modality P for the relation $R^{\smallsmile}$. The intended interpretation of a formula $F\varphi$ is 'φ is true at some *future* state', while that of $P\varphi$ is 'φ is true at some *past* state'. We can define box operators as usual by letting $G\varphi := \neg F \neg \varphi$ and $H\varphi := \neg P \neg \varphi$.

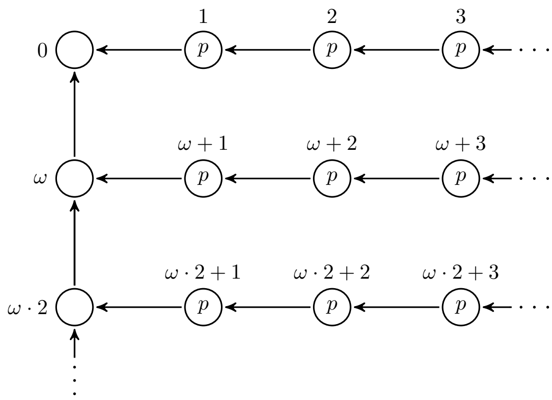

Figure 3.4: Bidirectional model corresponding to ω^2

In this setting we can rewrite the formula $\Box_a \bot \vee (\Diamond_a x \wedge p) \vee (\neg p \wedge \Diamond_a (\nu y. \Diamond_b (y \wedge x \wedge p)))$ from the previous section as $\varphi := G\bot \vee (Fx \wedge p) \vee (\neg p \wedge F(\nu y. P(y \wedge x \wedge p)))$, where we have identified R_a with R, and R_b with its converse. We can also redraw the model of Figure 3.2 of the previous section as in Figure 3.4, where we have only represented the arrows for the relation R.

We observe that the steps of the iteration of φ in this model are exactly as in the previous section: *(i)* first the state 0 is added to the iteration through the formula $G\bot$; *(ii)* then every finite ordinal, one by one, through the formula $(Fx \wedge p)$; *(iii)* once every finite ordinal is added to the iteration, an infinite $R^{\smile}$-path is formed witnessing the truth of $(\neg p \wedge F(\nu y.P(y \wedge x \wedge p)))$ at the state ω, which will then also be added to the iteration; *(iv)* the process continues with the states $\omega + 1, \omega + 2, \ldots$, which will be included in the iteration, one at a time, through the disjunct $(Fx \wedge p)$, and so on.

On the other hand, in this setting the model of Figure 3.3 is not a counterexample anymore: as seen in Figure 3.5, the path $t_0t_1t_2\ldots$ is now such that $t_i \in R[t_{i+1}]$ for every $i \geq 0$, so that, as soon as the state t_0 is added to the iteration of φ through the disjunct $(Fx \wedge p)$ at step $\omega + 2$, every other state in the path will be added in a finite number of steps also through $(Fx \wedge p)$, and the state ω^2 will be included at step $\omega \cdot 2 + 1$.

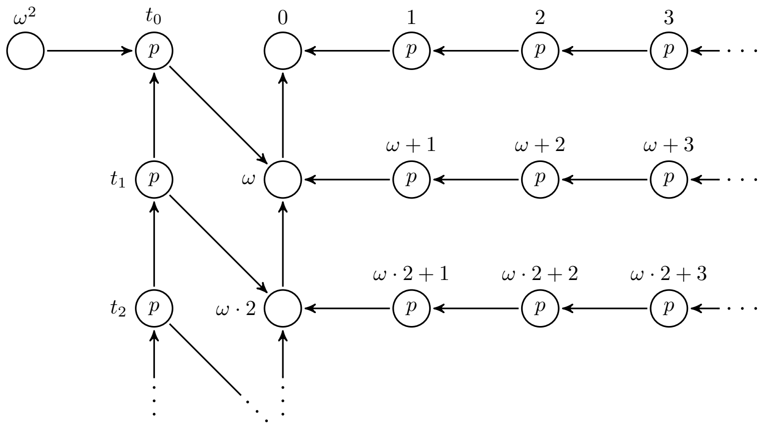

Figure 3.5: The counterexample of Figure 3.3 in this setting

We claim then that the formula

$$G\bot \vee (Fx \wedge p) \vee (\neg p \wedge F(\nu y.P(y \wedge x \wedge p)))$$

has closure ordinal ω^2 on bidirectional models (we will soon give a proof for a similar formula). Looking back at the model of Figure 3.4, we can observe that after step ω^2 in the iteration of φ an infinite $R^{\smile}$-path of points where $(x \wedge \neg p)$ is always true has been formed: we can almost use this fact in order to define a formula that allows the addition of new states to the model that can only be included after ω^2 many iterations. To actually do this, however, we are going to need more colours: fix a set $\{q_i \mid i \in \omega\} \subseteq \texttt{PROP}$ of propositional variables that is disjoint from the set $\{p_i \mid i \in \omega\} \subseteq \texttt{PROP}$ that we used to define fuses in Definition 3.1.1.

Definition 3.3.1 For every $0 < n < \omega$ we define the *colour* c_n as the conjunction of literals $c_n := \bigwedge_{0<i<n} \neg q_i \wedge q_n$.

Note that this definition is slightly different from that of fuses, not only for the variables involved: for example, $c_1 = q_1$ and $c_3 = \neg q_1 \wedge \neg q_2 \wedge q_3$, while $f_1 = \neg p_1 \wedge p_2$ and $f_3 = \neg p_1 \wedge \neg p_2 \wedge \neg p_3 \wedge p_4$. The reason is that the definition of fuses is more consistent with the original in [6], and that of colours is slightly more convenient for presentation, but in the end this is just a matter of convention. Clearly $c_i \wedge c_j \equiv \bot$ for every $i \neq j$ in this case too.

Now that we have more colours, we can consider the formula $\varphi := G\bot \vee (Fx \wedge c_1) \vee (c_2 \wedge F(\nu y.P(y \wedge x \wedge c_1))) \vee (c_3 \wedge F(\nu y.P(y \wedge x \wedge c_2)))$: the last disjunct allows the addition of a new point ω^2 to the model of Figure 3.4 that will be included in the least fixed point iteration of φ after ω^2 steps, as the following picture shows.

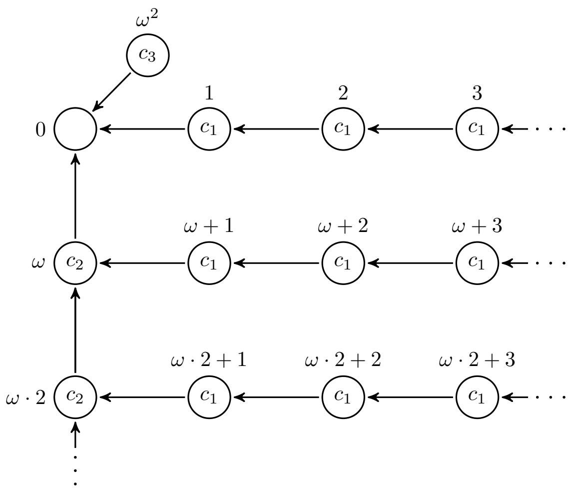

Figure 3.6: Adding the point ω^2

In order to construct a formula φ_n with closure ordinal ω^n for all $n \in \omega$ we now define a formula π_i^∞ that makes sure that models of φ_n are, in a sense, well-behaved: we need that whenever a state s in a model makes φ_n true and has colour c_i, then from this state begins an infinite $R^{\smile}$-path of points where c_{i-1} is always true, and from every point in this path starts an infinite $R^{\smile}$-path of points where c_{i-2} is always true, and so on.

Definition 3.3.2 For every $i \in \omega$ we define a formula π_i^∞ as follows:

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

Example 3.3.3 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 $R^{\smile}$-path $t_0t_1t_2\ldots$ 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 $R^{\smile}$-path $u_0u_1u_2\ldots$ 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 $R^{\smile}$-path $v_0v_1\ldots$, 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 Figure 3.7 makes π_3^∞ true (as does every state of the form $\omega^2 \cdot n$ for $n \in \omega$).

Proposition 3.3.4 For all $m, n \in \omega$, if $m \geq n$, then $\pi_m^\infty \models \pi_n^\infty$. Moreover, if $(\mathbb{S}, s)$ is a pointed bidirectional model, $\mathbb{S}, s \Vdash \pi_m^\infty$ and $t_0 t_1 \ldots$ is an $R^{\smile}$-path witnessing the truth of π_m^∞ at s, then $t_j \Vdash \pi_m^\infty$ for all $j \in \omega$.

We now have all the ingredients to define, for all $n \in \omega$, a formula φ_n that we will show has closure ordinal ω^n on bidirectional model.

Definition 3.3.5 For all $n \in \omega$ let the formula φ_n be

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

Before proving that φ_n has closure ordinal ω^n, we conclude the example regarding ω^3. To build a model where

$$\varphi_3 := G\bot \vee (c_1 \wedge Fx) \vee (c_2 \wedge \pi_1^\infty \wedge F(\nu y.P(y \wedge x \wedge c_1))) \vee (c_3 \wedge \pi_2^\infty \wedge F(\nu y.P(y \wedge x \wedge c_2)))$$

converges in exactly ω^3 steps, after we have added the point ω^2 to the model corresponding to the ordinal ω^2 like in Figure 3.6, we can start an infinite $R^{\smile}$-path from this new point where c_3 is always true and, at each point in this path, append a copy of ω^2, as is shown in Figure 3.7.

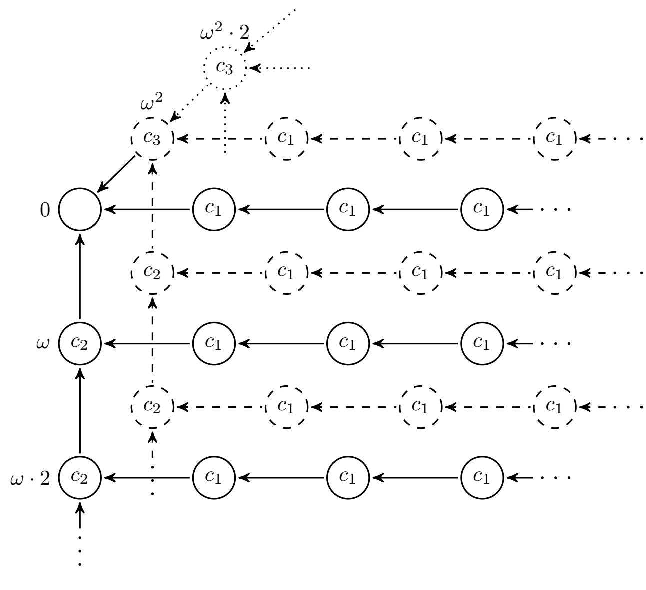

Figure 3.7: Model corresponding to ω^3

The iteration of φ_3 in this model behaves as follows: *(i)* first the state 0 is added through $G\bot$; *(ii)* then, every finite ordinal, one by one, through $(Fx \wedge c_1)$; *(iii)* after that, the point ω is added to the iteration at step $\omega + 1$ through the disjunct $(c_2 \wedge \pi_1^\infty \wedge F(\nu y.P(y \wedge x \wedge c_1)))$; *(iv)* after ω, every point of the form $\omega + n$ will be included, one at a time, to the iteration through $(Fx \wedge c_1)$; *(v)* at step $\omega \cdot 2 + 1$ the state $\omega \cdot 2$ will be added, and then every state of the form

$\omega \cdot 2 + n$, and so on; *(vi)* when every state of shape $\omega \cdot n$ has been added to the iteration, the state ω^2 will satisfy $(c_3 \wedge \pi_2^\infty \wedge F(\nu y.P(y \wedge x \wedge c_2)))$ and will be part of the iteration at step $\omega^2 + 1$; *(vii)* then it is the turn of states of the form $\omega^2 + n$, followed by $\omega^2 + \omega$, and so forth, building up the second copy of ω^2 appended to the state ω^2; *(viii)* when the entire second copy of ω^2 is part of the iteration, the state $\omega^2 \cdot 2$ will be added, and the process continues with $\omega^2 + 1, \omega^2 + 2, \ldots$. We thus see that the iteration of φ_3 in this model builds up ω copies of ω^2 and for every $\alpha < \omega^3$ it holds that $\varphi_3^\alpha = \alpha$, so that the least fixed point of φ_3 in this model is $\varphi_3^{\omega^3}$.

As a last observation before the formal proof, note that in the last model, after ω^3 many steps in the iteration of φ_3, an infinite $R^{\smile}$-path of points where $(c_3 \wedge x)$ is always true has been formed. The formula φ_4 contains the disjunct $(c_4 \wedge \pi_3^\infty \wedge F(\nu y.P(y \wedge x \wedge c_3)))$ which allows the addition of a new state ω^3 to the iteration of φ_4 after ω^3 many steps: if we build a model with an infinite $R^{\smile}$-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 Figure 3.7) we obtain a model where φ_4 converges in exactly ω^4 many steps. Then, since this model has an infinite $R^{\smile}$-chain where $(c_4 \wedge x)$ is always true, we could consider the formula φ_5 and similarly obtain a model where it converges in exactly ω^5 many steps, and so on.

We now finally prove that, for all $n \in \omega$, the closure ordinal of φ_n on bidirectional models is ω^n. We start with the following lemma.

Lemma 3.3.6 Let $\mathbb{S} = (S, R, V)$ be a bidirectional model, let $n \in \omega$ and φ_n be the formula from Definition 3.3.5. Let, for $1 \leq i \leq n$, $t_0 t_1 t_2 \ldots$ be an infinite $R^{\smile}$-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 $R^{\smile}$-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 $R^{\smile}$-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 $R^{\smile}$-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 3.3.4, since $u_0 \Vdash \pi_i^\infty$, then $u_0 \Vdash \pi_{i-1}^\infty$, and for every $k > 0$,

$u_k \Vdash c_i \wedge \pi^{\infty}_{i-1}$. 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

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

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

We now prove that φ_n converges to its least fixed point in at most ω^n many steps on every bidirectional model.

Lemma 3.3.7 Let $\mathbb{S} = (S, R, V)$ be a bidirectional model and $0 < n < \omega$ be a finite ordinal. Then $\varphi_n^{\omega^n+1} = \varphi_n^{\omega^n}$.

Proof. 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 $s \in (\varphi_n)^{\mathbb{S}}_x(\emptyset) \subseteq \varphi_n^{\omega^n}$, 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 $s \Vdash c_i \wedge \pi^{\infty}_{i-1} \wedge F(\nu y.P(y \wedge x \wedge c_{i-1}))$ for some $2 \leq i \leq n$. Then in particular there is a point $t \in R[s]$ and a $R^{\smallsmile}$-path $t_0t_1\ldots$ 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 $\varphi_n \wedge c_{i-1} \wedge F\top \models \pi^{\infty}_{i-2}$: this implies that $t_j \Vdash \pi^{\infty}_{i-2}$ 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 3.3.6 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 $\alpha + \omega^{i-1} + 1 < \omega^n$). □

Finally, we construct a model where φ_n converges to its least fixed point in exactly ω^n many steps.

Lemma 3.3.8 Let $0 < n < \omega$ be a finite ordinal. Then there is a bidirectional model $\mathbb{S}$ where $\mathrm{LFP}.(\varphi_n)^{\mathbb{S}}_x = \varphi_n^{\omega^n} \neq \varphi_n^{\beta}$ for all $\beta < \omega^n$.

Proof. For the rest of the proof we adopt the following notation: since every $\alpha < \omega^n$ is of the form $\omega^{n-1} \cdot k_1 + \ldots + \omega \cdot k_{n-1} + k_n$, 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 bidirectional 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 $1 \leq i \leq n$, $V(q_i) := \{(k_1, \ldots, k_{n-i+1} + 1, 0, \ldots, 0) \mid k_j \in \omega\}$.

Note that $R[(0, \ldots, 0)] = \emptyset$ and that $(0, \ldots, 0)$ 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, 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 $\varphi_x^{\mathbb{S}}(S_\beta) = S_{\beta+1}$.

For the $\supseteq$ inclusion 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 $\beta = 0 = (0, \ldots, 0)$ we are done. If $\beta = (k_1, \ldots, k_n + 1)$, then $\beta \in V(q_1)$ and $(k_1, \ldots, k_n) \in S_\beta \cap R[\beta]$, so $\mathbb{S}[x \mapsto S_\beta], \beta \Vdash c_1 \wedge Fx$ and $\beta \in \varphi_x^{\mathbb{S}}(S_\beta)$.

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

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

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

Now we move to the $\subseteq$ inclusion. 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 $\gamma = (k_1, \ldots, k_n + 1)$ and $\gamma' = (k_1, \ldots, k_n) \in S_\beta$: as $\gamma' \in S_\beta$, then $\gamma = \gamma' + 1 \in S_{\beta+1}$.

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

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

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

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

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

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

Theorem 3.3.9 *For all $0 < n < \omega$, the closure ordinal of $\varphi_n(x)$ on bidirectional models is ω^n.*

3.4 Bidirectional models: ordinals below ω^ω

In this section we combine fuses and colours to construct a formula with closure ordinal α for every ordinal $\alpha < \omega^\omega$. Note that every such ordinal has the form

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

so that we want a formula φ whose iteration builds up k_1 copies of ω^n, then k_2 copies of ω^{n-1}, and so on. We will use fuses to distinguish between each of these copies, so that we will need $k_1 + \ldots + k_n + 1$ fuses (we only need one additional fuse for the last k_{n+1} iterations, see for instance the example tied to Figure 3.1). The formula φ_α will then combine fuses with the formulas φ_n from the previous section, so that it will involve, for instance, k_1 subformulas that are essentially φ_n, each however containing a different fuse.

It will be helpful to recall the definitions of fuses and colours.

Definition 3.4.1 Fix disjoint sets $\{p_i \mid i \in \omega\} \subseteq \texttt{PROP}$ and $\{q_i \mid i \in \omega\} \subseteq \texttt{PROP}$ of propositional variables. For every $n \in \omega$ we define the *fuse* f_n as the conjunction of literals

$$f_n := \bigwedge_{0<i\leq n} \neg p_i \wedge p_{n+1}.$$

For every $0 < n < \omega$ we define the *colour* c_n as the conjunction of literals

$$c_n := \bigwedge_{0<i<n} \neg q_i \wedge q_n.$$

We also need the following definition and proposition, which are analogous to Definition 3.3.2 and Proposition 3.3.4.

Definition 3.4.2 For every $i, k \in \omega$ we define a formula $\pi^\infty_{i,k}$ inductively on i as follows:

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

Proposition 3.4.3 For all $m, n, k \in \omega$, if $m \geq n$, then $\pi^\infty_{m,k} \models \pi^\infty_{n,k}$. Moreover, if $(\mathbb{S}, s)$ is a pointed bidirectional model, $\mathbb{S}, s \Vdash \pi^\infty_{m,k}$ and $t_0 t_1 \ldots$ is an $R^\smallsmile$-path witnessing the truth of $\pi^\infty_{m,k}$ at s, then $t_j \Vdash \pi^\infty_{m,k}$ for all $j \in \omega$.

We finally state the definition of the formula φ_α.

Definition 3.4.4 For $n, k \in \omega$ define the formulas

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

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

$$\begin{array}{rcl}\psi & := & \bigvee_{i=0}^{k_{n+1}-1}(Gx \wedge \bigwedge_{j=0}^{i} G^j f_{k(\overrightarrow{n})} \wedge G^{i+1} f_{k(\overrightarrow{n})-1}),\\ \varphi_\alpha & := & G\bot \vee \bigvee_{k=0}^{k(\overrightarrow{n})-2} \chi_k \vee \bigvee_{m=0}^{n-1}\left(\bigvee_{k=k(\overrightarrow{m})}^{k(\overrightarrow{m+1})-1} \varphi_{(n-m,k)}\right) \vee \psi.\end{array}$$

To motivate this definition we first present a simple example and then provide some intuitions behind it.

Example 3.4.5 Let $\alpha = \omega^3 \cdot 2 + \omega^2 + \omega \cdot 3 + 4$: then $n = 3$, $k_1 = 2$, $k_2 = 1$, $k_3 = 3$ and $k_4 = 4$. We start by computing the disjuncts of φ_α: these are $G\bot$ and

$$\begin{array}{rcl}\bigvee_{k=0}^{4} \chi_k & = & (Gx \wedge f_1 \wedge Gf_0) \vee (Gx \wedge f_2 \wedge Gf_1) \vee \\ & & (Gx \wedge f_3 \wedge Gf_2) \vee (Gx \wedge f_4 \wedge Gf_3) \vee (Gx \wedge f_5 \wedge Gf_4),\\ (m=0) & \Rightarrow & \bigvee_{k=0}^{1} \varphi_{(3,k)} = \varphi_{(3,0)} \vee \varphi_{(3,1)},\\ (m=1) & \Rightarrow & \bigvee_{k=2}^{2} \varphi_{(2,k)} = \varphi_{(2,2)},\\ (m=2) & \Rightarrow & \bigvee_{k=3}^{5} \varphi_{(1,k)} = \varphi_{(1,3)} \vee \varphi_{(1,4)} \vee \varphi_{(1,5)},\\ \psi & = & \bigvee_{i=0}^{3}(Gx \wedge \bigwedge_{j=0}^{i} G^j f_6 \wedge G^{i+1} f_5).\end{array}$$

We now describe a model where φ_α converges in exactly α steps: this will consist of two copies of ω^3, one copy of ω^2, three copies of ω and four points. We let all the points of the first copy of ω^3 satisfy the fuse f_0, those of the second copy the fuse f_1, the states of the copy of ω^2 make the fuse f_2 true, and so on, while the last four points all satisfy the fuse f_6. We consider the domains of all these copies to be disjoint, so that for instance we index all points of the first copy of ω^3 with $(3, 0)$, where 0 refers to the fuse, and similarly for the rest. We call the last four points s_1, s_2, s_3 and s_4 for simplicity. We let the point $0_{(3,0)}$, be a dead end, the point $0_{(3,1)}$ see every point in the first copy of ω^3, the point $0_{(2,2)}$ see every point in the second copy of ω^3 and so forth; the state s_1 sees every point in the third copy of ω, and, for $1 \leq i \leq 3$, the state s_{i+1} sees s_i. This is all depicted in Figure 3.8, where two arrows going from a state to two vertices of a rectangle indicate that the said state has an arrow towards every point inside that rectangle.

We thus see that the iteration of φ_α will start from the point $0_{(3,0)}$ through the disjunct $\Box\bot$ and continue inside the first copy of ω^3 through

$$\begin{array}{rcl}\varphi_{(3,0)} & = & (Fx \wedge c_1 \wedge f_0 \wedge Gf_0) \vee \\ & & (c_2 \wedge f_0 \wedge Gf_0 \wedge \pi^\infty_{1,0} \wedge F(\nu y.f_0 \wedge P(y \wedge x \wedge Gf_0 \wedge c_1))) \vee \\ & & (c_3 \wedge f_0 \wedge Gf_0 \wedge \pi^\infty_{2,0} \wedge F(\nu y.f_0 \wedge P(y \wedge x \wedge Gf_0 \wedge c_2))).\end{array}$$

Once every point in the first copy of ω^3 is part of the iteration of φ_α after ω^3 many steps, the point $0_{(3,1)}$ will satisfy $\chi_0 = Gx \wedge f_1 \wedge Gf_0$, so that it will be added to the iteration. Note that for instance this point cannot be added before

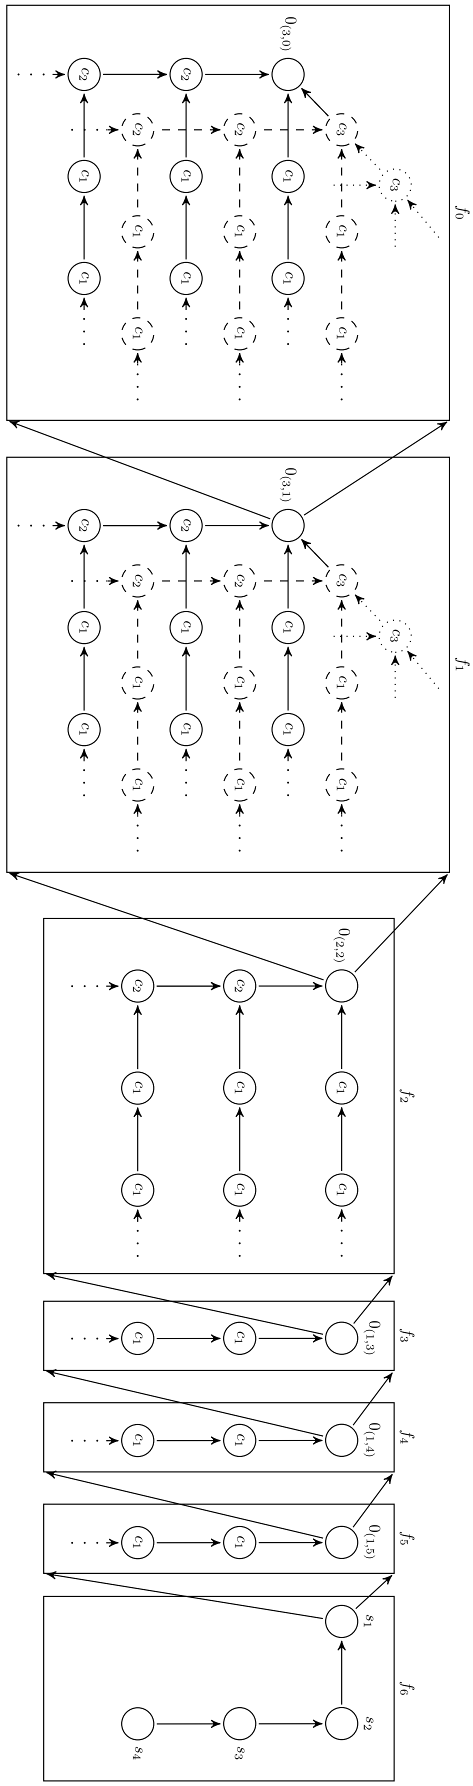

Figure 3.8: Model corresponding to $\omega^3 \cdot 2 + \omega^2 + \omega \cdot 3 + 4$.

through the disjunct $(Fx \wedge c_1 \wedge f_0 \wedge Gf_0)$ of $\varphi_{(3,0)}$, as it falsifies both c_1 and f_0, and neither through the disjunct $(Fx \wedge c_1 \wedge f_1 \wedge Gf_1)$ of $\varphi_{(3,1)}$, as it falsifies c_1 and Gf_1. After $0_{(3,1)}$ is added, the process carries on inside the second copy of ω^3 though $\varphi_{(3,1)}$, and every point in this copy will be included in the iteration in another ω^3 many steps. After $\omega^3 \cdot 2$ many steps, then, the point $0_{(2,2)}$ will satisfy $\chi_1 = Gx \wedge f_2 \wedge Gf_1$ and will be added to the iteration, followed by every point in the copy of ω^2 through the disjunct

$$\begin{aligned}\varphi_{(2,2)} \quad = \quad & (Fx \wedge c_1 \wedge f_2 \wedge Gf_2) \vee \\ & (c_2 \wedge f_2 \wedge Gf_2 \wedge \pi^{\infty}_{1,2} \wedge F(\nu y. f_2 \wedge P(y \wedge x \wedge Gf_2 \wedge c_1))).\end{aligned}$$

Then, after $\omega^3 \cdot 2 + \omega^2$ many steps, it is the turn of the state $0_{(1,3)}$ to be added to the iteration through $\chi_2 = Gx \wedge f_3 \wedge Gf_2$, and so every point in the first copy of ω through $\varphi_{(1,3)} = (Fx \wedge c_1 \wedge f_3 \wedge Gf_3)$. When the third copy of ω is part of the iteration after $\omega^3 \cdot 2 + \omega^2 + \omega \cdot 3$ many steps of the iteration, the state s_1 will make the first disjunct $Gx \wedge f_6 \wedge Gf_5$ of ψ true, then s_2 will satisfy $Gx \wedge f_6 \wedge Gf_6 \wedge GGf_5$ after s_1 is included in the iteration, followed similarly by s_3 and s_4. We conclude that in this model the formula φ^α converges to its least fixed point in exactly $\omega^3 \cdot 2 + \omega^2 + \omega \cdot 3 + 4$ many steps.

The intuition behind Definition 3.4.4 is the following. With $\alpha = \omega^n \cdot k_1 + \omega^{n-1} \cdot k_2 + \ldots + \omega \cdot k_n + k_{n+1}$:

- for each $0 \leq m \leq n-1$, the formula $\varphi_{(n-m,k)}$ witnesses the phase of the iteration that takes place inside the copy of the model ω^{n-m} corresponding to the fuse f_k;
- for each $0 \leq k < k(\overrightarrow{n}) - 2$, the formula χ_k allows to move from a copy of some ω^{n-m} with fuse f_k to a copy of some $\omega^{n-m'}$ with fuse f_{k+1}, with $m' \geq m$;
- the formula ψ allows to compute the last k_{n+1} many iterations.

Note that the formula $\varphi(n-m,k)$ has a very similar shape to the formula φ_{n-m} from Definition 3.3.5, but involves the fuse f_k in order to distinguish between copies of ω^{n-m} with different fuses.

Remark We observe that the definition of φ_α works also when some k_i is equal to 0 in $\omega^n \cdot k_1 + \ldots + \omega \cdot k_n + k_{n+1}$. For instance, let $\alpha = \omega^4 \cdot 2 + \omega \cdot 3$: then $n = 4$, $k_1 = 2$, $k_4 = 3$ and $k_2 = k_3 = k_5 = 0$, so that

$$\begin{aligned}\bigvee_{k=0}^{3} \chi_k \quad & = \quad (Gx \wedge f_1 \wedge Gf_0) \vee (Gx \wedge f_2 \wedge Gf_1) \vee \\ & \quad\quad (Gx \wedge f_3 \wedge Gf_2) \vee (Gx \wedge f_4 \wedge Gf_3), \\ (m=0) \quad & \Rightarrow \quad \bigvee_{k=0}^{1} \varphi_{(4,k)} = \varphi_{(4,0)} \vee \varphi_{(4,1)}, \\ (m=1) \quad & \Rightarrow \quad \bigvee_{k=2}^{1} \varphi_{(3,k)} = \bigvee \emptyset = \bot, \\ (m=2) \quad & \Rightarrow \quad \bigvee_{k=2}^{1} \varphi_{(2,k)} = \bigvee \emptyset = \bot, \\ (m=3) \quad & \Rightarrow \quad \bigvee_{k=2}^{4} \varphi_{(1,k)} = \varphi_{(1,2)} \vee \varphi_{(1,3)} \vee \varphi_{(1,4)},\end{aligned}$$

$$\psi \quad = \quad \bigvee_{i=0}^{-1} (Gx \wedge \bigwedge_{j=0}^{i} G^j f_5 \wedge G^{i+1} f_4) = \bigvee \emptyset = \bot.$$

We see then that the whole formula φ_α will behave as expected, as it contains the disjunct $G\bot$ to start the iteration, the disjuncts $\varphi_{(4,0)}$ and $\varphi_{(4,1)}$ to build the two copies of ω^4, the disjuncts $\varphi_{(1,2)}$, $\varphi_{(1,3)}$ and $\varphi_{(1,4)}$ to build the three copies of ω and the disjuncts $\chi_0, \ldots, \chi_3$ to move the iteration between these copies, while clearly the disjunct $\bot$ does not contribute to the iteration.

We now prove that α is the closure ordinal of φ_α on bidirectional models. We start with the following lemma, whose proof is essentially the same as that of Lemma 3.3.6.

Lemma 3.4.6 Let $\mathbb{S} = (S, R, V)$ be a bidirectional model, let $\alpha = \omega^n \cdot k_1 + \ldots + \omega \cdot k_n + k_{n+1}$ and φ_α be the formula from Definition 3.4.4 for some $n > 0$. For all $0 \leq m \leq n-1$, for all $k(\overrightarrow{m}) \leq k \leq k(\overrightarrow{m+1}) - 1$, for all $1 \leq i \leq n-m$, let $t_0 t_1 t_2 \ldots$ be an infinite $R^{\smile}$-path such that

$$\mathbb{S}, t_0 \Vdash \pi^\infty_{i-1,k} \wedge f_k \text{ and, for all } j > 0,\ \mathbb{S}, t_j \Vdash c_i \wedge \pi^\infty_{i-1,k} \wedge f_k \wedge G f_k.$$

Then, for any ordinal β: if $t_0 \in \varphi^\beta_\alpha$ then $t_j \in \varphi_\alpha^{\beta + \omega^{i-1} \cdot j + 1}$ for all $j \in \omega$.

Proof. Fix $0 \leq m \leq n-1$ and $k(\overrightarrow{m}) \leq k \leq k(\overrightarrow{m+1}) - 1$. We prove the statement by induction on $1 \leq i \leq n-m$.

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

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

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

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

Now suppose that $t_j \in \varphi_\alpha^{\beta + \omega^i \cdot j + 1}$: we show that $t_{j+1} \in \varphi_\alpha^{\beta + \omega^i \cdot (j+1) + 1}$. By assumption $t_j \in R[t_{j+1}]$ and $t_j \Vdash \pi^\infty_{i,k}$, which in particular means that there is an infinite $R^{\smile}$-path $u_0 u_1 \ldots$ (with $u_0 = t_j$) such that, for all $l > 0$, $u_l \Vdash c_i \wedge f_k \wedge G f_k$. But then this path satisfies the conditions of the inductive hypothesis: by Proposition 3.4.3, since $u_0 \Vdash \pi^\infty_{i,k} \wedge f_k$, then $u_0 \Vdash \pi^\infty_{i-1,k} \wedge f_k$, and for every $l > 0$, $u_l \Vdash c_i \wedge \pi^\infty_{i-1,k} \wedge f_k \wedge G f_k$. Then, by inductive hypothesis, since $u_0 = t_j \in \varphi_\alpha^{\beta + \omega^i \cdot j + 1}$ it follows that, for every $l \in \omega$, $u_l \in \varphi_\alpha^{\beta + \omega^i \cdot j + 1 + \omega^{i-1} \cdot l + 1}$. Since for all $l \in \omega$ it holds that

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

then also

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

It follows that $u_l \in \varphi_\alpha^{\beta+\omega^i\cdot(j+1)}$ for all $l \in \omega$, so that

$$\mathbb{S}[x \mapsto \varphi_\alpha^{\beta+\omega^i\cdot(j+1)}], t_{j+1} \Vdash c_{i+1} \wedge f_k \wedge G f_k \wedge \pi^\infty_{i,k} \wedge F(\nu y. f_k \wedge P(y \wedge x \wedge G f_k \wedge c_i)).$$

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

The next lemma involves fuses: with $\alpha = \omega^n \cdot k_1 + \ldots + \omega \cdot k_n + k_{n+1}$, it is shown that if a state belongs to some iteration of φ_α and makes $p_{k(\overrightarrow{m})+l}$ true for some $1 \leq l \leq k_{m+1}$ and $0 \leq m \leq n-1$, then it has to belong to $\varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-m}\cdot l}$.

Lemma 3.4.7 For $n \geq 1$ let $\alpha = \omega^n \cdot k_1 + \ldots + \omega \cdot k_n + k_{n+1}$. For all $0 \leq m \leq n-1$ and all pointed bidirectional models $(\mathbb{S}, s)$, if $\omega^n \cdot k_1 + \ldots + \omega^{n-m} \cdot k_{m+1} \leq \beta < \omega^\omega$, if $s \in \varphi_\alpha^\beta$ and $s \Vdash p_{k(\overrightarrow{m})+l}$ for some $1 \leq l \leq k_{m+1}$, then $s \in \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-m}\cdot l}$.

Proof. We prove the statement by induction on m, where $0 \leq m \leq n-1$.

As the base case take $m = 0$, so that by assumption we have that $\omega^n \cdot k_1 \leq \beta < \omega^\omega$, $s \in \varphi_\alpha^\beta$ and $s \Vdash p_l$ for some $1 \leq l \leq k_1$: we want to prove that $s \in \varphi_\alpha^{\omega^n\cdot l}$. We now proceed by induction on β and only consider the case where $\beta = \gamma + 1$, since the limit case is immediate.

Let $s \in (\varphi_\alpha)^{\mathbb{S}}_x(\varphi_\alpha^\gamma)$. We continue by case distinction as to which disjunct of φ_α is satisfied by s to prove $s \in \varphi_\alpha^{\omega^n\cdot l}$. If $s \Vdash G\bot$ the result follows immediately, and $\mathbb{S}[x \mapsto \varphi_\alpha^\gamma], s \not\Vdash \psi$ since $s \Vdash p_l$.

Next suppose $\mathbb{S}[x \mapsto \varphi_\alpha^\gamma], s \Vdash Gx \wedge f_{k+1} \wedge G f_k$ for some $k < l-1$. Then $R[s] \subseteq \varphi_\alpha^\gamma$ and $R[s] \subseteq [\![p_{k+1}]\!]^{\mathbb{S}}$. By inductive hypothesis on $\gamma < \beta$ then $R[s] \subseteq \varphi_\alpha^{\omega^n\cdot(k+1)}$, so that $s \in \varphi_\alpha^{\omega^n\cdot(k+1)+1} \subseteq \omega^n \cdot l$, since $k+1 < l$.

Now suppose $\mathbb{S}[x \mapsto \varphi_\alpha^\gamma], s \Vdash \varphi_{(n,k)}$ for some $k < l$. There are two cases to consider. First, if $\mathbb{S}[x \mapsto \varphi_\alpha^\gamma], s \Vdash Fx \wedge c_1 \wedge f_k \wedge G f_k$, then there is a $t \in R[s]$ such that $t \in \varphi_\alpha^\gamma$ and $t \Vdash p_{k+1}$, so that by inductive hypothesis on $\gamma < \beta$ it follows that $t \in \varphi_\alpha^{\omega^n\cdot(k+1)}$. Since $\omega^n \cdot (k+1) = \omega^n \cdot k + \omega^n$ is a limit ordinal, then there is a $\delta < \omega^n$ such that $t \in \varphi_\alpha^{\omega^n\cdot k+\delta}$. Because $\delta < \omega^{n-1} \cdot k'$ for some k', then $t \in \varphi_\alpha^{\omega^n\cdot k+\omega^{n-1}\cdot k'}$ and $s \in \varphi_\alpha^{\omega^n\cdot k+\omega^{n-1}\cdot k'+1} \subseteq \varphi_\alpha^{\omega^n\cdot l}$, since $k < l$.

Otherwise, let

$$\mathbb{S}[x \mapsto \varphi_\alpha^\gamma], s \Vdash c_i \wedge f_k \wedge G f_k \wedge \pi^\infty_{i-1,k} \wedge F(\nu y. f_k \wedge P(y \wedge x \wedge G f_k \wedge c_{i-1}))$$

for some $2 \leq i \leq n$. Then there is a state $t \in R[s]$ and a $R^{\smile}$-path $t_0 t_1 t_2 \ldots$ such that $t \in R[t_0]$, $t \Vdash f_k$ and, for all $j \in \omega$, $t_j \in \varphi_\alpha^\gamma$ and $t_j \Vdash c_{i-1} \wedge f_k \wedge G f_k$. Since $t_0 \Vdash f_k$, then $t_0 \Vdash p_{k+1}$ and, by inductive hypothesis on $\gamma < \beta$, $t_0 \in \varphi_\alpha^{\omega^n\cdot(k+1)}$, so that $t_0 \in \varphi_\alpha^{\omega^n\cdot k+\omega^{n-1}\cdot k'}$ for some k'. Moreover, since $(\varphi_\alpha \wedge c_{i-1} \wedge F\top \wedge f_k \wedge G f_k) \models \pi^\infty_{i-2,k}$, then $t_j \Vdash \pi^\infty_{i-2,k}$ for all $j \in \omega$. Hence we can apply Lemma 3.4.6 and obtain that, for all $j \in \omega$, $t_j \in \varphi_\alpha^{\omega^n\cdot k+\omega^{n-1}\cdot k'+\omega^{i-2}\cdot j+1}$. We thus observe that, for all $j \in \omega$ (recall that $i \leq n$),

$$\begin{aligned} \omega^n \cdot k + \omega^{n-1} \cdot k' + \omega^{i-2} \cdot j + 1 &< \omega^n \cdot k + \omega^{n-1} \cdot k' + \omega^{i-1} \\ &\leq \omega^n \cdot k + \omega^{n-1} \cdot k' + \omega^{n-1} \end{aligned}$$

$$= \omega^n \cdot k + \omega^{n-1} \cdot (k'+1),$$

so $t_j \in \varphi_\alpha^{\omega^n \cdot k + \omega^{n-1}\cdot(k'+1)}$ for all $j \in \omega$. Hence $\mathbb{S}[x \mapsto \varphi_\alpha^{\omega^n\cdot k+\omega^{n-1}\cdot(k'+1)}], s \Vdash \varphi_\alpha$, and it follows that

$$s \in \varphi_\alpha^{\omega^n\cdot k+\omega^{n-1}\cdot(k'+1)+1} \subseteq \varphi_\alpha^{\omega^n\cdot l}$$

in this case too, since $k < l$.

This concludes the base case of the induction on $0 \leq m \leq n-1$. To proceed, inductively assume that the statement of the lemma holds for all $m' \leq m < n-1$: we want to prove it for $m+1$. Assume that

$$\omega^n \cdot k_1 + \ldots + \omega^{n-m} \cdot k_{m+1} + \omega^{n-(m+1)} \cdot k_{m+2} \leq \beta < \omega^\omega,$$

$s \in \varphi_\alpha^\beta$ and $s \Vdash p_{k(\overrightarrow{m+1})+l}$, for some $1 \leq l \leq m+2$. We want to prove that $s \in \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-m}\cdot k_{m+1}+\omega^{n-(m+1)}\cdot l}$. Again, we proceed by induction on β and only consider the case where $\beta = \gamma+1$, so that we let $s \in (\varphi_\alpha)_x^{\mathbb{S}}(\varphi_\alpha^\gamma)$ and continue by case distinction as to which disjunct of φ_α is satisfied by s. The cases for $G\bot$ and ψ are again immediate, so that we focus on the other disjuncts of φ_α.

Suppose $\mathbb{S}[x \mapsto \varphi_\alpha^\gamma], s \Vdash Gx \wedge f_{k+1} \wedge Gf_k$ for some $k < k(\overrightarrow{m+1}) + l - 1$. We distinguish between two cases. If $k = k(\overrightarrow{m'}) + l'$ for some $0 \leq m' \leq m$ and $l' < k_{m'+1}$, then $s \Vdash Gp_{k(\overrightarrow{m'})+l'+1}$. Since $R[s] \subseteq \varphi_\alpha^\gamma$, by inductive hypothesis on $m' \leq m$ it follows that $R[s] \subseteq \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-m'}\cdot(l'+1)}$, so that $s \in \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-m'}\cdot(l'+1)+1} \subseteq \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-(m+1)}\cdot l}$ as desired.

Alternatively, $k = k(\overrightarrow{m+1})+l'$ for some $l' < l-1$, so $s \Vdash Gp_{k(\overrightarrow{m+1})+l'+1}$ and $R[s] \subseteq \varphi_\alpha^\gamma$, so that $R[s] \subseteq \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-(m+1)}\cdot(l'+1)}$ follows by inductive hypothesis on $\gamma < \beta$. Then $s \in \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-(m+1)}\cdot(l'+1)+1} \subseteq \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-(m+1)}\cdot l}$, since $l'+2 \leq l$.

Now suppose $\mathbb{S}[x \mapsto \varphi_\alpha^\gamma], s \Vdash \varphi_{(n-(m+1),k)}$ for some $k(\overrightarrow{m+1}) \leq k < k(\overrightarrow{m+1})+l$, so that $k = k(\overrightarrow{m+1}) + k'$ for some $k' < l$. There are two cases to consider. First, if $\mathbb{S}[x \mapsto \varphi_\alpha^\gamma], s \Vdash Fx \wedge c_1 \wedge f_k \wedge Gf_k$, then there is a $t \in R[s]$ such that $t \in \varphi_\alpha^\gamma$ and $t \Vdash p_{k(\overrightarrow{m+1})+k'+1}$, so that $t \in \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-(m+1)}\cdot(k'+1)}$ by inductive hypothesis on $\gamma < \beta$. Since $\omega^n \cdot k_1 + \ldots + \omega^{n-(m+1)} \cdot (k'+1)$ is a limit ordinal, then $t \in \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-(m+1)}\cdot k'+\omega^{n-(m+2)}\cdot k''}$ for some k'', and, since $k' < l$, then

$$s \in \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-(m+1)}\cdot k'+\omega^{n-(m+2)}\cdot k''+1} \subseteq \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-(m+1)}\cdot l}.$$

Otherwise, let

$$\mathbb{S}[x \mapsto \varphi_\alpha^\gamma], s \Vdash c_i \wedge f_k \wedge Gf_k \wedge \pi_{i-1,k}^\infty \wedge F(\nu y. f_k \wedge P(y \wedge x \wedge Gf_k \wedge c_{i-1}))$$

for some $2 \leq i \leq n-(m+1)$. Then there is a state $t \in R[s]$ and a $R^{\smile}$-path $t_0t_1t_2\ldots$ such that $t \in R[t_0]$, $t \Vdash f_k$ and, for all $j \in \omega$, $t_j \in \varphi_\alpha^\gamma$ and $t_j \Vdash c_{i-1} \wedge f_k \wedge Gf_k$. Since $t_0 \Vdash f_k$, then $t_0 \Vdash p_{k+1}$, that is, $t_0 \Vdash p_{k(\overrightarrow{m+1})+k'+1}$, and by inductive hypothesis on $\gamma < \beta$ we obtain $t_0 \in \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-(m+1)}\cdot(k'+1)}$, so that $t_0 \in \varphi_\alpha^{\omega^n\cdot k_1+\ldots+\omega^{n-(m+1)}\cdot k'+\omega^{n-(m+2)}\cdot k''}$ for some k''. Moreover, since $(\varphi_\alpha \wedge c_{i-1} \wedge F\top \wedge f_k \wedge Gf_k) \models \pi_{i-2,k}^\infty$, then $t_j \Vdash \pi_{i-2,k}^\infty$ for all $j \in \omega$. Hence we can apply Lemma 3.4.6 and obtain that, for all $j \in \omega$,

$$t_j \in \varphi_\alpha^{\omega^n \cdot k_1 + \ldots + \omega^{n-(m+1)} \cdot k' + \omega^{n-(m+2)} \cdot k'' + \omega^{i-2} \cdot j + 1}.$$

We thus observe that, for all $j \in \omega$ (recall that $i \leq n - (m+1)$, so that $i - 1 \leq n - (m+2)$),

$$\begin{array}{ll} \omega^n \cdot k_1 + \ldots + \omega^{n-(m+1)} \cdot k' + \omega^{n-(m+2)} \cdot k'' + \omega^{i-2} \cdot j + 1 & < \\ \omega^n \cdot k_1 + \ldots + \omega^{n-(m+1)} \cdot k' + \omega^{n-(m+2)} \cdot k'' + \omega^{i-1} & \leq \\ \omega^n \cdot k_1 + \ldots + \omega^{n-(m+1)} \cdot k' + \omega^{n-(m+2)} \cdot k'' + \omega^{n-(m+2)} & = \\ \omega^n \cdot k_1 + \ldots + \omega^{n-(m+1)} \cdot k' + \omega^{n-(m+2)} \cdot (k'' + 1), & \end{array}$$

so that $t_j \in \varphi_\alpha^{\omega^n \cdot k_1 + \ldots + \omega^{n-(m+1)} \cdot k' + \omega^{n-(m+2)} \cdot (k''+1)}$ for all $j \in \omega$.

It follows that $\mathbb{S}[x \mapsto \varphi_\alpha^{\omega^n \cdot k_1 + \ldots + \omega^{n-(m+1)} \cdot k' + \omega^{n-(m+2)} \cdot (k''+1)}], s \Vdash \varphi_\alpha$ and we conclude that also in this case

$$s \in \varphi_\alpha^{\omega^n \cdot k_1 + \ldots + \omega^{n-(m+1)} \cdot k' + \omega^{n-(m+2)} \cdot (k''+1) + 1} \subseteq \varphi_\alpha^{\omega^n \cdot k_1 + \ldots + \omega^{n-(m+1)} \cdot l}$$

since $k' < l$. □

We now show that φ_α converges to its least fixed point in at most α steps on every bidirectional model.

Lemma 3.4.8 Let $\mathbb{S}$ be a bidirectional model, $\omega \leq \alpha < \omega^\omega$ be an ordinal and φ_α be the formula from Definition 3.4.4. Then $\varphi_\alpha^{\alpha+1} = \varphi_\alpha^\alpha$.

Proof. Let $\alpha = \omega^n \cdot k_1 + \ldots + \omega \cdot k_n + k_{n+1}$ and $s \in \varphi_\alpha^{\alpha+1}$: we want to show that $s \in \varphi_\alpha^\alpha$. If $s \Vdash p_k$ for some $1 \leq k \leq k(\overrightarrow{n})$ then $s \in \varphi_\alpha^\alpha$ follows by Lemma 3.4.7. Hence suppose otherwise: $s \Vdash \neg p_k$ for all $1 \leq k \leq k(\overrightarrow{n})$.

By assumption $s \in (\varphi_\alpha)_x^{\mathbb{S}}(\varphi_\alpha^\alpha)$, so that we proceed by case distinction as to which disjunct of φ_α is satisfied by s to prove that $s \in \varphi_\alpha^\alpha$. It cannot be the case that

$$\mathbb{S}[x \mapsto \varphi_\alpha^\alpha], s \Vdash \bigvee_{k=0}^{k(\overrightarrow{n})-2} \chi_k \vee \bigvee_{m=0}^{n-1} \left(\bigvee_{k=k(\overrightarrow{m})}^{k(\overrightarrow{m+1})-1} \varphi_{(n-m,k)} \right),$$

since otherwise $s \Vdash p_k$ for some $1 \leq k \leq k(\overrightarrow{n})$, against assumption. We thus consider the cases where $s \Vdash G\bot$ and $\mathbb{S}[x \mapsto \varphi_\alpha^\alpha] \Vdash \psi$, where we recall that

$$\psi := \bigvee_{i=0}^{k_{n+1}-1} (Gx \wedge \bigwedge_{j=0}^{i} G^j f_{k(\overrightarrow{n})} \wedge G^{i+1} f_{k(\overrightarrow{n})-1}).$$

Since the first case is trivial we focus on the second one.

We mention that the formula ψ here is essentially the same as the second disjunct of the formula $\psi_{\omega \cdot n+m}$ from Definition 3.1.2, and also the following induction is an adaptation to our setting of the same induction in [6].

Let $\mathbb{S}[x \mapsto \varphi_\alpha^\alpha] \Vdash Gx \wedge \bigwedge_{j=0}^{i} G^j f_{k(\overrightarrow{n})} \wedge G^{i+1} f_{k(\overrightarrow{n})-1}$ for some $0 \leq i < k_{n+1}$: in particular, $R[s] \subseteq \varphi_\alpha^\alpha$. By induction on i we prove that

$$\text{if } s \Vdash \bigwedge_{j=0}^{i} G^j f_{k(\overrightarrow{n})} \wedge G^{i+1} f_{k(\overrightarrow{n})-1}, \text{ then } s \in \varphi_\alpha^{\omega^n \cdot k_1 + \ldots + \omega \cdot k_n + i + 1}.$$

For $i = 0$, assume $s \Vdash f_{k(\overrightarrow{n})} \wedge G f_{k(\overrightarrow{n})-1}$, so every $t \in R[s]$ is such that $t \Vdash p_{k(\overrightarrow{n})}$, that is, $t \Vdash p_{k(\overrightarrow{n-1})+k_n}$, and by Lemma 3.4.7 it holds that $R[s] \subseteq \varphi_\alpha^{\omega^n \cdot k_1 + \ldots + \omega^{n-(n-1)} \cdot k_n} = \varphi_\alpha^{\omega^n \cdot k_1 + \ldots + \omega \cdot k_n}$, hence $s \in \varphi_\alpha^{\omega^n \cdot k_1 + \ldots + \omega \cdot k_n + 1}$ as desired.

Now suppose that the statement of the induction holds for all i' such that $0 \leq i' < i \leq k_{n+1}$: we prove it for i. Let $s \Vdash \bigwedge_{j=0}^{i} G^j f_{k(\overrightarrow{n})} \wedge G^{i+1} f_{k(\overrightarrow{n})-1}$: then, for all $t \in R[s]$, $t \Vdash \bigwedge_{j=0}^{i-1} G^j f_{k(\overrightarrow{n})} \wedge G^i f_{k(\overrightarrow{n})-1}$, which by induction hypothesis implies $R[s] \subseteq \varphi_\alpha^{\omega^n \cdot k_1 + \ldots + \omega \cdot k_n + i}$. We conclude that $s \in \varphi_\alpha^{\omega^n \cdot k_1 + \ldots + \omega \cdot k_n + i + 1}$. □

We finally construct a bidirectional model where φ_α converges to its least fixed point in exactly α steps.

Lemma 3.4.9 For every $\omega \leq \alpha < \omega^\omega$, there is a bidirectional model where the formula φ_α from Definition 3.4.4 converges to its least fixed point in exactly α steps.

Proof. Let $\alpha = \omega^n \cdot k_1 + \ldots + \omega \cdot k_n + k_{n+1}$ for $n > 0$. We define a model $\mathbb{S}$ where we argue that φ_α converges in exactly α many steps to its least fixed point. To simplify the definition we assume that $k_i \neq 0$ for all $1 \leq i \leq k_{n+1}$: when this is not the case a similar definition with very obvious variations is possible. For each $0 \leq m \leq n-1$ and $k(\overrightarrow{m}) \leq k \leq k(\overrightarrow{m+1}) - 1$ we take a copy of the bidirectional model corresponding to the ordinal ω^{n-m} as defined in the proof of Lemma 3.3.8, where we let f_k be true everywhere by making all states satisfy p_{k+1} and falsify every p_i for $i \neq k+1$ (but leaving the valuation of other propositional variables unchanged, so that every point has the same colour as before): we call this model $\mathbb{S}_{(n-m,k)}$. We let the domains of all these models be disjoint, for instance by indexing all the elements of $\mathbb{S}_{(n-m,k)}$ with $(n-m,k)$. Note that every model $\mathbb{S}_{(n-m,k)}$ has a blind state $0_{(n-m,k)}$. Define $\mathbb{S}$ to be the disjoint union of all these models, enriched with the following relations:

- the point $0_{(n,0)}$, which is the blind state in the copy of $\mathbb{S}_{(n,0)}$ that makes the fuse f_0 true, is also a blind state in $\mathbb{S}$;
- for $m \geq 0$ and $k(\overrightarrow{m}) < k \leq k(\overrightarrow{m+1}) - 1$, we let $0_{(n-m,k)}$ have an arrow towards every point of the copy of $\mathbb{S}_{(n-m,k-1)}$ in the disjoint union;
- For $m > 0$ and $k = k(\overrightarrow{m})$, we let $0_{(n-m,k)}$ have an arrow towards every point of the copy of $\mathbb{S}_{(n-m+1,k(\overrightarrow{m})-1)}$ in the disjoint union.

We also let the model $\mathbb{S}$ have k_{n+1} new points that we call $s_1, \ldots s_{k_{n+1}}$ and let the state s_1 have an arrow towards every point of the copy of $\mathbb{S}_{1,k(\overrightarrow{n})-1}$ in the disjoint union and let s_{i+1} see s_i for every $1 \leq i < k_{n+1}$. We also let every point s_i make $f_{k(\overrightarrow{n})}$ true.

We now argue that in this model the least fixed point of φ_α is reached in exactly α many iterations. Note that by construction every point of the form $0_{(n-m,k)}$ has no colour, so that it can be added to the iteration only through $G\bot$ or through χ_{k-1}. The iteration of φ_α starts through the disjunct $G\bot$ at state $0_{(n,0)}$ in the submodel $\mathbb{S}_{(n,0)}$ which corresponds to the first copy of ω^n making the fuse f_0 true, then proceeds to include every point of $\mathbb{S}_{(n,0)}$ through the disjunct $\varphi_{(n,0)}$ in ω^n many steps. After that, the point $0_{(n,1)} \in \mathbb{S}_{(n,1)}$ will make true $\chi_0 = (Gx \wedge f_1 \wedge G f_0)$ and will be included in the iteration, followed by every point of $\mathbb{S}_{(n,1)}$ through the disjunct $\varphi_{(n,1)}$, again in ω^n many steps. Then

$0_{(n,2)}$ is added and so on, until every point of $\mathbb{S}_{(n,k_1-1)}$ is part of the iteration, in a process that so far has taken $\omega^n \cdot k_1$ many steps. Then, the point $0_{(n-1,k_1)}$ in $\mathbb{S}_{(n-1,k_1)}$ will satisfy $\chi_{k_1-1} = Gx \wedge f_{k_1} \wedge Gf_{k_1-1}$ and the process continues by adding every point of $\mathbb{S}_{(n-1,k_1)}$ in ω^{n-1} many steps, followed by $\mathbb{S}_{(n-1,k_1+1)}$, $\mathbb{S}_{(n-1,k_1+2)}$, $\ldots$, $\mathbb{S}_{(n-1,k_1+k_2-1)}$, each taking ω^{n-1} many steps in order to be fully added to the iteration, which so far has taken $\omega^n \cdot k_1 + \omega^{n-1} \cdot k_2$ many steps. The iteration then moves to $\mathbb{S}_{(n-2,k_1+k_2)}$ and so on, until every point in $\mathbb{S}_{(1,k(\overrightarrow{n})-1)}$ is added to the iteration in $\omega^n \cdot k_1 + \ldots + \omega \cdot k_n$ many steps. At this point the state s_1 will make true the disjunct $(G_x \wedge f_{k(\overrightarrow{n})} \wedge Gf_{k(\overrightarrow{n})-1})$ of ψ, and one by one the states $s_2, s_3, \ldots, s_{k_{n+1}}$ will make true the corresponding disjuncts of ψ and the formula φ_α will converge to its least fixed point in exactly α steps. □

Theorem 3.4.10 *For every $\omega \leq \alpha < \omega^\omega$, the closure ordinal of φ_α on bidirectional models is α.*

3.5 Bidirectional models: sum of ordinals

We mentioned already that Gouveia and Santocanale proved that closure ordinals are closed under ordinal sum [12] and that this result in particular entails that every ordinal strictly less than ω^2 is a closure ordinal. In this section we show that the same can be achieved in the setting of bidirectional models, which together with the theorem from Section 3.3 in particular implies that every ordinal strictly below ω^ω is a closure ordinal on bidirectional models. Since we are going to follow closely Section 6 and Section 8 from [12], we will not provide proofs, which are variations of the original ones; we will try however to offer some intuitions.

Suppose that φ_0 and φ_1 are formulas with, respectively, α and β as their closure ordinals on bidirectional models, and we want to define some formula ψ that has closure ordinal $\alpha + \beta$ on bidirectional models. Essentially, we want ψ to depend on φ_0 and φ_1, and be such that its iteration first focuses on φ_0 and then moves to φ_1. One way to look at this is letting $\mathbb{S}_0 = (S_0, R_0, V_0)$ and $\mathbb{S}_1 = (S_1, R_1, V_1)$ be two bidirectional models where φ_0 and φ_1 converge in exactly α and β steps, respectively, and try to define a model $\mathbb{S}$ and a formula ψ where ψ converges in exactly $\alpha + \beta$ steps. Thanks to the following proposition [12] we can also assume, without loss of generality, that $S_i = [\![\mu x.\varphi_i]\!]^{\mathbb{S}_i}$.

Proposition 3.5.1 Let α be a closure ordinal of the modal μ-calculus on bidirectional models. Then there is a formula $\varphi(x)$ with closure ordinal α on bidirectional models and a bidirectional model $\mathbb{S} = (S, R, V)$ such that:

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

We now consider the disjoint union of the two models $\mathbb{S}_0$ and $\mathbb{S}_1$ and try to come up with a formula ψ whose iteration first focuses on the formula φ_0 and the model $\mathbb{S}_0$ for α many steps, and then moves to the model $\mathbb{S}_1$, iterating φ_1 for β many steps. To make sure that this process does not accidentally involve φ_1 and $\mathbb{S}_1$ during the first part of the iteration (and viceversa), we can, in some sense, strengthen the link between φ_0 and $\mathbb{S}_0$, and between φ_1 and $\mathbb{S}_1$. We consider a fresh variable p occurring in neither φ_0 or φ_1, make p false everywhere in the

$\mathbb{S}_0$ part and true everywhere in $\mathbb{S}_1$ part of the disjoint union and define, for $i \in \{0, 1\}$, a formula $\mathtt{tr}_i(\varphi_i)$ that involves the variable p and whose meaning in the disjoint union so defined is the same as the meaning of φ_i in the original model $\mathbb{S}_i$: most importantly, each iteration of $\mathtt{tr}_i(\varphi_i)$ in the disjoint union will be the same as the corresponding iteration of φ_i in $\mathbb{S}_i$. This is made precise in the following definition and proposition.

Definition 3.5.2 Let $p \notin \mathtt{PROP}$ be a fresh variable and define $\mathtt{p}_0 := \neg p$ and $\mathtt{p}_1 := p$. For $i \in \{0, 1\}$ the formula $\mathtt{tr}_i(\varphi)$ is defined by induction as follows.

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

The proofs of the items in the following proposition are either particular cases or variations of the proofs of Proposition 41 and Proposition 45 from [12]. Recall that a subset $S' \subseteq S$ of some model $\mathbb{S}$ is closed if whenever $s \in S'$ and $t \in R[s]$ then also $t \in S'$.

Proposition 3.5.3 Let $\varphi(x)$ be a formula in the basic temporal language (with fixed point operators) and $\mathbb{S} = (S, R, V)$ be an arbitrary bidirectional model. With $S' \subseteq S$, let $\mathbb{S}' = (S', R', V')$ be the submodel of $\mathbb{S}$ induced by S'. Then:

1. $[\![\mathtt{tr}_0(\varphi)]\!]^{\mathbb{S}[p \mapsto S \setminus S']} = [\![\varphi]\!]^{\mathbb{S}'}$;
2. $[\![\mathtt{tr}_1(\varphi)]\!]^{\mathbb{S}[p \mapsto S']} = [\![\varphi]\!]^{\mathbb{S}'}$;
3. $\mathtt{tr}_0(\varphi)^{\alpha}_{\mathbb{S}[p \mapsto S \setminus S']} = \varphi^{\alpha}_{\mathbb{S}'}$;
4. $\mathtt{tr}_1(\varphi)^{\alpha}_{\mathbb{S}[p \mapsto S']} = \varphi^{\alpha}_{\mathbb{S}'}$;
5. If S' is closed, then $(p \wedge \varphi)^{\alpha}_{\mathbb{S}[p \mapsto S']} = \varphi^{\alpha}_{\mathbb{S}'}$.

Returning to our formulas φ_0 and φ_1, models $\mathbb{S}_0$, $\mathbb{S}_1$ and their disjoint union as discussed above (with valuation of p as described), now we have defined formulas $\mathtt{tr}_0(\varphi_0)$ and $\mathtt{tr}_1(\varphi_1)$ such that the iteration of $\mathtt{tr}_i(\varphi_i)$ in the disjoint union behaves exactly as the iteration of φ_i in $\mathbb{S}_i$. What is missing is a way to first compute the α many steps of the iteration of $\mathtt{tr}_0(\varphi_0)$, and then move to the β many steps of the iteration of $\mathtt{tr}_1(\varphi_1)$. This last ingredient is given by the formula in the next definition, whose formulation is an adaptation to our setting of Theorem 53 in [12].

Definition 3.5.4 Let $\varphi_0(x)$ and $\varphi_1(x)$ be two formulas, let p be a variable occurring neither in φ_0 nor in φ_1 and define:

$$\begin{array}{lll}
\chi_0 & := & \neg p \to (G\neg p \wedge \mu z.\mathtt{tr}_0(\varphi_0)(z)), \\
\chi_1 & := & p \to (G(\neg p \to \mu z.\mathtt{tr}_0(\varphi_0)(z)) \wedge \mu z.\mathtt{tr}(\varphi_1)(z)), \\
\chi & := & \chi_0 \wedge \chi_1, \\
\psi(x) & := & (\neg p \wedge \mathtt{tr}_0(\varphi_0)(x)) \vee (p \wedge \mathtt{tr}_1(\varphi_1)(x) \wedge G(\neg p \to x)), \\
\Psi(x) & := & [\mathcal{U}]\chi \wedge \psi(x).
\end{array}$$

where $[\mathcal{U}]\chi := \nu z.(\chi \wedge Gz \wedge Hz)$.

Theorem 3.5.5 *Suppose $\varphi_0(x)$ and $\varphi_1(x)$ are such that their closure ordinals on bidirectional models are, respectively, α and β. Let p be a variable occurring neither in φ_0 nor in φ_1. Then the closure ordinal of $\Psi(x)$ on bidirectional models is $\alpha + \beta$.*

Before explaining the meaning of the formula $\Psi(x)$, define an *acceptable model* to be a bidirectional model where $[\mathcal{U}]\chi$ is true at every state. As the next lemma states, to prove that $\alpha+\beta$ is the closure ordinal of $\Psi(x)$ on bidirectional models, it is enough to prove that $\alpha + \beta$ is the closure ordinal of $\psi(x)$ on acceptable models: we refer to [12] for its proof.

Lemma 3.5.6 An ordinal γ is the closure ordinal of $\Psi(x)$ on bidirectional models if and only if *(i)* $\psi(x)$ converges to its least fixed point in at most γ steps on every acceptable bidirectional model, and *(ii)* there exists and acceptable bidirectional model on which the formula $\psi(x)$ converges to its least fixed point in exactly γ steps.

We now explain the meaning of the formula $\Psi(x)$. The first conjunct $[\mathcal{U}]\chi$ involves the *master modality* $[\mathcal{U}]$: the formula $[\mathcal{U}]\chi := \nu z.(\chi \wedge Gz \wedge Hz)$ is true at a state s in a bidirectional model if χ is true at every state that is reachable (in any finite amount of steps) from s, both through the accessibility relation R and its converse. An acceptable model then is a model where χ is true everywhere, so that its universe is divided in two sections: one section, where p is false, is dedicated to computing the least fixed point of $\mathtt{tr}_0(\varphi_0)$, while the second section, where p is true, is where the computation of the least fixed point of $\mathtt{tr}_1(\varphi_1)$ takes place. The variable p can be considered a colour, so that for instance we can say that the states that make p false are red, while the others are blue, so that $[\mathcal{U}]\chi$ also expresses that from the red part there are no forward transitions to blue states, but the other way around is possible. The formula $\psi(x)$ consists of two disjuncts that witness the two phases of the fixed point computation: the first disjunct $(\neg p \wedge \mathtt{tr}_0(\varphi_0)(x))$ allows the addition of points from the red part of the model to the iteration, while the second disjunct those of the blue part. Note that the states of the blue part of the model can only be added after every red point that is accessible to them through R is already included in the iteration, as they must satisfy $G(\neg p \to x)$.

Consider again the models $\mathbb{S}_0$ and $\mathbb{S}_1$ from before and their disjoint union where every point in S_0 makes p false, while every point in S_1 makes p true. To obtain an acceptable model where $\psi(x)$ converges in exactly $\alpha + \beta$ steps, we join the two submodels of the disjoint union by adding an arrow from every state of S_1 to every state of S_0: call this model $\mathbb{S}$. The computation of the least fixed point of ψ on $\mathbb{S}$ will then behave as follows. Since no point in S_1 can make $G(\neg p \to x)$ true at the beginning, the iteration will start by including points in S_0 through the disjunct $(\neg p \wedge \mathtt{tr}_0(\varphi_0)(x))$: as we argued before, this will take exactly α steps and will cover the whole red part of $\mathbb{S}$. Then, after every point in S_0 has been added, the states in S_1 will satisfy $G(\neg p \to x)$, so that they can be added through the disjunct $(p \wedge \mathtt{tr}_1(\varphi_1)(x) \wedge G(\neg p \to x))$: this process will take exactly β steps before converging.

We have thus argued that there exists an acceptable model where convergence of $\psi(x)$ happens in exactly $\alpha + \beta$ steps: a precise proof of this statement can be obtained as an adaptation of Proposition 56 in [12]. To conclude that $\alpha+\beta$ is the closure ordinal of $\Psi(x)$ it remains to show that $\psi(x)$ converges to its least fixed point in at most $\alpha + \beta$ steps on every acceptable bidirectional model: this can be proved by a variation of Proposition 55 in [12]. By Lemma 3.5.6 we conclude that Theorem 3.5.5 holds.

3.6 Uncountable closure ordinals

In this last section we discuss ω_1, the first uncountable ordinal, as a closure ordinal. Similarly to how the closure ordinal ω is linked to the semantic property of continuity, the ordinal ω_1 has been studied in [12] in relation to the semantic property of $\aleph_1$-continuity. In the following, let κ be an infinite regular cardinal.

Definition 3.6.1 A formula $\varphi \in \mu\mathtt{ML}$ is *κ-continuous* in x if

$$\mathbb{S}, s \Vdash \varphi \text{ iff } \mathbb{S}[x{\upharpoonright}C], s \Vdash \varphi, \text{ for some subset } C \subseteq S \text{ such that } |C| < \kappa$$

for every pointed model $(\mathbb{S}, s)$.

In other words, a formula φ is $\aleph_1$-continuous in x if whenever φ is true in a pointed model $(\mathbb{S}, s)$, it is enough to restrict the valuation of x to a countable subset of $V(x)$ to satisfy φ at the same state s of the model. Equivalently, the definition of κ-continuity can be stated in terms of distributivity over κ-directed sets.

Definition 3.6.2 A subset $D \subseteq \wp(S)$ is *κ-directed* if every $C \subseteq D$ such that $|C| < \kappa$ has an upper bound in D. A function $f : \wp(S) \to \wp(S)$ is κ-continuous if $f(\bigcup D) = \bigcup f[D]$ whenever $D \subseteq \wp(S)$ is κ-directed.

Proposition 3.6.3 A formula $\varphi \in \mu\mathtt{ML}$ is κ-continuous in x if and only if $\varphi_x^{\mathbb{S}}$ is a κ-continuous function for every model $\mathbb{S}$.

We observe that the notion of continuity we have already defined coincides with $\aleph_0$-continuity. Gouveia and Santocanale [12] have found a syntactic fragment of the μ-calculus that characterises $\aleph_1$-continuity.

Definition 3.6.4 Given a finite set $X \subseteq \mathtt{PROP}$, define the fragment $\mu\mathtt{ML}_X^{\aleph_1}$ by the following grammar:

$$\varphi ::= p \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Diamond\varphi \mid \mu z.\varphi' \mid \nu z.\varphi'$$

where $p \in X$, ψ is a X-free formula and $\varphi' \in \mu\mathtt{ML}_{X\cup\{z\}}^{\aleph_1}$. In case X is a singleton, say, $X = \{x\}$ we will write $\mu\mathtt{ML}_x^{\aleph_1}$ rather than $\mu\mathtt{ML}_{\{x\}}^{\aleph_1}$.

Theorem 3.6.5 (Gouveia & Santocanale) *Every formula in $\mu\mathtt{ML}_x^{\aleph_1}$ is $\aleph_1$-continuous in x. Moreover, there is an effective translation which, given a $\mu\mathtt{ML}$-formula φ, computes a formula $\varphi^{\aleph_1} \in \mu\mathtt{ML}_x^{\aleph_1}$ such that*

$$\varphi \text{ is } \aleph_1\text{-continuous in } x \text{ iff } \varphi \equiv \varphi^{\aleph_1},$$

and it is decidable whether a given formula φ is $\aleph_1$-continuous in x.

In [12] Gouveia and Santocanale actually prove that if φ is κ-continuous in x for some regular cardinal κ, then $\varphi \equiv \varphi^{\aleph_1}$ for some $\varphi^{\aleph_1} \in \mu\mathtt{ML}_x^{\aleph_1}$. This leads to the following theorem.

Theorem 3.6.6 (Gouveia & Santocanale) *There are only two fragments of the modal μ-calculus determined by continuity conditions: the fragment $\mu\mathtt{ML}_x^C$ and the fragment $\mu\mathtt{ML}_x^{\aleph_1}$.*

We also mention that the fragment of Definition 3.6.4 conincides with the fragment characterising the *finite width property* of [11, 10].

Definition 3.6.7 A formula $\varphi \in \mu\mathtt{ML}$ has the *finite width property* for $x \in \mathtt{PROP}$ if φ is monotone in x and, for every tree model $(\mathbb{S}, s)$,

$$\mathbb{S}, s \Vdash \varphi \text{ iff } \mathbb{S}[x \upharpoonright U], s \Vdash \varphi, \text{ for some finitely branching subtree } U \subseteq S,$$

where a subset $U \subseteq S$ is a *finitely branching subtree* if U is downward closed and the set $R[u] \cap U$ is finite for every $u \in U$.

Theorem 3.6.8 (Fontaine & Venema) *Every formula in $\mu\mathtt{ML}_x^{\aleph_1}$ has the finite width property for $x \in \mathtt{PROP}$. Moreover, there is an effective translation which, given a $\mu\mathtt{ML}$-formula φ, computes a formula $\varphi^{\aleph_1} \in \mu\mathtt{ML}_x^{\aleph_1}$ such that*

$$\varphi \textit{ has the finite width property for } x \textit{ iff } \varphi \equiv \varphi^{\aleph_1},$$

and it is decidable whether a given formula φ has the finite width property for $x \in \mathtt{PROP}$.

The following corollary immediately follows.

Corollary 3.6.9 A formula $\varphi \in \mu\mathtt{ML}$ has the finite width property for x if and only if it is $\aleph_1$-continuous in x.

In relation to closure ordinals, it can be proved that every κ-continuous formula converges to its least fixed point in at most κ steps [12]: this fact immediately gives the following proposition.

Proposition 3.6.10 If a formula $\varphi \in \mu\mathtt{ML}$ is $\aleph_1$-continuous in x, then its closure ordinal is at most ω_1.

An example of a formula that has closure ordinal ω_1 is provided in [12]: the bimodal formula

$$(\nu y.\Diamond_a x \wedge \Diamond_b y) \vee \Box_a \bot.$$

Recall that in Section 3.2, while trying to come up with a formula with closure ordinal ω^2, we ended up with the formula

$$\varphi := \Box_a \bot \vee (c_1 \wedge \Diamond_a x) \vee (c_2 \wedge \Diamond_a (\nu y.\Diamond_b (y \wedge x \wedge c_1)))$$

that we claimed has closure ordinal ω_1 (the formula of Section 3.2 actually involved different propositional variables, but of course this is not important). We will soon prove that this is indeed the case as motivation for a question that will be formulated at the end of the chapter. Before providing the proof we mention the result by Gouveia and Santocanale that the bimodal language of the μ-calculus is not needed for ω_1 to be a closure ordinal: the proof of the following proposition can be found in [12].

Proposition 3.6.11 For each bimodal formula φ there is a monomodal formula φ' with the same free variables of φ such that, if $\varphi \in \mu\mathtt{ML}_x^{\aleph_1}$, then $\varphi' \in \mu\mathtt{ML}_x^{\aleph_1}$, and with the following property: for each bimodal model $\mathbb{S} = (S, R_a, R_b, V)$ there is a monomodal model $\mathbb{S}' = (S', R', V')$ such that *(i)* $S \subseteq S'$, *(ii)* $\mathbb{S}, s \Vdash \varphi$ if and only if $\mathbb{S}', s \Vdash \varphi'$ for each $s \in S$, and *(iii)* $(\varphi')^\alpha_{\mathbb{S}'} = \varphi^\alpha_{\mathbb{S}}$ for every ordinal α.

We now prove that the formula φ has closure ordinal ω_1. Since $\varphi \in \mu\mathtt{ML}_x^{\aleph_1}$ it is enough to find a model where φ converges to its least fixed point in exactly ω_1 steps. To define this model we exploit the fact that for every countable limit ordinal λ there exists an ω-chain $(\gamma_{\lambda,i})_{i\in\omega}$ that is cofinal in λ, which means that for every $\beta < \lambda$ there is an $i \in \omega$ such that $\beta \leq \gamma_{\lambda,i}$. The model that we will define will have, for every countable limit ordinal λ, a state s_λ that makes c_2 true and with an R_a-transition to an infinite R_b-path $(t_{\lambda,i})_{i\in\omega}$ where c_1 is always true. Every point $t_{\lambda,i}$ in this path will in turn have an R_a-transition to a point $s_{\lambda,i}$: since the chain $(\gamma_{\lambda,i})_{i\in\omega}$ is cofinal in λ, before the state s_λ can be included in the iteration of φ through the disjunct $(c_2 \wedge \Diamond_a(\nu y.\Diamond_b(y \wedge x \wedge c_1)))$, all states of the form $s_{\lambda,i}$ and $t_{\lambda,i}$ need to be added first. This is depicted in the next picture.

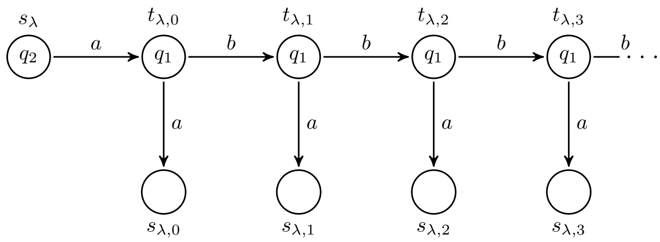

Observe that a similar situation is shown in Figure 3.3 for the state ω^2.

Proposition 3.6.12 There is a model where the formula

$$\varphi := \Box_a\bot \vee (c_1 \wedge \Diamond_a x) \vee (c_2 \wedge \Diamond_a(\nu y.\Diamond_b(y \wedge x \wedge c_1)))$$

converges to its least fixed point in exactly ω_1 many steps.

Proof. We construct a model where φ converges in exactly ω_1 many steps. Define

$$\Lambda := \{\lambda < \omega_1 \mid \lambda \text{ is a limit ordinal}\}$$

and for an ordinal α define

$$S_\alpha := \{s_\beta \mid \beta < \alpha\}.$$

For every $\lambda \in \Lambda$ fix a chain $(\gamma_{\lambda,i})_{i\in\omega}$ cofinal in λ and define the set $T_\lambda := \{t_{\lambda,i} \mid i \in \omega\}$; let

$$T := \bigcup\{T_\lambda \mid \lambda \in \Lambda\}.$$

We may choose the sets T_λ such that $T_\lambda \cap T_{\lambda'} = \emptyset$ for every distinct $\lambda, \lambda' \in \Lambda$ and $S_{\omega_1} \cap T = \emptyset$, so that in the model defined below the paths $(t_{\lambda,i})_{i\in\omega}$ and $(t_{\lambda',i})_{i\in\omega}$ will not cross. For every $\lambda \in \Lambda$ and $i \in \omega$ the set S_{ω_1} contains an element $s_{\gamma_{\lambda,i}}$ that we will denote by $s_{\lambda,i}$ to avoid notational clutter.

Let $\mathbb{S} = (S, R_a, R_b, V)$ be the model where $S := S_{\omega_1} \cup T$, the accessibility relations are defined by letting

$$\begin{aligned} R_a &:= \{(s_{\alpha+1}, s_\alpha) \mid \alpha < \omega_1\} \cup \{(s_\lambda, t_{\lambda,0}) \mid \lambda \in \Lambda\} \\ &\quad \cup \{(t_{\lambda,i}, s_{\lambda,i}) \mid \lambda \in \Lambda \text{ and } i \in \omega\}, \\ R_b &:= \{(t_{\lambda,i}, t_{\lambda,i+1}) \mid \lambda \in \Lambda \text{ and } i \in \omega\}, \end{aligned}$$

and the valuation by letting (recall that $c_1 = q_1$ and $c_2 = \neg q_1 \wedge q_2$):

$$V(q_1) := \{s_\beta \mid \beta < \omega_1 \text{ is a successor ordinal}\} \cup T,$$
$$V(q_2) := \{s_\lambda \mid \lambda \in \Lambda\}.$$

Finally, for two ordinals α and β define the relation $\alpha \prec \beta$ as follows:

- if β is a limit ordinal, then $\alpha \prec \beta$ if $\alpha < \beta$;
- if $\beta = \gamma + 1$ is a successor ordinal, then $\alpha \prec \beta$ if $\alpha < \gamma$.

From now on let λ always indicate some (non fixed) countable limit ordinal.

Claim. For each ordinal $\alpha < \omega_1$:

$$\varphi^\alpha = S_\alpha \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \alpha\}.$$

Proof of Claim. The proof goes by induction on α. The case for $\alpha = 0$ is immediate.

If α is a limit, then $\varphi^\alpha = \bigcup_{\beta<\alpha} \varphi^\beta =_{IH} \bigcup_{\beta<\alpha}(S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\})$, so that we need to show

$$\bigcup_{\beta<\alpha} (S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\}) = S_\alpha \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \alpha\}.$$

Since $S_\alpha = \bigcup_{\beta<\alpha} S_\beta$ is obvious, we only show

$$\bigcup_{\beta<\alpha} \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\} = \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \alpha\}.$$

For the $\subseteq$ inclusion, consider $t_{\lambda,i}$ with $\gamma_{\lambda,i} \prec \beta$: note then that $\gamma_{\lambda,i} < \beta < \alpha$. For the converse inclusion consider $t_{\lambda,i}$ for $\gamma_{\lambda,i} \prec \alpha$. Since α is a limit then by definition of $\prec$ it follows that $\gamma_{\lambda,i} < \alpha$. The desired result follows because $\gamma_{\lambda,i} < \gamma_{\lambda,i} + 2 < \alpha$ implies $\gamma_{\lambda,i} \prec \gamma_{\lambda,i} + 1 < \alpha$.

Now let $\alpha = \beta + 1$: then $\varphi^\alpha = \varphi_x^{\mathbb{S}}(\varphi^\beta) =_{IH} \varphi_x^{\mathbb{S}}(S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\})$. We want to show

$$\varphi_x^{\mathbb{S}}(S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\}) = S_{\beta+1} \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta + 1\}.$$

For the $\subseteq$ inclusion, suppose $u \in \varphi_x^{\mathbb{S}}(S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\})$, that is, $\mathbb{S}[x \mapsto S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\}], u \Vdash \varphi$. We proceed by case distinction as to which disjunct of φ is satisfied by u to prove that $u \in S_{\beta+1} \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta + 1\}$. First, if $u \Vdash \Box_a \bot$, then $u = s_0 \in S_{\beta+1}$.

Second, suppose $u \Vdash c_1 \wedge \Diamond_a x$. Then $u \in V(q_1)$, so $u \in \{s_\beta \mid \beta < \omega_1 \text{ is a successor ordinal}\} \cup T$, and there is a $v \in R_a[u]$ such that $v \in S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\}$. Since $(u, v) \in R_a$ and $u \neq s_\lambda$ for any $\lambda \in \Lambda$, by definition of R_a we have either *(i)* $u = s_{\gamma+1}$ and $v = s_\gamma$ for some $\gamma < \beta$, or *(ii)* $u = t_{\lambda,i}$ and $v = s_{\lambda,i}$ for some $\gamma_{\lambda,i} \prec \beta + 1$. In the former case $\gamma < \beta$ implies $\gamma + 1 < \beta + 1$ and $s_{\gamma+1} \in S_{\beta+1}$; in the latter $u = t_{\lambda,i} \in \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta + 1\}$.

Finally, suppose that u satisfies the third disjunct of φ, i.e. $u \Vdash c_2 \wedge \Diamond_a(\nu y.\Diamond_b(y \wedge x \wedge c_1))$: then there is a state $u_0 \in R_a(u)$ and an infinite R_b-path $u_0u_1u_2\dots$ such that, for all $j \in \omega$, $u_j \in S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\}$. Since $u \in V(q_2)$, then $u = s_\lambda$ for some countable limit ordinal λ, and by construction of our model then it must be that for all $i \in \omega$, $u_i = t_{\lambda,i}$ Since $\gamma_{\lambda,i} \prec \beta$ for all $i \in \omega$ and $(\gamma_{\lambda,i})_{i\in\omega}$ is cofinal in λ, then $\lambda \leq \beta$, so that $\lambda < \beta + 1$ and $u = s_\lambda \in S_{\beta+1} \subseteq S_{\beta+1} \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta + 1\}$.

Now we move to the $\supseteq$ inclusion. Suppose $u \in S_{\beta+1} \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta + 1\}$, and we want to show that $u \in \varphi^{\mathbb{S}}_x(S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\})$. First suppose $u = t_{\lambda,i}$ for some $\gamma_{\lambda,i} \prec \beta + 1$, that is, $\gamma_{\lambda,i} < \beta$. Then $s_{\lambda,i} \in S_\beta \cap R_a[u]$, so $\mathbb{S}[x \mapsto S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\}], u \Vdash c_1 \wedge \Diamond_a x$.

Otherwise, let $u \in S_{\beta+1}$. Since $S_{\beta+1} = S_\beta \cup \{s_\beta\}$ and $S_\beta \subseteq_{IH} \varphi^\beta \subseteq \varphi^{\beta+1} = \varphi^{\mathbb{S}}_x(\varphi^\beta) =_{IH} \varphi^{\mathbb{S}}_x(S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\})$, we only need to consider the case where $u = s_\beta$. We want to show $\mathbb{S}[x \mapsto S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\}], s_\beta \Vdash \varphi$. This is obvious if $\beta = 0$. If $\beta = \gamma + 1$ is a successor ordinal, then $s_\gamma \in S_\beta \cap R_a[s_\beta]$ and the result follows since $\mathbb{S}[x \mapsto S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\}], s_\beta \Vdash c_1 \wedge \Diamond_a x$. If β is a limit, then by construction $s_\beta \Vdash q_2$ and there is an infinite R_b-path $(t_{\beta,i})_{i\in\omega}$ such that $t_{\beta,0} \in R_a[s_\beta]$ and $t_{\beta,i} \Vdash q_1$ for all $i \in \omega$. Moreover, $\gamma_{\beta,i} \prec \beta$ for all $i \in \omega$ implies that $\{t_{\beta,i} \mid i \in \omega\} \subseteq \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\}$. We conclude that $\mathbb{S}[x \mapsto S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\}], s_\beta \Vdash c_2 \wedge \Diamond_a(\nu y.\Diamond_b(y \wedge x \wedge c_1))$. $\lhd$

Now we want to show that for each $\beta < \omega_1$ there is a $t \in \varphi^{\omega_1}$ such that $t \notin \varphi^\beta$. Fix $\beta < \omega_1$: then $s_\beta \notin \varphi^\beta$, since $s_\beta \notin S_\beta \cup \{t_{\lambda,i} \mid \gamma_{\lambda,i} \prec \beta\}$, but $s_\beta \in S_{\beta+1} \subseteq \varphi^{\beta+1} \subseteq \varphi^{\omega_1}$. $\square$

We conclude the chapter by posing a question. Looking for a formula with closure ordinal ω^2 in Section 3.2, we ended up with a very natural candidate that, as we have just proved, in fact has closure ordinal ω_1. On the other hand, the same formula interpreted in the setting of bidirectional models has closure ordinal ω^2 indeed, as we have shown in Section 3.3. We believe that the following question points to an interesting direction in the research for closure ordinals of the modal μ-calculus.

Question *Is there an $\aleph_1$-continuous formula $\varphi(x)$ such that $\omega^2 \leq \mathtt{cl}_x(\varphi) < \omega_1$?*

A negative answer to this question would constitute quite a surprising fact about closure ordinals of μML-formulas.

Conclusion

After presenting the modal μ-calculus in Chapter 1 we started our study of closure ordinals of its formulas in Chapter 2, where we provided a syntactic characterisation of formulas that have closure ordinal 0, we briefly discussed bounded formulas and then focused on the property of continuity by investigating its connection with the property of constructivity, and by formulating the syntactic fragment of the μ-calculus that characterises the property of continuity on finitely branching models. In Chapter 3 we presented Czarnecki's construction of a formula with closure ordinal α for every $\alpha < \omega^2$ [6] and our failed attempt to define a formula with closure ordinal ω^2, which eventually led to the definition of formula with closure ordinal α for every $\alpha < \omega^\omega$ on bidirectional models; we then adopted methodologies similar to Gouveia and Santocanale's [12] to prove that closure ordinals are closed under ordinal sum on bidirectional models, and finally considered the first uncountable closure ordinal ω_1.

In the last chapter we concluded our discussion about closure ordinals with an open question. In fact, there are many more: we close our thesis by pointing to some further research directions.

1. One open problem is, of course, the existence of a formula with a countable closure ordinal that is at least ω^2, which we discussed in Section 3.2.

2. We have also already mentioned in Section 2.4 the question formulated by Venema [9, 11] whether every formula that is constructive in x is μx-equivalent to some formula that is continuous in x, and the evidence supporting a positive answer.

3. As we have seen, Gouveia and Santocanale have proved that closure ordinals are closed under ordinal sum [12] and we have proved an analogous statement for closure ordinals in the setting of bidirectional models. A research direction proposed by Venema involves investigating whether similar statements can be formulated with different operations between ordinals: are closure ordinals closed under ordinal multiplication, exponentiation, …? Note that the question regarding closure under ordinal multiplication is closely related to the existence of a formula with closure ordinal ω^2. Furthermore, we could also ask ourselves whether $\varphi \vee \psi$, $\varphi \wedge \psi$, $\varphi[\psi/x]$, … have a closure ordinal whenever $\varphi(x)$ and $\psi(x)$ have a closure ordinal.

4. Another research direction involves bidirectional models and whether it is possible to define a formula with closure ordinal ω^ω or greater in this setting, including ω_1. One way to accomplish the latter in fact is by

taking advantage of Proposition 3.6.11 and applying it to the formula $\Box_a \bot \vee (c_1 \wedge \Diamond_a x) \vee (c_2 \wedge \Diamond_a(\nu y.\Diamond_b(y \wedge x \wedge c_1)))$ to obtain a monomodal formula φ with closure ordinal ω_1: by replacing $\Box$ with G and $\Diamond$ with F we get a formula φ' in the temporal language (but where the modalities P and H do not occur) that clearly has closure ordinal ω_1 on bidirectional models. Regarding the first part of the question, a formula with closure ordinal ω^ω on bidirectional models might involve a disjunct that generalises the infinite disjunction

$$\bigvee_{2 \leq i < \omega} (c_i \wedge \pi^{\infty}_{i-1} \wedge F(\nu y.P(y \wedge x \wedge c_{i-1})))$$

allowing to move the iteration between ω^n and ω^{n+1} for every $n \geq 1$: whether this is actually possible is left for future research.

5. Modal automata play an important role in the theory of the modal μ-calculus, and several fundamental results about this logic are proved by resorting to them. It is not very difficult to define what a closure ordinal of a modal automaton is, so that one further research direction might consist in understanding if modal automata can be a natural and useful tool for proving results about closure ordinals.

6. A game-theoretic definition of closure ordinal of a μ-calculus formula (or closure ordinal of a modal automaton) is also possible. Given a formula $\varphi(x)$ and a model $\mathbb{S}$ we can look at a variant of the evaluation game where the positions are of the form (ψ, s, β), with ψ a subformula of φ, $s \in S$ and β an ordinal: now we let a position with shape (x, s, β) belong to $\exists$, who has to move to $(\varphi(x), s, \gamma)$ for some $\gamma < \beta$, if possible. In a sense, there is a bound on how many times the variable x can be unfolded in the game. One can prove that $s \in \varphi^\alpha_{\mathbb{S}}$ if and only if $\exists$ has a winning strategy in this game from position (x, s, α), and consequently give an equivalent definition of closure ordinal in game-theoretic terms. In the setting of modal automata one could consider a similar variant of the acceptance game. We leave for future research the question whether and how this alternative definition can be taken advantage of.

7. Following Afshari and Leigh's result that formulas of the alternation-free fragment of the μ-calculus have closure ordinal strictly less than ω^2 [1], it would be interesting to further explore the connection between alternation of fixed points and closure ordinals, also in the bidirectional setting.

8. Finally, another research question involves decidability results. We have mentioned several such results in this thesis: Otto proved that it is decidable whether a modal logic formula is bounded [17]; whether a μ-calculus formula is normal (equivalently, has closure ordinal 0) is decidable; Fontaine proved that it is decidable whether a formula of the μ-calculus is continuous [9]; Gouveia and Santocanale showed that the property of $\aleph_1$-continuity is decidable [12]; by Theorem 2.4.14 the property of continuity on finitely branching models is decidable. The problems of deciding whether a formula is constructive or whether it converges to its least fixed point in at most ω_1 many steps are still open, as is the problem of deciding whether a given μ-calculus formula has a closure ordinal.

Bibliography

[1] B. Afshari and G. E. Leigh. On closure ordinals for the modal mu-calculus. In *Computer Science Logic 2013 (CSL 2013)*, volume 23 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 30–44. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2013.

[2] A. Arnold and D. Niwiński. *Rudiments of μ-calculus*, volume 146 of *Studies in Logic and the Foundations of Mathematics.* North-Holland Publishing Co., Amsterdam, 2001.

[3] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic.* Cambridge University Press, 2001.

[4] J. Bradfield and C. Stirling. Modal mu-calculi. In Patric Blackburn, Johan Van Benthem, and Frank Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical Reasoning*, pages 721 – 756. Elsevier, 2007.

[5] F. Carreiro. *Fragments of fixpoint logics.* PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, 2015.

[6] M. Czarnecki. How fast can the fixpoints in modal μ-calculus be reached? In *Fixed Points in Computer Science 2010 (FICS 2010)*, pages 35–39, August 2010.

[7] M. Czarnecki. Constructive fragment of the μ-calculus. presented at Brno for the Young Researchers' Forum, 2010. [Links can be found at `http://students.mimuw.edu.pl/~mc208417/`].

[8] E. A. Emerson and C. S. Jutla. The complexity of tree automata and logics of programs (extended abstract). In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, pages 328–337. IEEE Computer Society Press, 1988.

[9] G. Fontaine. Continuous fragment of the μ-calculus. In *Computer Science Logic*, volume 5213 of *Lecture Notes in Computer Science (LNCS)*, pages 139–153. Springer, 2008.

[10] G. Fontaine. *Modal fixpoint logic: some model theoretic questions.* PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, 2010.

[11] G. Fontaine and Y. Venema. Some model theory for the modal μ-calculus: syntactic characterisations of semantic properties. *Logical Methods in Computer Science*, Volume 14, Issue 1, February 2018.

[12] M. J. Gouveia and L. Santocanale. $\aleph_1$ and the modal mu-Calculus. In *26th EACSL Annual Conference on Computer Science Logic (CSL 2017)*, volume 82 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 38:1–38:16. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017. [An updated version can be found at `https://arxiv.org/abs/1704.03772v2`].

[13] D. Janin and I. Walukiewicz. Automata for the modal μ-calculus and related results. In *Proceedings of the Twentieth International Symposium on Mathematical Foundations of Computer Science 1995 (MFCS '95)*, volume 969 of *Lecture Notes in Computer Science (LNCS)*, pages 552–562. Springer, 1995.

[14] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ-calculus with respect to monadic second-order logic. In *Proceedings of the Seventh International Conference on Concurrency Theory, CONCUR '96*, volume 1119 of *Lecture Notes in Computer Science (LNCS)*, pages 263–277. Springer, 1996.

[15] D. Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27:333–354, 1983.

[16] D. Kozen. A finite model theorem for the propositional μ-calculus. *Studia Logica*, 47(3):233–241, 1988.

[17] M. Otto. Eliminating recursion in the μ-calculus. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS '99)*, volume 1563 of *Lecture Notes in Computer Science (LNCS)*, pages 531–540. Springer, 1999.

[18] J. van Benthem. *Modal correspondence theory.* PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, 1976.

[19] Y. Venema. Lectures on the modal μ-calculus. Lecture Notes, ILLC, University of Amsterdam, 2012.

[20] I. Walukiewicz. Completeness of Kozen's axiomatisation of the propositional μ-calculus. *Information and Computation*, 157(1):142 – 182, 2000.