

\aleph_1 AND THE MODAL μ -CALCULUS

In memory of Zoltán Ésik

MARIA JOÃO GOUVEIA AND LUIGI SANTOCANALE

Faculdade de Ciências, Universidade de Lisboa, Portugal
e-mail address: mjgouveia@fc.ul.pt

LIS, CNRS UMR 7020, Aix-Marseille Université, France
e-mail address: luigi.santocanale@lis-lab.fr

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 κ -directed sets. We define the fragment $\mathcal{C}_{\aleph_1}(x)$ of the modal μ -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 [21, 27] 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 [17]. Despite its strong expressive power, this logic is still considered as a tractable one. Its model checking problem, known to be in the class $UP \cap co-UP$ [20], has

Key words and phrases: Modal μ -calculus; regular cardinal; continuous function; \aleph_1 ; ω_1 ; closure ordinal; ordinal sum.

Partially supported by FCT under grant SFRH/BSAB/128039/2016.

recently been proved to be quasi-polynomial and fixed-parameter tractable [10]. Moreover, this problem becomes polynomial if some restricted classes of models are considered [31, 3, 7]. The widespread interest for this logic has triggered further researches that spread beyond the realm of verification: these concern the expressive power [8, 6], axiomatic bases [40], algebraic and order theoretic approaches [36], deductive systems [30, 37], and the semantics of functional programs [16].

The present paper lies at the intersection of two lines of research on the modal μ -calculus, on continuity [14] and on closure ordinals [11, 2]. Continuity of monotone functions is a fundamental phenomenon in modal logic, on which well-known uniform completeness theorems rely [32, 33, 18]. Fontaine [14] 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 [35] 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 $\phi(x) \in \mathcal{L}_\mu$, 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 [14], 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 $\mathcal{C}_{\aleph_1}(x)$. In [15], an extended journal version of the conference paper [14], the fragment $\mathcal{C}_{\aleph_1}(x)$ is also studied, yet the semantic property pinpointed there and corresponding to the syntactic fragment $\mathcal{C}_{\aleph_1}(x)$ is different from the property that we consider, κ -continuity. Say that a formula of the modal μ -calculus has the *finite width property* if, whenever it is satisfied in a tree model, it is satisfied in a finitely branching subtree of this model. It is proved in [15] that a formula has the finite width property if and only if it is equivalent to a formula in $\mathcal{C}_{\aleph_1}(x)$. Combining these results with ours, we deduce a quite surprising statement: a formula has the finite width property if and only if it is κ -continuous for some regular cardinal κ . While it is easy to guess why the finite width property implies \aleph_1 -continuity, the converse implication appears to be a non-obvious strong statement, whose potential consequences and applications need to be uncovered.

Our interest in \aleph_1 -continuity was wakened once more when researchers started investigating closure ordinals of formulas of the modal μ -calculus [11, 2]. The notion of closure ordinal was studied in the context of first order inductive definitions [29]. Closure ordinals for the modal μ -calculus are more directly related to *global* inductive definability, see [5], in that a class of structures is being tested, not a single structure. 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, in our opinion, naturally intertwined and the results we present in this paper are in support of this thesis. 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 $\text{cl}(\phi(x)) \leq \kappa$ —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 show 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 $\text{Ord}(\mathbf{L}_\mu)$* of all closure ordinals of formulas of the propositional modal μ -calculus. Secondly, we prove that *$\text{Ord}(\mathbf{L}_\mu)$ 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 $\text{Ord}(\mathbf{L}_\mu)$. Let us recall that Czarnecki [11] proved that all the ordinals $\alpha < \omega^2$ belong to $\text{Ord}(\mathbf{L}_\mu)$. 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, even considering that our work does not yield methods to exclude ordinals from $\text{Ord}(\mathbf{L}_\mu)$, 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 $\text{Ord}(\mathbf{L}_\mu)$ 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. The fragment $\mathcal{C}_{\aleph_1}(x)$ of the propositional modal μ -calculus has imposed itself by its robustness, which can be recognised in our work as well as in [15]. We believe $\mathcal{C}_{\aleph_1}(x)$ is worth investigating further in order to enlighten a hidden dimension (and thus new tools, new ideas, new perspectives, etc.) of the modal μ -calculus and of fixed-point logics. As an example, take the modal μ -calculus on deterministic models: states have at most one successor and it is immediate to conclude 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 research. We also believe that the scope of this work, as well as of the problems studied within, goes much beyond the pure theory of the modal μ -calculus. For example, our interest in closure ordinals stems from a previous proof-theoretic investigation of induction and coinduction [16, 35]. In these works ordinal notations are banned from the syntax because of an alleged non-constructiveness of the set theory needed to represent ordinals. However, also considering that elegant constructive theories of ordinals exist, see e.g. [19], the present work encourages us to develop alternative proof-theoretic frameworks based on ordinals.

The paper is structured as follows. In Section 2 we introduce the notion of κ -continuity. In the following Section 3 we illustrate the interactions between κ -continuity and least/greatest fixed-points of monotone maps. In Section 4 we present the modal μ -calculus and some of the related theory that we shall need in the following sections. Section 5 presents our results on the fragment $\mathcal{C}_{\aleph_1}(x)$. The following Section 6 presents a tool—roughly speaking the observation that various kind of submodels can be logically described modulo the introduction of a new propositional variable—that is repeatedly used in the rest of the paper to obtain results on closure ordinals. In Section 7 we argue that the least uncountable ordinal is a closure ordinal for the modal μ -calculus. In the final Section 8 we argue that $\text{Ord}(\mathbf{L}_\mu)$, the set of closure ordinal of formulas of the modal μ -calculus, is closed under ordinal sum.

Acknowledgment. The authors are thankful to the anonymous reviewers for their analytical reading and for their valuable comments by which the presentation of this research could substantially improve.

CONTENTS

1. Introduction	1
2. κ -continuous maps	4
3. Fixed-points of κ -continuous maps	6
4. The propositional modal μ -calculus	9
4.1. The closure of a formula	10
4.2. Game semantics	11
4.3. Bisimulations	11
5. \aleph_1 -continuous fragment of the modal μ -calculus	12
5.1. Syntactic considerations	12
5.2. The $\mathcal{C}_{\aleph_1}(x)$ -flattening of formulas	14
5.3. Comparing the closures of ϕ and ϕ^b	17
5.4. The continuous fragments	18
6. On p -definability	20
7. An uncountable closure ordinal	26
7.1. ω_1 is a closure ordinal	27
8. Closure under ordinal sum.	29
References	33

2. κ -CONTINUOUS MAPS

In this section we consider κ -continuity of monotone maps 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 found for example in [14, 15]. The reader might find further information in the monograph [1] where this notion is presented in the more general context of categories.

In the following κ is an infinite regular cardinal, A and B are sets, for which $P(A)$ and $P(B)$ denote the corresponding powerset Boolean algebras, and $f : P(A) \rightarrow P(B)$ is a monotone map. We shall say that a subset X of a set A is κ -small if $\text{card } X < \kappa$. 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. Regularity of the cardinal κ essentially amounts to the following property: if \mathcal{J} is a κ -small collection of κ -small subsets of A , then $\bigcup \mathcal{J}$ is κ -small.

Definition 2.1. A subset $\mathcal{I} \subseteq P(A)$ is a κ -directed set if every collection $\mathcal{J} \subseteq \mathcal{I}$ with $\text{card } \mathcal{J} < \kappa$ has an upper bound in \mathcal{I} . A map $f : P(A) \rightarrow P(B)$ is κ -continuous if $f(\bigcup \mathcal{I}) = \bigcup f(\mathcal{I})$, whenever $\mathcal{I} \subseteq P(A)$ is a κ -directed set.

Observe that if κ' is a regular cardinal and $\kappa < \kappa'$, then a κ' -directed set is also a κ -directed set. Therefore, if f is κ -continuous, then it also preserves unions of κ' -directed sets, thus it is κ' -continuous as well. Also, notice that the wording “monotone κ -continuous” is redundant: if f is κ -continuous, then it is monotone, since if $X \subseteq Y$, then $\{X, Y\}$ is κ -directed, so $f(Y) = f(Y \cup X) = f(X) \cup f(Y)$, so $f(X) \subseteq f(Y)$.

For each subset X of A , define

$$\mathcal{I}_\kappa(X) := \{X' \mid X' \subseteq X, X' \text{ is } \kappa\text{-small}\}.$$

Notice that $\bigcup \mathcal{I}_\kappa(X) = X$ and $\mathcal{I}_\kappa(X)$ is a κ -directed set. For this latter property, it is useful to note that if $\{X_i \subseteq X \mid i \in I\}$ is a κ -small set of κ -small subsets of X , then the union $\bigcup \{X_i \mid i \in I\}$ is still κ -small, so it belongs to $\mathcal{I}_\kappa(X)$.

Proposition 2.2. *A subset X of A 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. We firstly prove that if X is κ -small and $\mathcal{I} \subseteq P(A)$ is a κ -directed set such that $X \subseteq \bigcup \mathcal{I}$, 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 $\text{card } \mathcal{J} < \kappa$, 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 -directed set. Suppose therefore that, for every κ -directed set \mathcal{I} , $X \subseteq \bigcup \mathcal{I}$ implies $X \subseteq I$ for some $I \in \mathcal{I}$. Applying this property when $\mathcal{I} = \mathcal{I}_\kappa(X)$ yields $X \subseteq X'$ for some κ -small $X' \subseteq X$. Therefore $X' = X$ and X is κ -small. \square

Proposition 2.3. *A monotone map $f : P(A) \rightarrow P(B)$ 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 $f : P(A) \rightarrow P(B)$ be a κ -continuous monotone map. Notice that the equation above is $f(\bigcup \mathcal{I}_\kappa(X)) = \bigcup f(\mathcal{I}_\kappa(X))$, since $X = \bigcup \mathcal{I}_\kappa(X)$. The equation holds since $\mathcal{I}_\kappa(X)$ is κ -directed and we are supposing that f is κ -continuous.

Conversely suppose that $f : P(A) \rightarrow P(B)$ is a monotone map such that $f(X) = \bigcup f(\mathcal{I}_\kappa(X))$ for every $X \in P(A)$. Also let $\mathcal{I} \subseteq P(A)$ be a κ -directed set, so we aim to show that $f(\bigcup \mathcal{I}) = \bigcup f(\mathcal{I})$. Since f is κ -continuous, $f(\bigcup \mathcal{I}) = \bigcup f(\mathcal{I}_\kappa(\bigcup \mathcal{I}))$. Since f is monotone, we have $\bigcup f(\mathcal{I}) \subseteq f(\bigcup \mathcal{I})$ and therefore we only need to verify the opposite inclusion. Let Y be a κ -small set contained in $\bigcup \mathcal{I}$. By Proposition 2.2 there exists $Z \in \mathcal{I}$ such that $Y \subseteq Z$. Hence for every $Y \in \mathcal{I}_\kappa(\bigcup \mathcal{I})$ there exists $Z \in \mathcal{I}$ such that $Y \subseteq Z$ and so also $f(Y) \subseteq f(Z)$. Thus, $\bigcup f(\mathcal{I}_\kappa(\bigcup \mathcal{I})) \subseteq \bigcup f(\mathcal{I})$. Consequently, we have

$$f(\bigcup \mathcal{I}) = \bigcup f(\mathcal{I}_\kappa(\bigcup \mathcal{I})) \subseteq \bigcup f(\mathcal{I}),$$

proving the opposite inclusion. \square

Next we extend the notion of κ -continuity to functions of many variables, that is, to functions whose domain is a finite product of the form $P(A_1) \times \dots \times P(A_n)$, the ordering being coordinate-wise. To achieve this goal, we observe that there is a standard isomorphism $\psi : P(A_1 \cup \dots \cup A_n) \rightarrow P(A_1) \times \dots \times P(A_n)$, where \cup denotes the disjoint union. Therefore, we say that a monotone function $f : P(A_1) \times \dots \times P(A_n) \rightarrow P(B)$ is κ -continuous if the function of one variable $f \circ \psi : P(A_1 \cup \dots \cup A_n) \rightarrow P(B)$ is κ -continuous. The standard isomorphism associates to a subset $S \subseteq A_1 \cup \dots \cup A_n$ the tuple $\psi(S) = \langle S \cap A_1, \dots, S \cap A_n \rangle$.

The next Lemma (that, for simplicity, we state and prove for $n = 2$) states the expected property of κ -continuous functions of many variables: these functions are κ -continuous exactly when they are κ -continuous in each variable.

Lemma 2.4. *A monotone map $f : P(A_1) \times P(A_2) \rightarrow P(B)$ is κ -continuous w.r.t. the coordinate-wise order on $P(A_1) \times P(A_2)$ if and only if it is κ -continuous in every variable.*

Proof. Obviously if $f \circ \psi : P(A_1 \cup A_2) \rightarrow P(B)$ is κ -continuous, then it is κ -continuous when we fix a subset, say $X \subseteq A_1$. Indeed, a family of the form $\{X \cup Y_i \mid i \in I, Y_i \subseteq A_2\}$ is κ -directed if and only if $\{Y_i \subseteq A_2 \mid i \in I\}$ is κ -directed.

Conversely, suppose that $f \circ \psi : P(A_1 \cup A_2) \rightarrow P(B)$ is κ -continuous in every variable. First observe that, for any families $\mathcal{X} = \{X_i \subseteq A_1 \mid i \in I\}$ and $\mathcal{Y} = \{Y_i \subseteq A_2 \mid i \in I\}$, we have that

$$\bigcup_i \{X_i \cup Y_i \mid i \in I\} = \bigcup_{i,j} \{X_i \cup Y_j \mid i, j \in I\} = \bigcup_i (X_i \cup \bigcup_j Y_j)$$

and when $\{X_i \cup Y_i \mid i \in I\}$ is κ -directed also \mathcal{X} , \mathcal{Y} and $\{X_i \cup Y_j \mid i, j \in I\}$ are κ -directed. Consequently, given a κ -directed set $\{X_i \cup Y_i \mid i \in I\}$ with $X_i \subseteq A_1$ and $Y_i \subseteq A_2$, the following holds

$$\begin{aligned} (f \circ \psi)\left(\bigcup_i X_i \cup Y_i\right) &= (f \circ \psi)\left(\bigcup_{i,j} X_i \cup Y_j\right) = (f \circ \psi)\left(\bigcup_i (X_i \cup \bigcup_j Y_j)\right) \\ &= \bigcup_i (f \circ \psi)(X_i \cup \bigcup_j Y_j) = \bigcup_i \bigcup_j (f \circ \psi)(X_i \cup Y_j), \\ &\qquad\qquad\qquad \text{since } f \text{ is } \kappa\text{-continuous in each variable,} \\ &= \bigcup_{i,j} (f \circ \psi)(X_i \cup Y_j) = \bigcup_i (f \circ \psi)(X_i \cup Y_i). \end{aligned}$$

This concludes the proof of Lemma 2.4. □

3. FIXED-POINTS OF κ -CONTINUOUS MAPS

The interplay between κ -continuity of monotone maps (recall that κ is assumed to be an infinite regular cardinal) and their least and greatest fixed-points is the focus of the present section. On the one hand, the Knaster-Tarski theorem [38] states that the least fixed-point of a monotone map $f : P(A) \rightarrow P(A)$ is the set $\bigcap\{X \subseteq A \mid f(X) \subseteq X\}$. On the other hand, Kleene's fixed-point theorem states that the least fixed-point of an \aleph_0 -continuous map f is constructible by iterating ω -times f starting from the empty set, namely it is equal to $\bigcup_{n \geq 0} f^n(\emptyset)$. Generalisations of Kleene's theorem appeared later and give ways to build the least fixed-point of monotone maps by ordinal approximations; see [26] for an historical account of this family of theorems.

The first result we present in this section is a generalised Kleene's fixed-point theorem specifically suited to κ -continuous maps (we do not claim the authorship of Proposition 3.2, even if we could not find it stated as it is in the literature).

Definition 3.1. Let $f : P(A) \rightarrow P(A)$ be a monotone map. The *approximants* $f^\alpha(\emptyset)$, with α an ordinal, are inductively defined as follows:

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

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 3.2. *If $f : P(A) \rightarrow P(A)$ is a κ -continuous monotone map, then f converges to its least fixed-point in at most κ steps.*

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

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

since the regularity of κ implies that $\{f^\alpha(\emptyset) \mid \alpha < \kappa\}$ is a κ -directed set. Since the inclusion $f^\kappa(\emptyset) \subseteq f(f^\kappa(\emptyset))$ holds by monotonicity of f , $f^\kappa(\emptyset)$ is also a fixed-point of f . \square

Until now we have focused on least fixed-points of monotone maps. Greatest fixed-points are dual to least fixed-points: namely, for a monotone map $f : P(A) \rightarrow P(A)$, its greatest fixed-point is the largest subset Z of A such that $f(Z) = Z$; by Tarski's theorem, it is equal to $\bigcup\{Z \subseteq A \mid Z \subseteq f(Z)\}$. Propositions 3.3 and 3.4 relate both kind of (parametrized) fixed points to continuity; they are specific instances of a result stated for categories in [35]. To clarify their statements, let us recall that if $f : P(B) \times P(A) \rightarrow P(B)$ is a monotone map, then, for each $X \in P(A)$, the unary map $f(-, X) : P(B) \rightarrow P(B)$, $Z \mapsto f(Z, X)$, is also monotone. Hence, we may consider the map $P(A) \rightarrow P(A)$ 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. $\nu_z.f(z, -)$).¹ Let us also recall that f is κ -continuous w.r.t. the coordinatewise order on $P(B) \times P(A)$ if and only if it is κ -continuous in every variable (see Lemma 2.4).

Proposition 3.3. *Let $f : P(B) \times P(A) \rightarrow P(B)$ be a κ -continuous monotone map. If $\kappa > \aleph_0$ then $\nu_z.f(z, -) : P(A) \rightarrow P(B)$ 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 every $X \in P(A)$, if $b \in g(X)$, then $b \in g(X')$ for some κ -small X' contained in X . Having shown this, the continuity of g follows from Proposition 2.3. Let therefore $b \in g(X)$ and note that this condition implies that, for some $Z \subseteq B$, $b \in Z$ and $Z \subseteq f(Z, X)$; let us fix such Z . Aiming at constructing a κ -small subset $X' \subseteq A$ such that $b \in g(X')$, we recursively define a family $(X_n)_{n \geq 1}$ of κ -small subsets of X and a family $(Z_n)_{n \geq 0}$ 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 a κ -small set Z_n that satisfies $Z_n \subseteq f(Z, X)$. Let us consider

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

Since $Z_n \subseteq f(Z, X) = \bigcup \mathcal{I}$ and \mathcal{I} is a κ -directed set, by Proposition 2.2 there exist Z_{n+1}, X_{n+1} κ -small such that $Z_n \subseteq f(Z_{n+1}, X_{n+1})$. Moreover, $Z_{n+1} \subseteq Z \subseteq f(Z, X)$.

¹ Let us mention that later we shall emphasize the distinction syntax/semantics. Then, we shall use **lfp** and **gfp** in the semantics for the symbols μ and ν , respectively, and reserve these symbols for the syntax.

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

$$Z_\omega = \bigcup_{n \geq 0} Z_n \subseteq \bigcup_{n \geq 1} f(Z_n, X_n) \subseteq f\left(\bigcup_{n \geq 1} Z_n, \bigcup_{n \geq 1} X_n\right) \subseteq f(Z_\omega, X_\omega).$$

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. \square

Proposition 3.4. *Let $f : P(B) \times P(A) \rightarrow P(B)$ be a κ -continuous monotone map. If $\kappa \geq \aleph_0$ then $\mu_x.f(z, -) : P(A) \rightarrow P(B)$ is also κ -continuous.*

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

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

$$\begin{aligned} f\left(\bigcup_{i \in I} \mu_x.f(x, X_i), X\right) &= \bigcup_{i \in I} f(\mu_x.f(x, X_i), X) \quad \text{since } f \text{ is } \kappa\text{-continuous in its first argument} \\ &= \bigcup_{i \in I} f(\mu_x.f(x, X_i), \bigcup_{j \in I} X_j) \\ &= \bigcup_{i \in I, j \in I} f(\mu_x.f(x, X_i), X_j) \\ &\quad \text{since } f \text{ is } \kappa\text{-continuous in its second argument} \\ &= \bigcup_{i \in I} f(\mu_x.f(x, X_i), X_i) \quad \text{since } \{X_i \mid i \in I\} \text{ is } \kappa\text{-directed} \\ &= \bigcup_{i \in I} \mu_x.f(x, X_i). \end{aligned}$$

This concludes the proof of Proposition 3.4. \square

Notice that the statement of Proposition 3.4 holds for κ -continuous monotone maps $f : P \times Q \rightarrow P$, that is, we might only assume that P and Q are complete lattices, not powerset algebras. Indeed, the corresponding proof is obtained from the proof of Proposition 3.4 by replacing the set theoretic \bigcup with the supremum symbol \bigvee . Similarly, the statement of Proposition 3.3 is suitable to be generalized to posets P and Q satisfying appropriate conditions, see [35].

Let $\mathcal{F} = \{f_i : P(A)^{n_i} \rightarrow P(A) \mid i \in I\}$ be a collection of monotone operations on $P(A)$. We define *the μ -clone of \mathcal{F}* to be the least set of finitary operations on $P(A)$ that contains \mathcal{F} and the projections and which is closed under the following operations: substitution, taking parametrized least fixed-points and greatest fixed-points.

Corollary 3.5. *Let $\kappa > \aleph_0$ be a regular cardinal. If all the maps in \mathcal{F} are κ -continuous, then all the maps in the μ -clone of \mathcal{F} are also κ -continuous.*

Proof. We shall observe that projections are κ -continuous and that the set of κ -continuous functions is closed under substitution and under the operations of taking least and greatest fixed-points. Projections are lower and upper adjoints, so they actually preserve all unions and intersections, see [12, §7.23 and Proposition 7.31]. For substitution, argue first that the

composition of two κ -continuous maps is κ -continuous. Observe then that if $f_i : P(A) \rightarrow P(B_i)$ is κ -continuous, for $i = 1, \dots, n$, then the unique map $\langle f_1, \dots, f_n \rangle : P(A) \rightarrow \prod_i P(B_i)$ such that $\pi_i \circ \langle f_1, \dots, f_n \rangle = f_i$ for each $i = 1, \dots, n$, is κ -continuous; this is because suprema are computed coordinatewise in $\prod_i P(B_i)$. Therefore, if f_0, f_1, \dots, f_n are κ -continuous, then also the composite $f_0 \circ \langle f_1, \dots, f_n \rangle$ is κ -continuous. For least and greatest fixed-points use Propositions 3.3 and 3.4. \square

4. THE PROPOSITIONAL MODAL μ-CALCULUS

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

Hereinafter Act is a fixed finite set of actions and $Prop$ is a countable set of propositional variables. The set \mathbb{L}_μ of formulas of the propositional modal μ -calculus over Act is generated by the following grammar:

$$\phi := y \mid \neg y \mid \top \mid \phi \wedge \phi \mid \perp \mid \phi \vee \phi \mid \langle a \rangle \phi \mid [a] \phi \mid \mu_z. \phi \mid \nu_z. \phi, \quad (4.1)$$

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. In general, we shall use $x, x_1, \dots, x_n, \dots$ for variables that are never under the scope of a negation nor bound in a formula ϕ ; $y, y_1, \dots, y_n, \dots$ for variables that are free in formulas; $z, z_1, \dots, z_n, \dots$ for variables that are bound in formulas. However, this convention cannot be rigorously enforced, since we shall often consider the steps from a formula ϕ with a free occurrence of the variable z to the formula $\mu_z. \phi$, where z is bound. We think of the grammar (4.1) as a way of specifying the abstract syntax of a formula, as if it was the specification of an inductive type in a programming language such as Haskell. Nonetheless, we shall write formulas thus we need to be able to disambiguate them. To achieve this goal we use standard conventions: \wedge has higher priority than \vee , unary modal connectors have higher priority than binary logical connectors. The least and greatest fixed-points operators yield priority instead, the dot notation emphasizes this. For example, the formula $\mu_x. \phi \wedge \psi$ is implicitly parenthesised as $\mu_x. (\phi \wedge \psi)$ instead of $(\mu_x. \phi) \wedge \psi$.

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 (of worlds or states); for each $a \in Act$, $R_a \subseteq |\mathcal{M}| \times |\mathcal{M}|$ is a (accessibility or transition) relation; $v : Prop \rightarrow P(|\mathcal{M}|)$ is a valuation, i.e., 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 $\psi \in \mathbb{L}_\mu$ as subsets of $|\mathcal{M}|$ is recursively defined using the standard clauses from multimodal logic **K** (see e.g. [25]). For example, we have

$$\begin{aligned} \llbracket \langle a \rangle \psi \rrbracket_{\mathcal{M}} &:= \{ s \in |\mathcal{M}| \mid \exists s' (sR_a s' \ \& \ s' \in \llbracket \psi \rrbracket_{\mathcal{M}}) \}, \\ \llbracket [a] \psi \rrbracket_{\mathcal{M}} &:= \{ s \in |\mathcal{M}| \mid \forall s' (sR_a s' \Rightarrow s' \in \llbracket \psi \rrbracket_{\mathcal{M}}) \}. \end{aligned}$$

We present next the semantics of the least and greatest fixed-point constructors μ and ν . For this purpose, given a subset $Z \subseteq |\mathcal{M}|$, we define $\mathcal{M}[z \mapsto Z]$ to be the model that possibly 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:

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

A formula $\phi \in \mathbf{L}_\mu$ and a variable $x \in Prop$ determine on every model \mathcal{M} the correspondence $\phi_{\mathcal{M}}^x : P(|\mathcal{M}|) \rightarrow P(|\mathcal{M}|)$, that sends each $S \subseteq |\mathcal{M}|$ to $\llbracket \phi \rrbracket_{\mathcal{M}[x \mapsto S]} \subseteq |\mathcal{M}|$. We shall write in the following $\phi_{\mathcal{M}}$ for $\phi_{\mathcal{M}}^x$, when x is understood. Coming back to the clauses for the fixed-point constructors, the syntactic restrictions on the variable z in the productions of $\mu_z.\psi$ and $\nu_z.\psi$ (z must be positive in ψ) imply that the function $\psi_{\mathcal{M}}^z$ is monotone. By Tarski's theorem [38], the above clauses state that $\llbracket \mu_z.\psi \rrbracket_{\mathcal{M}}$ and $\llbracket \nu_z.\psi \rrbracket_{\mathcal{M}}$ are, respectively, the least and the greatest fixed-point of $\psi_{\mathcal{M}}^z$. As usual, we write $\mathcal{M}, s \Vdash \psi$ to mean that $s \in \llbracket \psi \rrbracket_{\mathcal{M}}$.

4.1. The closure of a formula. For $\phi \in \mathbf{L}_\mu$, we denote by $Sub(\phi)$ the set of subformulas of ϕ . A *substitution* is an expression of the form $[\psi_1/y_1, \dots, \psi_n/y_n]$ where, for $i = 1 \dots, n$, y_i is a propositional variable and $\psi_i \in \mathbf{L}_\mu$. We use $\phi[\psi_1/y_1, \dots, \psi_n/y_n]$ to denote *application* of the substitution $[\psi_1/y_1, \dots, \psi_n/y_n]$ to the formula ϕ —that is, the result of simultaneously replacing every free occurrence of the variable y_i in ϕ by the formula ψ_i , $i = 1, \dots, n$. As usual for formal systems with variable binders, we may assume that variable capture does not arise when applying substitutions to formulas. When we want to emphasize application (of a substitution to a formula) we use a dot: for example, $\phi \cdot [\psi_1/y_1, \dots, \psi_n/y_n]$ and $\phi[\psi_1/y_1, \dots, \psi_n/y_n]$ denote the same formula. We also use the symbol \cdot to denote composition of substitutions. For $\sigma_1 := [\phi_1/x_1, \dots, \phi_n/x_n]$ and $\sigma_2 := [\psi_1/y_1, \dots, \psi_m/y_m]$, the composite substitution $\sigma_1 \cdot \sigma_2$ is defined by

$$\sigma_1 \cdot \sigma_2 := [\phi_1[\psi_1/y_1, \dots, \psi_m/y_m]/x_1, \dots, \phi_n[\psi_1/y_1, \dots, \psi_m/y_m]/x_n].$$

A formula $\phi \in \mathbf{L}_\mu$ 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 occurrence ψ of ϕ of the form $Q_z.\psi'$, with $Q \in \{\mu, \nu\}$.

It is well-known that every formula $\phi \in \mathbf{L}_\mu$ is equivalent to a well-named formula. We shall use well-named formulas only to have an accurate description of the game semantics, see § 4.2.

For $\phi \in \mathbf{L}_\mu$ well-named and $\psi \in Sub(\phi)$, the *standard context* of ψ in ϕ is the composite substitution

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

uniquely determined by the following conditions:

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

The *closure* of a well-named $\phi \in \mathbf{L}_\mu$, see [21], is the set $CL(\phi)$ defined as follows:

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

Recall from [21] that $CL(\phi)$ can be characterised as the least subset of \mathbf{L}_μ such that

- $\phi \in CL(\phi)$,
- if $\psi_1 @ \psi_2 \in CL(\phi)$, then $\psi_1, \psi_2 \in CL(\phi)$, with $@ \in \{\wedge, \vee\}$,
- if $\langle a \rangle \psi \in CL(\phi)$ or $[a]\psi \in CL(\phi)$, then $\psi \in CL(\phi)$,
- if $Q_z.\psi \in CL(\phi)$, then $\psi[Q_z.\psi/z] \in CL(\phi)$, with $Q \in \{\mu, \nu\}$.

The definition of $CL(\phi)$ implies it is finite.

4.2. Game semantics. Given $\phi \in \mathbb{L}_\mu$ well-named and a model $\mathcal{M} = \langle |\mathcal{M}|, \{R_a \mid a \in \text{Act}\}, v \rangle$, the game $\mathcal{G}(\mathcal{M}, \phi)$ is the two player game of perfect information and possibly infinite duration—a *parity game*, see e.g. [4, Chapter 4]—defined as follows. Players of $\mathcal{G}(\mathcal{M}, \phi)$ are named Eva and Adam. The set of positions is the Cartesian product $|\mathcal{M}| \times \text{CL}(\phi)$. Moves are as in the table below:

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

From a position of the form (s, \top) Adam loses, and from a position of the form (s, \perp) 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 defining infinite winning plays. To achieve this goal, we choose a rank function $\rho : \text{CL}(\phi) \rightarrow \mathbb{N}$ 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. The winner of an infinite play $\{(s_n, \psi_n) \mid n \geq 0\}$ is determined by the *parity condition*: it is a win for Eva if and only if $\max\{n \geq 0 \mid \{i \mid \rho^{-1}(\psi_i) \text{ is infinite}\} \}$ is even.

Let us recall the following fundamental result (see for example [9, Theorem 6]):

Proposition 4.1. *For each model \mathcal{M} and each well-named formula $\phi \in \mathbb{L}_\mu$, $\mathcal{M}, s \Vdash \phi$ if and only if Eva has a winning strategy from position (s, ϕ) in the game $\mathcal{G}(\mathcal{M}, \phi)$.*

4.3. Bisimulations. Let $P \subseteq \text{Prop}$ be a subset of variables and let $B \subseteq \text{Act}$ 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 $p \in P$,
- for each $b \in B$,
 - $xR_b y$ implies $x'R_b y'$ for some y' such that $(y, y') \in \mathcal{B}$,
 - $x'R_b y'$ implies $xR_b y$ for some y such that $(y, y') \in \mathcal{B}$.

A pointed model is a pair $\langle \mathcal{M}, s \rangle$ with $\mathcal{M} = \langle |\mathcal{M}|, \{R_a \mid a \in \text{Act}\}, v \rangle$ 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 $(\text{Prop}, \text{Act})$ -bisimilar.

Let us denote by $\mathbb{L}_\mu[P, B]$ 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 [9, Theorem 10].

Proposition 4.2. *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 $\phi \in \mathbb{L}_\mu[P, B]$.*

5. \aleph_1 -CONTINUOUS FRAGMENT OF THE MODAL μ -CALCULUS

In this section we introduce a fragment of the modal μ -calculus which we name $\mathcal{C}_{\aleph_1}(x)$. Formulas in this fragment give rise to \aleph_1 -continuous maps when interpreted as monotone maps of the variable x . We show how to construct a formula $\phi' \in \mathcal{C}_{\aleph_1}(x)$ from a given arbitrary formula ϕ in order to satisfy the following property: ϕ is κ -continuous for some infinite regular cardinal κ if and only if ϕ and ϕ' are equivalent formulas. Our conclusions are twofold. Firstly, we deduce the decidability of the problem whether a formula is κ -continuous for some κ is decidable. Decidability relies on the effectiveness of the construction and on the well-known fact that equivalence for the modal μ -calculus is elementary [13]. Secondly, we observe that if a formula is κ -continuous, then it is already \aleph_1 -continuous or even \aleph_0 -continuous. Thus, there are no interesting notions of κ -continuity for the modal μ -calculus besides those for the cardinals \aleph_0 and \aleph_1 .

Definition 5.1. A formula $\phi \in \mathsf{L}_\mu$ is κ -continuous in x if $\phi_{\mathcal{M}}$ is κ -continuous, for each model \mathcal{M} . If $X \subseteq \text{Prop}$, then we say that ϕ is κ -continuous in X if ϕ is κ -continuous in x for each $x \in X$.

Definition 5.2. We 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:

$$\phi := x \mid \psi \mid \top \mid \perp \mid \phi \wedge \phi \mid \phi \vee \phi \mid \langle a \rangle \phi \mid \mu_z \cdot \chi \mid \nu_z \cdot \chi, \quad (5.1)$$

where $x \in X$, $\psi \in \mathsf{L}_\mu$ is a μ -calculus formula not containing any variable $x \in X$, and $\chi \in \mathcal{C}_{\aleph_1}(X \cup \{z\})$.

If we omit the last production from the above grammar, we obtain a grammar for the continuous fragment of the modal μ -calculus, see [14], which we denote here by $\mathcal{C}_{\aleph_0}(X)$. For $i = 0, 1$, we shall write $\mathcal{C}_{\aleph_i}(x)$ for $\mathcal{C}_{\aleph_i}(\{x\})$. The main result of [14] is that a formula $\phi \in \mathsf{L}_\mu$ is \aleph_0 -continuous in x if and only if it is equivalent to a formula in $\mathcal{C}_{\aleph_0}(x)$. It must be observed that the fragment presented above is the same as the one presented in [15] under the name of finite width fragment.

Let $X = \{x_1, \dots, x_n\}$; a straightforward induction shows that, for each $\phi \in \mathcal{C}_{\aleph_1}(X)$, the map that sends a tuple $(S_1, \dots, S_n) \in P(|\mathcal{M}|)^n$ to $\llbracket \phi \rrbracket_{\mathcal{M}[x_1 \mapsto S_1, \dots, x_n \mapsto S_n]}$ belongs to the μ -clone generated by intersections, unions, the modal operators $\langle a \rangle_{\mathcal{M}}$ and the constants $\llbracket \psi \rrbracket_{\mathcal{M}}$. Since all these generating operations are \aleph_1 -continuous maps (actually, they are \aleph_0 -continuous) we can use Corollary 3.5 to derive the following statement.

Proposition 5.3. *Every formula in the fragment $\mathcal{C}_{\aleph_1}(X)$ is \aleph_1 -continuous in X .*

5.1. Syntactic considerations.

Definition 5.4. The *digraph* $G(\phi)$ of a formula $\phi \in \mathsf{L}_\mu$ 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 5.5. 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 the following property: for every node, there exists a unique simple path from the root to this node.

Definition 5.6. 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) *slightly-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 5.7. Figure 1 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 slightly-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.




Figure 1: The digraph of a formula in L_μ .

Lemma 5.8. 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. Let X be a set of variables and $\phi \in L_\mu$.

(1) implies (2). The proof is by induction on the structure of formulas. Consider a formula $\phi \in \mathcal{C}_{\aleph_1}(X)$ and observe that the only way to introduce a bad path from the root of $G(\phi)$ to an occurrence of some variable $x \in X$ is either by using a modal operator $[a]$ —which, however, is excluded by the grammar defining the fragment $\mathcal{C}_{\aleph_1}(X)$ —or by a fixed-point formation rule. Therefore, we focus on the case where ϕ is of the form $Q_z.\chi$, for $Q \in \{\mu, \nu\}$ and $\chi \in \mathcal{C}_{\aleph_1}(X \cup \{z\})$, inductively assuming that no occurrence of a variable $x \in X \cup \{z\}$ is bad in χ . Suppose that there is an occurrence of a variable $x \in X$ and a bad path from the root of $G(Q_z.\chi)$ to this occurrence. Since this occurrence of x is not bad in χ , this path necessarily crosses an edge from an occurrence of the variable z to the root of $G(Q_z.\chi)$. But then this occurrence of z is bad in $G(\chi)$, contradicting the inductive hypothesis.

(2) implies (1). Suppose there exist pairs of the form (X, ϕ) where X is a finite set of variables, $\phi \notin \mathcal{C}_{\aleph_1}(X)$ and, for each $x \in X$, ϕ has no bad occurrence of x . Among these pairs, consider (X, ϕ) with ϕ of least complexity, where we define the complexity of a formula ϕ as the number of vertices in $G(\phi)$. Clearly, ϕ has to be of the form $Q_z.\chi$. Moreover, by the second production of the grammar (5.1), it must contain a free occurrence of a variable $x \in X$. Observe that χ has no bad occurrence of any $x \in X$, since such a bad occurrence yields a bad occurrence of x in $Q_z.\chi$. Also, if an occurrence of z is bad in χ , then any occurrence of some $x \in X$ is bad in $Q_z.\chi$. Therefore, χ has no bad occurrence of any variable in $X \cup \{z\}$. By the minimality assumption on (X, ϕ) , χ belongs to $\mathcal{C}_{\aleph_1}(X \cup \{z\})$ and so $\phi \in \mathcal{C}_{\aleph_1}(X)$, a contradiction. \square

5.2. The $\mathcal{C}_{\aleph_1}(x)$ -flattening of formulas. We aim at defining the $\mathcal{C}_{\aleph_1}(x)$ -flattening ϕ^{bx} of any formula ϕ of the modal μ -calculus. This will go through the definition of the intermediate formula $\phi^{\#x}$ which has one more new free variable \bar{x} . The formula $\phi^{\#x}$ is obtained from ϕ by renaming to \bar{x} all the boxed occurrences of the variable x . In the definition of $\phi^{\#x}$ below, we assume that x has no bound occurrences in ϕ . The formal definition is given by induction as follows:

$$\begin{aligned} y^{\#x} &= y & (\neg y)^{\#x} &= \neg y \\ \top^{\#x} &= \top & \perp^{\#x} &= \perp \\ (\psi_0 @ \psi_1)^{\#x} &= \psi_0^{\#x} @ \psi_1^{\#x} & \text{with } @ \in \{ \wedge, \vee \}, \\ (\langle a \rangle \psi)^{\#x} &= \langle a \rangle \psi^{\#x} & ([a] \psi)^{\#x} &= [a] \psi[\bar{x}/x] \\ (Q_z.\psi)^{\#x} &= Q_z.\psi^{\#x} & \text{with } Q \in \{ \mu, \nu \}. \end{aligned}$$

The following fact is proved by a straightforward induction.

Lemma 5.9. *For each $\phi \in \mathbf{L}_\mu$, we have*

$$\phi^{\#x} \cdot [x/\bar{x}] = \phi. \quad (5.2)$$

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

$$\phi^{bx} := \phi^{\#x} \cdot [\perp/\bar{x}]$$

and henceforward we shorten it to ϕ^b .

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

Definition 5.10. 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. $\{x\}$.

Remark 5.11. Let ψ be a well-named variant of a formula ϕ , so ψ is obtained from ϕ by renaming some bound variables. The digraphs $G(\psi)$ and $G(\phi)$ differ only for the labelling of some pairs of nodes lying on a back edge from an occurrence of a bound variable to its binding fixed-point quantifier. Now let P be a property of formulas defined by means of the digraphs $G(\phi)$ without mentioning the labels of nodes on any of those back-edges. Then a formula ϕ has the property P if and only if any of its well-named variant has the property P . One such P is the property of being almost-good. Therefore, if ϕ is almost-good, then so it is any of its well-named variants.

Lemma 5.12. *If ϕ is an almost-good formula, then both $\phi^{\sharp x}$ and ϕ^{\flat} belong to $\mathcal{C}_{N_1}(x)$.*

Proof. We prove the result for $\phi^{\sharp x}$. Consider a bad occurrence of x in $\phi^{\sharp x}$. After substituting \bar{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 slightly-bad, that is, under the scope of a necessity modal operator $[a]$. But then this same occurrence of x in ϕ would correspond to an occurrence of \bar{x} in $\phi^{\sharp x}$ and not to an occurrence of x as assumed. \square

We aim to transform a formula ϕ into an equivalent formula in which there are no very-bad occurrences of the variable x . The transformation that we define next achieves this goal. For $\phi \in \mathbf{L}_\mu$ and a finite set X of variables not bound in ϕ , we define a formula $\psi^{\square X}$, with all the occurrences of a bad variable $x \in X$ boxed (aka slightly-bad). We let

$$\psi^{\square X} := \psi, \quad \text{if no occurrence of a variable } x \in X \text{ is very-bad in } \psi,$$

and, otherwise,

$$\begin{aligned} \langle a \rangle \psi^{\square X} &:= \langle a \rangle (\psi)^{\square X}, \\ (\psi_1 @ \psi_2)^{\square X} &:= (\psi_1)^{\square X} @ (\psi_2)^{\square X}, \quad \text{with } @ \in \{ \wedge, \vee \}, \\ (Q_z. \psi)^{\square X} &:= \psi_0[\psi_1/\bar{z}], \quad \text{where} \\ \psi_0 &:= Q_z. \psi_2, \quad \psi_2 := (\psi^{\square X \cup \{z\}})^{\sharp z}, \quad \text{and } \psi_1 := Q_{\bar{z}}. \psi_0, \end{aligned}$$

with $Q \in \{ \mu, \nu \}$. That is, in the last clause, ψ_2 is obtained from $\psi^{\square X \cup \{z\}}$ by renaming all the boxed occurrences of z to \bar{z} . A key point of the definition of $(Q_z. \psi)^{\square X}$ is that, when we split, with ψ_2 , the fixed-point variable z into its boxed/unboxed parts, we also split, with ψ_1 and ψ_0 , the respective fixed-point bindings, see Figure 2. Observe that the first defining clause implies that

$$\begin{aligned} x^{\square X} &= x \text{ if } x \in X, \\ \psi^{\square X} &= \psi \text{ if } \psi \text{ contains no variable } x \in X, \\ ([a]\psi)^{\square X} &= [a]\psi. \end{aligned}$$

Example 5.13. Consider the formula $\psi := x \vee \mu_z. x \vee z \vee [a](x \wedge z)$, where only the second occurrence of x is very-bad. For $X = \{ x \}$ we have

$$\psi^{\square X} = x \vee \mu_z. x \vee z \vee [a](x \wedge \mu_{\bar{z}}. \mu_z. x \vee z \vee [a](x \wedge \bar{z}))$$

where no occurrence of x is very-bad and so the formula $\psi^{\square X}$ is almost-good.

Proposition 5.14. *The formula $\phi^{\square X}$ is almost-good w.r.t. X and it is equivalent to the formula ϕ .*

We split the proof of the proposition in two lemmas.

Lemma 5.15. *The formula $\phi^{\square X}$ 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 $@ \in \{ \wedge, \vee \}$, the statement is an immediate consequence of the inductive hypothesis. In

case of a formula of the form $(Q_z.\psi)^{\square X}$ with $Q \in \{\mu, \nu\}$, we argue as follows:

$$\begin{aligned}
(Q_z.\psi)^{\square X} &= \psi_0[Q_{\bar{z}}.\psi_0/\bar{z}] \equiv Q_{\bar{z}}.\psi_0, && \text{by the fixed-point equation,} \\
&= Q_{\bar{z}}.Q_z.\psi_2 \equiv Q_z.\psi_2[z/\bar{z}], && \text{by the equational properties of fixed-points,} \\
&= Q_z.((\psi^{\square X \cup \{z\}})^{\sharp z}[z/\bar{z}]) = Q_z.(\psi^{\square X \cup \{z\}}), && \text{by equation (5.2),} \\
&\equiv Q_z.\psi, && \text{by the inductive hypothesis.}
\end{aligned}$$

□

Lemma 5.16. *The formula $\phi^{\square X}$ is almost-good, that is, it has no very-bad occurrence of a variable $x \in X$.*

Figure 2 illustrates the proof of this lemma.

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 $@ \in \{\wedge, \vee\}$, the statement is an immediate consequence of the inductive hypothesis. The only non-trivial case is that of a formula of the form $(Q_z.\psi)^{\square X}$ with $Q \in \{\mu, \nu\}$.

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

We argue next that every occurrence of a variable $x \in X$ within $\psi_0 = Q_z.\psi_2$ 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 $Q_z.\psi_2$. 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 $Q_w.\chi$, with this subformula lying on the path from ψ_2 to the occurrence of z . Yet (a) is not possible: recall that $\psi_2 = (\psi^{\square X \cup \{z\}})^{\sharp z}$, thus all the occurrences of z within ψ_2 are not boxed (such an occurrence in $\psi^{\square X \cup \{z\}}$ has been renamed to \bar{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 $\psi^{\square X \cup \{z\}}$, and renaming the boxed occurrences of z to \bar{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 $\psi_0[\psi_1/\bar{z}]$. 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 \bar{z} within ψ_0 is boxed, all the variable occurrences of x within ψ_1 become boxed in the formula $\psi_0[\psi_1/\bar{z}]$.

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

We can finally state our first main result.

Theorem 5.17. *Every formula ϕ is equivalent to a formula ψ with $\psi^{\sharp x}$ and ψ^{\flat} in $\mathcal{C}_{\aleph_1}(x)$. Moreover, we can choose ψ well-named.*

In the theorem we can take ψ to be a well-named variant of the almost-good formula $\phi^{\square \{x\}}$. Then, by Remark 5.11, ψ is almost-good and therefore, by Lemma 5.12, $\psi^{\sharp x}$ and ψ^{\flat} belong to $\mathcal{C}_{\aleph_1}(x)$.




Figure 2: Illustration of the proof of Lemma 5.16

5.3. Comparing the closures of ϕ and ϕ^b . In the following, let ϕ be a well-named formula. Observe that both $\phi^{\sharp x}$ and ϕ^b are also well-named—we verify this for $\phi^{\sharp x}$, the argument for ϕ^b is similar. Indeed, $\phi^{\sharp x}$ has the same bound variables as ϕ and therefore \bar{x} , assumed to be new, cannot be bound in $\phi^{\sharp x}$. Next, if z is bound in $\phi^{\sharp x}$, then it is bound in ϕ and there is a unique subformula occurrence of ϕ of the form $Q_z.\psi$ and therefore a unique subformula occurrence of $\phi^{\sharp x}$ of the form $Q_z.\psi'$, the latter being either $Q_z.\psi^{\sharp x}$ or $Q_z.\psi[\bar{x}/x]$.

We develop here some syntactic considerations that allow us to relate the closures of ϕ and ϕ^b . In turn, that will make it possible to relate the positions of the games $\mathcal{G}(\mathcal{M}, \phi)$ and $\mathcal{G}(\mathcal{M}, \phi^b)$, and so to construct, in the proof of Proposition 5.21, a winning strategy in the latter game from a winning strategy in the former.

Recall that we use $Sub(\phi)$ for the set of subformulas of ϕ .

Remark 5.18. If \bar{x} and y are distinct variables and χ is a formula that does not contain the variable y , then

$$[\psi/y] \cdot [\chi/\bar{x}] = [\chi/\bar{x}] \cdot [\psi[\chi/\bar{x}]/y]. \quad (5.3)$$

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

The above remark is easily justified considering that for terms t, s over an arbitrary signature we have $Sub(t[s/\bar{x}]) = \{ t'[s/\bar{x}] \mid t' \in Sub(t) \} \cup Sub(s)$, whenever \bar{x} is a variable occurring free in t , where now $Sub(t)$ denotes the set of subterms of t .

Lemma 5.19. *If \bar{x} is a free variable of ϕ and γ is either a variable not bound in ϕ or a constant, then*

$$CL(\phi \cdot [\gamma/\bar{x}]) = \{ \psi \cdot [\gamma/\bar{x}] \mid \psi \in CL(\phi) \}.$$

In particular, we have

$$CL(\phi) = \{ \phi' \cdot [x/\bar{x}] \mid \phi' \in CL(\phi^{\sharp x}) \}, \quad CL(\phi^b) = \{ \phi' \cdot [\perp/\bar{x}] \mid \phi' \in CL(\phi^{\sharp x}) \}.$$

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

Proof. By repeatedly using equation (5.3) with $\chi = \gamma$, we have

$$\begin{aligned} \sigma_\psi^\phi \cdot [\gamma/\bar{x}] &= [Q_n y_n \cdot \psi_n / y_n] \cdot \dots \cdot [Q_1 y_1 \cdot \psi_1 / y_1] \cdot [\gamma/\bar{x}] \\ &= [\gamma/\bar{x}] \cdot [Q_n y_n \cdot \psi_n \cdot [\gamma/\bar{x}] / y_n] \cdot \dots \cdot [Q_1 y_1 \cdot \psi_1 \cdot [\gamma/\bar{x}] / y_1]. \end{aligned}$$

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

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

holds. From this we deduce

$$(\psi \cdot [\gamma/\bar{x}]) \cdot \sigma_{\psi \cdot [\gamma/\bar{x}]}^{\phi \cdot [\gamma/\bar{x}]} = (\psi \cdot \sigma_\psi^\phi) \cdot [\gamma/\bar{x}]. \quad (5.4)$$

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

This concludes the proof of Lemma 5.19. \square

5.4. The continuous fragments. Now we aim to prove some sort of converse of Proposition 5.3, namely that every κ -continuous formula ϕ of the propositional modal μ -calculus is equivalent to ϕ^b , where κ is still assumed to be an infinite regular cardinal.

A pointed model $\langle \mathcal{M}, s \rangle$ is a *tree model* if the rooted digraph $\langle |\mathcal{M}|, \bigcup_{a \in \text{Act}} R_a, s \rangle$ is a tree. Let κ be a cardinal. A tree model $\langle \mathcal{M}, s \rangle$ is κ -*expanded* if, for each $a \in \text{Act}$, whenever $x R_a x'$, there are at least κ a -successors of x that are bisimilar to x' . The following lemma is straightforward, see e.g. [14, Proposition 1] for the case where $\kappa = \aleph_0$.

Lemma 5.20. *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 5.21. *If $\mathcal{M}, s \Vdash \phi$ and ϕ is κ -continuous in x , then $\mathcal{M}, s \Vdash \phi^b$.*

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 $s_0 \Vdash \phi^b$. Notice first that, by Lemma 5.20, we can assume that $\langle \mathcal{M}, s_0 \rangle$ is a 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 $\mathcal{M}[x \mapsto U], s_0 \Vdash \phi$. We shall argue that $\mathcal{M}[x \mapsto U], s_0 \Vdash \phi^b$, from which it follows that $s_0 \in \phi_{\mathcal{M}}^b(U) \subseteq \phi_{\mathcal{M}}^b(v(x))$ —since $\phi_{\mathcal{M}}^b$ is monotone—thus $\mathcal{M}, s_0 \Vdash \phi^b$.

In the following let $\mathcal{N} = \mathcal{M}[x \mapsto U]$ (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 $\mathcal{G}(\mathcal{N}, \phi^b)$ from position (s_0, ϕ^b) . Observe first that, by Lemma 5.19, positions in $\mathcal{G}(\mathcal{N}, \phi)$ (respectively, $\mathcal{G}(\mathcal{N}, \phi^b)$) are of the form $(s, \psi[x/\bar{x}])$ (resp., $(s, \psi[\perp/\bar{x}])$) for a formula $\psi \in \text{CL}(\phi^{\sharp x})$. Therefore, at the beginning of the play, Eva plays in $\mathcal{G}(\mathcal{N}, \phi^b)$ 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 $p := (s, [a]\chi\sigma_{[a]\chi}^{\phi^{\sharp x}} \cdot [x/\bar{x}])$ and $p' := (s, [a]\chi\sigma_{[a]\chi}^{\phi^{\sharp x}} \cdot [\perp/\bar{x}])$, for some subformula $[a]\chi$ of $\phi^{\sharp x}$, where $\chi = \chi'[\bar{x}/x]$ for some subformula χ' of ϕ .

Claim 5.22. The positions p and p' are respectively of the form $(s, [a]\psi) \in \mathcal{G}(\mathcal{N}, \phi)$ and $(s, [a]\psi') \in \mathcal{G}(\mathcal{N}, \phi^b)$ for some ψ and ψ' such that $\psi[\perp/x] \rightarrow \psi'$ is a tautology.

Proof of Claim. In the computations that follows we use the notation $\phi \geq \phi'$ (for $\phi, \phi' \in \mathbf{L}_\mu$) to mean that $\llbracket \phi \rrbracket_{\mathcal{M}} \supseteq \llbracket \phi' \rrbracket_{\mathcal{M}}$ for every \mathcal{M} (i.e., $\phi' \rightarrow \phi$ is a tautology).

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

$$\begin{aligned} \psi &= \chi\sigma_{\chi}^{\phi^{\sharp x}} \cdot [x/\bar{x}] = \chi'[\bar{x}/x] \cdot \sigma_{\chi'[\bar{x}/x]}^{\phi^{\sharp x}} \cdot [x/\bar{x}] \\ &= \chi'[\bar{x}/x] \cdot [x/\bar{x}] \cdot \sigma_{\chi'[\bar{x}/x] \cdot [x/\bar{x}]}^{\phi^{\sharp x} \cdot [x/\bar{x}]}, && \text{by equation (5.4),} \\ &= \chi' \cdot \sigma_{\chi'}^{\phi}, \end{aligned}$$

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

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

Thus, Eva needs to continue playing in the game $\mathcal{G}(\mathcal{N}, \phi^b)$ from a position of the form $(s, [a]\psi')$ where $\psi[\perp/x] \rightarrow \psi'$ is a tautology. We construct a winning strategy 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 $\mathcal{G}(\mathcal{N}, \phi^b)$, which she shall use to continue the play.

Claim 5.23. $s \in \llbracket [a]\psi \rrbracket_{\mathcal{N}}$ implies $s \in \llbracket [a]\psi[\perp/x] \rrbracket_{\mathcal{N}}$.

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. $sR_a s'$), so $\mathcal{N}, s' \Vdash \psi$; we want to show that $\mathcal{N}, s' \Vdash \psi[\perp/x]$. To this goal, recalling that $\psi[\perp/x] \in \mathbf{L}_\mu[\text{Prop} \setminus \{x\}, \text{Act}]$ and using Proposition 4.2, it is enough to prove that $\langle \mathcal{N}, s' \rangle$ is $(\text{Prop} \setminus \{x\}, \text{Act})$ -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 $\{t \mid sR_a t, \langle \mathcal{M}, t \rangle \text{ is bisimilar to } \langle \mathcal{M}, s' \rangle\}$ 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 and $\mathcal{N}, s'' \Vdash \psi[\perp/x]$ imply $\mathcal{N}, s' \Vdash \psi[\perp/x]$. \square *Claim.*

To complete the proof of Proposition 5.21 we need to argue that the strategy so defined for Eva to play in the game $\mathcal{G}(\mathcal{M}, \phi^b)$ 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 $(s, \bar{x}[x/\bar{x}])$ and $(s, \bar{x}[\perp/\bar{x}])$. This is however excluded since in $\phi^{\#x}$ all the occurrences of \bar{x} are boxed, so we are enforced to go through the second step of the strategy. \square

Proposition 5.24. *If, for some regular cardinal κ , $\phi \in \mathsf{L}_\mu$ is κ -continuous, then ϕ is equivalent to ϕ^b .*

Proof. Notice that, by monotonicity in the variable x , $\phi^b \rightarrow \phi$ is a tautology. Proposition 5.21 exhibits the converse implication as another tautology. \square

Theorem 5.25. *If for some regular cardinal κ , $\phi \in \mathsf{L}_\mu$ is a κ -continuous formula, then ϕ is equivalent to a formula $\phi' \in \mathcal{C}_{\aleph_1}(x)$.*

Proof. Suppose that ϕ is κ -continuous. By Corollary 5.17, ϕ is equivalent to a formula ψ with $\psi^b \in \mathcal{C}_{\aleph_1}(x)$. Clearly, ψ is κ -continuous as well, so it is equivalent to ψ^b by Proposition 5.24. It follows that ϕ is equivalent to $\psi^b \in \mathcal{C}_{\aleph_1}(x)$. \square

A fragment of the modal μ -calculus is a subset of L_μ . For an infinite regular cardinal κ , we let $\mathcal{C}_\kappa(x)$ be the set of κ -continuous formulas $\phi(x) \in \mathsf{L}_\mu$, cf. Definition 5.1. We say that a fragment \mathcal{F} of the modal μ -calculus is determined by a continuity condition if, for some infinite regular cardinal κ , $\mathcal{F} = \mathcal{C}_\kappa(x)$. Combining the main result of [14] and Theorem 5.25, we immediately obtain the following result.

Theorem 5.26. *There are only two fragments of the modal μ -calculus determined by continuity conditions: the fragment $\mathcal{C}_{\aleph_0}(x)$ and the fragment $\mathcal{C}_{\aleph_1}(x)$.*

Theorem 5.27. *The following problem is decidable: given a formula $\phi(x) \in \mathsf{L}_\mu$, 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 $\phi' \in \mathcal{C}_{\aleph_1}(x)$, where $\phi' = (\phi^{\square x})^b$. 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' . \square

6. ON p -DEFINABILITY

We collect in this section some technical results, mainly on relating different types of submodels via formulas, that we shall use later to prove two main results on closure ordinals of the modal μ -calculus, Theorem 7.6 and Theorem 8.1.

We start recalling the usual notion of Kripke frame (hereinafter referred to as frame). An *Act-frame* (or simply, a frame, if *Act* is understood) is a pair $\mathcal{F} = \langle |\mathcal{F}|, \{R_a \mid a \in Act\} \rangle$ where $|\mathcal{F}|$ is a set and $R_a \subseteq |\mathcal{F}| \times |\mathcal{F}|$, for each $a \in Act$ – in other words, a frame is a model without a valuation of propositional variables. If $v : Prop \rightarrow P(|\mathcal{F}|)$ is a valuation, then we denote by \mathcal{F}_v the model $\langle \mathcal{F}, v \rangle$. The complex algebra $\mathcal{F}^\#$ of a frame \mathcal{F} is the Boolean

algebra of subsets of $|\mathcal{F}|$ endowed with (the interpretation of) the modal operators $\langle a \rangle_{\mathcal{F}^\sharp}$, $a \in Act$, defined by

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

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 $G \subseteq F$.

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

Definition 6.1. Let $\Psi = \{\psi_b \in \mathbf{L}_\mu[p, q] \mid b \in B\}$ 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 $\langle \mathcal{F}, [G/p, S/q] \rangle$.

Example 6.2. 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 $\{p \wedge \langle a \rangle(p \wedge q) \mid a \in A\}$. \square *Example 6.2.*

The two examples we present below illustrate the notion of p -definability. Moreover, they both shall allow (in conjunction with Proposition 6.7) to transfer results from a bimodal setting (that is, when $\text{card}(Act) = 2$) to a monomodal one ($\text{card}(Act) = 1$). In particular, the second example shall be used to prove Theorem 7.6.

In the following $B := \{h, v\}$ and A is a singleton. The choice of the letters is suggested by the construction in Section 7.1 where the actions h and v are interpreted respectively as horizontal and vertical transitions.

Example 6.3. We are thankful to an anonymous referee for suggesting the following construction. Given a bimodal frame \mathcal{G} , we define a monomodal frame \mathcal{F} on the disjoint union of the sets $|\mathcal{G}|$ and R_v by letting the accessibility relation be as follows:

$$\begin{aligned} xRy, & \quad \text{when } xR_h y, \\ xR(x, y) \text{ and } (x, y)Ry, & \quad \text{when } xR_v y. \end{aligned}$$

Clearly $|\mathcal{G}|$ embeds into $|\mathcal{F}|$. By identifying $|\mathcal{G}|$ with its image in $|\mathcal{F}|$, \mathcal{G} is p -defined in \mathcal{F} by $\Psi = \{\psi_v, \psi_h\}$, where

$$\begin{aligned} \psi_h(p, q) &= p \wedge \langle \rangle(p \wedge q), \\ \psi_v(p, q) &= p \wedge \langle \rangle(\neg p \wedge \langle \rangle(p \wedge q)). \end{aligned} \quad \square \text{ Example 6.3.}$$

Example 6.4 (Thomason's coding of bimodal logic into monomodal logic). In [39], see also [24, Section 4], Thomason constructs:

- (i) a monomodal formula ϕ^{sim} , for each (fixed-point free) bimodal formula ϕ ;
- (ii) a monomodal model \mathcal{M}^{sim} and an injective function $(-)^{\circ} : |\mathcal{M}| \rightarrow |\mathcal{M}^{sim}|$, for each bimodal model \mathcal{M} .

These data have the following property:

Fact 6.5. For each $s \in |\mathcal{M}|$, $\mathcal{M}, s \Vdash \phi$ if and only if $M^{sim}, s^{\circ} \Vdash \phi^{sim}$.

We recall how \mathcal{M}^{sim} is defined: for a $\{h, v\}$ -model \mathcal{M} , \mathcal{M}^{sim} is the monomodal model with $|\mathcal{M}^{sim}| = |\mathcal{M}| \times \{h, v\} \cup \{p_0\}$, such that $v(x, i) = v(x)$ and whose accessibility relation R is described as follows:

$$\begin{aligned} (x, h)R(y, h), & \quad \text{when } xR_h y, \\ (x, v)R(y, v), & \quad \text{when } xR_v y, \\ (x, v)R(x, h), & \quad (x, h)R(x, v), \quad \text{and } (x, h)R p_0, \end{aligned}$$

for each $x, y \in |\mathcal{M}|$. Since the function sending $x \in |\mathcal{M}|$ to $x^\circ := (x, h) \in |\mathcal{M}^{sim}|$ is injective, we can identify $|\mathcal{M}|$ with a subset of $|\mathcal{M}^{sim}|$. 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 $|\mathcal{M}^{sim}|$. Fact 6.5 relies on \mathcal{G} being p -defined in \mathcal{F} by $\Psi = \{\psi_h, \psi_v\}$, where

$$\begin{aligned} \psi_h(p, q) &= p \wedge \langle \rangle (p \wedge q), \\ \psi_v(p, q) &= p \wedge \langle \rangle (\neg p \wedge \langle \rangle (\neg p \wedge \langle \rangle (p \wedge q))). \end{aligned}$$

The reader has remarked the similarity with the previous example. Thomason's construction is slightly more subtle: by adding the pit p_0 to \mathcal{M}^{sim} and transitions as in the third line of the above display, the image of \mathcal{M} under the embedding becomes definable by the formula $\langle \rangle [\] \perp$. Consequently, the monomodal formula ϕ^{sim} does not contain p as an additional propositional variable. \square *Example 6.4.*

We tackle next the proof of the main technical result of this section, Proposition 6.7. This proposition allows lifting standard simulation results (such as Thomason's one) from modal logic to the modal μ -calculus.

Definition 6.6. Let $p \notin Prop$ be a fresh variable and let $\Psi := \{\psi_b \in \mathsf{L}_\mu[p, q] \mid b \in B\}$. The formula $\mathbf{tr}^\Psi(\phi)$ is defined by induction as follows:

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

In the above definition, ψ_b^{op} is a formula dual to ψ_b , thus semantically behaving as $\neg\psi_b[\neg q/q]$. We need this since in the grammar (4.1) we allowed negation only on propositional variables.

Aiming at a proof of the next Proposition, let us introduce/recall some notation: we let $\pi : P(F) \rightarrow P(G)$ be defined by $\pi(S) := S \cap G$; if $v : Prop \rightarrow P(F)$, then $\pi \circ v : Prop \rightarrow P(G)$ is the valuation in G such that $(\pi \circ v)(y) := G \cap v(y)$, for each $y \in Prop$.

Proposition 6.7. *Let p, Ψ , and \mathbf{tr}^Ψ be as in Definition 6.6. If \mathcal{G} is p -defined in \mathcal{F} by Ψ , then, for each valuation $v : Prop \rightarrow P(F)$,*

$$\llbracket \phi \rrbracket_{\mathcal{G}_{\pi \circ v}} = \llbracket \mathbf{tr}^\Psi(\phi) \rrbracket_{\mathcal{F}_v[p \mapsto G]}. \quad (6.1)$$

Remark 6.8. For a formula ϕ , let us denote by $\llbracket \mathbf{tr}^\Psi(\phi) \rrbracket_{\mathcal{F}[p \mapsto G]}$ the mapping from $P(F)^{Prop}$ to $P(F)$ sending a valuation $v \in P(F)^{Prop}$ to $\llbracket \mathbf{tr}^\Psi(\phi) \rrbracket_{\mathcal{F}_v[p \mapsto G]} \in P(F)$; let us denote by

$\llbracket \phi \rrbracket_G$ the mapping sending a valuation $v' \in P(G)^{Prop}$ to $\llbracket \phi \rrbracket_{\mathcal{G}_{v'}}$ in $P(G)$. The statement of Proposition 6.7 implies that $\llbracket \mathbf{tr}^\Psi(\phi) \rrbracket_{\mathcal{F}[G/p]}$ takes values in $P(G)$ and, moreover, that the following diagram commutes:

$$\begin{array}{ccc} P(F)^{Prop} & & \\ \downarrow \pi_{o_-} & \searrow \llbracket \mathbf{tr}^\Psi(\phi) \rrbracket_{\mathcal{F}[G/p]} & \\ P(G)^{Prop} & \xrightarrow{\llbracket \phi \rrbracket_G} & P(G) . \end{array}$$

Proof of Proposition 6.7. The proof that equation (6.1) holds is by induction on formulas. The basic cases are treated below:

$$\begin{aligned} \llbracket \mathbf{tr}^\Psi(y) \rrbracket_{\mathcal{F}_v[p \mapsto G]} &= \llbracket p \wedge y \rrbracket_{\mathcal{F}_v[p \mapsto G]} = G \cap v(y) = \llbracket y \rrbracket_{\mathcal{G}_{\pi \circ v}} , \\ \llbracket \mathbf{tr}^\Psi(\neg y) \rrbracket_{\mathcal{F}_v[p \mapsto G]} &= \llbracket p \wedge \neg y \rrbracket_{\mathcal{F}_v[p \mapsto G]} = G \cap v(\neg y) = G \cap v(y)^c \\ &= G \cap (G \cap v(y))^c = \llbracket \neg y \rrbracket_{\mathcal{G}_{\pi \circ v}} , \\ \llbracket \mathbf{tr}^\Psi(\perp) \rrbracket_{\mathcal{F}_v[p \mapsto G]} &= \llbracket \perp \rrbracket_{\mathcal{F}_v[p \mapsto G]} = \emptyset = \llbracket \perp \rrbracket_{\mathcal{G}_{\pi \circ v}} , \\ \llbracket \mathbf{tr}^\Psi(\top) \rrbracket_{\mathcal{F}_v[p \mapsto G]} &= \llbracket p \rrbracket_{\mathcal{F}_v[p \mapsto G]} = G = \llbracket \top \rrbracket_{\mathcal{G}_{\pi \circ v}} . \end{aligned}$$

For formulas of the form $\psi_0 @ \psi_1$ with $@ \in \{\wedge, \vee\}$, the result is immediate by induction. We give below explicit computations for formulas whose main logical connector is a modal operator:

$$\begin{aligned} \llbracket \mathbf{tr}^\Psi(\langle b \rangle \psi) \rrbracket_{\mathcal{F}_v[p \mapsto G]} &= \llbracket \psi_b[\mathbf{tr}^\Psi(\psi)/q] \rrbracket_{\mathcal{F}_v[p \mapsto G]} \\ &= \llbracket \psi_b \rrbracket_{\mathcal{F}_v[p \mapsto G, q \mapsto \llbracket \mathbf{tr}^\Psi(\psi) \rrbracket_{\mathcal{F}_v[p \mapsto G]}]} \\ &= \llbracket \langle b \rangle q \rrbracket_{\mathcal{G}_{\pi \circ v}[q \mapsto \llbracket \psi \rrbracket_{\mathcal{G}_{\pi \circ v}}]} = \llbracket \langle b \rangle \psi \rrbracket_{\mathcal{G}_{\pi \circ v}} , \\ \llbracket \mathbf{tr}^\Psi([b] \psi) \rrbracket_{\mathcal{F}_v[p \mapsto G]} &= \llbracket p \wedge \neg \psi_b[\neg \mathbf{tr}^\Psi(\psi)/q] \rrbracket_{\mathcal{F}_v[p \mapsto G]} \\ &= G \cap (\llbracket \psi_b[\neg \mathbf{tr}^\Psi(\psi)/q] \rrbracket_{\mathcal{F}_v[p \mapsto G]})^c \\ &= G \cap (\llbracket \psi_b \rrbracket_{\mathcal{F}_v[p \mapsto G, q \mapsto S^c]})^c \quad \text{with } S = \llbracket \mathbf{tr}^\Psi(\psi) \rrbracket_{\mathcal{F}_v[p \mapsto G]} = \llbracket \phi \rrbracket_{\mathcal{G}_{\pi \circ v}} \\ &= G \cap (\llbracket \langle b \rangle q \rrbracket_{\mathcal{G}_{\pi \circ v}[G \cap S^c/q]})^c \\ &= \llbracket \neg \langle b \rangle \neg q \rrbracket_{\mathcal{G}_{\pi \circ v}[q \mapsto S]} \\ &= \llbracket [b] q \rrbracket_{\mathcal{G}_{\pi \circ v}[q \mapsto S]} \\ &= \llbracket [b] q \rrbracket_{\mathcal{G}_{\pi \circ v}[q \mapsto \llbracket \psi \rrbracket_{\mathcal{G}_{\pi \circ v}}]} \\ &= \llbracket [b] \psi \rrbracket_{\mathcal{G}_{\pi \circ v}} . \end{aligned}$$

We finally consider least and the greatest fixed-point formulas of the form $\mu_z.\phi$ and $\nu_z.\phi$. Consider the two functions defined by

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

and remark firstly their typing, that is we have $f : P(F) \rightarrow P(F)$ and $g : P(G) \rightarrow P(G)$. Since by the inductive hypothesis we have

$$\llbracket \mathbf{tr}^\Psi(\phi) \rrbracket_{\mathcal{F}_w[p \mapsto G]} = \llbracket \phi \rrbracket_{\mathcal{G}_{\pi \circ w}}$$

for each valuation w , this in particular holds for the valuation $v[z \mapsto S]$, with $S \subseteq F$; that is, we have

$$f(S) = g(S \cap G), \quad (6.2)$$

for each $S \subseteq F$. Let us denote by \mathbf{Pre}_h the set of prefixed-points of a monotone function h and by $\mathbf{lfp}.h$ its least element.² It immediately follows from equation (6.2) that \mathbf{Pre}_g is included in \mathbf{Pre}_f and that $S \in \mathbf{Pre}_f$ implies $\pi(S) \in \mathbf{Pre}_g$. Therefore the inclusion of \mathbf{Pre}_g into \mathbf{Pre}_f has π as an upper adjoint, so it is a lower adjoint and therefore (as usual for lower adjoints) it preserves the least element: $\mathbf{lfp}.g = \mathbf{lfp}.f$. We obtain

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

For the greatest fixed-point, denote by \mathbf{Pos}_h the set of postfixes-points of some monotone function h and by $\mathbf{gfp}.h$ its greatest element. Using equation 6.2, observe that $S \subseteq f(S)$ implies $S \subseteq G$. It immediately follows that $\mathbf{Pos}_f = \mathbf{Pos}_g$, so

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

This concludes the proof of Proposition 6.7. \square

It has been easier for us to expose the proof of Proposition 6.7 using frames. Next, we recast our previous observations using models, for the particular cases of submodels (Example 6.2) and of bimodal models (Examples 6.3 and 6.4).

If $\mathcal{M} = \langle |\mathcal{M}|, \{R_a^{\mathcal{M}} \mid a \in \mathit{Act}\}, v \rangle$ and $\mathcal{N} = \langle |\mathcal{N}|, \{R_a^{\mathcal{N}} \mid a \in \mathit{Act}\}, v_{\mathcal{N}} \rangle$ 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 $y \in \mathit{Prop}$ and each $a \in \mathit{Act}$,

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

Thus, \mathcal{N} is a submodel of \mathcal{M} if and only if, for some frame \mathcal{F} , for a valuation $v : \mathit{Prop} \rightarrow P(|\mathcal{F}|)$, 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 $\mathcal{M}_{\upharpoonright S}$ of \mathcal{M} defined as follows:

$$\mathcal{M}_{\upharpoonright S} := \langle S, \{R_a \cap S \times S \mid a \in \mathit{Act}\}, v' \rangle \quad (6.3)$$

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* . We write $\mathbf{tr}(\phi)$ in place of $\mathbf{tr}^\Psi(\phi)$ if Ψ is the collection of formulas given in Example 6.2. Proposition 6.7 instantiates then to models and submodels as follows:

Proposition 6.9. *For each formula $\phi \in \mathbf{L}_\mu$, the formula $\mathbf{tr}(\phi) \in \mathbf{L}_\mu$ (which contains p as a new propositional variable) has the following property: for each model \mathcal{M} , each subset $S \subseteq |\mathcal{M}|$, and each $s \in |\mathcal{M}|$,*

$$\mathcal{M}[p \mapsto S], s \models \mathbf{tr}(\phi) \text{ iff } s \in S \text{ and } \mathcal{M}_{\upharpoonright S}, s \models \phi.$$

A subset S of $|\mathcal{M}|$ is *closed* if $s \in S$ and $sR_a s'$ imply $s' \in S$, for every $a \in \mathit{Act}$. A submodel \mathcal{N} of \mathcal{M} is *closed* if $|\mathcal{N}|$ is a closed subset of $|\mathcal{M}|$. The attentive reader might have already observed that if S is a closed subset of \mathcal{M} , then the statement of Proposition 6.9 holds with the simpler $p \wedge \phi$ in place of the recursively defined $\mathbf{tr}(\phi)$.

Let us fix Ψ from one of Example 6.3 or 6.4. The translating function \mathbf{tr}^Ψ has now the following properties:

²We prefer to use here the notation \mathbf{lfp} in place of μ so to reserve the symbol μ for the syntax and to emphasize the gap between semantics and syntax that we are trying to fill.

- (i) it associates to each bimodal formula ϕ of the modal μ -calculus a monomodal formula $\mathbf{tr}^\Psi(\phi)$ of the modal μ -calculus,
- (ii) the formula $\mathbf{tr}^\Psi(\phi)$ contains a new propositional variable,
- (iii) the formula $\mathbf{tr}^\Psi(\phi)$ belongs to $\mathcal{C}_{\aleph_1}(x)$ if ϕ does.

Moreover, in case Ψ comes from Example 6.4, then (ii) can be strengthened to the statement that $\mathbf{tr}^\Psi(\phi)$ has exactly the same propositional variables as ϕ . Proposition 6.7 then yields the following result.

Proposition 6.10. *For each bimodal model \mathcal{M} there is a monomodal model \mathcal{M}^{sim} and an injective function $(-)^{\circ} : |\mathcal{M}| \rightarrow |\mathcal{M}^{sim}|$ such that, for each $s \in |\mathcal{M}|$, $\mathcal{M}, s \Vdash \phi$ if and only if $\mathcal{M}^{sim}[p \mapsto S], s^{\circ} \Vdash \mathbf{tr}^\Psi(\phi)$, where S is the image of s under the injective function $(-)^{\circ}$.*

Proposition 6.7 also yields the following result, needed to transfer results on closure ordinals:

Proposition 6.11. *Let $\phi \in \mathcal{L}_{\mu}$ with x occurring positively in ϕ .*

- (i) *If \mathcal{M} is a model and $S \subseteq |\mathcal{M}|$, then*

$$\mathbf{tr}(\phi)_{\mathcal{M}[p \mapsto S]}^{\alpha}(\emptyset) = \phi_{\mathcal{M}_{\uparrow S}}^{\alpha}(\emptyset).$$

- (ii) *If \mathcal{M} is a model and $S \subseteq |\mathcal{M}|$ is closed, then*

$$(p \wedge \phi)_{\mathcal{M}[p \mapsto S]}^{\alpha}(\emptyset) = \phi_{\mathcal{M}_{\uparrow S}}^{\alpha}(\emptyset).$$

- (iii) *If \mathcal{M} is a bimodal model and both Ψ and the construction \mathcal{M}^{sim} come from one of the Examples 6.3 or 6.4, then*

$$\mathbf{tr}^\Psi(\phi)_{\mathcal{M}^{sim}[p \mapsto S]}^{\alpha}(\emptyset) = [\phi_{\mathcal{M}}^{\alpha}(\emptyset)]^{\circ},$$

where S is the image of $|\mathcal{M}|$ under the injective function $(-)^{\circ}$.

Proof. Let $\mathcal{F}, \mathcal{G}, F, G$ and $v : Prop \rightarrow P(F)$ be as in the statement of Proposition 6.7. If S is a subset of G , then

$$\begin{aligned} \mathbf{tr}^\Psi(\phi)_{\mathcal{F}_v[p \mapsto G]}(S) &= \llbracket \mathbf{tr}^\Psi(\phi) \rrbracket_{\mathcal{F}_v[p \mapsto G][x \mapsto S]} = \llbracket \mathbf{tr}^\Psi(\phi) \rrbracket_{\mathcal{F}_v[x \mapsto S][p \mapsto G]} \\ &= \llbracket \phi \rrbracket_{\mathcal{G}_{\pi \circ v}[x \mapsto S]} = \phi_{\mathcal{G}_{\pi \circ v}}(S). \end{aligned}$$

Then, by induction, we easily derive

$$(\mathbf{tr}^\Psi(\phi)_{\mathcal{F}_v[p \mapsto G]})^{\alpha}(\emptyset) = \phi_{\mathcal{G}_{\pi \circ v}}^{\alpha}(\emptyset),$$

for each ordinal α . The three statements above follow considering Examples 6.2, 6.3, and 6.4. \square

Finally, consider again Example 6.4 and formulas (resp. models) ϕ' (resp. \mathcal{M}') defined by

$$\phi' := \mathbf{tr}^\Psi(\psi)[\langle \rangle[\]_{\perp}/p], \quad \mathcal{M}' := \mathcal{M}^{sim}[p \mapsto \llbracket \langle \rangle[\]_{\perp}/p \rrbracket_{\mathcal{M}}].$$

Let us identify the injective function $(-)^{\circ} : |\mathcal{M}| \rightarrow |\mathcal{M}^{sim}|$ with an inclusion (so that, instead of embedding \mathcal{M} into \mathcal{M}^{sim} , we are actually extending it into some bigger model). We derive henceforth the following simpler statement that we shall use in the next section to argue that ω_1 is the closure ordinal of a monomodal formula. In the statement the role of the special variable p is not transparent anymore.

Proposition 6.12. *For each bimodal formula ϕ there is a monomodal formula ϕ' (with the same free variables of ϕ) such that if $\phi \in \mathcal{C}_{\aleph_1}(x)$, then $\phi' \in \mathcal{C}_{\aleph_1}(x)$, and with the following property: for each bimodal model \mathcal{M} there is a monomodal model \mathcal{M}' (that does not depend on ϕ) such that,*

- (i) $|\mathcal{M}| \subseteq |\mathcal{M}'|$,
- (ii) $\mathcal{M}, s \Vdash \phi$ if and only if $\mathcal{M}', s \Vdash \phi'$, for each $s \in |\mathcal{M}|$,
- (iii) $(\phi'_{\mathcal{M}'})^\alpha(\emptyset) = \phi_{\mathcal{M}}^\alpha(\emptyset)$, for each ordinal α .

7. AN UNCOUNTABLE CLOSURE ORDINAL

In this section we firstly formally define the notion of closure ordinal, present some tools required later, here and in the next section, and then we prove that ω_1 , the least uncountable ordinal, is a closure ordinal of a formula of the modal μ -calculus. We firstly prove it in a bimodal setting and then, using the tools developed in the previous section, we argue that ω_1 is also the closure ordinal of a monomodal μ -formula.

For a formula $\phi(x)$ of the modal μ -calculus and a Kripke model \mathcal{M} , let $\text{cl}_{\mathcal{M}}(\phi)$ be the least ordinal β for which $\phi_{\mathcal{M}}^\beta(\emptyset) = \phi_{\mathcal{M}}^{\beta+1}(\emptyset)$. Recall from Definition 3.1 that we say that $\phi_{\mathcal{M}}$ converges to its least fixed-point in exactly α steps when $\text{cl}_{\mathcal{M}}(\phi) = \alpha$.

Definition 7.1. 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 $\text{cl}(\phi) = \alpha$) 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.

Elsewhere in the literature, see e.g. [2], the closure ordinal of a formula $\phi(x)$ w.r.t. a class of models \mathcal{K} is defined as the supremum of the ordinals $\text{cl}_{\mathcal{M}}(\phi)$ for $\mathcal{M} \in \mathcal{K}$. If \mathcal{K} is the class of Kripke models, then this definition coincides with the one given above. This is a consequence of the class of Kripke models being closed under disjoint unions: consider a family $\{\mathcal{M}_i \mid i \in I\}$ such that $\alpha = \sup\{\text{cl}_{\mathcal{M}_i}(\phi)\}$; then the disjoint union $\bigcup_{i \in I} \mathcal{M}_i$ carries a canonical structure of a Kripke model, call it \mathcal{M} , and it is easily seen that $\text{cl}_{\mathcal{M}}(\phi) = \alpha$.

The notions of *closure ordinal of a formula on a structure* and of *closure ordinal of a structure* appear in the monograph [29, Chapter 2B]. The notion of closure ordinal presented here is on the other hand strictly related to *global* inductive definability, see [5]. Indeed, it is well-known that each fixed-point-free modal formula ψ can be transformed into some equivalent first order logic sentence $ST_y(\psi)$, known as the *standard translation* of ψ . The formula $ST_y(\psi)$ contains y as the only free-variable and is related to ψ by the equivalence $\mathcal{M}, s \Vdash \psi$ if and only if $\mathcal{M} \models ST_y(\psi)(s)$, where \mathcal{M} is considered as a relational structure for first-order logic. The closure ordinal of a fixed-point-free modal formula $\phi(x)$, as defined here and when it exists, coincides with the global closure ordinal of the first-order inductive definition given by $ST_y(\phi(x))$.

Let us recall that formulas may have no closure ordinal. For example $\phi(x) := []x$ has no closure ordinal. Indeed, it is not difficult to construct, for each ordinal α , a model \mathcal{M}_α such that $\phi_{\mathcal{M}_\alpha}^\alpha(\emptyset)$ is strictly included in $\phi_{\mathcal{M}_\alpha}^{\alpha+1}(\emptyset)$. We collect with the following Proposition the observations developed in the course of the paper that are relevant to closure ordinals.

Proposition 7.2. *If a formula $\phi(x)$ belongs to the syntactic fragment $\mathcal{C}_{\aleph_1}(x)$, then it has a closure ordinal $\text{cl}(\phi(x))$ and ω_1 is an upper bound for $\text{cl}(\phi(x))$.*

Proof. The formula ϕ belongs to the syntactic fragment $\mathcal{C}_{\aleph_1}(x)$, thus it is \aleph_1 -continuous and, for every model \mathcal{M} , $\phi_{\mathcal{M}}$ is \aleph_1 -continuous. It follows then from Proposition 3.2 that $\phi_{\mathcal{M}}$ converges to its least fixed-point in at most ω_1 steps. Therefore, for such ϕ , $\sup\{\text{cl}_{\mathcal{M}}(\phi) \mid \mathcal{M} \text{ a Kripke model}\} \leq \omega_1$. As we have seen at the beginning of this section, there exists a model \mathcal{M} such that $\text{cl}_{\mathcal{M}}(\phi) = \sup\{\text{cl}_{\mathcal{M}}(\phi) \mid \mathcal{M} \text{ a Kripke model}\}$. \square

The following Lemma will be useful in the next section, when we shall show that closure ordinals of the modal μ -calculus are closed under ordinal sum.

Lemma 7.3. *Let $\alpha \neq 0$ be a closure ordinal of the modal μ -calculus. Among the formulas 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,*

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

Proof. For a formula $\psi(x)$, let $(\mu_x.\psi(x))^{op}$ be a formula semantically equivalent to the negation of $\mu_x.\psi(x)$ and define then

$$\phi(x) := (\mu_x.\psi(x))^{op} \vee \psi(x \wedge \mu_x.\psi(x)).$$

Observe that $\phi(x)$ is not well-named, yet this will not be a concern here. For the sake of readability, let $\mu := \mathbf{lfp}.\psi_{\mathcal{M}}$. We verify next that

$$\phi_{\mathcal{M}}^{\gamma}(\emptyset) = \mu \rightarrow \psi_{\mathcal{M}}^{\gamma}(\emptyset), \quad \text{for each ordinal } \gamma \geq 1. \quad (7.1)$$

The symbol \rightarrow used above stands for the Heyting implication of the Boolean algebra $P(|\mathcal{M}|)$. Equation (7.1) clearly holds if $\gamma = 1$. Assuming the equation holds for γ , then

$$\begin{aligned} \phi_{\mathcal{M}}^{\gamma+1}(\emptyset) &= \mu \rightarrow \psi_{\mathcal{M}}(\mu \rightarrow \psi_{\mathcal{M}}^{\gamma}(\emptyset)) \cap \mu \\ &= \mu \rightarrow \psi_{\mathcal{M}}(\psi_{\mathcal{M}}^{\gamma}(\emptyset) \cap \mu) \\ &= \mu \rightarrow \psi_{\mathcal{M}}(\psi_{\mathcal{M}}^{\gamma}(\emptyset)), & \text{since } \psi_{\mathcal{M}}^{\gamma}(\emptyset) \subseteq \mathbf{lfp}.\psi_{\mathcal{M}} = \mu, \\ &= \mu \rightarrow \psi_{\mathcal{M}}^{\gamma+1}(\emptyset). \end{aligned}$$

The inductive step to a limit ordinal is obvious. From equation (7.1) it follows that, for each $\gamma \neq 0$, $\phi_{\mathcal{M}}^{\gamma+1}(\emptyset) \subseteq \phi_{\mathcal{M}}^{\gamma}(\emptyset)$ if and only if $\psi_{\mathcal{M}}^{\gamma+1}(\emptyset) \subseteq \psi_{\mathcal{M}}^{\gamma}(\emptyset)$, so $\text{cl}_{\mathcal{M}}(\phi) = \text{cl}_{\mathcal{M}}(\psi)$ provided that $\text{cl}_{\mathcal{M}}(\psi) > 0$. Finally, $\llbracket \mu_x.\phi(x) \rrbracket_{\mathcal{M}} = \mathbf{lfp}.\phi_{\mathcal{M}} = \mu \rightarrow \mu = |\mathcal{M}|$. \square

7.1. ω_1 is a closure ordinal. We are going to prove that ω_1 is the closure ordinal of the following bimodal formula:

$$\Phi(x) := (\nu_z.\langle v \rangle x \wedge \langle h \rangle z) \vee [v]\perp. \quad (7.2)$$

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

The construction relies on a 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

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

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 $(n, \beta)R_h(n+1, \beta)$, for each $n < \omega$ and each ordinal β , and vertical transitions from a state $(n, \beta) \in S_{\omega_1}$ are as follows:

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

Lemma 7.4. *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 $\phi_{\mathcal{M}_{\omega_1}}^\alpha(\emptyset) = S_\alpha$.

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

$$\begin{aligned} \phi_{\mathcal{M}_{\omega_1}}(S_0) &= \phi_{\mathcal{M}_{\omega_1}}(\emptyset) = \llbracket \nu_z.(\langle h \rangle z \wedge \langle v \rangle x) \vee [v] \perp \rrbracket_{\mathcal{M}_{\omega_1}[x \mapsto \emptyset]} \\ &= \llbracket [v] \perp \rrbracket_{\mathcal{M}_{\omega_1}} = \{(n, 0) \mid n < \omega\} = S_1. \end{aligned}$$

Consider now an ordinal $\alpha > 0$.

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

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

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

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

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

If α is a limit ordinal, then

$$\phi_{\mathcal{M}_{\omega_1}}^\alpha(\emptyset) = \bigcup_{\beta < \alpha} \phi_{\mathcal{M}_{\omega_1}}^\beta(\emptyset) = \bigcup_{\beta < \alpha} S_\beta = S_\alpha.$$

This concludes the proof of Lemma 7.4. \square

We conclude the section by stating its main result.

Theorem 7.5. *The closure ordinal of $\Phi(x)$ is ω_1 .*

Proof. As we mentioned before the formula $\Phi(x)$ has a closure ordinal and $\text{cl}(\Phi(x)) \leq \omega_1$, by Proposition 7.2. We claim that $\Phi_{\mathcal{M}_{\omega_1}}$ converges to its least fixed-point in exactly ω_1 steps, that is, we have $\Phi_{\mathcal{M}_{\omega_1}}^{\omega_1}(\emptyset) \not\subseteq \Phi_{\mathcal{M}_{\omega_1}}^\alpha(\emptyset)$ for each $\alpha < \omega_1$. Our claim is verified as follows. By Lemma 7.4, the claim is equivalent to $S_{\omega_1} \not\subseteq S_\alpha$, for each $\alpha < \omega_1$. The latter relation holds since if $\alpha < \omega_1$, then we can find an ordinal β with $\alpha < \beta < \omega_1$, so the states (n, β) , $n \geq 0$, belong to $S_{\omega_1} \setminus S_\alpha$. \square

Finally, we argue that a bimodal language is not needed for ω_1 to be a closure ordinal. To this goal, let Ψ be as in Example 6.4 and let

$$\Phi' := \text{tr}^\Psi(\Phi)[\langle \rangle[\]\perp/p], \quad \mathcal{M}'_{\omega_1} := \mathcal{M}_{\omega_1}^{\text{sim}}[p \mapsto \llbracket \langle \rangle[\]\perp \rrbracket_{\mathcal{M}_{\omega_1}}],$$

where Φ is the bimodal formula defined in equation (7.2). As in the statement of Proposition 6.12, we consider $|\mathcal{M}'_{\omega_1}|$ as a superset of $|\mathcal{M}_{\omega_1}|$.

Theorem 7.6. *The monomodal formula Φ' has closure ordinal ω_1 .*

Proof. Consider the statement of Proposition 6.12. Since the correspondence $\phi \mapsto \phi'$ sends formulas in $\mathcal{C}_{\aleph_1}(x)$ to formulas in $\mathcal{C}_{\aleph_1}(x)$, Φ' is \aleph_1 -continuous and therefore it has a closure ordinal bounded by ω_1 . To argue that the closure ordinal of Φ' is equal to ω_1 it is enough to consider the model \mathcal{M}'_{ω_1} and rely on item (iii) of Proposition 6.12. \square

8. CLOSURE UNDER ORDINAL SUM.

In this section we prove that the ordinal sum of two closure ordinals of the modal μ -calculus is again a closure ordinal of this logic, as stated in the next theorem.

Theorem 8.1. *Suppose $\phi_0(x)$ and $\phi_1(x)$ are monomodal formulas that have, respectively, α and β as closure ordinals. Then there is a monomodal formula $\Psi(x)$, constructible from ϕ_0 and ϕ_1 , whose closure ordinal is $\alpha + \beta$.*

We prove the theorem through a series of observations. With the first one, Lemma 8.2, we make use of the master modality $[\mathcal{U}]$ of the propositional modal μ -calculus. In principle, the use master modality in the proof of Theorem 8.1 may be avoided, at the cost of reducing its readability. Given a monomodal formula χ this modality is defined as follows:

$$[\mathcal{U}]\chi := \nu_z.(\chi \wedge [\]z).$$

The master modality allows us to focus on those models of a fixed shape since they satisfy, globally, a given formula. Indeed, the semantics of this modality is the following:

$$\mathcal{M}, s \Vdash [\mathcal{U}]\chi \text{ if and only if } \mathcal{M}, s' \Vdash \chi, \text{ for each } s' \text{ reachable from } s.$$

In particular, if \mathcal{M} is a tree model, then $\mathcal{M} \Vdash \chi$ if and only if $\mathcal{M}, r \Vdash [\mathcal{U}]\chi$, where r is the root of the tree. Let us mention that the modality $[\mathcal{U}]$ satisfies all the axioms (reflexivity

and transitivity) of the modal system **S4**, see e.g. [22, § 2.5], and yields a deduction theorem for the modal μ -calculus, see [23, 34].

When $\mathcal{M} \Vdash [\mathcal{U}]_\chi$ (that is, $\mathcal{M}, s \Vdash [\mathcal{U}]_\chi$, for each $s \in |\mathcal{M}|$), we say that \mathcal{M} is χ -*acceptable*.

Lemma 8.2. *Let χ and $\psi(x)$ be monomodal formulas and define $\Psi(x) := [\mathcal{U}]_\chi \wedge \psi(x)$. 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 χ -acceptable models, and (ii) there exists an χ -acceptable model on which the formula $\psi(x)$ converges to its least fixed point in exactly γ steps.*

Proof. If \mathcal{N} is an χ -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 $\llbracket [\mathcal{U}]_\chi \rrbracket_{\mathcal{M}}$ is closed and χ -acceptable. Call \mathcal{N} such a submodel of \mathcal{M} . Thus, by Proposition 6.11.(ii), for any ordinal $\gamma \geq 0$, we have

$$\Psi_{\mathcal{M}}^\gamma(\emptyset) = \psi_{\mathcal{N}}^\gamma(\emptyset). \quad (8.1)$$

The statement of the lemma immediately follows. \square

Next, recall that we write $\mathbf{tr}(\phi)$ in place of $\mathbf{tr}^\Psi(\phi)$ if Ψ is the collection of formulas given in Example 6.2. Let $\phi_0(x)$ and $\phi_1(x)$ be monomodal formulas as in the statement of Theorem 8.1. For a variable p occurring neither in ϕ_0 nor in ϕ_1 , we define

$$\chi := \chi_0 \wedge \chi_1 \text{ with } \chi_0 := p \vee ([]\neg p \wedge \mu_z.\phi_0(z)) \text{ and } \chi_1 := \neg p \vee \mu_z.\mathbf{tr}(\phi_1(z)), \quad (8.2)$$

$$\psi(x) := (\neg p \wedge \phi_0(x)) \vee (\mathbf{tr}(\phi_1)(x) \wedge [](p \vee x)), \quad (8.3)$$

$$\Psi(x) := [\mathcal{U}]_\chi \wedge \psi(x). \quad (8.4)$$

From now on, we shall say that a model \mathcal{N} is *acceptable* if it is χ -acceptable, where χ is the formula given in equation (8.2). We shall argue that $\Psi(x)$ defined in (8.4) has closure ordinal $\alpha + \beta$ using Lemma 8.2.

Next, we continue by studying the structure of an acceptable model \mathcal{N} and how $\psi_{\mathcal{N}}$ acts on it—where ψ is the formula defined in (8.3). 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 $N_0 := v(\neg p)$, and $N_1 := v(p)$. A model \mathcal{N} is acceptable if and only if N_0 is a closed subset of $|\mathcal{N}|$ (since $\mathcal{N} \Vdash p \vee []\neg p \equiv \neg p \rightarrow []\neg p$) and moreover

$$N_0 \subseteq \llbracket \mu_z.\phi_0(z) \rrbracket_{\mathcal{N}}, \quad N_1 \subseteq \llbracket \mu_z.\mathbf{tr}(\phi_1(z)) \rrbracket_{\mathcal{N}}.$$

Let also $\phi_{\mathcal{N}_0} := (\phi_0)_{\mathcal{N}_0}$ and $\phi_{\mathcal{N}_1} := (\phi_1)_{\mathcal{N}_1}$, so $\phi_{\mathcal{N}_0} : P(N_0) \rightarrow P(N_0)$ and $\phi_{\mathcal{N}_1} : P(N_1) \rightarrow P(N_1)$. We claim that $\psi_{\mathcal{N}}$ is of the form

$$\psi_{\mathcal{N}}(X) = \phi_{\mathcal{N}_0}(X \cap N_0) \cup (\phi_{\mathcal{N}_1}(X \cap N_1) \cap \nabla(X \cap N_0)), \quad (8.5)$$

with

$$\nabla(X) := N_1 \cap []_{\mathcal{N}}(N_0 \rightarrow X). \quad (8.6)$$

This is because, for each $X \subseteq |\mathcal{N}|$,

$$\begin{aligned}\psi_{\mathcal{N}}(X) &= (\psi_{\mathcal{N}}(X) \cap N_0) \cup (\psi_{\mathcal{N}}(X) \cap N_1), \\ \psi_{\mathcal{N}}(X) \cap N_0 &= \phi_{\mathcal{N}_0}(X \cap N_0), \\ \psi_{\mathcal{N}}(X) \cap N_1 &= \mathbf{tr}(\phi_1)_{\mathcal{N}_1}(X) \cap []_{\mathcal{N}}(N_0 \rightarrow X) = N_1 \cap \phi_{\mathcal{N}_1}(X \cap N_1) \cap []_{\mathcal{N}}(N_0 \rightarrow X) \\ &= \phi_{\mathcal{N}_1}(X \cap N_1) \cap N_1 \cap []_{\mathcal{N}}(N_0 \rightarrow (X \cap N_0)).\end{aligned}$$

We notice now that if \mathcal{N} is acceptable, then

$$N_0 = \llbracket \mu_z.\phi_0(z) \rrbracket_{\mathcal{N}} \cap N_0 = \llbracket \mu_z.\phi_0(z) \rrbracket_{\mathcal{N}_0} = \phi_{\mathcal{N}_0}^{\alpha}(\emptyset) \quad (8.7)$$

and

$$N_1 = \llbracket \mu_z.\mathbf{tr}(\phi_1(z)) \rrbracket_{\mathcal{N}} \cap N_1 = \llbracket \mu_z.\phi_1(z) \rrbracket_{\mathcal{N}_1} = \phi_{\mathcal{N}_1}^{\beta}(\emptyset). \quad (8.8)$$

Observe that $\nabla(X) = N_1$ whenever $N_0 \subseteq X$ and therefore, using $\phi_{\mathcal{N}_0}^{\alpha}(\emptyset) = N_0$, we have

$$\nabla(X) = N_1, \quad \text{whenever } X \supseteq \phi_{\mathcal{N}_0}^{\alpha}(\emptyset). \quad (8.9)$$

Lemma 8.3. *On every acceptable model \mathcal{N} the equality $\psi_{\mathcal{N}}^{\alpha+\beta}(\emptyset) = |\mathcal{N}|$ holds and, consequently, the formula $\psi(x)$ converges within $\alpha + \beta$ steps.*

Proof. Since N_0 is a closed subset of $|\mathcal{N}|$, by Proposition 6.11, we have

$$\psi_{\mathcal{N}}^{\delta}(\emptyset) \cap N_0 = \psi_{\mathcal{N}_0}^{\delta}(\emptyset) = \phi_{\mathcal{N}_0}^{\delta}(\emptyset) \quad (8.10)$$

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 .

Claim 8.4. The following relation holds for every ordinal $\gamma \geq 0$:

$$\phi_{\mathcal{N}_1}^{\gamma}(\emptyset) \subseteq \psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_1. \quad (8.11)$$

Proof of Claim. 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 (the inductive step to a limit ordinal is obvious). We have

$$\begin{aligned}\psi_{\mathcal{N}}^{\alpha+\gamma+1}(\emptyset) \cap N_1 &= \phi_{\mathcal{N}_1}(\psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_1) \cap \nabla(\psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_0) \\ &= \phi_{\mathcal{N}_1}(\psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_1) \cap \nabla(\phi_{\mathcal{N}_0}^{\alpha+\gamma}(\emptyset)), && \text{by equation (8.10),} \\ &= \phi_{\mathcal{N}_1}(\psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_1), && \text{by equation (8.9),} \\ &\supseteq \phi_{\mathcal{N}_1}(\phi_{\mathcal{N}_1}^{\gamma}(\emptyset)), && \text{by the IH,} \\ &= \phi_{\mathcal{N}_1}^{\gamma+1}(\emptyset). && \square \text{ Claim.}\end{aligned}$$

Therefore

$$\begin{aligned}|\mathcal{N}| = N_0 \cup N_1 &= \phi_{\mathcal{N}_0}^{\alpha}(\emptyset) \cup \phi_{\mathcal{N}_1}^{\beta}(\emptyset) && \text{using (8.7) and (8.8)} \\ &\subseteq (\psi_{\mathcal{N}}^{\alpha+\beta}(\emptyset) \cap N_0) \cup (\psi_{\mathcal{N}}^{\alpha+\beta}(\emptyset) \cap N_1) = \psi_{\mathcal{N}}^{\alpha+\beta}(\emptyset).\end{aligned}$$

This terminates the proof of Lemma 8.3. \square

Lemma 8.5. *There exists an acceptable model \mathcal{N} on which $\psi(x)$ converges in exactly $\alpha + \beta$ steps.*

Proof. Since the formulas $\phi_0(x)$ and $\phi_1(x)$ have, respectively, α and β as closure ordinals, by Lemma 7.3 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 $\beta' < \beta$ $\llbracket \mu_x.\phi_0(x) \rrbracket_{\mathcal{M}_\alpha} = |\mathcal{M}_\alpha| = \phi_{0_{\mathcal{M}_\alpha}}^\alpha(\emptyset) \neq \phi_{0_{\mathcal{M}_\alpha}}^{\alpha'}(\emptyset)$ and $\llbracket \mu_x.\phi_1(x) \rrbracket_{\mathcal{M}_\beta} = |\mathcal{M}_\beta| = \phi_{1_{\mathcal{M}_\beta}}^\beta(\emptyset) \neq \phi_{1_{\mathcal{M}_\beta}}^{\beta'}(\emptyset)$.

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 (8.11) is actually an equality, as stated and proved next.

Claim 8.6. 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 (8.11) is an equality, for each ordinal $\gamma \geq 0$:

$$\phi_{\mathcal{N}_1}^\gamma(\emptyset) = \psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_1. \quad (8.12)$$

Proof of Claim. It is enough to verify that the above equality holds for $\gamma = 0$. Indeed, for $\gamma > 0$, we can use the same computations as in the proof of the claim in Lemma 8.3, by substituting an equality for the inclusion in the inductive hypothesis.

If $\delta < \alpha$, then

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

since by assumption $\phi_{\mathcal{N}_0}^\delta(\emptyset)$ is strictly included in N_0 . In particular, if α is a successor ordinal, we have $\psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) \cap N_1 = \emptyset$. If α is a limit ordinal, then

$$\psi_{\mathcal{N}}^\alpha(\emptyset) \cap N_1 \subseteq \bigcup_{\delta < \alpha} \psi_{\mathcal{N}}^{\delta+1}(\emptyset) \cap N_1 \subseteq \bigcup_{\delta < \alpha} \nabla(\psi_{\mathcal{N}}^\delta(\emptyset) \cap N_0) = \emptyset. \quad \square \text{ Claim.}$$

We can then use equations (8.10) and (8.12) to obtain

$$\psi_{\mathcal{N}}^\alpha(\emptyset) = \phi_{\mathcal{N}_0}^\alpha(\emptyset) \supsetneq \phi_{\mathcal{N}_0}^\delta(\emptyset) = \psi_{\mathcal{N}}^\delta(\emptyset) \text{ and } \psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) = N_0 \cup \phi_{\mathcal{N}_1}^\gamma(\emptyset)$$

for ordinals γ, δ such that $\delta < \alpha$. Finally,

$$\psi_{\mathcal{N}}^{\alpha+\beta}(\emptyset) = |\mathcal{N}| = N_0 \cup \phi_{\mathcal{N}_1}^\beta(\emptyset) \supsetneq N_0 \cup \phi_{\mathcal{N}_1}^\gamma(\emptyset) = \psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset), \quad \text{for } \gamma < \beta.$$

This shows that ψ converges in exactly $\alpha + \beta$ steps in $\mathcal{M}_{\alpha+\beta}$ and therefore terminates the proof of Lemma 8.5. \square

Now Theorem 8.1 immediately follows from Lemmas 8.2, 8.3 and 8.5 when applied to the formulas χ, ψ and Ψ defined in 8.2, 8.3 and 8.4 respectively.

In the introduction we used $\text{Ord}(\mathbf{L}_\mu)$ to denote the set of closure ordinals of formulas of the modal μ -calculus. This section yields an insight on Czarnecki's work [11] by proving the closure of $\text{Ord}(\mathbf{L}_\mu)$ under the ordinal sum. The general problem of characterizing $\text{Ord}(\mathbf{L}_\mu)$ is open. At the time of writing this paper, it is our opinion that still a few ordinals are known to belong to $\text{Ord}(\mathbf{L}_\mu)$ —all of them can be constructed from the cardinals $1, \omega$ and ω_1 by iterating the binary ordinal sum. Our results from Section 5 show that no other infinite regular cardinal κ (apart from ω and ω_1) can be proved to belong to $\text{Ord}(\mathbf{L}_\mu)$ in a straightforward way, that is, by relying on the κ -continuity of some formula in \mathbf{L}_μ and on the generalized Kleene theorem (Proposition 3.2). Therefore, any other membership of

$\text{Ord}(\mathbf{L}_\mu)$ requires a very different justification from the known ones. New questions about $\text{Ord}(\mathbf{L}_\mu)$ need to be raised, such as whether this set is closed under other ordinal operations. Let us mention that a recent work [28] exhibits a rich structure for closure ordinals of the modal μ -calculus on bidirectional models. It is conceivable that studying closure ordinals on restricted classes of models will eventually yield a finer understanding of the structure of $\text{Ord}(\mathbf{L}_\mu)$.

REFERENCES

- [1] 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] A. Arnold and D. Niwiński. *Rudiments of μ -calculus*, volume 146 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, Amsterdam, 2001.
- [5] J. Barwise and Y. N. Moschovakis. Global inductive definability. *J. Symbolic Logic*, 43(3):521–534, 09 1978.
- [6] D. Berwanger, E. Grädel, and G. Lenzi. On the variable hierarchy of the modal μ -calculus. In J. Bradfield, editor, *CSL 2002, Proceedings*, volume 2471 of *Lecture Notes in Computer Science*, pages 352–366. Springer, 2002.
- [7] 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.
- [8] J. Bradfield. The modal mu-calculus alternation hierarchy is strict. *Theoretical Computer Science*, 195(2):133–153, 1998.
- [9] J. Bradfield and I. Walukiewicz. The mu-calculus and model-checking. In E. Clarke, T. Henzinger, and H. Veith, editors, *Handbook of Model Checking*. Springer-Verlag, 2015.
- [10] C. S. Calude, S. Jain, B. Khossainov, W. Li, and F. Stephan. Deciding parity games in quasipolynomial time. In H. Hatami, P. McKenzie, and V. King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 252–263. ACM, 2017.
- [11] 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>.
- [12] B. A. Davey and H. A. Priestley. *Introduction to lattices and order*. Cambridge University Press, New York, second edition, 2002.
- [13] E. A. Emerson and C. S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991*, pages 368–377. IEEE Computer Society, 1991.
- [14] 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.
- [15] G. Fontaine and Y. Venema. Some model theory for the modal μ -calculus: syntactic characterisations of semantic properties. *Logical Methods in Computer Science*, Volume 14, Issue 1, Feb. 2018.
- [16] 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.
- [17] 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, *CONCUR '96, Proceedings*, volume 1119 of *Lecture Notes in Computer Science*, pages 263–277. Springer, 1996.
- [18] B. Jónsson. On the canonicity of Sahlqvist identities. *Studia Logica*, 53(4):473–491, 1994.
- [19] A. Joyal and I. Moerdijk. *Algebraic Set Theory*. Lecture note series. Cambridge University Press, 1995.

- [20] M. Jurdzinski. Deciding the winner in parity games is in $UP \cap co-UP$. *Information Processing Letters*, 68(3):119–124, 1998.
- [21] D. Kozen. Results on the propositional μ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
- [22] M. Kracht. *Tools and Techniques in Modal Logic*. Elsevier, 1999.
- [23] M. Kracht. Modal consequence relations. In J. van Benthem, P. Blackburn, and F. Wolter, editors, *Handbook of Modal Logic*, pages 495–549. Elsevier, 2006.
- [24] M. Kracht and F. Wolter. Simulation and transfer results in modal logic - A survey. *Studia Logica*, 59(1):149–177, 1997.
- [25] A. Kurucz, F. Wolter, M. Zakharyashev, and D. Gabbay. *Many-dimensional Modal Logics: Theory and Applications*. Studies in logic and the foundations of mathematics. North Holland Publishing Company, 2003.
- [26] J. Lassez, V. L. Nguyen, and L. Sonenberg. Fixed point theorems and semantics: A folk tale. *Inf. Process. Lett.*, 14(3):112–116, 1982.
- [27] G. Lenzi. The modal μ -calculus: a survey. *Task Quarterly*, 9(3):29–316, 2005.
- [28] G. C. Milanese. An exploration of closure ordinals in the modal μ -calculus. Master’s thesis, University of Amsterdam, 2018. Available from <https://eprints.iillc.uva.nl/1648/>.
- [29] Y. Moschovakis. *Elementary Induction on Abstract Structures*. Studies in logic and the foundations of mathematics. North-Holland, 1974.
- [30] D. Niwiński and I. Walukiewicz. Games for the μ -calculus. *Theoretical Computer Science*, 163(1):99–116, 1996.
- [31] J. Obdržálek. Fast μ -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.
- [32] 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.
- [33] G. Sambin and V. Vaccaro. A new proof of Sahlqvist’s theorem on modal definability and completeness. *Journal of Symbolic Logic*, 54:992–999, 1989.
- [34] L. Santocanale. Congruences of modal μ -algebras. In Z. Ésik and A. Ingólfssdóttir, editors, *FICS02*, volume NS-02-02 of *BRICS Notes Series*, pages 83–87, June 2002. Available from: <http://www.brics.dk/NS/02/Ref/BRICS-NS-02-Ref/BRICS-NS-02-Ref.html#BRICS-NS-02-2>.
- [35] 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>.
- [36] L. Santocanale. Completions of μ -algebras. *Ann. Pure Appl. Logic*, 154(1):27–50, 2008.
- [37] T. Studer. On the proof theory of the modal μ -calculus. *Studia Logica*, 89(3):343–363, 2008.
- [38] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5:285–309, 1955.
- [39] S. K. Thomason. Independent propositional modal logics. *Studia Logica*, 39(2):143–144, 1980.
- [40] I. Walukiewicz. Completeness of Kozen’s axiomatisation of the propositional μ -calculus. *Information and Computation*, 157(1):142–182, 2000.