

$\ 1 and the Modal \mu$ -Calculus

Maria Joao Gouveia, Luigi Santocanale

▶ To cite this version:

Maria Joao Gouveia, Luigi Santocanale. \aleph_1 and the Modal μ -Calculus. 26th EACSL Annual Conference on Computer Science Logic (CSL 2017), Aug 2017, Stockholm, Sweden. hal-01503091

HAL Id: hal-01503091 https://hal.science/hal-01503091

Submitted on 6 Apr 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

\aleph_1 AND THE MODAL μ -CALCULUS

MARIA JOÃO GOUVEIA

CEMAT-CIÊNCIAS Faculdade de Ciências, Universidade de Lisboa

LUIGI SANTOCANALE

Laboratoire d'Informatique Fondamentale de Marseille Aix-Marseille Université

ABSTRACT. For a regular cardinal κ , a formula of the modal μ -calculus is κ -continuous in a variable x if, on every model, its interpretation as a unary function of x is monotone and preserves unions of $\kappa\text{-directed sets}.$ We define the fragment $\mathcal{C}_{\aleph_1}(x)$ of the modal $\mu\text{-calculus}$ and prove that all the formulas in this fragment are \aleph_1 -continuous. For each formula $\phi(x)$ of the modal μ -calculus, we construct a formula $\psi(x) \in \mathcal{C}_{\aleph_1}(x)$ such that $\phi(x)$ is κ -continuous, for some κ , if and only if $\phi(x)$ is equivalent to $\psi(x)$. Consequently, we prove that (i) the problem whether a formula is κ -continuous for some κ is decidable, (ii) up to equivalence, there are only two fragments determined by continuity at some regular cardinal: the fragment $\mathcal{C}_{\aleph_0}(x)$ studied by Fontaine and the fragment $\mathcal{C}_{\aleph_1}(x)$. We apply our considerations to the problem of characterizing closure ordinals of formulas of the modal μ -calculus. An ordinal α is the closure ordinal of a formula $\phi(x)$ if its interpretation on every model converges to its least fixed-point in at most α steps and if there is a model where the convergence occurs exactly in α steps. We prove that ω_1 , the least uncountable ordinal, is such a closure ordinal. Moreover we prove that closure ordinals are closed under ordinal sum. Thus, any formal expression built from $0, 1, \omega, \omega_1$ by using the binary operator symbol + gives rise to a closure ordinal.

1. INTRODUCTION

The propositional modal μ -calculus [16, 19] is a well established logic in theoretical computer science, mainly due to its convenient properties for the verification of computational systems. It includes as fragments many other computational logics, PDL, CTL, CTL^{*}, its expressive power is therefore highly appreciated. Also, being capable to express all the bisimulation invariant properties of transition systems that are definable in monadic second order logic, the modal μ -calculus can itself be considered as a robust fragment of an already very expressive logic [13]. Despite its strong expressive power, this logic is still considered as a tractable one: its model checking problem, even if in the class UP \cap co-UP [15], becomes polynomial as soon

E-mail addresses: mjgouveia@fc.ul.pt, luigi.santocanale@lif.univ-mrs.fr.

The first author acknowledges partial support by FCT under grant SFRH/BSAB/128039/2016.

as some critical parameters are fixed or restricted classes of models are considered [21, 3, 5]. The widespread interest for this logic has triggered further researches that spread beyond the realm of verification: these concern the expressive power [6, 4], axiomatic bases [28], algebraic and order theoretic approaches [24], deductive systems [20, 25] and the semantics of functional programs [11].

The present paper lies at the intersection of two lines of research on the modal μ -calculus, on continuity [10] and on closure ordinals [9, 2]. Continuity of monotone functions is a fundamental phenomenon in modal logic, on which well known uniform completeness theorems rely [22, 12, 14]. Fontaine [10] characterized the formulas of the modal μ -calculus that give rise to continuous functions on Kripke models. It is well known, for example in categorical approaches to model theory [1], that the notion of continuity of monotone functions (and of functors) can be generalized to κ -continuity, where the parameter κ is an infinite regular cardinal. In the work [23] one of the authors proved that \aleph_1 -continuous functors are closed under their greatest fixed-points. Guided by this result, we present in this paper a natural syntactic fragment $\mathcal{C}_{\aleph_1}(x)$ of the modal μ -calculus whose formulas are \aleph_1 -continuous—that is, they give rise to \aleph_1 -continuous monotone unary functions of the variable x on arbitrary models. A first result that we present here is that the fragment $\mathcal{C}_{\aleph_1}(x)$ is decidable: for each $\phi(x) \in \mathsf{L}_{\mu}$, we construct a formula $\psi(x) \in \mathcal{C}_{\aleph_1}(x)$ such that $\phi(x)$ is \aleph_1 -continuous on every model if and only if $\phi(x)$ and $\psi(x)$ are semantically equivalent formulas. We borrow some techniques from [10], yet the construction of the formula $\psi(x)$ relies on a new notion of normal form for formulas of the modal μ -calculus. A closer inspection of our proof uncovers a stronger fact: the formulas $\phi(x)$ and $\psi(x)$ are equivalent if and only if, for some regular cardinal κ , $\phi(x)$ is κ -continuous on every model. The stronger statement implies that we cannot find a fragment $\mathcal{C}_{\kappa}(x)$ of κ -continuous formulas for some cardinal κ strictly larger than \aleph_1 ; any such hypothetical fragment collapses, semantically, to the fragment $\mathcal{C}_{\aleph_1}(x)$.

Our interest in \aleph_1 -continuity was wakened once more when researchers started investigating closure ordinals of formulas of the modal μ -calculus [9, 2]. Indeed, we consider closure ordinals as a wide field where the notion of κ -continuity can be exemplified and applied; the two notions, κ -continuity and closure ordinals, are naturally intertwined. An ordinal α is the closure ordinal of a formula $\phi(x)$ if (the interpretation of) this formula (as a monotone unary function of the variable x) converges to its least fixed-point $\mu_x \phi(x)$ in at most α steps in every model and, moreover, there exists at least one model in which the formula converges exactly in α steps. Not every formula has a closure ordinal. For example, the simple formula []x has no closure ordinal; more can be said, this formula is not κ -continuous for any κ . As a matter of fact, if a formula $\phi(x)$ is κ -continuous (that is, if its interpretation on every model is κ -continuous), then it has a closure ordinal $\mathsf{cl}(\phi(x)) \leq \kappa$ —here we use the fact that, using the axiom of choice, a cardinal can be identified with a particular ordinal, for instance $\aleph_0 = \omega$ and $\aleph_1 = \omega_1$. Our results on \aleph_1 -continuity shows that all the formulas in $\mathcal{C}_{\aleph_1}(x)$ have a closure ordinal bounded by ω_1 . For closure ordinals, our results are threefold. Firstly we prove that the least uncountable ordinal ω_1 belongs to the set $\operatorname{Ord}(\mathsf{L}_{\mu})$ of all closure ordinals of formulas of the propositional modal μ -calculus. Secondly, we prove that $\operatorname{Ord}(\mathsf{L}_{\mu})$ is closed under ordinal sum. It readily follows that any formal expression built from $0, 1, \omega, \omega_1$ by using the binary operator symbol + gives rise to an ordinal in $Ord(L_{\mu})$. Let us recall that Czarnecki [9] proved that all the ordinals $\alpha < \omega^2$ belong to $\operatorname{Ord}(\mathsf{L}_{\mu})$. Our results generalize Czarnecki's construction of closure ordinals and give it a rational reconstruction—every ordinal strictly smaller than ω^2 can be generated by 0, 1 and ω by repeatedly using the sum operation. Finally, the fact that there are no relevant fragments of the modal μ -calculus determined by continuity at some regular cardinal other than \aleph_0 and \aleph_1 implies that the methodology (adding regular cardinals to $\operatorname{Ord}(\mathsf{L}_{\mu})$ and closing them under ordinal sum) used until now to construct new closure ordinals for the modal μ -calculus cannot be further exploited.

Let us add some final considerations. In our view, the discovery of the fragment $\mathcal{C}_{\aleph_1}(x)$ opens an unsuspected new dimension (thus new tools, new ideas, new perspectives, etc.) in the theory of the modal μ -calculus and of fixed-point logics. Consider for example the modal μ -calculus on deterministic models, where states have at most one successor; we immediately obtain that every formula is \aleph_1 continuous on these models. Whether this and other observations can be exploited (towards understanding alternation hierarchies or reasoning using axiomatic bases, for example) is part of future researches. Yet we believe that the scopes of this work and of the problems studied here go well beyond the pure theory of the modal μ -calculus. Our interest in closure ordinals stems from a proof-theoretic work on induction and coinduction [11, 23]. There we banned ordinal notations from the syntax, as we considered the theory of ordinals too strong for our constructive goals. Yet our judgement might have gone too far, since the theory needed to deal with ordinals is not that strong; for example, many statements on ordinals do not need the axiom of choice. This makes reasonable to devise syntaxes based on ordinals. With respect to these problems, related to the semantics of programming languages, the closure ordinal problem becomes an optimal playground where to develop and test intuitions.

The paper is structured as follows. In Section 2 we introduce the notion of κ continuity and illustrate its interactions with fixed-points. In Section 3 we present the modal μ -calculus and some tools that shall be needed in the following sections. Section 4 presents our results on the fragment $C_{\aleph_1}(x)$. In Section 5 we argue that the least countable ordinal is a closure ordinal for the modal μ -calculus and that $\operatorname{Ord}(\mathsf{L}_{\mu})$ is closed under ordinal sum.

2. κ -continuous mappings and their extremal fixed-points

In this section we consider κ -continuity of mappings between powerset Boolean algebras, where the parameter κ is an infinite regular cardinal. If $\kappa = \aleph_0$, then κ -continuity coincides with the usual notion of continuity as known, for example, from [10]. The interested reader might find further informations in the monograph [1]. In the second part of this section we recall how κ -continuity interacts with least and greatest fixed-points.

In the following, κ is a fixed infinite regular cardinal, P(A) and P(B) are the powerset Boolean algebras, for some sets A and B, and $f : P(A) \to P(B)$ is a monotone mapping.

Definition 1. A subset $\mathcal{I} \subseteq P(A)$ is a κ -directed set if every collection $\mathcal{J} \subseteq \mathcal{I}$ with card $\mathcal{J} < \kappa$ has an upper bound in \mathcal{I} . A mapping $f : P(A) \to P(B)$ is κ -continuous if $f(\bigcup \mathcal{I}) = \bigcup f(\mathcal{I})$, whenever \mathcal{I} is a κ -directed set.

Remark 2. If κ' is a regular cardinal and $\kappa < \kappa'$, then a κ' -directed set is also a κ -directed set. Whence, if f is κ -continuous, then it also preserves unions of κ' -directed sets, thus it is also κ' -continuous.

We shall say that a subset X of A is κ -small if card $X < \kappa$. For example, a set X is \aleph_0 -small if and only if it is finite, and it is \aleph_1 -small if and only if it is countable.

Proposition 3. For each $X \subseteq A$, X is κ -small if and only if, for every κ -directed set $\mathcal{I}, X \subseteq \bigcup \mathcal{I}$ implies $X \subseteq I$, for some $I \in \mathcal{I}$.

Proposition 4. A monotone mapping $f : P(A) \to P(B)$ is κ -continuous if and only if, for every $X \in P(A)$,

$$f(X) = \bigcup \{ f(X') \mid X' \subseteq X, X' \text{ is } \kappa\text{-small} \}.$$

Proof. Let $f : P(A) \to P(B)$ be a monotone mapping and suppose that f is κ continuous. In P(A) every element X is the union of the set $\mathcal{I}_{\kappa}(X) := \{X' \mid X' \subseteq X, X' \text{ is } \kappa\text{-small}\}$ which is a $\kappa\text{-directed set:}$ just observe that if $\mathcal{J} \subseteq P(A)$ is a $\kappa\text{-small collection of } \kappa\text{-small subsets of } A$, then $\bigcup \mathcal{J}$ is $\kappa\text{-small, since the cardinal}$ κ is regular. Then $f(X) = f(\bigcup \mathcal{I}_{\kappa}(X)) = \bigcup f(\mathcal{I}_{\kappa}(X)) = \bigcup f(\mathcal{I}_{\kappa}(X)).$

Conversely suppose that $f : P(A) \to P(B)$ is a monotone mapping such that $f(X) = \bigcup f(\mathcal{I}_{\kappa}(X))$ for every $X \in P(A)$. Let \mathcal{I} be a κ -ideal and let X' be a κ -small set contained in $\bigcup \mathcal{I}$. By Proposition 3 there exists $I \in \mathcal{I}$ such that $X' \subseteq I$. But then $X' \in \mathcal{I}$ since \mathcal{I} is a downward closed set. Thus $\mathcal{I}_{\kappa}(\bigcup \mathcal{I}) \subseteq \mathcal{I}$ and consequently $f(\bigcup \mathcal{I}) = \bigcup f(\mathcal{I}_{\kappa}(\bigcup \mathcal{I})) \subseteq \bigcup f(\mathcal{I})$. Since $\bigcup f(\mathcal{I}) \subseteq f(\bigcup \mathcal{I})$ we obtain $f(\bigcup \mathcal{I}) = \bigcup f(\mathcal{I})$.

Fixed-points of κ -continuous mappings. The Knaster-Tarski theorem [26] states that if $f: P(A) \to P(A)$ is monotone, then the set $\bigcap \{X \subseteq A \mid f(X) \subseteq X\}$ is the least fixed-point of f. On the other hand, Kleene's fixed-point theorem states that least fixed-point of an \aleph_0 -continuous mapping f is constructible by iterating ω_0 -times f starting from the empty set, namely it is equal to $\bigcup_{n\geq 0} f^n(\emptyset)$. Generalizations of Kleene's theorem, constructing the least fixed-point of a monotone f by ordinal approximations, appeared later, see for example [8, 18]. The following Proposition 6 generalizes Kleene's theorem to κ -continuous mappings. To state it, we firstly introduce the notions of approximant and convergence.

Definition 5. If $f : P(A) \to P(A)$ is a monotone function, then the *approximants* $f^{\alpha}(\emptyset)$, α an ordinal, are inductively defined as follows:

 $f^{\alpha+1}(\emptyset):=f(f^\alpha(\emptyset))\,,\qquad f^\alpha(\emptyset):=\bigcup_{\beta<\alpha}f^\beta(\emptyset)\quad\text{when α is a limit ordinal.}$

We say that f converges to its least fixed-point in at most α steps if $f^{\alpha}(\emptyset)$ is a fixed-point (necessarily the least one) of f. We say that f converges to its least fixed-point in exactly α steps if $f^{\alpha}(\emptyset)$ is a fixed-point of f and $f^{\beta}(\emptyset) \subsetneq f^{\beta+1}(\emptyset)$, for each ordinal $\beta < \alpha$.

Let us recall that in set theory a cardinal κ is identified with the least ordinal of cardinality equal to κ . We exploit this, notationally, in the next proposition.

Proposition 6. If $f : P(A) \to P(A)$ is a κ -continuous monotone function, then it converges to its least fixed-point in at most κ steps.

Proof. Let us argue that $f^{\kappa}(\emptyset)$ is a fixed-point of f:

$$f(f^{\kappa}(\emptyset)) = f(\bigcup_{\alpha < \kappa} f^{\alpha}(\emptyset)) = \bigcup_{\alpha < \kappa} f(f^{\alpha}(\emptyset)) \subseteq \bigcup_{\alpha < \kappa} f^{\alpha}(\emptyset) = f^{\kappa}(\emptyset)$$

since the regularity of κ implies that $\{ f^{\alpha}(\emptyset) \mid \alpha < \kappa \}$ is a κ -directed set.

Propositions 7 and 8 are specific instances of a result stated for categories [23]. In order to clarify their statements, we first observe that if $f: P(B) \times P(A) \to P(B)$ is a monotone mapping, then the unary mapping $f(-, X): P(B) \to P(B), Z \mapsto f(Z, X)$, is also monotone. Hence we may consider the mapping $P(A) \to P(A)$ that sends X to the least (resp. greatest) fixed-point of f(-, X); by using the standard μ -calculus notation, we denote it by $\mu_z.f(z, -)$ (resp. $\nu_z.f(z, -)$). We also recall that f is κ -continuous w.r.t. the coordinate-wise order on $P(B) \times P(A)$ if and only if it is κ -continuous in every variable.

Proposition 7. Let $f: P(B) \times P(A) \to P(B)$ be a κ -continuous monotone mapping. If $\kappa > \aleph_0$ then $\nu_z.f(z, _): P(A) \to P(B)$ is also κ -continuous.

Proof. Let us write $g(x) := \nu_z f(z, x)$. We shall show that, for every $b \in B$ and for $X \in P(A)$, if $b \in g(X)$, then $b \in g(X')$ for some κ -small X' contained in X. Having shown this, it follows by Proposition 4 that g is continuous. Note that the condition $b \in g(X)$ holds when there exists $Z \subseteq B$ such that $b \in Z$ and $Z \subseteq f(Z, X)$. Aiming to find such a set Z we recursively obtain a family $(X_n)_{n\geq 1}$ of κ -small subsets of X and a family $(Z_n)_{n>0}$ of κ -small subsets of Z satisfying $Z_n \subseteq f(Z_{n+1}, X_{n+1})$.

For n = 0 we take $Z_0 := \{b\}$ which is a κ -small subset of f(Z, X). Now suppose we have already constructed Z_n which is κ -small and satisfies $Z_n \subseteq f(Z, X)$. Let us consider

$$\mathcal{I} := \left\{ f(Z', X') \mid X' \subseteq X, Z' \subseteq Z \text{ and } X', Z' \kappa \text{-small} \right\}.$$

Since $Z_n \subseteq f(Z, X) = \bigcup \mathcal{I}$ and \mathcal{I} is a κ -directed set, by Proposition 3 there exist $Z_{n+1}, X_{n+1} \kappa$ -small such that $Z_n \subseteq f(Z_{n+1}, X_{n+1})$. Moreover, $Z_{n+1} \subseteq Z \subseteq f(Z, X)$.

Let now $X_{\omega} := \bigcup_{n \ge 1} X_n$ and $Z_{\omega} := \bigcup_{n \ge 0} Z_n$. Notice that Z_{ω} and X_{ω} are κ -small, since we assume that $\kappa > \aleph_0$. We have therefore

$$Z_{\omega} = \bigcup_{n \ge 0} Z_n \subseteq \bigcup_{n \ge 1} f(Z_n, X_n) \subseteq f(\bigcup_{n \ge 1} Z_n, \bigcup_{n \ge 1} X_n) \subseteq f(Z_{\omega}, X_{\omega})$$

Whence $b \in Z_{\omega} \subseteq \nu_z f(z, X_{\omega})$, with $X_{\omega} \subseteq X$ and X_{ω} κ -small, proving that $\nu_z f(z, -)$ is κ -continuous.

Proposition 8. Suppose that $\kappa \geq \aleph_0$ and let $f : P(B) \times P(A) \to P(B)$ be a κ -continuous monotone mapping. Then $\mu_z f(z, -) : P(A) \to P(B)$ is also κ -continuous. (See the appendix).

3. The propositional modal μ -calculus

In this section we present the propositional modal μ -calculus and some known results on this logic that we shall need later.

Henceforward Act is a fixed finite set of actions and Prop is a countable set of propositional variables. The set L_{μ} of formulas of the propositional modal μ -calculus over Act is generated by the following grammar:

$$\phi := y \mid \neg y \mid \top \mid \phi \land \phi \mid \bot \mid \phi \lor \phi \mid \langle a \rangle \phi \mid [a] \phi \mid \mu_z.\phi \mid \nu_z.\phi,$$

where $a \in Act$, $y \in Prop$, and $z \in Prop$ is a positive variable in the formula ϕ , i.e. no occurrence of z is under the scope of a negation. We assume that Prop contains variables $x, x_1, \ldots, x_n, \ldots$ that are never under the scope of a negation nor bound in a formula ϕ . In general, we shall use $y, y_1, \ldots, y_n, \ldots$ for variables that are free in formulas, and $z, z_1, \ldots, z_n, \ldots$ for variables that are bound in formulas.

An Act-model (hereinafter referred to as model) is a triple $\mathcal{M} = \langle |\mathcal{M}|, \{R_a \mid a \in Act\}, v \rangle$ where $|\mathcal{M}|$ is a set, $R_a \subseteq |\mathcal{M}| \times |\mathcal{M}|$ for each $a \in Act$, and $v : Prop \to P(|\mathcal{M}|)$ is an interpretation of the propositional variables as subsets of $|\mathcal{M}|$. Given a model \mathcal{M} , the semantics $\llbracket \psi \rrbracket_{\mathcal{M}}$ of formulas $\psi \in L_{\mu}$ as subsets of $|\mathcal{M}|$ is recursively defined using the standard clauses from polymodal logic **K**. For example, we have

$$\llbracket \langle a \rangle \psi \rrbracket_{\mathcal{M}} = \{ s \in |\mathcal{M}| \mid \exists s' \text{ s.t. } sR_as' \text{ and } \in \llbracket \psi \rrbracket_{\mathcal{M}} \}, \\ \llbracket [a] \psi \rrbracket_{\mathcal{M}} = \{ s \in |\mathcal{M}| \mid \forall s'sR_as' \text{ implies } s' \in \llbracket \psi \rrbracket_{\mathcal{M}} \}.$$

Here we only define the semantics of the least and greatest fixed-point constructors μ and ν . To this goal, given a subset $Z \subseteq |\mathcal{M}|$, we define $\mathcal{M}[Z/z]$ to be the model that differs from \mathcal{M} only on the value Z that its valuation takes on z. The clauses for the fixed-point constructors are the following:

$$\llbracket \mu_z . \psi \rrbracket_{\mathcal{M}} := \bigcap \{ Z \subseteq |\mathcal{M}| \mid \llbracket \psi \rrbracket_{\mathcal{M}[Z/z]} \subseteq Z \}, \\ \llbracket \nu_z . \psi \rrbracket_{\mathcal{M}} := \bigcup \{ Z \subseteq |\mathcal{M}| \mid Z \subseteq \llbracket \psi \rrbracket_{\mathcal{M}[Z/z]} \}.$$

A formula $\phi \in \mathsf{L}_{\mu}$ and a variable $x \in Prop$ determine on every model \mathcal{M} the correspondence $\phi_{\mathcal{M}}^{x}: P(|\mathcal{M}|) \to P(|\mathcal{M}|)$, that sends each $S \subseteq |\mathcal{M}|$ to $\llbracket \phi \rrbracket_{\mathcal{M}[S/x]} \subseteq |\mathcal{M}|$ —in the following we shall write $\phi_{\mathcal{M}}$ for $\phi_{\mathcal{M}}^{x}$, when x is understood. Due to the syntactic restriction on the variable z in the productions of $\mu_{z}.\phi$ and $\nu_{z}.\phi$, the function $\phi_{\mathcal{M}}^{z}$ is monotone. By Tarski's theorem [26], the above clauses state that $\llbracket \mu_{z}.\phi \rrbracket_{\mathcal{M}}$ and $\llbracket \nu_{z}.\phi \rrbracket_{\mathcal{M}}$ are, respectively, the least and the greatest fixed-point of $\phi_{\mathcal{M}}^{z}$. As usual, we write $\mathcal{M}, s \Vdash \psi$ to mean that $s \in \llbracket \psi \rrbracket_{\mathcal{M}}$.

The closure of a formula. For $\phi \in L_{\mu}$, we denote by $Sub(\phi)$ the set of subformulas of ϕ . For $\psi \in Sub(\phi)$, the standard context of ψ in ϕ is the (composed) substitution

$$\sigma_{\psi}^{\phi} := [Q_{z_n}^n . \psi_n / z_n] \cdot \cdots \cdot [Q_{z_1}^1 . \psi_1 / z_1]$$

uniquely determined by the following conditions:

- (1) $\{z_1, \ldots, z_n\}$ is the set of variables bound in ϕ and free in ψ ,
- (2) for each i = 1, ..., n, $Q_{z_i}^i . \psi_i$ is the unique subformula of ϕ such that $Q^i \in \{\mu, \nu\},\$
- (3) $Q_{z_i}^j \psi_j$ is a subformula of ψ_i , for i < j.

For $\phi \in L_{\mu}$, the *closure* of ϕ , see [16], is the set $CL(\phi)$ defined as follows:

$$\operatorname{CL}(\phi) := \left\{ \psi \cdot \sigma_{\psi}^{\phi} \mid \psi \in Sub(\phi) \right\}.$$

By definition, $CL(\phi)$ is finite. The characterization of $CL(\phi)$ as the least subset satisfying some conditions yields the following observation: if $\psi \in CL(\phi)$, then $CL(\psi) \subseteq CL(\phi)$.

Game semantics. Given $\phi \in L_{\mu}$ and a model $\mathcal{M} = \langle |\mathcal{M}|, \{R_a \mid a \in Act\}, v \rangle$, the game $\mathcal{G}(\mathcal{M}, \phi)$ has $|\mathcal{M}| \times CL(\phi)$ as its set of positions. Moves are as in the table below:

| Adam's moves | Eva's moves |
|--|---|
| $(s, \psi_1 \wedge \psi_2) \rightarrow (s, \psi_i), i = 1, 2$ | $(s, \psi_1 \lor \psi_2) \to (s, \psi_i), i = 1, 2,$ |
| $(s, [a]\psi) \rightarrow (s', \psi), sR_as'$ | $(s, \langle a \rangle \psi) \rightarrow (s', \psi), sR_a s',$ |
| $(s,\nu_z.\psi) \to (s,\psi[\nu_z.\psi/z])$ | $(s, \mu_z.\psi) \rightarrow (s, \psi[\mu_z.\psi/z]).$ |

From a position of the form (s, \top) Adam loses, and from a position of the form (s, \bot) Eva loses. Also, from a position of the form (s, p) with p a propositional variable, Eva wins if and only if $s \in v(p)$; from a position of the form $(s, \neg p)$ with p a propositional variable, Eva wins if and only if $s \notin v(p)$. The definition of the game is completed by wins on infinite plays. To this goal we choose a rank function ρ : $\operatorname{CL}(\phi) \to \mathbb{N}$ such that, when ψ_1 is a subformula of ψ_2 , then $\rho(\psi_1 \cdot \sigma_{\psi_1}^{\phi}) \leq \rho(\psi_2 \cdot \sigma_{\psi_2}^{\phi})$, and such that $\rho(\mu_z.\psi)$ is odd and $\rho(\nu_z.\psi)$ is even. An infinite play $\{(s_n, \psi_n) \mid n \ge 0\}$ is a win for Eva if and only if $\max\{n \ge 0 \mid \{i \mid \rho^{-1}(\psi_i) \text{ is infinite }\}$ is even. Let us recall the following fundamental result (see for example [7, Theorem 6]):

Proposition 9. For each model \mathcal{M} and each formula $\phi \in L_{\mu}$, $\mathcal{M}, s \Vdash \phi$ if and only if Eva has a winning strategy from position (s, ϕ) in the game $\mathcal{G}(\mathcal{M}, \phi)$.

Bisimulations. Let $P \subseteq Prop$ be a subset of variables and let $B \subseteq Act$ be a subset of actions. Let \mathcal{M} and \mathcal{M}' be two models. A (P, B)-bisimulation is a relation $\mathcal{B} \subseteq |\mathcal{M}| \times |\mathcal{M}'|$ such that, for all $(x, x') \in \mathcal{B}$, we have

- $x \in v(p)$ if and only if $x' \in v'(p)$, for all $p \in P$,
- xR_by implies $x'R_by'$, for some y' such that $(y, y') \in \mathcal{B}$, and for each $b \in B$;
- $x'R_by'$ implies xR_by , for some y such that $(y, y') \in \mathcal{B}$, and for each $b \in B$.

A pointed model is a pair $\langle \mathcal{M}, s \rangle$ with $\mathcal{M} = \langle |\mathcal{M}|, \{ R_a \mid a \in Act \}, v \rangle$ a model and $s \in |\mathcal{M}|$. We say that two pointed models $\langle \mathcal{M}, s \rangle$ and $\langle \mathcal{M}', s' \rangle$ are (P, B)-bisimilar if there exists a (P, B)-bisimulation $\mathcal{B} \subseteq |\mathcal{M}| \times |\mathcal{M}|'$ with $(s, s') \in \mathcal{B}$; we say that they are bisimilar if they are (Prop, Act)-bisimilar.

Let us denote by $L_{\mu}[P, B]$ the set of formulas whose free variables are in P and whose modalities are only indexed by actions in B. The following statement is a straightforward refinement of [7, Theorem 10].

Proposition 10. If $\langle \mathcal{M}, s \rangle$ and $\langle \mathcal{M}', s' \rangle$ are (P, B)-bisimilar, then $\mathcal{M}, s \Vdash \phi$ if and only if $\mathcal{M}', s' \Vdash \phi$, for each $\phi \in \mathsf{L}_{\mu}[P, B]$.

Submodels. If $\mathcal{M} = \langle |\mathcal{M}|, \{ R_a^{\mathcal{M}} \mid a \in Act \}, v \rangle$ is a model, then a subset S of $|\mathcal{M}|$ determines the model $\mathcal{M}_{\uparrow S} := \langle S, \{ R_a \cap S \times S \mid a \in Act \}, v' \rangle$ where $v'(y) = v(y) \cap S$. We call the submodel of \mathcal{M} induced by S. A subset S of $|\mathcal{M}|$ is closed if $s \in S$ and sR_as' imply $s' \in S$, for every $a \in Act$.

Proposition 11. For each formula $\phi \in L_{\mu}$, there exists a formula $\operatorname{tr}(\phi) \in L_{\mu}$, containing a new propositional variable p, with the following property: for each model \mathcal{M} , each subset $S \subseteq |\mathcal{M}|$, and each $s \in |\mathcal{M}|$,

 $\mathcal{M}[S/p], s \models \operatorname{tr}(\phi) \text{ iff } s \in S \text{ and } \mathcal{M}_{\restriction S}, s \models \phi.$

Moreover, for each ordinal α , $\operatorname{tr}(\phi)^{\alpha}_{\mathcal{M}}(\emptyset) = \phi^{\alpha}_{\mathcal{M}_{\uparrow S}}(\emptyset)$.

Remark 12. In the statement of the previous proposition, the formula $tr(\phi)$ is, in general, defined by induction. (See Section A in the Appendix.) Yet, if S is a closed subset of \mathcal{M} , then we can simply let $tr(\phi) := p \wedge \phi$.

4. \aleph_1 -continuous fragment of the modal μ -calculus

We introduce in this section the fragment $C_{\aleph_1}(x)$ of the modal μ -calculus whose formulas, when interpreted as monotone functions of the variable x, are all \aleph_1 continuous. We show how to construct a formula $\phi' \in C_{\aleph_1}(x)$ from a given formula ϕ such that ϕ is κ -continuous, for some κ , if and only if ϕ and ϕ' are equivalent formulas. We argue therefore that the problem whether a formula is κ -continuous for some κ is decidable and, moreover, that there are no interesting notions of κ -continuity, for the modal μ -calculus, besides those for the cardinals \aleph_0 and \aleph_1 .

A formula $\phi \in L_{\mu}$ is κ -continuous in x if $\phi_{\mathcal{M}}$ is κ -continuous, for each model \mathcal{M} . If $X \subseteq Prop$, then we say that ϕ is κ -continuous in X if ϕ is κ -continuous in x, for each $x \in X$.

Define $\mathcal{C}_{\aleph_1}(X)$ to be the set of formulas of the modal μ -calculus that can be generated by the following grammar:

$$\phi := x \mid \psi \mid \top \mid \bot \mid \phi \land \phi \mid \phi \lor \phi \mid \langle a \rangle \phi \mid \mu_z . \chi \mid \nu_z . \chi,$$

where $x \in X$, $\psi \in L_{\mu}$ is a μ -calculus formula not containing any variable $x \in X$, and $\chi \in C_{\aleph_1}(X \cup \{z\})$. If we omit the last production from the above grammar, we obtain a grammar for the continuous fragment of the modal μ -calculus, see [10], which we denote here by $C_{\aleph_0}(X)$. For i = 0, 1, we shall write $C_{\aleph_i}(x)$ for $C_{\aleph_i}(\{x\})$. The main achievement of [10] is that a formula $\phi \in L_{\mu}$ is \aleph_0 -continuous in x if and only if it is equivalent to a formula in $C_{\aleph_0}(x)$.

Observe that the set of κ -continuous functions from $P(|\mathcal{M}|)^n$ to $P(|\mathcal{M}|)$, with $n \ge 1$, contains constants, projections, intersections and unions, as well as the unary functions $\phi_{\mathcal{M}}$ with $\phi = \langle a \rangle x$ for some $a \in Act$. Moreover, this set is closed under composition and diagonalisation, and so Propositions 7 and 8 immediately yield the following result:

Proposition 13. Every formula in the fragment $\mathcal{C}_{\aleph_1}(X)$ is \aleph_1 -continuous in X.

Syntactic considerations.

Definition 14. The digraph $G(\phi)$ of a formula $\phi \in L_{\mu}$ is obtained from the syntax tree of ϕ by adding an edge from each occurrence of a bound variable to its binding fixed-point quantifier. The root of $G(\phi)$ is ϕ .

Definition 15. A path in $G(\phi)$ is *bad* if one of its nodes corresponds to a subformula occurrence of the form $[a]\psi$. A bad cycle in $G(\phi)$ is a bad path starting and ending at the same vertex.

Recall that a path in a digraph is simple if it does not visit twice the same vertex. The rooted digraph $G(\phi)$ is a tree with back-edges; in particular, it has this property: for every node, there exists a unique simple path from the root to this node.

Definition 16. We say that an occurrence of a free variable x of ϕ is

- (1) bad if there is a bad path in $G(\phi)$ from the root to it;
- (2) not-so-bad (or boxed) if the unique simple path in $G(\phi)$ from the root to it is bad;

(3) very bad if it is bad and not boxed.

Example 17. It is natural to expect that a presence of the operator [] in a formula might have implications on the continuity of the formula since []x is not κ -continuous in x for every κ . However a bad occurrence of a variable in a formula can be not-so-bad (aka boxed) or very bad and the distinction between these two kind of bad occurrences can be illustrated as follows: Consider the formula $\phi(x_0, x_1) := \mu_z .[a]x_0 \wedge x_1 \wedge [a]z$. The occurrence of x_0 is not-so-bad; the word notso-bad is chosen since it is immediate from the syntax tree to guess that it is not continuous in x_0 . On the other hand, the occurrence of x_1 is called very bad since in order to recognize that ϕ is not continuous in x_1 , it is necessary to go through its dependency on the bound variable z.

Example 18. The digraph of a formula in L_{μ} .

The figure on the right represents the digraph of the formula

 $\mu_{z_1}.(y_0 \wedge \nu_{z_0}.(z_0 \wedge []z_1)) \vee (\langle \rangle y_0 \wedge y_1).$

From the figure we observe that:

The free occurrence of z_1 in the digraph of $\nu_{z_0} (z_0 \wedge []z_1)$ (in dashed) is bad but not-so-bad.

The free occurrence of y_0 in the left branch of the digraph (in bold) is very bad. The other occurrence of y_0 is not bad.

The unique free occurrence of y_1 in ϕ is not bad.

Lemma 19. For every set X of variables and every $\phi \in L_{\mu}$, the following are equivalent:

- (1) $\phi \in \mathcal{C}_{\aleph_1}(X),$
- (2) no occurrence of a variable $x \in X$ is bad in ϕ .

The $C_{\aleph_1}(x)$ -flattening of formulas. We aim at defining the $C_{\aleph_1}(x)$ -flattening $\phi^{\flat x}$ of any formula ϕ of the modal μ -calculus. This will go through the definition of the intermediate formula $\phi^{\sharp x}$ which has one more new free variable \overline{x} . The formula $\phi^{\sharp x}$ is obtained from ϕ by renaming to \overline{x} all the boxed occurrences of the variable x. The formal definition is given by induction as follows:

$$y^{\sharp x} = y \qquad (\neg y)^{\sharp x} = \neg y \qquad \top^{\sharp x} = \top \qquad \bot^{\sharp x} = \bot (\psi_0 @\psi_1)^{\sharp x} = \psi_0^{\sharp x} @\psi_1^{\sharp x} \qquad \text{with } @\in \{\land,\lor\}, \quad (\langle a\rangle\psi)^{\sharp x} = \langle a\rangle\psi^{\sharp x} \qquad ([a]\psi)^{\sharp x} = [a]\psi[\overline{x}/x] (Q_z.\psi)^{\sharp x} = Q_z.\psi^{\sharp x} \qquad \text{with } Q \in \{\mu,\nu\}.$$

In the definition of $\phi^{\sharp x}$ above, we assume that x has no bound occurrences in ϕ . The following fact is proved by a straightforward induction.

Lemma 20. Let $\phi \in L_{\mu}$. We have $\phi^{\sharp x} \cdot [x/\overline{x}] = \phi$.

The $\mathcal{C}_{\aleph_1}(x)$ -flattening $\phi^{\flat x}$ of formula $\phi \in \mathsf{L}_{\mu}$ is then defined by:

$$\phi^{\flat x} := \phi^{\sharp x} \cdot [\bot/\overline{x}]$$

and henceforward we shorten it up to ϕ^{\flat} .



Let us notice that $\phi^{\sharp x}$ (or ϕ^{\flat}) does not in general belong to $\mathcal{C}_{\aleph_1}(x)$. For example, $(\mu_z.x \vee [a]z)^{\flat} = \mu_z.x \vee [a]z \notin \mathcal{C}_{\aleph_1}(x)$ since $x \vee [a]z \notin \mathcal{C}_{\aleph_1}(X \cup \{z\})$. Yet, the following definition and lemma partially justify the choice of naming.

Definition 21. A formula ϕ is *almost-good* w.r.t. a set X of variables if no occurrence of a variable $x \in X$ is very bad. A formula ϕ is *almost-good* if it is almost-good w.r.t. $\{x\}$.

Lemma 22. If ϕ is an almost-good formula, then both $\phi^{\sharp x}$ and ϕ^{\flat} belong to $\mathcal{C}_{\aleph_1}(x)$.

We aim therefore to transform a formula ϕ into an equivalent formula in which there are not very bad occurrences of the variable x. The transformation that we define next achieves this goal. For $\phi \in \mathsf{L}_{\mu}$ and a finite set X of variables not bound in ϕ , we define $\psi^{\Box X}$ as follows. When in ψ no occurrence of a variable $x \in X$ is very bad, we take $\psi^{\Box X} := \psi$. Otherwise:

 $\begin{array}{ll} (\langle a \rangle \psi)^{\Box X} := \langle a \rangle (\psi)^{\Box X} & (\psi_1 @ \psi_2)^{\Box X} := (\psi_1)^{\Box X} @ (\psi_2)^{\Box X} & (Q_z.\psi)^{\Box X} := \psi_0 [\psi_1/\overline{z}] \\ \text{with } @ \in \{ \land, \lor \}, \ Q \in \{ \mu, \nu \} \text{ and where } \psi_0 = Q_z.\psi_2, \psi_1 = Q_{\overline{z}}.\psi_0 \text{ and } \psi_2 = (\psi^{\Box X \cup \{z\}})^{\sharp z}. \\ \text{That is, in the last clause, } \psi_2 \text{ is obtained from } \psi^{\Box X \cup \{z\}} \text{ by renaming all the boxed occurrences of } z \text{ to } \overline{z}. \\ \end{array}$

$$x^{\Box X} = x$$
 if $x \in X$, $\psi^{\Box X} = \psi$ if ψ contains no variable $x \in X$, and
 $([a]\psi)^{\Box X} = [a]\psi$.

Proposition 23. The formula $\phi^{\Box X}$ is almost-good w.r.t. X and it is equivalent to ϕ .

We can finally state the main result up to now.

Theorem 24. Every formula ϕ is equivalent to a formula ψ with $\psi^{\sharp x}$ and ψ^{\flat} in $C_{\aleph_1}(x)$.

Comparing the closures of ϕ and ϕ^{\flat} . We develop here some syntactic considerations allowing us to relate the closures of ϕ and ϕ^{\flat} . In turn, this will make it possible to relate the positions of the games $\mathcal{G}(\mathcal{M}, \phi)$ and $\mathcal{G}(\mathcal{M}, \phi^{\flat})$, so to construct, in the proof of Proposition 27, a winning strategy in the latter game from a winning strategy in the former.

Lemma 25. If \overline{x} is a free variable of ϕ and κ is either a variable not bound in ϕ or a constant, then

 $\operatorname{CL}(\phi \cdot [\kappa/\overline{x}]) = \left\{ \psi \cdot [\kappa/\overline{x}] \mid \psi \in \operatorname{CL}(\phi) \right\}.$

In particular, we have

 $\operatorname{CL}(\phi) = \left\{ \phi' \cdot [x/\overline{x}] \mid \phi' \in \operatorname{CL}(\phi^{\sharp x}) \right\}, \quad \operatorname{CL}(\phi^{\flat}) = \left\{ \phi' \cdot [\bot/\overline{x}] \mid \phi' \in \operatorname{CL}(\phi^{\sharp x}) \right\}.$

The second statement of the lemma is an immediate of the first, considering that $\phi = \phi^{\sharp x} \cdot [x/\overline{x}]$ and $\phi^{\flat} = \phi^{\sharp x} \cdot [\perp/\overline{x}]$.

The continuous fragments. Our next goal is to prove some sort of converse to Proposition 13.

A pointed model $\langle \mathcal{M}, s \rangle$ is a *tree model* if the rooted digraph $\langle |\mathcal{M}|, \bigcup_{a \in Act} R_a, s \rangle$ is a tree. Let κ be a cardinal. A tree model $\langle \mathcal{M}, s \rangle$ is κ -expanded if, for each $a \in Act$,

10

whenever xR_ax' , there are at least κ a-successors of x that are bisimilar to x'. The following lemma is a straightforward generalization of [10, Proposition 1].

Lemma 26. For each pointed model $\langle \mathcal{M}, s \rangle$ there exists a κ -expanded tree model $\langle \mathcal{T}, t \rangle$ bisimilar to $\langle \mathcal{M}, s \rangle$.

Proposition 27. If $\mathcal{M}, s \Vdash \phi$ and ϕ is κ -continuous in x, then $\mathcal{M}, s \Vdash \phi^{\flat}$.

Proof. Suppose that $\mathcal{M} = (|\mathcal{M}|, \{R_a \mid a \in A\}, v)$ is a model and that $s_0 \Vdash \phi$. We want to prove that $s_0 \Vdash \phi^{\flat}$. Notice first that, by Lemma 26, we can assume that $\langle \mathcal{M}, s_0 \rangle$ is κ -expanded tree model.

Since ϕ is κ -continuous in x and $s_0 \in \phi_{\mathcal{M}}(v(x))$, there exists $U \subseteq v(x)$, with cardinality of U strictly smaller than κ , such that $s_0 \in \phi_{\mathcal{M}}(U)$, so $\mathcal{M}[U/x], s_0 \Vdash \phi$. We shall argue that $\mathcal{M}[U/x], s_0 \Vdash \phi^{\flat}$, from which it follows that $s_0 \in \phi^{\flat}_{\mathcal{M}}(U) \subseteq \phi^{\flat}_{\mathcal{M}}(v(x))$ —since $\phi^{\flat}_{\mathcal{M}}$ is monotonic—thus $\mathcal{M}, s_0 \Vdash \phi^{\flat}$.

In the following let $\mathcal{N} = \mathcal{M}[U/x]$ (notice that \mathcal{N} is not anymore κ -expanded). Since $\mathcal{N}, s_0 \Vdash \phi$, let us fix a winning strategy for Eva in the game $\mathcal{G}(\mathcal{N}, \phi)$ from position (s_0, ϕ) . We define next a strategy for Eva in the game $\mathcal{G}(\mathcal{N}, \phi^{\flat})$ from position (s_0, ϕ^{\flat}) . Observe first that, by Lemma 25, positions in $\mathcal{G}(\mathcal{N}, \phi)$ (respectively, $\mathcal{G}(\mathcal{N}, \phi^{\flat})$) are of the form $(s, \psi[x/\overline{x}])$ (resp., $(s, \psi[\perp/\overline{x}]))$ for a formula $\psi \in \mathrm{CL}(\phi^{\sharp x})$. Therefore, at the beginning of the play, Eva plays in $\mathcal{G}(\mathcal{N}, \phi^{\flat})$ simulating the moves of the given winning strategy for the game $\mathcal{G}(\mathcal{N}, \phi)$. The simulation goes on until the play reaches a pair of positions $p = (s, [a]\chi \sigma_{[a]\chi}^{\phi^{\sharp x}} \cdot [x/\overline{x}])$ and $p' = (s, [a]\chi \sigma_{[a]\chi}^{\phi^{\sharp x}} \cdot [\perp/\overline{x}])$, for some subformula $[a]\chi$ of $\phi^{\sharp x}$, where $\chi = \chi'[\overline{x}/x]$ for some subformula χ' of ϕ .

Claim. The positions p and p' are respectively of the form $(s, [a]\psi) \in \mathcal{G}(\mathcal{N}, \phi)$ and $(s, [a]\psi') \in \mathcal{G}(\mathcal{N}, \phi^{\flat})$ for some ψ and ψ' such that $\psi[\perp/x] \to \psi'$ is a tautology. (See Lemma 48 in the Appendix.)

Thus Eva needs to continue playing in the game $\mathcal{G}(\mathcal{N}, \phi^{\flat})$ from a position of the form $(s, [a]\psi')$ where $\psi[\perp/x] \to \psi'$ is a tautology. We construct a winning stategy for Eva from this position as follows. Since the play has reached the position $(s, [a]\psi)$ of $\mathcal{G}(\mathcal{N}, \phi)$ we also know that $s \in \llbracket[a]\psi\rrbracket_{\mathcal{N}}$. We argue then that $s \in \llbracket[a]\psi\rrbracket_{\mathcal{N}}$ implies $s \in \llbracket[a]\psi[\perp/x]\rrbracket_{\mathcal{N}}$. Since $\llbracket[a]\psi[\perp/x]\rrbracket_{\mathcal{N}} \subseteq \llbracket[a]\psi'\rrbracket_{\mathcal{N}}$, Eva also has a winning strategy from position $(s, [a]\psi')$ of the game $\mathcal{G}(\mathcal{N}, \phi^{\flat})$, which she shall use to continue the play.

Claim. $s \in \llbracket [a]\psi \rrbracket_{\mathcal{N}}$ implies $s \in \llbracket [a]\psi [\bot/x] \rrbracket_{\mathcal{N}}$.

Proof of Claim. The statement of the claim trivially holds if s has no successors. Let s' be a fixed a-successor of s (i.e. sR_as'), so $\mathcal{N}, s' \Vdash \psi$; we want to show that $\mathcal{N}, s' \Vdash \psi[\perp/x]$. To this goal, recalling that $\psi[\perp/x] \in \mathsf{L}_{\mu}[Prop \setminus \{x\}, Act]$ and using Proposition 10, it is enough to prove that $\langle \mathcal{N}, s' \rangle$ is $(Prop \setminus \{x\}, Act)$ -bisimilar to some $\langle \mathcal{N}, s'' \rangle$ such that $\mathcal{N}, s'' \Vdash \psi[\perp/x]$.

Let S be the set

 $\{t \mid sR_a t, \langle \mathcal{M}, t \rangle \text{ is bisimilar to } \langle \mathcal{M}, s' \rangle, \text{ and } \downarrow t \cap U \neq \emptyset \},\$

where we have used $\downarrow t$ to denote the subtree of $\langle \mathcal{M}, s_0 \rangle$ rooted at t. Recall that the cardinality of U is strictly smaller than κ and so is the cardinality of S once it is at most equal to the cardinality of U. But the cardinality of $\{t \mid sR_at, \langle \mathcal{M}, t \rangle$ is bisimilar to $\langle \mathcal{M}, s' \rangle$ is at least κ (recall $\langle \mathcal{M}, s_0 \rangle$ is a κ -expanded tree model). Consequently there must be a successor s'' of s such that $\langle \mathcal{M}, s'' \rangle$ is

bisimilar to $\langle \mathcal{M}, s' \rangle$ and which does not belong to S, that is $\downarrow s'' \cap U = \emptyset$ (i.e. no states in U are reachable from s''). Since $\mathcal{N}, s'' \Vdash \psi$ and $\downarrow s'' \cap U = \emptyset$, we have $\mathcal{N}, s'' \Vdash \psi[\perp/x]$. Yet $\langle \mathcal{M}, s'' \rangle$ and $\langle \mathcal{M}, s' \rangle$ are bisimilar and since \mathcal{N} is obtained from \mathcal{M} just by modifying the value of the variable $x, \langle \mathcal{N}, s'' \rangle$ and $\langle \mathcal{N}, s' \rangle$ are $(Prop \setminus \{x\}, Act)$ -bisimilar. As stated before, this together with $\mathcal{N}, s'' \Vdash \psi[\perp/x]$ imply that $\mathcal{N}, s' \Vdash \psi[\perp/x]$. \Box Claim.

To complete the proof of Proposition 27 we need to argue that the strategy so defined for Eva to play in the game $\mathcal{G}(\mathcal{M}, \phi^{\flat})$ is winning. The only difficulty in asserting this is to exclude the case where the initial simulation leads to a pair of positions of the form $(s, \overline{x}[x/\overline{x}])$ and $(s, \overline{x}[\perp/\overline{x}])$. This is however excluded since in $\phi^{\sharp x}$ all the occurrences of \overline{x} are boxed, so we are enforced to go through the second step of the strategy.

Proposition 28. If, for some regular cardinal κ , $\phi \in L_{\mu}$ is κ -continuous, then ϕ is equivalent to ϕ^{\flat} .

Proof. Notice that, by monotonicity in the variable $x, \phi^{\flat} \to \phi$ is a tautology. Proposition 27 exhibits the converse implication as another tautology. \Box

Theorem 29. If for some regular cardinal κ , $\phi \in L_{\mu}$ is a κ -continuous formula, then ϕ is equivalent to a formula $\phi' \in C_{\aleph_1}(x)$.

Proof. Suppose that ϕ is κ -continuous. By Corollary 24, ϕ is equivalent to a formula ψ with $\psi^{\flat} \in \mathcal{C}_{\aleph_1}(x)$. Clearly, ψ is κ -continuous as well, so it is equivalent to ψ^{\flat} by Proposition 28. It follows that ϕ is equivalent to $\psi^{\flat} \in \mathcal{C}_{\aleph_1}(x)$.

As a consequence of the previous Theorem 29, we obtain the following result.

Theorem 30. There are only two fragments of the modal μ -calculus determined by continuity conditions: the fragment $C_{\aleph_0}(x)$ and the fragment $C_{\aleph_1}(x)$.

Theorem 31. The following problem is decidable: given a formula $\phi(x) \in L_{\mu}$, is $\phi(x)$ κ -continuous for some regular cardinal κ ?

Proof. From what has been exposed above, ϕ is κ -continuous if and only if it equivalent to the formula $\phi' \in C_{\aleph_1}(x)$, where $\phi' = (\phi^{\Box x})^{\flat}$. It is then enough to observe that there are effective processes to construct the formula ϕ' and to check whether ϕ is equivalent to ϕ' .

5. Large closure ordinals

We start by presenting some of the tools required for the two subsections in which this section is organized. Then, we prove that ω_1 , the least countable ordinal, is a closure ordinal for the modal μ -calculus. Finally, in the second subsection, we show that the set of closure ordinals is closed under the ordinal sum.

Definition 32. Let $\phi(x)$ be a formula of the modal μ -calculus. We say that an ordinal α is the *closure ordinal* of ϕ (and write $cl(\phi) = \alpha$) if, for each model \mathcal{M} , the function $\phi_{\mathcal{M}}$ converges to its least fixed-point in at most α steps, and there exists a model \mathcal{M} in which $\phi_{\mathcal{M}}$ converges to its least fixed-point in exactly α steps.

Lemma 33. If α is a closure ordinal, then there is a formula $\phi(x)$ such that $\operatorname{cl}(\phi(x)) = \alpha$ and that is total, meaning that $\llbracket \mu_x . \phi(x) \rrbracket_{\mathcal{M}} = |\mathcal{M}|$, for each model \mathcal{M} .

Proposition 34. If a formula $\phi(x)$ belongs to the syntactic fragment $C_{\aleph_1}(x)$, then it has a closure ordinal $cl(\phi(x))$ and ω_1 is an upper bound for $cl(\phi(x))$.

Proof. The formula ϕ belongs to the syntactic fragment $C_{\aleph_1}(x)$, thus it is \aleph_1 continuous and, for every model \mathcal{M} , $\phi_{\mathcal{M}}$ is \aleph_1 -continuous. It follows then from
Proposition 6 that $\phi_{\mathcal{M}}$ converges to its least fixed-point in at most ω_1 steps. \Box

 ω_1 is a closure ordinal. Let us recall that ω_1 denotes the least uncountable ordinal. In set theory cardinals are identified with particular ordinals so, using this convention, we have $\aleph_0 = \omega$ and $\aleph_1 = \omega_1$. We are going to prove that ω_1 is the closure ordinal of the following bimodal formula:

$$\Phi(x) := \nu_z . (\langle v \rangle x \land \langle h \rangle z) \lor [v] \bot .$$
⁽¹⁾

Later we shall also argue that ω_1 is the closure ordinal of a monomodal formula.

For the time being, consider $Act = \{h, v\}$; if $\mathcal{M} = \langle |\mathcal{M}|, R_h, R_v, v \rangle$ is a model, we think of R_h as a set of horizontal transitions and of R_v as a set of vertical transitions. Thus, for $s \in |\mathcal{M}|, \mathcal{M}, s \Vdash \Phi(x)$ if either (i) there are no vertical transitions from s, or (ii) there exists an infinite horizontal path from s such that each state on this path has a vertical transition to a state s' such that $\mathcal{M}, s' \Vdash x$.

By Proposition 34, the formula $\Phi(x)$ has a closure ordinal and $\mathsf{cl}(\Phi(x)) \leq \omega_1$. In order to prove that $\mathsf{cl}(\Phi(x)) = \omega_1$, we are going to construct a model \mathcal{M}_{ω_1} where $\Phi^{\omega_1}_{\mathcal{M}_{\omega_1}}(\emptyset) \not\subseteq \Phi^{\alpha}_{\mathcal{M}_{\omega_1}}(\emptyset)$ for each $\alpha < \omega_1$.

The construction relies on few combinatorial properties of posets and ordinals that we recall here. For a poset P and an ordinal α , an α -chain in P is a subset $\{p_{\beta} \mid \beta < \alpha\} \subseteq P$, with $p_{\beta} \leq p_{\gamma}$ whenever $\beta \leq \gamma < \alpha$. An α -chain $\{p_{\beta} \mid \beta < \alpha\} \subseteq P$ is cofinal in P if, for every $p \in P$ there exists $\beta < \alpha$ with $p \leq p_{\beta}$. The cofinality κ_P of a poset P is the least ordinal α for which there exists an α -chain cofinal in P. Recall that an ordinal α might be identified with the poset $\{\beta \mid \beta \text{ is an ordinal}, \beta < \alpha\}$ and so $\kappa_{\alpha} = \omega$, whenever α is a countable infinite limit ordinal; this means that, for such an α , it is always possible to pick an ω -chain cofinal in α .

For a given ordinal $\alpha \leq \omega_1$, let

$$S_{\alpha} := \{ (\beta, n) \mid \beta \text{ is an ordinal}, \beta < \alpha, 0 \le n < \omega \}.$$

We define \mathcal{M}_{ω_1} to be the model $\langle S_{\omega_1}, R_h, R_v, v \rangle$ where $v(y) = \emptyset$, for each $y \in Prop$, horizontal transitions are of the form $(\beta, n)R_h(\beta, n+1)$, for each ordinal β and each $n < \omega$, and vertical transitions from a state $(\beta, n) \in S_{\omega_1}$ are as follows:

- if $\beta = 0$, then there are no vertical transitions outgoing from (0, n);
- if β = γ + 1 is a successor ordinal, then the only vertical transitions are of the form (γ + 1, n)R_v(γ, 0);
- for β a countable limit ordinal distinct from 0, the vertical transitions are of the form $(\beta, n)R_v(\beta_n, 0)$, where the set $\{\beta_n \mid n < \omega\}$ is an ω -chain cofinal in β .

We prove that, we have $\Phi_{\mathcal{M}_{\omega_1}}(S_{\alpha}) = S_{\alpha+1}$, for each countable ordinal α , and, consequently, $\Phi^{\alpha}_{\mathcal{M}_{\omega_1}}(\emptyset) = S_{\alpha}$, for each ordinal $\alpha \leq \omega_1$. (See Lemma 49 in the Appendix.) To conclude the proof, it is enough to observe that $S_{\omega_1} \not\subseteq S_{\alpha}$, for each $\alpha < \omega_1$. Indeed, if $\alpha < \omega_1$, then we can find an ordinal β with $\alpha < \beta < \omega_1$, so the states $(\beta, n), n \geq 0$, do not belong to S_{α} .

Theorem 35. The closure ordinal of $\Phi(x)$ is ω_1 .

From a bimodal language to a monomodal language. The following statement generalizes to the modal μ -calculus a well known coding of polymodal logic to monomodal logic, see [27] and [17, Section 4].

Proposition 36. For each bimodal formula ϕ of the modal μ -calculus, we construct a monomodal formula $\operatorname{tr}^{\Psi}(\phi)$; if ϕ belongs to $\mathcal{C}_{\aleph_1}(x)$, then so does $\operatorname{tr}^{\Psi}(\phi)$. Moreover, for each bimodal model \mathcal{M} we can also construct a monomodal model \mathcal{M}^{sim} , together with an injective function $(-)^\circ : |\mathcal{M}| \to |\mathcal{M}^{sim}|$ such that, for each $s \in |\mathcal{M}|, \mathcal{M}, s \Vdash \phi$ if and only if $M^{sim}, s^\circ \Vdash \operatorname{tr}^{\Psi}(\phi)$.

Theorem 37. The monomodal formula $tr^{\Psi}(\Phi)$ has closure ordinal ω_1 .

Proof. Since the translation $\phi \mapsto \operatorname{tr}^{\Psi}(\phi)$ sends formulas in $\mathcal{C}_{\aleph_1}(x)$ to formulas in $\mathcal{C}_{\aleph_1}(x)$, $\operatorname{tr}^{\Psi}(\Phi)$ is \aleph_1 -continuous and therefore it has a closure ordinal bounded by ω_1 . To argue that the closure ordinal of $\operatorname{tr}^{\Psi}(\Phi)$ is equal to ω_1 it is enough to consider the model $\mathcal{M}_{\omega_1}^{sim}$ and rely on Proposition 36.

Closure under ordinal sum. Here we prove that the sum of any two closure ordinals is again a closure ordinal. To ease the exposition, we shall make use of the universal modality $[\mathcal{U}]$ of the μ -calculus which, in case of a monomodal language, is defined as $[\mathcal{U}]\chi := \nu_z.(\chi \wedge []z)$. The modal operator $[\mathcal{U}]$ does not satisfy the Euclidean axiom 5, yet it is satisfies all the axioms of an S4 modality.

Theorem 38. Suppose $\phi_0(x)$ and $\phi_1(x)$ are monomodal formulas that have, respectively, α and β as closure ordinals. For a variable p occurring neither in ϕ_0 nor in ϕ_1 , for $\chi := \chi_0 \wedge \chi_1$ with $\chi_0 = \neg p \rightarrow ([] \neg p \wedge (\neg p \wedge \mu_z.\phi_0(z)))$ and $\chi_1 := p \rightarrow ([] (\neg p \rightarrow \mu_z.\phi_0(z)) \wedge \mu_z.tr(\phi_1(z)))$, and for

$$\psi(x) := (\neg p \land \phi_0(x)) \lor (\operatorname{tr}(\phi_1)(x) \land [](\neg p \to x)),$$

the formula $\Psi(x) := [\mathcal{U}]\chi \wedge \psi(x)$ has closure ordinal $\alpha + \beta$.

We prove the theorem through a series of observations. Say that a model \mathcal{N} is acceptable if $\mathcal{N} \models [\mathcal{U}]\chi$. The first observation is the following: an ordinal γ is the closure ordinal of the formula $\Psi(x)$ if and only if (i) the formula $\psi(x)$ converges to its least fixed point in at most γ steps on all the acceptable models, and (ii) there exists an acceptable model on which the formula $\psi(x)$ converges to its least fixed point in exactly γ steps. (See Lemma 50 in the Appendix.)

We continue by understanding how $\psi_{\mathcal{N}}$ acts on an acceptable model \mathcal{N} . To this goal, let \mathcal{N}_0 and \mathcal{N}_1 be the submodels of \mathcal{N} induced by $v(\neg p)$ and v(p), respectively. To ease the reading, let also $N_0 := v(\neg p), N_1 := v(p), \phi_{\mathcal{N}_0} := \phi_{0\mathcal{N}_0}$ and $\phi_{\mathcal{N}_1} := \phi_{1\mathcal{N}_1}$, so $\phi_{\mathcal{N}_0} : P(N_0) \to P(N_0)$ and $\phi_{\mathcal{N}_1} : P(N_1) \to P(N_1)$. Observe that since $\mathcal{N}, s \models \neg p \to [] \neg p$ for every $s \in |\mathcal{N}|, N_0$ is a closed subset of $|\mathcal{N}|$. Then, by Proposition 11, we have $\psi_{\mathcal{N}}(X) \cap N_0 = \phi_{\mathcal{N}_0}(X \cap N_0)$ and $\operatorname{tr}(\phi_1)_{\mathcal{N}_1}(X) = \phi_{\mathcal{N}_1}(X \cap N_1)$ for each $X \subseteq |\mathcal{N}|$. Now let $\nabla(X) := N_1 \cap []_{\mathcal{N}}(N_0 \to X)$. We consider that the domain of ∇ is $P(N_0)$ while its codomain is $P(N_1)$. Therefore, $\psi_{\mathcal{N}}$ is of the form

$$\psi_{\mathcal{N}}(X) = \phi_{\mathcal{N}_0}(X \cap N_0) \cup (\phi_{\mathcal{N}_1}(X \cap N_1) \cap \nabla(X \cap N_0)).$$
⁽²⁾

We notice that if \mathcal{N} is an acceptable model, then $N_0 = \llbracket \mu_z . \phi_0(z) \rrbracket_{\mathcal{M}} = \phi_{\mathcal{N}_0}^{\alpha}(\emptyset)$ and $N_1 = \llbracket \mu_z . \phi_1(z) \rrbracket_{\mathcal{M}} = \phi_{\mathcal{N}_1}^{\beta}(\emptyset)$. Moreover, $\mathcal{N}, s \models p \to [](\neg p \to \mu_x . \phi_0(x))$, for each

 $s \in |\mathcal{N}|, \text{ so}$

$$\nabla(X) = N_1, \quad \text{whenever } X \supseteq \phi^{\alpha}_{N_0}(\emptyset) (= N_0).$$
 (3)

Proposition 39. On every acceptable model \mathcal{N} the equality $\psi_{\mathcal{N}}^{\alpha+\beta}(\emptyset) = |\mathcal{N}|$ holds and, consequently, the formula $\psi(x)$ converges before $\alpha + \beta$ steps.

Proof. Since N_0 is a closed subset of $|\mathcal{N}|$, by Proposition 11, we have

$$\psi^{\delta}_{\mathcal{N}}(\emptyset) \cap N_0 = \psi^{\delta}_{\mathcal{N}_0}(\emptyset) = \phi^{\delta}_{\mathcal{N}_0}(\emptyset) \tag{4}$$

for each ordinal δ . Consequently, $\psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_0 \supseteq \psi_{\mathcal{N}}^{\alpha}(\emptyset) \cap N_0 = \phi_{\mathcal{N}_0}^{\alpha}(\emptyset)$, for every ordinal γ . By a straightforward induction (See Lemma 51 in the Appendix) we also prove that, for each ordinal γ ,

$$\phi_{\mathcal{N}_1}^{\gamma}(\emptyset) \subseteq \psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_1.$$
(5)

Therefore $|\mathcal{N}| = N_0 \cup N_1 = \phi_{\mathcal{N}_0}^{\alpha}(\emptyset) \cup \phi_{\mathcal{N}_1}^{\beta}(\emptyset) \subseteq (\psi_{\mathcal{N}}^{\alpha+\beta}(\emptyset) \cap N_0) \cup (\psi_{\mathcal{N}}^{\alpha+\beta}(\emptyset) \cap N_1) = \psi_{\mathcal{N}}^{\alpha+\beta}(\emptyset)$.

Proposition 40. There exists an acceptable model \mathcal{N} on which $\psi(x)$ converges exactly after $\alpha + \beta$ steps.

Proof. Since the formulas $\phi_0(x)$ and $\phi_1(x)$ have, respectively, α and β as closure ordinals, by Proposition 33 there exist models $\mathcal{M}_{\gamma} = \langle |\mathcal{M}_{\gamma}|, R_{\gamma}, v_{\gamma} \rangle, \ \gamma \in \{\alpha, \beta\}$, such that for every $\alpha' < \alpha$ and $\beta' < \beta \ [\![\mu_x.\phi_0(x)]\!]_{\mathcal{M}_{\alpha}} = |\mathcal{M}_{\alpha}| = \phi_0^{\alpha}_{\mathcal{M}_{\alpha}}(\emptyset) \neq \phi_0^{\alpha'}_{\mathcal{M}_{\alpha}}(\emptyset)$ and $[\![\mu_x.\phi_1(x)]\!]_{\mathcal{M}_{\beta}} = |\mathcal{M}_{\beta}| = \phi_1^{\beta}_{\mathcal{M}_{\beta}}(\emptyset) \neq \phi_1^{\beta'}_{\mathcal{M}_{\beta}}(\emptyset).$

We construct now the model $\mathcal{M}_{\alpha+\beta}$ by making the disjoint union of the sets $|\mathcal{M}_{\alpha}|$ and $|\mathcal{M}_{\beta}|$, endowed with $R_{\alpha} \cup R_{\beta} \cup \{(s,s') \mid s \in |\mathcal{M}_{\beta}|, s' \in |\mathcal{M}_{\alpha}|\}$ and the valuation v defined by $v(q) := |\mathcal{M}|_{\beta}$, if q = p, and $v(q) := v_{\alpha}(q) \cup v_{\beta}(q)$ otherwise. Let us put $\mathcal{N} = \mathcal{M}_{\alpha+\beta}$. Observe now that $\mathcal{M}_{\alpha+\beta}$ is an acceptable model and that $\nabla(X) = \emptyset$ for every $X \subseteq |\mathcal{N}|$ such that $X \cap N_0 \subsetneq \phi_{\mathcal{N}_0}^{\alpha}(\emptyset)$. Because of this, the inclusion (5) is actually an equality. (See Lemma 52 in the Appendix.) But then we apply equations (4) and (5) to obtain $\psi_{\mathcal{N}}^{\alpha}(\emptyset) = \phi_{\mathcal{N}_0}^{\alpha}(\emptyset) = \psi_{\mathcal{N}}^{\delta}(\emptyset)$ and $\psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) = N_0 \cup \phi_{\mathcal{N}_1}^{\gamma}(\emptyset)$, for ordinals $\delta < \alpha$ and γ . Finally, $\psi_{\mathcal{N}}^{\alpha+\beta}(\emptyset) = |\mathcal{N}| = N_0 \cup \phi_{\mathcal{N}_1}^{\beta}(\emptyset) = \psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset)$, for $\gamma < \beta$.

References

- J. Adamek and J. Rosicky. Locally Presentable and Accessible Categories:. Cambridge University Press, Cambridge, 1994.
- [2] B. Afshari and G. E. Leigh. On closure ordinals for the modal mu-calculus. In S. Ronchi Della Rocca, editor, *Computer Science Logic 2013 (CSL 2013), CSL 2013, September 2-5,* 2013, Torino, Italy, volume 23 of *LIPIcs*, pages 30–44. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.
- [3] L. Alberucci and A. Facchini. The Modal μ-Calculus Hierarchy on Restricted Classes of Transition Systems. The Journal of Symbolic Logic, 74(4):1367–1400, Dec. 2009.
- [4] D. Berwanger, E. Grädel, and G. Lenzi. On the variable hierarchy of the modal μ-calculus. In J. Bradfield, editor, CSL 2002, Proceedings, pages 352–366, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
- [5] M. Bojanczyk, C. Dittmann, and S. Kreutzer. Decomposition theorems and model-checking for the modal μ-calculus. In CSL-LICS '14, pages 17:1–17:10, New York, NY, USA, 2014. ACM.
- [6] J. Bradfield. The modal mu-calculus alternation hierarchy is strict. Theoretical Computer Science, 195(2):133 – 153, 1998.

- [7] J. Bradfield and I. Walukiewicz. The mu-calculus and model-checking. In H. V. E. Clarke, T. Henzinger, editor, *Handbook of Model Checking*. Springer-Verlag, 2015.
- [8] R. Cousot and P. Cousot. Constructive versions of Tarski's fixed point theorems. Pacific Journal of Mathematics, 82(1):43-57, 1979.
- [9] M. Czarnecki. How fast can the fixpoints in modal μ-calculus be reached? In L. Santocanale, editor, 7th Workshop on Fixed Points in Computer Science, FICS 2010, page 89, Brno, Czech Republic, Aug. 2010. Available from Hal: https://hal.archives-ouvertes.fr/hal-00512377.
- [10] G. Fontaine. Continuous fragment of the mu-calculus. In M. Kaminski and S. Martini, editors, *CSL 2008. Proceedings*, volume 5213 of *Lecture Notes in Computer Science*, pages 139–153. Springer, 2008.
- [11] J. Fortier and L. Santocanale. Cuts for circular proofs: semantics and cut-elimination. In S. Ronchi Della Rocca, editor, *Computer Science Logic 2013 (CSL 2013), CSL 2013, September 2-5, 2013, Torino, Italy*, volume 23 of *LIPIcs*, pages 248–262. Schloss Dagstuhl -Leibniz-Zentrum fuer Informatik, 2013.
- [12] V. V. G. Sambin. A new proof of sahlqvists theorem on modal definability and completeness. Journal of Symbolic Logi, 54:992–999, 1989.
- [13] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic. In U. Montanari and V. Sassone, editors, CON-CUR '96, Proceedings, pages 263–277, Berlin, Heidelberg, 1996. Springer Berlin Heidelberg.
- [14] B. Jónsson. On the canonicity of Sahlqvist identities. Studia Logica: An International Journal for Symbolic Logic, 53(4):473–491, 1994.
- [15] M. Jurdziski. Deciding the winner in parity games is in UP ∩ co-UP. Information Processing Letters, 68(3):119 – 124, 1998.
- [16] D. Kozen. Results on the propositional mu-calculus. Theor. Comput. Sci., 27:333–354, 1983.
- [17] M. Kracht and F. Wolter. Simulation and transfer results in modal logic A survey. Studia Logica, 59(1):149–177, 1997.
- [18] J. Lassez, V. L. Nguyen, and L. Sonenberg. Fixed point theorems and semantics: A folk tale. Inf. Process. Lett., 14(3):112–116, 1982.
- [19] G. Lenzi. The modal $\mu\text{-calculus:}$ a survey. Task Quarterly, 9(3):29–316, 2005.
- [20] D. Niwiski and I. Walukiewicz. Games for the μ-calculus. Theoretical Computer Science, 163(1):99 – 116, 1996.
- [21] J. Obdržálek. Fast mu-calculus model checking when tree-width is bounded. In W. A. Hunt and F. Somenzi, editors, CAV 2003, Proceedings, pages 80–92, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.
- [22] H. Sahlqvist. Completeness and correspondence in the first and second order semantics for modal logic. In S. Kanger, editor, *Proceedings of the Third Scandinavian Logic Symposium*, volume 82 of *Studies in Logic and the Foundations of Mathematics*, pages 110 – 143. Elsevier, 1975.
- [23] L. Santocanale. μ -Bicomplete Categories and Parity Games. *ITA*, 36(2):195–227, 2002. Extended version appeared as LaBRI report RR-1281-02 is available from Hal: https://hal.archives-ouvertes.fr/hal-01376731.
- [24] L. Santocanale. Completions of µ-algebras. Ann. Pure Appl. Logic, 154(1):27–50, 2008.
- [25] T. Studer. On the proof theory of the modal mu-calculus. *Studia Logica*, 89(3):343–363, 2008.
- [26] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. Pacific Journal of Mathematics, 5:285–309, 1955.
- [27] S. K. Thomason. Independent propositional modal logics. Studia Logica, 39(2):143-144, 1980.
- [28] I. Walukiewicz. Completeness of Kozen's axiomatisation of the propositional μ -calculus. Information and Computation, 157(1):142 – 182, 2000.

APPENDIX A. SUBFRAMES AND SUBMODELS

Here we collect some considerations about different types of submodels that are needed to prove Propositions 11 and Proposition 36. To this goal le us recall the usual notion of Kripke frame (hereafter, frame). An Act-frame (or simply, a frame, if Act is understood) is a pair $\mathcal{F} = \langle |\mathcal{F}|, \{R_a \mid a \in Act\} \rangle$ with $|\mathcal{F}|$ a set and $R_a \subseteq |\mathcal{F}| \times |\mathcal{F}|$, for each $a \in Act$. That is, a frame is a model without a valuation of propositional variables. If $v : Prop \to P(F)$ is a valuation, then we denote by \mathcal{F}_v the model $\langle \mathcal{F}, v \rangle$. The complex algebra \mathcal{F}^{\sharp} of a frame \mathcal{F} is the Boolean algebra of subsets of $|\mathcal{F}|$ endowed with the modal operators $\langle a \rangle_{\mathcal{F}^{\sharp}}, a \in Act$, defined by

 $\langle a \rangle_{\mathcal{F}^{\sharp}}(S) := \{ s \in |\mathcal{F}| \mid \exists s' \in S \text{ s.t. } sR_as' \}, \text{ for } S \subseteq |\mathcal{F}|.$

We consider next two frames \mathcal{F} and \mathcal{G} such that $|\mathcal{G}| \subseteq |\mathcal{F}|$. \mathcal{F} and \mathcal{G} might have different sets of actions: say that \mathcal{F} is an A-frame, \mathcal{G} is a B-frame, while we do not suppose that A = B. To ease the reading, we let $F := |\mathcal{F}|$ and $G := |\mathcal{G}|$, so $G \subseteq F$.

Definition 41. Let $\Psi = \{ \psi_b \in \mathsf{L}_\mu[p,q] \mid b \in B \}$ be a collection of formulas containing only the free variables p, q in positive position. If \mathcal{F} and \mathcal{G} are frames as above, then we say that \mathcal{G} is *p*-defined in \mathcal{F} by Ψ if, for each $b \in B$ and each $S \subseteq F$,

$$\langle b \rangle_{\mathcal{G}^{\sharp}}(G \cap S) = \llbracket \psi_b(p,q) \rrbracket_{\mathcal{F}_{[G/p,S/q]}}.$$

Above [G/p, S/q] is the valuation that sends p to G and q to S (and, say, any other propositional variable to \emptyset). In this sense, $\mathcal{F}_{[G/p,S/q]}$ denotes the model $\langle \mathcal{F}, [G/p, S/q] \rangle$.

The previous definition means that each modal operator $\langle b \rangle$ of the algebra \mathcal{G}^{\sharp} is described using a term of the algebra \mathcal{F}^{\sharp} .

Example 42. Suppose that \mathcal{G} is a *subframe* of $\mathcal{F} = \langle F, \{R_a \mid a \in A\}\rangle$, by which we mean that $A = B, \mathcal{G} = \langle G, \{R'_a \mid a \in A\}\rangle$ with $R'_a = R_a \cap G \times G$, for each $a \in A$. Then \mathcal{G} is *p*-defined in \mathcal{F} by the collection of formulas $\{p \land \langle a \rangle (p \land q) \mid a \in A\}$. Example 42

Example 43. A second important example of *p*-definability comes from Thomason's coding of bimodal logic into monomodal logic, see [27] and [17, Section 4]. For each (fixed-point free) bimodal formula ϕ , Thomason constructs a monomodal formula ϕ^{sim} ; for each bimodal model \mathcal{M} Thomason constructs a monomodal model M^{sim} and an injective function $(-)^{\circ} : |\mathcal{M}| \to |\mathcal{M}^{sim}|$. These data have the following property:

Fact. For each $s \in |\mathcal{M}|$, $\mathcal{M}, s \Vdash \phi$ if and only if $M^{sim}, s^{\circ} \Vdash \phi^{sim}$.

Here we give an approximate description of the construction of \mathcal{M}^{sim} that, however, will be enough for our goals (namely proving Proposition 36). For a $\{h, v\}$ model \mathcal{M} , let \mathcal{M}^{sim} be the monomodal model with $|\mathcal{M}^{sim}| = |\mathcal{M}| \times \{h, v\}$, such that v(x, i) = v(x) and with accessibility relation described as follows:

$$(x,h)R(y,h)$$
, when xR_hy ,
 $(x,v)R(y,v)$, when xR_vy ,
 $(x,v)R(x,h)$, and $(x,h)R(x,v)$

Since the function sending $x \in |\mathcal{M}|$ to $x^{\circ} := (x, h) \in |M^{sim}|$ is injective, so we can identify $|\mathcal{M}|$ as a subset of $|M^{sim}|$. Call \mathcal{N} the image of \mathcal{M} within $|\mathcal{M}^{sim}|$, call \mathcal{G}

the underlying frame of \mathcal{N} and \mathcal{F} the underlying frame of $|\mathcal{M}^{sim}|$. The above Fact relies on \mathcal{G} being *p*-defined in \mathcal{F} by $\Psi = \{\psi_v, \psi_h\}$, where

$$\begin{split} \psi_h(p,q) &= p \land \langle \ \rangle (p \land q) \,, \\ \psi_v(p,q) &= p \land \langle \ \rangle (\neg p \land \langle \ \rangle (\neg p \land \langle \ \rangle (p \land q))) \,. \end{split} \tag{Example 43} \ \Box \end{split}$$

Example 44. We are thankful to an anonymous reader to suggest that a coding of bimodal logic into monomodal logic can be realized in a more straighforward way. Namely, for a given bimodal frame \mathcal{M} we can define \mathcal{M}^{sim} as having underlying set $|\mathcal{M}| \cup R$, and let accessibility relations as follows:

$$xRy$$
, when xR_hy ,
 $xR(x,y)$ and $(x,y)Ry$, when xR_vy .

Clearly $|\mathcal{M}|$ embeds into $|M^{sim}|$, so we can identify $|\mathcal{M}|$ as a subset of $|M^{sim}|$. Call \mathcal{N} the image of \mathcal{M} within $|\mathcal{M}^{sim}|$, call \mathcal{G} the underlying frame of \mathcal{N} and \mathcal{F} the underlying frame of $|\mathcal{M}^{sim}|$. Then \mathcal{G} is *p*-defined in \mathcal{F} by $\Psi = \{\psi_v, \psi_h\}$, where

$$egin{aligned} &\psi_h(p,q) = p \land \langle \
angle(p\land q) \,, \ &\psi_v(p,q) = p \land \langle \
angle(\neg p \land \langle \
angle(p\land q)) \,. \end{aligned}$$
 Example 44 \Box

Let $\pi : P(F) \to P(G)$ be defined by $\pi(S) = S \cap G$. If $v : Prop \to P(F)$, then $\pi \circ v : Prop \to P(G)$ is the valuation in G such that $(\pi \circ v)(y) = G \cap v(y)$ for each $y \in Prop$.

Proposition 45. Let $\Psi = \{ \psi_b \in \mathsf{L}_\mu[p,q] \mid b \in B \}$ and suppose that p is a new variable (i.e., $p \notin Prop$). For each formula $\phi \in \mathsf{L}_\mu[Prop, B]$ there exists a formula $\mathsf{tr}^\Psi(\phi) \in \mathsf{L}_\mu[Prop \cup \{p\}, A]$ such that, if \mathcal{G} is p-defined in \mathcal{F} by Ψ , then, for each valuation $v : Prop \to P(F)$,

$$\llbracket \phi \rrbracket_{\mathcal{G}_{\pi \circ v}} = \llbracket \operatorname{tr}^{\Psi}(\phi) \rrbracket_{\mathcal{F}_{v}[G/p]}.$$
(6)

The formula $tr^{\Psi}(\phi)$ is defined by induction as follows:

$$\begin{split} \mathbf{tr}^{\Psi}(y) &= p \wedge y & \mathbf{tr}^{\Psi}(\neg y) = p \wedge \neg y \\ \mathbf{tr}^{\Psi}(\bot) &= \bot & \mathbf{tr}^{\Psi}(\top) = p \\ \mathbf{tr}^{\Psi}(\psi_0 @ \psi_1) &= \mathbf{tr}^{\Psi}(\psi_0) @ \mathbf{tr}^{\Psi}(\psi_1) , & @ \in \{ \wedge, \lor \} \\ \mathbf{tr}^{\Psi}(\langle b \rangle \psi) &= \psi_b [\mathbf{tr}^{\Psi}(\psi)/q] \\ \mathbf{tr}^{\Psi}([b]\psi) &= p \wedge \psi_b^{op} [\mathbf{tr}^{\Psi}(\psi)/q] \\ \mathbf{tr}^{\Psi}(\mu_z.\psi) &= \mu_z.\mathbf{tr}^{\Psi}(\psi) & \mathbf{tr}^{\Psi}(\nu_z.\psi) = \nu_z.\mathbf{tr}^{\Psi}(\psi) . \end{split}$$

In the above definition ψ_b^{op} is a formula dual of ψ_b , thus semantically behaving as $\neg \psi_b[\neg q/q]$. In case \mathcal{G} is a subframe of \mathcal{F} , as in Example 42, then we shall simply write $\operatorname{tr}(\phi)$ for $\operatorname{tr}^{\Psi}(\phi)$.

Remark 46. For a formula ϕ , let us denote by $[\![tr^{\Psi}(\phi)]\!]_{\mathcal{F}[G/p]}$ the mapping sending a valuation $v \in P(F)^{Prop}$ to $[\![tr^{\Psi}(\phi)]\!]_{\mathcal{F}_v[G/p]} \in P(F)$; let us denote by $[\![\phi]\!]_{\mathcal{G}}$ the mapping sending a valuation $v' \in P(G)^{Prop}$ to $[\![\phi]\!]_{\mathcal{G}_{v'}} \in P(G)$. The statement of Proposition 45 implies that $[\![tr^{\Psi}(\phi)]\!]_{\mathcal{F}[G/p]}$ takes values in P(G) and, moreover, that the following diagram commutes:



Proof of Proposition 45. The proof that equation (6) holds is by induction on formulas. The base cases as well as the case for the logical operators \land and \lor are straightforward.

$$\begin{split} \llbracket \mathbf{tr}^{\Psi}(\langle b \rangle \psi) \rrbracket_{\mathcal{F}_{v}[G/p]} &= \llbracket \psi_{b} \llbracket \mathbf{tr}^{\Psi}(\psi) / q \rrbracket_{\mathcal{F}_{v}[G/p]} \\ &= \llbracket \psi_{b} \rrbracket_{\mathcal{F}_{v}[G/p, \llbracket \mathbf{tr}^{\Psi}(\psi) \rrbracket_{\mathcal{F}_{v}[G/p]} / q] \\ &= \llbracket \langle b \rangle q \rrbracket_{\mathcal{G}_{\pi \circ v}} \llbracket \psi \rrbracket_{\mathcal{G}_{\pi \circ v}} / q] = \llbracket \langle b \rangle \psi \rrbracket_{\mathcal{G}_{\pi \circ v}}, \\ \llbracket \mathbb{tr}^{\Psi}([b] \psi) \rrbracket_{\mathcal{F}_{v}[G/p]} &= \llbracket p \land \neg \psi_{b} [\neg \mathbf{tr}^{\Psi}(\psi) / q] \rrbracket_{\mathcal{F}_{v}[G/p]} \\ &= G \cap (\llbracket \psi_{b} [\neg \mathbf{tr}^{\Psi}(\psi) / q] \rrbracket_{\mathcal{F}_{v}[G/p]})^{c} \\ &= G \cap (\llbracket \psi_{b} \llbracket \mathcal{F}_{v}[G/p, S^{c}/q])^{c} \\ &\qquad \text{with } S = \llbracket \mathbf{tr}^{\Psi}(\psi) \rrbracket_{\mathcal{F}_{v}[G/p]} = \llbracket \phi \rrbracket_{\mathcal{G}_{\pi \circ v}} \\ &= \llbracket \cap \langle [[\langle b \rangle q] \rrbracket_{\mathcal{G}_{\pi \circ v}} [G \cap S^{c}/q] \rangle^{c} \\ &= \llbracket [b] q \rrbracket_{\mathcal{G}_{\pi \circ v}} [S/q] \\ &= \llbracket [b] q \rrbracket_{\mathcal{G}_{\pi \circ v}} [W \rrbracket_{\mathcal{G}_{\pi \circ v}} / q] \\ &= \llbracket [b] q \rrbracket_{\mathcal{G}_{\pi \circ v}} [W \rrbracket_{\mathcal{G}_{\pi \circ v}} / q] \\ &= \llbracket [b] \psi \rrbracket_{\mathcal{G}_{\pi \circ v}} . \end{split}$$

For the least and greatest fixed-points, let us consider the two functions

$$f(S) = \llbracket \operatorname{tr}^{\Psi}(\phi) \rrbracket_{\mathcal{F}_{v}[G/p, S/z]} \quad \text{and} \quad g(T) = \llbracket \phi \rrbracket_{\mathcal{G}_{\pi \circ v}[T/z]}$$

Firstly notice that $f: P(F) \to P(F)$ and that $g: P(G) \to P(G)$. Yet, the inductive hypothesis is that

$$\llbracket \operatorname{tr}^{\Psi}(\phi) \rrbracket_{\mathcal{F}_{v}[G/p]} = \llbracket \phi \rrbracket_{\mathcal{G}_{\pi \circ v}}$$

for each vaualtion v, in particular for the valuation v[S/z], with $S \subseteq F$. That is, we have

$$f(S) = g(S \cap G),$$

for each $S \subseteq F$.

Let us denote by Pre_h the set of prefixed-points of some monotone function h and by $\operatorname{lfp}.h$ its least element. It follows from our previous considerations that Pre_g is included in Pre_f and that if $S \in \operatorname{Pre}_f$, then $\pi(S) \in \operatorname{Pre}_g$. As is it readily seen, π is right adjoint of the inclusion of Pre_g is into Pre_f , and as usual for adjoint pair of maps, the inclusion of Pre_f necessarily preserves the least element (i.e. $\operatorname{lfp}.g = \operatorname{lfp}.f$). We obtain

$$\llbracket \texttt{tr}^{\Psi}(\mu_z.\psi) \rrbracket_{\mathcal{F}_v[G/p]} = \texttt{lfp}.f = \texttt{lfp}.g = \llbracket \mu_z.\psi \rrbracket_{\mathcal{G}_{\pi \circ v}}$$

For the greatest fixed-point, denote by Pos_h the set of postfixed-points of some monotone function h and by gfp.h its greatest element. Observe that if $S \subseteq f(S)$ then $S \subseteq G$. It immediately follows that $\mathsf{Pos}_f = \mathsf{Pos}_g$, so

$$\llbracket \operatorname{tr}^{\Psi}(\mu_{z}.\psi) \rrbracket_{\mathcal{F}_{v}[G/p]} = \operatorname{gfp}.f = \operatorname{gfp}.g = \llbracket \mu_{z}.\psi \rrbracket_{\mathcal{G}_{\pi\circ v}}.$$

We recast now our previous observations in the language of models. If $\mathcal{M} = \langle |\mathcal{M}|, \{R_a^{\mathcal{M}} \mid a \in Act\}, v \rangle$ and $\mathcal{N} = \langle |\mathcal{N}|, \{R_a^{\mathcal{N}} \mid a \in Act\}, v_{\mathcal{N}} \rangle$ are models, then we say that \mathcal{N} is a *submodel* of \mathcal{M} if $|\mathcal{N}|$ is a subset of $|\mathcal{M}|$ and, for each $y \in Prop$ and each $a \in Act$,

$$v_{\mathcal{N}}(y) = v_{\mathcal{M}}(y) \cap |\mathcal{N}| \text{ and } R_a^{\mathcal{N}} = |\mathcal{N}| \times |\mathcal{N}| \cap R_a^{\mathcal{M}}.$$

Thus \mathcal{N} is a submodel of \mathcal{M} if and only if, for some frame \mathcal{F} , for a valuation $v: Prop \to P(|\mathcal{F}|)$, and for a subframe \mathcal{G} of \mathcal{F} , $\mathcal{M} = \mathcal{F}_v$ and $\mathcal{N} = \mathcal{G}_{\pi \circ v}$. Every subset S of $|\mathcal{M}|$ induces the submodel $\mathcal{M}_{\uparrow S}$ of \mathcal{M} defined as follows:

$$\mathcal{M}_{\uparrow S} := \langle S, \{ R_a \cap S \times S \mid a \in Act \}, v' \rangle \tag{7}$$

where $v'(y) = v(y) \cap S$, is a submodel of \mathcal{M} and it is called the *submodel of* \mathcal{M} induced by S. A subset S of $|\mathcal{M}|$ is closed if $s \in S$ and sR_as' imply $s' \in S$, for every $a \in Act$. A submodel \mathcal{N} of \mathcal{M} is closed if $|\mathcal{N}|$ is a closed subset of $|\mathcal{M}|$. Proposition 45 instantiates then to models and submodels as follows:

Proposition 11. For each formula $\phi \in L_{\mu}$, there exists a formula $\operatorname{tr}(\phi) \in L_{\mu}$, containing a new propositional variable p, with the following property: for each model \mathcal{M} , each subset $S \subseteq |\mathcal{M}|$, and each $s \in |\mathcal{M}|$,

$$\mathcal{M}[S/p], s \models \mathsf{tr}(\phi) \quad iff \ s \in S \ and \ \mathcal{M}_{\upharpoonright S}, s \models \phi.$$

Moreover, for each ordinal α , $\operatorname{tr}(\phi)^{\alpha}_{\mathcal{M}}(\emptyset) = \phi^{\alpha}_{\mathcal{M} \upharpoonright S}(\emptyset)$.

Proposition 36. For each bimodal formula ϕ of the modal μ -calculus, we construct a monomodal formula $\operatorname{tr}^{\Psi}(\phi)$; if ϕ belongs to $\mathcal{C}_{\aleph_1}(x)$, then so does $\operatorname{tr}^{\Psi}(\phi)$. Moreover, for each bimodal model \mathcal{M} we can also construct a monomodal model \mathcal{M}^{sim} , together with an injective function $(_)^{\circ} : |\mathcal{M}| \to |\mathcal{M}^{sim}|$ such that, for each $s \in |\mathcal{M}|, \mathcal{M}, s \Vdash \phi$ if and only if $M^{sim}, s^{\circ} \Vdash \operatorname{tr}^{\Psi}(\phi)$.

Proof. The modal formula $tr^{\Psi}(\phi)$ is the one inductively defined in this section, by either taking Ψ as the set of formulas from Example 43 or the one from the following Example 44.

APPENDIX B. PROOFS FROM SECTION 2

Proposition 3. For each $X \subseteq A$, X is κ -small if and only if, for every κ -directed set $\mathcal{I}, X \subseteq \bigcup \mathcal{I}$ implies $X \subseteq I$, for some $I \in \mathcal{I}$.

Proof. Let us firstly prove that if $\mathcal{I} \subseteq P(A)$ is a κ -directed set and $X \subseteq \bigcup \mathcal{I}$ is κ -small, then there exists $I \in \mathcal{I}$ with $X \subseteq I$. For each $a \in X$, let $I_a \in \mathcal{I}$ such that $a \in I_a$. Then $\mathcal{J} = \{I_a \mid a \in X\}$ is a subfamily of \mathcal{I} with card $\mathcal{J} < \kappa$, whence there exists $I \in \mathcal{I}$ with $I_a \subseteq I$, for each $a \in X$; whence $X \subseteq I$.

For the converse, recall that $X = \bigcup \mathcal{I}_{\kappa}(X)$ and that $\mathcal{I}_{\kappa}(X)$ is a κ -ideal. If, for every κ -directed set $\mathcal{I}, X \subseteq \bigcup \mathcal{I}$ implies $X \subseteq I$ for some $I \in \mathcal{I}$, then $X \subseteq \bigcup \mathcal{I}_{\kappa}(X)$ yields $X \subseteq X'$ for some κ -small $X' \subseteq X$. Since $X' \subseteq X$ and X' is κ -small, we obtain X' = X and X is κ -small. \Box

Proposition 8. Suppose that $\kappa \geq \aleph_0$ and let $f : P(B) \times P(A) \to P(B)$ be a κ -continuous monotone mapping. Then $\mu_x f(x, -) : P(A) \to P(B)$ is also κ -continuous.

Since the above statement holds for any κ -continuous $f : P \times Q \to P$ (that is, we do not need P and Q be power set Booelan algebras) we use the supremum symbol \bigvee in place of the set theoretic \bigcup .

Proof. We suppose that f is κ -continuous, that $\{X_i \mid i \in I\}$ is a κ -directed set of elements of P(A) and that $X = \bigvee_{i \in I} X_i$. Let us show that $\mu_x f(x, X) = \bigvee_{i \in I} \mu_x f(x, X_i)$.

Firstly, notice that the relation $\mu_x f(x, X) \ge \bigvee_{i \in I} \mu_x f(x, X_i)$ follows from monotonicity; thus we only need to prove the coverse relation and, to this end, it is enough to show that $\bigvee_{i \in I} \mu_x f(x, X_i)$ is a fixed-point of f(x, X). This goes as follows:

$$f(\bigvee_{i \in I} \mu_x . f(x, X_i), X) = \bigvee_{i \in I} f(\mu_x . f(x, X_i), X)$$

since f is κ -continuous

since f is κ -continuous in its first argument

$$= \bigvee_{i \in I} f(\mu_x.f(x, X_i), \bigvee_{j \in I} X_j)$$
$$= \bigvee_{i \in I, j \in I} f(\mu_x.f(x, X_i), X_j)$$

since f is κ -continuous in its second argument

$$= \bigvee_{i \in I} f(\mu_x.f(x, X_i), X_i) \quad \text{since } \{ X_i \mid i \in I \} \text{ is directed}$$
$$= \bigvee_{i \in I} \mu_x.f(x, X_i)$$
by the fixed point equation for $\mu_x.f(x, X_i).$

APPENDIX C. PROOFS FROM SECTION 4

Lemma 19. For every set X of variables and every $\phi \in L_{\mu}$, the following are equivalent:

- (1) $\phi \in \mathcal{C}_{\aleph_1}(X),$
- (2) no occurrence of a variable $x \in X$ is bad in ϕ .

Proof. (1) implies (2). Let $\phi \in L_{\mu}$ and suppose that $\phi \in \mathcal{C}_{\aleph_1}(X)$ for some set X of variables. The only way to introduce a bad path to an occurrence of some variable $x \in X$ is either by using some modal operator [a], which is excluded by the grammar, or by a fixed-point formation rule. In the latter case, a bad path to x could only be introduced in $Q_z \cdot \chi$, for $Q \in \{\mu, \nu\}$, through some edge from occurrence variable z to $Q_z \cdot \chi$. But such a free occurrence of z in χ would be bad and so $\chi \notin \mathcal{C}_{\aleph_1}(X \cup \{z\})$ contradicting the hypothesis.

(2) implies (1). Suppose there exist formulas without bad occurrences of any of the variables in X but which do not belong to $\mathcal{C}_{\aleph_1}(X)$, for some set X of variables. Among those formulas consider a formula ϕ of least complexity. Clearly, this formula has to be of the form $Q_z \cdot \chi$ and, moreover, there must be some occurrences of variables $x \in X$. Observe that χ has no bad occurrences of any $x \in X$, since such a bad occurrence would also be a bad occurrence of x in $Q_z \cdot \chi$. Also, if z is bad in χ , then any occurrence of some $x \in X$ is bad in $Q_z \cdot \chi$. Thus, χ has no bad occurrences of variables in $X \cup \{y\}$. Therefore, by the minimality assumption, χ belongs to $C(X \cup \{z\})$ and so $\phi \in \mathcal{C}_{\aleph_1}(X)$, against the assumptions.

Lemma 22. If ϕ is an almost-good formula, then both $\phi^{\sharp x}$ and ϕ^{\flat} belong to $\mathcal{C}_{\aleph_1}(x)$.

Proof. We prove the result for $\phi^{\sharp x}$. Consider a bad occurrence of x in $\phi^{\sharp x}$. After substituting \overline{x} for x, such an occurrence yields a bad occurrence of x in ϕ . Since there are no very bad occurrences of x in ϕ , then this occurrence should be not-sobad, that is, under the scope of a necessity modal operator [a], in ϕ as well as in $\phi^{\sharp x}$. Yet, there are no not-so-bad occurrences of x in $\phi^{\sharp x}$, as they have been replaced by occurrences of \overline{x} .

Proposition 23.I. The formula $\phi^{\Box X}$ is equivalent to ϕ .

Proof. The statement of the proposition is obvious if a formula matches the base case of the definition. Also, in the cases of a modal formula $\langle a \rangle \psi$ and of a formula $\psi_1 @ \psi_2$ with $@ \in \{ \land, \lor \}$, the statement is an immediate consequence of the inductive hypothesis. In case of a formula of the form $(Q_z.\psi)^{\Box X}$ with $Q \in \{ \mu, \nu \}$, we argue as follows:

$$(Q_z.\psi)^{\Box X} \equiv \psi_0[Q_{\overline{z}}.\psi_0/\overline{z}] \equiv Q_{\overline{z}}.\psi_0, \qquad \text{by the fixed-point equation,} = Q_{\overline{z}}.Q_z.\psi_2 \equiv Q_z.\psi_2[z/\overline{z}],$$

by the equational properties of fixed-points,

$$= Q_z.((\psi^{\Box X \cup \{z\}})^{\sharp z}[z/\overline{z}]) \equiv Q_z.(\psi^{\Box X \cup \{z\}}), \qquad \text{by equation (20)},$$
$$\equiv Q_z.\psi, \qquad \qquad \text{by the inductive hypothesis.} \quad \Box$$

Proposition 23.II. The formula $\phi^{\Box X}$ is almost-good, that is, it has no very bad occurrence of a variable $x \in X$.

Proof. The statement of the proposition is obvious if a formula matches the base case of the definition. Also, in the cases of a modal formula $\langle a \rangle \psi$ and of a formula $\psi_1 @ \psi_2$ with $@ \in \{ \land, \lor \}$, the statement is an immediate consequence of the inductive hypothesis. The only non-trivial case is that of a formula of the form $(Q_z, \psi)^{\Box X}$ with $Q \in \{ \mu, \nu \}$.

Let us firstly recall that $(Q_z.\psi)^{\Box X}$ is of the form $\psi_0[Q_{\overline{z}}.\psi_0/\overline{z}]$ with $\psi_0 = Q_z.\psi_2$ and $\psi_2 = (\psi^{\Box X \cup \{z\}})^{\sharp z}$. Also, for the sake of readability, we have let in the definition $\psi_1 := Q_{\overline{z}}.\psi_0$, so $(Q_z.\psi)^{\Box X} = \psi_0[\psi_1/\overline{z}]$. In particular, every occurrence of a variable is located within ψ_0 , or it is located in some subtree of $\psi_0[\psi_1/\overline{z}]$ rooted at some occurrence of the subformula ψ_1 .

We argue next that every occurrence of a variable $x \in X$ within $\psi_0 = Q_z \cdot \psi_2$ is not very bad. By the induction hypothesis, such an occurrence of x is not very bad within ψ_2 ; the only reason for becoming very bad in ψ_0 is then the existence of a cycle going through an edge from some occurrence of the variable z to the formula $Q_z \cdot \psi_2$. Such a bad cycle can arise for two reasons: either (a) there is a necessity modal operator [a] from ψ_2 to this occurrence of z, or (b) there is a bad cycle in some subformula of ψ_2 of the form $Q_w \cdot \chi$, with this subformula lying on the path from ψ_2 to the occurrence of z. Yet (a) is not possible: recall that $\psi_2 = (\psi^{\Box X \cup \{z\}})^{\sharp z}$, thus all the occurrences of z within ψ_2 are not boxed (such an occurrence in $\psi^{\Box X \cup \{z\}}$ has been renamed to \overline{z} in ψ_2). Also (b) is not possible, since otherwise the occurrence of z in ψ_2 is very bad. Yet we know that the same occurrence of z is not very bad in $\psi^{\Box X \cup \{z\}}$, and renaming the boxed occurrences of z to \overline{z} in this formula cannot transform another occurrence of z into a very bad occurrence. Finally, we argue that there is no very bad occurrence of some variable $x \in X$ in $\psi_0[\psi_1/\overline{z}]$. Suppose there is such an occurrence of x. If this occurrence is located within ψ_0 , then this would also be a bad occurrence for ψ_0 , which we have excluded. Thus, such an occurrence is located within some occurrence of the subformula ψ_1 . But since every occurrence of the variable \overline{z} within ψ_0 is boxed, all the variable occurences of x within ψ_1 become boxed in the formula $\psi_0[\psi_1/\overline{z}]$.

Therefore, no occurrence of $x \in X$ is very bad in $\psi_0[\psi_1/\overline{z}]$.

A formula ϕ is well-named if no bound variable of ϕ is also free in ϕ and, for each bound variable z of ϕ , there is a unique subformula ψ of ϕ of the form $Q_z.\psi'$, with $Q \in \{\mu, \nu\}$. It is well known that every formula $\phi \in \mathsf{L}_{\mu}$ is equivalent to a well-named formula. Therefore, we have tacitly assumed that all the formulas are well-named.

Let ϕ be a well-named formula; recall that we use $Sub(\phi)$ for the set of subformulas of ϕ .

Remark 47. If \overline{x} and y are distinct variables and χ is a formula that does not contain the variable y, then

$$[\psi/y] \cdot [\chi/\overline{x}] = [\chi/\overline{x}] \cdot [\psi[\chi/\overline{x}]/y].$$
(8)

Also, if \overline{x} is variable occurring free in ϕ and κ is either a variable or a constant, then $Sub(\phi \cdot [\kappa/\overline{x}]) = \{ \psi \cdot [\kappa/\overline{x}] \mid \psi \in Sub(\phi) \}.$

Lemma 25. If \overline{x} is a free variable of ϕ and κ is either a variable not bound in ϕ or a constant, then

$$\operatorname{CL}(\phi \cdot [\kappa/\overline{x}]) = \{ \psi \cdot [\kappa/\overline{x}] \mid \psi \in \operatorname{CL}(\phi) \}$$

Proof. By repeatedly using equation (8) with $\chi = \kappa$, we have

$$\sigma_{\psi}^{\varphi} \cdot [\kappa/\overline{x}] = [Q_n y_n \cdot \psi_n / y_n] \cdot \dots \cdot [Q_1 y_1 \cdot \psi_1 / y_1] \cdot [\kappa/\overline{x}]$$

= $[\kappa/\overline{x}] \cdot [Q_n y_n \cdot \psi_n \cdot [\kappa/\overline{x}] / y_n] \cdot \dots \cdot [Q_1 y_1 \cdot \psi_1 \cdot [\kappa/\overline{x}] / y_1].$

Inspection of the three properties defining the standard context σ_ψ^ϕ shows that the equality

$$\sigma_{\psi\cdot[\kappa/\overline{x}]}^{\phi\cdot[\kappa/\overline{x}]} = [Q_n y_n \cdot \psi_n \cdot [\kappa/\overline{x}]/y_n] \cdot \dots \cdot [Q_1 y_1 \cdot \psi_1 \cdot [\kappa/\overline{x}]/y_1]$$

holds. From this we deduce

$$(\psi \cdot [\kappa/\overline{x}]) \cdot \sigma_{\psi \cdot [\kappa/\overline{x}]}^{\phi \cdot [\kappa/\overline{x}]} = (\psi \cdot \sigma_{\psi}^{\phi}) \cdot [\kappa/\overline{x}].$$
(9)

Thus
$$\phi' \in \operatorname{CL}(\phi \cdot [\kappa/\overline{x}])$$
 iff $\phi' = \psi \cdot \sigma_{\psi}^{\phi \cdot [\kappa/\overline{x}]}$ for some $\psi \in \operatorname{Sub}(\phi \cdot [\kappa/\overline{x}])$
iff $\phi' = \psi \cdot [\kappa/\overline{x}] \cdot \sigma_{\psi \cdot [\kappa/\overline{x}]}^{\phi \cdot [\kappa/\overline{x}]}$ for some $\psi \in \operatorname{Sub}(\phi)$
iff $\phi' = \psi \cdot \sigma_{\psi}^{\phi} \cdot [\kappa/\overline{x}]$ for some $\psi \in \operatorname{Sub}(\phi)$
iff $\phi' = \phi'' \cdot [\kappa/\overline{x}]$ for some $\phi'' \in \operatorname{CL}(\phi)$.

Lemma 48. The positions p and p' are respectively of the form $(s, [a]\psi) \in \mathcal{G}(\mathcal{N}, \phi)$ and $(s, [a]\psi') \in \mathcal{G}(\mathcal{N}, \phi^{\flat})$ for some ψ and ψ' such that $\psi[\perp/x] \to \psi'$ is a tautology.

Proof. We let $\psi = \chi \sigma_{\chi}^{\phi^{\sharp x}} \cdot [x/\overline{x}]$ and observe that

$$\psi = \chi \sigma_{\chi}^{\phi^{\sharp x}} \cdot [x/\overline{x}] = \chi'[\overline{x}/x] \cdot \sigma_{\chi'[\overline{x}/x]}^{\phi^{\sharp x}} \cdot [x/\overline{x}]$$
$$= \chi'[\overline{x}/x] \cdot [x/\overline{x}] \cdot \sigma_{\chi'[\overline{x}/x] \cdot [x/\overline{x}]}^{\phi^{\sharp x} \cdot [x/\overline{x}]}, = \chi' \cdot \sigma_{\chi'}^{\phi}.$$
by equation (9)

On the other hand, we let $\psi' = \chi \sigma_{\chi}^{\phi^{\sharp x}} \cdot [\perp/\overline{x}]$, so that

$$\begin{split} \psi' &= \chi \sigma_{\chi}^{\phi^{\sharp x}} \cdot [\perp/\overline{x}] = \chi'[\overline{x}/x] \cdot \sigma_{\chi'[\overline{x}/x]}^{\phi^{\sharp x}} \cdot [\perp/\overline{x}] = \chi'[\overline{x}/x] \cdot [\perp/\overline{x}] \cdot \sigma_{\chi'[\overline{x}/x] \cdot [\perp/\overline{x}]}^{\phi^{\sharp x}} \\ &= \chi'[\perp/x] \cdot \sigma_{\chi'[\perp/x]}^{\phi^{\sharp x} \cdot [\perp/\overline{x}]}, \qquad \text{since } \chi' \text{ does not contain the variable } \overline{x}, \\ &\geq \chi'[\perp/x] \cdot \sigma_{\chi'[\perp/x]}^{\phi^{\sharp x} \cdot [\perp/\overline{x}, \perp/x]}, \qquad \text{since } [\perp/\overline{x}, \perp/x] \text{ and } \phi^{\sharp x} \text{ is monotone in } x \text{ and } \overline{x}, \\ &= \chi'[\perp/x] \cdot \sigma_{\chi'[\perp/x]}^{\phi^{\sharp x} \cdot [x/\overline{x}] \cdot [\perp/x]} = \chi'[\perp/x] \cdot \sigma_{\chi'[\perp/x]}^{\phi^{[\perp/x]}} \\ &= \chi' \cdot \sigma_{\chi'}^{\phi} \cdot [\perp/x] = \psi[\perp/x], \qquad \text{by the previous computations,} \end{split}$$

where for $\phi, \phi' \in \mathsf{L}_{\mu}$, the notation $\phi \geq \phi'$ means that $\llbracket \phi \rrbracket_{\mathcal{M}} \supseteq \llbracket \phi' \rrbracket_{\mathcal{M}}$ for every \mathcal{M} .

Appendix D. Proofs from Section 5

Lemma 33. Let α be a closure ordinal. Among the formulas that that have α as its closure ordinal there exists one formula $\phi(x)$ such that $\mu_x . \phi(x)$ is total in some model \mathcal{M} where the convergence occurs in exactly α steps, that is,

$$|\mathcal{M}| = \llbracket \mu_x . \phi(x) \rrbracket_{\mathcal{M}} = \phi^{\alpha}_{\mathcal{M}}(\emptyset) \neq \phi^{\alpha'}_{\mathcal{M}}(\emptyset), \qquad \text{for every } \alpha' < \alpha$$

Proof. For a fomula $\psi(x)$, let

$$\phi(x) := \mu_x . \psi(x) \to \psi(x \land \mu_x . \psi(x))$$

Let us verify that, for each ordinal $\gamma \geq 1$, we have

$$\phi^{\gamma}_{\mathcal{M}}(\emptyset) = \mu.\psi_{\mathcal{M}} \to \psi^{\gamma}_{\mathcal{M}}(\emptyset)$$

This is clear for $\alpha = 1$. Assuming it holds for γ , then

$$\phi_{\mathcal{M}}^{\gamma+1}(\emptyset) = \mu.\psi_{\mathcal{M}} \to \psi_{\mathcal{M}}(\mu.\psi_{\mathcal{M}} \to \psi_{\mathcal{M}}^{\gamma}(\emptyset) \cap \mu.\psi_{\mathcal{M}})$$
$$= \mu.\psi_{\mathcal{M}} \to \psi_{\mathcal{M}}(\psi_{\mathcal{M}}^{\gamma}(\emptyset) \cap \mu.\psi_{\mathcal{M}})$$
$$= \mu.\psi_{\mathcal{M}} \to \psi_{\mathcal{M}}(\psi_{\mathcal{M}}^{\gamma}(\emptyset)) = \mu.\psi_{\mathcal{M}} \to \psi_{\mathcal{M}}^{\gamma+1}(\emptyset)$$

The inductive step is obvious. Clearly, $\mu . \phi_{\mathcal{M}}$ is always the total set $|\mathcal{M}|$.

Lemma 49. For each countable ordinal α , we have

$$\phi_{\mathcal{M}_{\omega_1}}(S_\alpha) = S_{\alpha+1}.$$

Consequently, for each ordinal $\alpha \leq \omega_1$, we have $\phi^{\alpha}_{\mathcal{M}_{\omega_1}}(\emptyset) = S_{\alpha}$.

Proof. If $\alpha = 0$, then $S_{\alpha} = \emptyset$ and

$$\begin{split} \phi_{\mathcal{M}_{\omega_1}}(S_0) &= \phi_{\mathcal{M}_{\omega_1}}(\emptyset) = \llbracket \nu_z . (\langle h \rangle z \land \langle v \rangle x) \lor [v] \bot \rrbracket_{\mathcal{M}_{\omega_1}[\emptyset/x]} \\ &= \llbracket [v] \bot \rrbracket_{\mathcal{M}_{\omega_1}} = \{ (0,n) \mid n < \omega \} = S_1 \,. \end{split}$$

Consider now an ordinal $\alpha > 0$.

24

Let us argue firstly that $S_{\alpha+1} \subseteq \phi_{\mathcal{M}_{\omega_1}}(S_{\alpha})$. Let $(\beta, n) \in S_{\alpha+1}$, so $\beta < \alpha + 1$ implies $\beta \leq \alpha$. From (β, n) , there is the infinite horizontal path $\{(\beta, m) \mid n \leq m < \omega\}$ and each vertex on this path has a vertical transition to a vertex $(\beta', 0)$ with $\beta' < \beta \leq \alpha$, in particular $(\beta', 0) \in S_{\alpha}$. So $(\beta, n) \in \phi_{\mathcal{M}_{\omega_1}}(S_{\alpha})$.

Next, we argue that the converse inclusion, $\phi_{\mathcal{M}_{\omega_1}}(S_\alpha) \subseteq S_{\alpha+1}$, holds. Suppose $(\beta, n) \in \phi(S_\alpha)$. If there are no vertical transitions from (β, n) then $\beta = 0$ and $(\beta, n) \in S_1 \subseteq S_{\alpha+1}$, since $S_\beta \subseteq S_\gamma$ for $\beta \leq \gamma$. Otherwise $\beta > 0$, there is an infinite horizontal path from (β, n) and each vertex on this path has a transition to some vertex in S_α . Notice that such an infinite horizontal path is, necessarily, the path $\pi := \{ (\beta, m) \mid n \leq m < \omega \}.$

If $\beta = \gamma + 1$ is a successor ordinal then the unique outgoing vertical transition from (β, n) is to $(\gamma, 0)$. Hence $(\gamma, 0) \in S_{\alpha}$, thus $\gamma < \alpha, \beta = \gamma + 1 < \alpha + 1$ and $(\beta, n) \in S_{\alpha+1}$. Otherwise β is a limit ordinal distinct from 0 and, for each $m \ge n$, there is a vertical transition $(\beta, m)R_v(\beta_m, 0)$ with $(\beta_m, 0) \in S_{\alpha}$, so $\beta_m < \alpha$. If $\alpha + 1 \le \beta$, then $\alpha < \beta$, that is, $\alpha \in \beta$. Since the ω -chain $\{\beta_k \mid k \in \omega\}$ is cofinal in β , we can find $k \in \omega$ such that $\alpha \le \beta_k$. Since $\beta_k \le \beta_{k'}$ for $k \le k' \in \omega$, we can also suppose that $n \le k$. But we obtain here a contradiction, since we mentioned before that $\beta_m < \alpha$ for $m \ge n$, in particular $\beta_k < \alpha$.

The proof of the second statement is now a straightforward induction on α . If $\alpha = \beta + 1$ is a successor ordinal, then

$$\phi^{\alpha}_{\mathcal{M}_{\omega_1}}(\emptyset) = \phi_{\mathcal{M}_{\omega_1}}(\phi^{\beta}_{\mathcal{M}_{\omega_1}}(\emptyset)) = \phi_{\mathcal{M}_{\omega_1}}(S_{\beta}) = S_{\beta+1}.$$

If α is a limit ordinal, then

$$\phi^{\alpha}_{\mathcal{M}_{\omega_1}}(\emptyset) = \bigcup_{\beta < \alpha} \phi^{\beta}_{\mathcal{M}_{\omega_1}}(\emptyset) = \bigcup_{\beta < \alpha} S_{\beta} = S_{\alpha} \,. \qquad \Box$$

Lemma 50. An ordinal γ is the closure ordinal of the formula $\Psi(x)$ if and only if (i) the formula $\psi(x)$ converges to its least fixed point in at most γ steps on all the acceptable models, and (ii) there exists an acceptable model on which the formula $\psi(x)$ converges to its least fixed point in exactly γ steps.

Proof. If \mathcal{N} is an acceptable model, then $\llbracket [\mathcal{U}]\chi \rrbracket_{\mathcal{N}} = |\mathcal{N}|$, so that $\Psi_{\mathcal{N}} = \psi_{\mathcal{N}}$.

On the other hand, if \mathcal{M} is any model, then the submodel of \mathcal{M} induced by $[[\mathcal{U}]\chi]_{\mathcal{M}}$ —call it \mathcal{N} —is closed. Thus, by Proposition 11, for any ordinal $\gamma \geq 0$, we have

$$\Psi_{\mathcal{M}}^{\gamma}(\emptyset) = \psi_{\mathcal{N}}^{\gamma}(\emptyset) \,. \tag{10}$$

The statement of the lemma immediately follows.

Lemma 51. The following relation holds for every ordinal $\gamma \geq 0$:

$$\phi_{N_1}^{\gamma}(\emptyset) \subseteq \psi_{N}^{\alpha+\gamma}(\emptyset) \cap N_1.$$
(5)

Proof. Clearly the relation holds for $\gamma = 0$. In order to prove the above inclusion, it will be enough to prove that it holds at a successor ordinal $\gamma + 1$, assuming it

holds at γ . We have

$$\begin{split} \psi_{\mathcal{N}}^{\alpha+\gamma+1}(\emptyset) \cap N_{1} &= \phi_{\mathcal{N}_{1}} \left(\psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_{1} \right) \cap \nabla(\psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_{0}) \\ &= \phi_{\mathcal{N}_{1}} \left(\psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_{1} \right) \cap \nabla(\psi_{\mathcal{N}}^{\alpha}(\emptyset) \cap N_{0}), \\ \text{since by equation (4) } \psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_{0} &= \phi_{\mathcal{N}_{0}}^{\alpha}(\emptyset) = \psi_{\mathcal{N}}^{\alpha}(\emptyset) \cap N_{0} \\ &= \phi_{\mathcal{N}_{1}} \left(\psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_{1} \right), \qquad \text{by equation (3),} \\ &\supseteq \phi_{\mathcal{N}_{1}} \left(\phi_{\mathcal{N}_{1}}^{\gamma}(\emptyset) \right) = \phi_{\mathcal{N}_{1}}^{\gamma+1}(\emptyset), \qquad \text{by the IH.} \quad \Box \end{split}$$

Lemma 52. Suppose that $\phi_{\mathcal{N}_0}^{\delta}(\emptyset)$ is strictly included in N_0 for $\delta < \alpha$ and that $\nabla(X) = \emptyset$ whenever X is a proper subset of N_0 . Then, the inclusion (5) is an equality, for each ordinal $\gamma \geq 0$:

$$\phi_{\mathcal{N}_1}^{\gamma}(\emptyset) = \psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_1$$

Proof. It is enough to verify that the equality holds at $\gamma = 0$. Then all the inductive computations are as in the proof of the previous Lemma. If $\delta < \alpha$, we have

$$\psi_{\mathcal{N}}^{\delta+1}(\emptyset) \cap N_1 \subseteq \nabla(\psi_{\mathcal{N}}^{\delta}(\emptyset) \cap N_0) = \nabla(\phi_{\mathcal{N}_0}^{\delta}(\emptyset)) = \emptyset,$$

since by assumption $\phi_{N_0}^{\delta}(\emptyset)$ is strictly included in N_0 . In particular, if α is a successor ordinal, we have $\psi_{N}^{\alpha+\gamma}(\emptyset) \cap N_1 = \emptyset$. If α is a limit ordinal, then

$$\psi^{\alpha}_{\mathcal{N}}(\emptyset) \cap N_1 \subseteq \bigcup_{\delta < \alpha} \psi^{\delta+1}_{\mathcal{N}}(\emptyset) \cap N_1 \subseteq \bigcup_{\delta < \alpha} \nabla(\psi^{\delta}_{\mathcal{N}}(\emptyset) \cap N_0) = \emptyset. \qquad \Box$$