$\aleph_1$ and the Modal μ-Calculus[*]

Maria João Gouveia[†1] and Luigi Santocanale[2]

1   **CEMAT-CIÊNCIAS (UID/Multi/04621/2013), Faculdade de Ciências, Universidade de Lisboa, Lisbon, Portugal**
    `mjgouveia@fc.ul.pt`
2   **Laboratoire d'Informatique Fondamentale de Marseille, UMR 7279, CNRS AMU, Marseille, France**
    `luigi.santocanale@lif.univ-mrs.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 [17, 20] 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 [14]. Despite its strong expressive power, this logic is still considered as a tractable one: its model checking problem, even if in the class UP ∩ co-UP [15], becomes polynomial as soon as some critical parameters are fixed or restricted classes of models are considered [22, 3, 5]. The widespread interest for this logic has triggered further researches

---

that spread beyond the realm of verification: these concern the expressive power [7, 4], axiomatic bases [30], algebraic and order theoretic approaches [26], deductive systems [21, 27] and the semantics of functional programs [11].

The present paper lies at the intersection of two lines of research on the modal μ-calculus, on continuity [10] and on closure ordinals [9, 2]. Continuity of monotone functions is a fundamental phenomenon in modal logic, on which well known uniform completeness theorems rely [24, 12, 16]. Fontaine [10] characterized the formulas of the modal μ-calculus that give rise to continuous functions on Kripke models. It is well known, for example in categorical approaches to model theory [1], that the notion of continuity of monotone functions (and of functors) can be generalized to κ-continuity, where the parameter κ is an infinite regular cardinal. In the work [25] one of the authors proved that $\aleph_1$-continuous functors are closed under their greatest fixed-points. Guided by this result, we present in this paper a natural syntactic fragment $\mathcal{C}_{\aleph_1}(x)$ of the modal μ-calculus whose formulas are $\aleph_1$-continuous—that is, they give rise to $\aleph_1$-continuous monotone unary functions of the variable x on arbitrary models. A first result that we present here is that *the fragment $\mathcal{C}_{\aleph_1}(x)$ is decidable*: for each $\phi(x) \in \mathsf{L}_\mu$, we construct a formula $\psi(x) \in \mathcal{C}_{\aleph_1}(x)$ such that $\phi(x)$ is $\aleph_1$-continuous on every model if and only if $\phi(x)$ and $\psi(x)$ are semantically equivalent formulas. We borrow some techniques from [10], yet the construction of the formula $\psi(x)$ relies on a new notion of normal form for formulas of the modal μ-calculus. A closer inspection of our proof uncovers a stronger fact: the formulas $\phi(x)$ and $\psi(x)$ are equivalent if and only if, for some regular cardinal κ, $\phi(x)$ is κ-continuous on every model. The stronger statement implies that we cannot find a fragment $\mathcal{C}_\kappa(x)$ of κ-continuous formulas for some cardinal κ strictly larger than $\aleph_1$; any such hypothetical fragment collapses, semantically, to the fragment $\mathcal{C}_{\aleph_1}(x)$.

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

Let us add some final considerations. In our view, the discovery of the fragment $\mathcal{C}_{\aleph_1}(x)$ opens an unsuspected new dimension (thus new tools, new ideas, new perspectives, etc.) in the theory of the modal μ-calculus and of fixed-point logics. Consider for example the modal μ-calculus on deterministic models, where states have at most one successor; we immediately obtain that every formula is $\aleph_1$-continuous on these models. Whether this and other observations can be exploited (towards understanding alternation hierarchies or reasoning using axiomatic bases, for example) is part of future researches. Yet we believe that the scopes of this work and of the problems studied here go well beyond the pure theory of the modal μ-calculus. Our interest in closure ordinals stems from a proof-theoretic work on induction and coinduction [11, 25]. There we banned ordinal notations from the syntax, as we considered the theory of ordinals too strong for our constructive goals. Yet our judgement might have gone too far, since the theory needed to deal with ordinals is not that strong; for example, many statements on ordinals do not need the axiom of choice. This makes reasonable to devise syntaxes based on ordinals. With respect to these problems, related to the semantics of programming languages, the closure ordinal problem becomes an optimal playground where to develop and test intuitions.

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

Proof of all the statements can be found in the preprint [13].

2 κ-continuous mappings and their extremal fixed-points

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

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

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

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

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

▶ **Proposition 2.** *For each $X \subseteq A$, X is κ-small if and only if, for every κ-directed set $\mathcal{I}$, $X \subseteq \bigcup \mathcal{I}$ implies $X \subseteq I$, for some $I \in \mathcal{I}$.*

▶ **Proposition 3.** *A monotone mapping $f : P(A) \to P(B)$ is κ-continuous if and only if, for every $X \in P(A)$,*

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

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

Conversely suppose that $f : P(A) \to P(B)$ is a monotone mapping such that $f(X) = \bigcup f(\mathcal{I}_\kappa(X))$ for every $X \in P(A)$. Let $\mathcal{I}$ be a κ-ideal and let X' be a κ-small set contained in $\bigcup \mathcal{I}$. By Proposition 2 there exists $I \in \mathcal{I}$ such that $X' \subseteq I$. But then $X' \in \mathcal{I}$ since $\mathcal{I}$ is a downward closed set. Thus $\mathcal{I}_\kappa(\bigcup \mathcal{I}) \subseteq \mathcal{I}$ and consequently $f(\bigcup \mathcal{I}) = \bigcup f(\mathcal{I}_\kappa(\bigcup \mathcal{I})) \subseteq \bigcup f(\mathcal{I})$. Since $\bigcup f(\mathcal{I}) \subseteq f(\bigcup \mathcal{I})$ we obtain $f(\bigcup \mathcal{I}) = \bigcup f(\mathcal{I})$.   ◀

2.1   Fixed-points of κ-continuous mappings

The Knaster-Tarski theorem [28] states that if $f : P(A) \to P(A)$ is monotone, then the set $\bigcap \{ X \subseteq A \mid f(X) \subseteq X \}$ is the least fixed-point of f. On the other hand, Kleene's fixed-point theorem states that least fixed-point of an $\aleph_0$-continuous mapping f is constructible by iterating ω_0-times f starting from the empty set, namely it is equal to $\bigcup_{n \geq 0} f^n(\emptyset)$. Generalizations of Kleene's theorem, constructing the least fixed-point of a monotone f by ordinal approximations, appeared later, see for example [8], [19]. The following Proposition 5 generalizes Kleene's theorem to κ-continuous mappings. To state it, we firstly introduce the notions of approximant and convergence.

▶ **Definition 4.** If $f : P(A) \to P(A)$ is a monotone function, then the *approximants* $f^\alpha(\emptyset)$, α an ordinal, are inductively defined as follows:

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

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

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

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

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

Propositions 6 and 7 are specific instances of a result stated for categories [25]. In order to clarify their statements, we first observe that if $f : P(B) \times P(A) \to P(B)$ is a monotone mapping, then the unary mapping $f(-, X) : P(B) \to P(B), Z \mapsto f(Z, X)$, is also monotone. Hence we may consider the mapping $P(A) \to P(A)$ that sends X to the least (resp. greatest) fixed-point of $f(-, X)$; by using the standard μ-calculus notation, we denote it by $\mu_z.f(z, -)$ (resp. $\nu_z.f(z, -)$). We also recall that f is κ-continuous w.r.t. the coordinate-wise order on $P(B) \times P(A)$ if and only if it is κ-continuous in every variable.

▶ **Proposition 6.** *Let $f : P(B) \times P(A) \to P(B)$ be a κ-continuous monotone mapping. If $\kappa > \aleph_0$ then $\nu_z.f(z, \_) : P(A) \to P(B)$ is also κ-continuous.*

Proof. Let us write $g(x) := \nu_z.f(z,x)$. We shall show that, for every $b \in B$ and for $X \in P(A)$, if $b \in g(X)$, then $b \in g(X')$ for some κ-small X' contained in X. Having shown this, it follows by Proposition 3 that g is continuous. Note that the condition $b \in g(X)$ holds when there exists $Z \subseteq B$ such that $b \in Z$ and $Z \subseteq f(Z,X)$. Aiming to find such a set Z we recursively obtain a family $(X_n)_{n\geq 1}$ of κ-small subsets of X and a family $(Z_n)_{n\geq 0}$ 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 Z_n which is κ-small and satisfies $Z_n \subseteq f(Z,X)$. Let us consider

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

Since $Z_n \subseteq f(Z,X) = \bigcup \mathcal{I}$ and $\mathcal{I}$ is a κ-directed set, by Proposition 2 there exist Z_{n+1}, X_{n+1} κ-small such that $Z_n \subseteq f(Z_{n+1}, X_{n+1})$. Moreover, $Z_{n+1} \subseteq Z \subseteq f(Z,X)$.

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

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

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

▶ **Proposition 7.** *Suppose that $\kappa \geq \aleph_0$ and let $f : P(B) \times P(A) \to P(B)$ be a κ-continuous monotone mapping. Then $\mu_z.f(z,\_) : P(A) \to P(B)$ is also κ-continuous.*

3 The propositional modal μ-calculus

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

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

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

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

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

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

Here we only define the semantics of the least and greatest fixed-point constructors μ and ν. To this goal, given a subset $Z \subseteq |\mathcal{M}|$, we define $\mathcal{M}[Z/z]$ to be the model that differs

from $\mathcal{M}$ only on the value Z that its valuation takes on z. The clauses for the fixed-point constructors are the following:

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

A formula $\phi \in \mathsf{L}_\mu$ and a variable $x \in Prop$ determine on every model $\mathcal{M}$ the correspondence $\phi^x_{\mathcal{M}} : P(|\mathcal{M}|) \to P(|\mathcal{M}|)$, that sends each $S \subseteq |\mathcal{M}|$ to $[\![\phi]\!]_{\mathcal{M}[S/x]} \subseteq |\mathcal{M}|$—in the following we shall write $\phi_{\mathcal{M}}$ for $\phi^x_{\mathcal{M}}$, when x is understood. Due to the syntactic restriction on the variable z in the productions of $\mu_z.\phi$ and $\nu_z.\phi$, the function $\phi^z_{\mathcal{M}}$ is monotone. By Tarski's theorem [28], the above clauses state that $[\![\mu_z.\phi]\!]_{\mathcal{M}}$ and $[\![\nu_z.\phi]\!]_{\mathcal{M}}$ are, respectively, the least and the greatest fixed-point of $\phi^z_{\mathcal{M}}$. As usual, we write $\mathcal{M}, s \Vdash \psi$ to mean that $s \in [\![\psi]\!]_{\mathcal{M}}$.

The closure of a formula

For $\phi \in \mathsf{L}_\mu$, we denote by $Sub(\phi)$ the set of subformulas of ϕ. For $\psi \in Sub(\phi)$, the *standard context* of ψ in ϕ is the (composed) substitution

$$\sigma^\phi_\psi := [Q^n_{z_n}.\psi_n/z_n] \cdot \cdots \cdot [Q^1_{z_1}.\psi_1/z_1]$$

uniquely determined by the following conditions:
1. $\{\, z_1, \ldots, z_n \,\}$ is the set of variables bound in ϕ and free in ψ,
2. for each $i = 1, \ldots, n$, $Q^i_{z_i}.\psi_i$ is the unique subformula of ϕ such that $Q^i \in \{\, \mu, \nu \,\}$,
3. $Q^j_{z_j}.\psi_j$ is a subformula of ψ_i, for $i < j$.

For $\phi \in \mathsf{L}_\mu$, the *closure* of ϕ, see [17], is the set $\mathrm{CL}(\phi)$ defined as follows:

$$\mathrm{CL}(\phi) := \{\, \psi \cdot \sigma^\phi_\psi \mid \psi \in Sub(\phi) \,\}.$$

Recall from [17] that $\mathrm{CL}(\phi)$ is the least subset of L_μ such that
- $\phi \in \mathrm{CL}(\phi)$,
- if $\psi_1 @ \psi_2 \in \mathrm{CL}(\phi)$, then $\psi_1, \psi_2 \in \mathrm{CL}(\phi)$, with $@ \in \{\, \wedge, \vee \,\}$,
- if $\langle a \rangle \psi \in \mathrm{CL}(\phi)$ or $[a]\psi \in \mathrm{CL}(\phi)$, then $\psi \in \mathrm{CL}(\phi)$,
- if $Q_z.\psi \in \mathrm{CL}(\phi)$, then $\psi[\mu_z.\psi/z] \in \mathrm{CL}(\phi)$, with $Q \in \{\, \mu, \nu \,\}$.

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

Game semantics

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

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

From a position of the form $(s, \top)$ Adam loses, and from a position of the form $(s, \bot)$ Eva loses. Also, from a position of the form (s, p) with p a propositional variable, Eva wins if and

only if $s \in v(p)$; from a position of the form $(s, \neg p)$ with p a propositional variable, Eva wins if and only if $s \notin v(p)$. The definition of the game is completed by wins on infinite plays. To this goal we choose a rank function $\rho : \mathrm{CL}(\phi) \to \mathbb{N}$ such that, when ψ_1 is a subformula of ψ_2, then $\rho(\psi_1 \cdot \sigma_{\psi_1}^\phi) \leq \rho(\psi_2 \cdot \sigma_{\psi_2}^\phi)$, and such that $\rho(\mu_z.\psi)$ is odd and $\rho(\nu_z.\psi)$ is even. An infinite play $\{ (s_n, \psi_n) \mid n \geq 0 \}$ 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 [6, Theorem 6]):

▶ **Proposition 8.** *For each model $\mathcal{M}$ and each formula $\phi \in \mathsf{L}_\mu$, $\mathcal{M}, s \Vdash \phi$ if and only if Eva has a winning strategy from position (s, ϕ) in the game $\mathcal{G}(\mathcal{M}, \phi)$.*

Bisimulations

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

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

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

Let us denote by $\mathsf{L}_\mu[P, B]$ 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 [6, Theorem 10].

▶ **Proposition 9.** *If $\langle \mathcal{M}, s \rangle$ and $\langle \mathcal{M}', s' \rangle$ are (P, B)-bisimilar, then $\mathcal{M}, s \Vdash \phi$ if and only if $\mathcal{M}', s' \Vdash \phi$, for each $\phi \in \mathsf{L}_\mu[P, B]$.*

Submodels

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

▶ **Proposition 10.** *For each formula $\phi \in \mathsf{L}_\mu$, there exists a formula $\mathtt{tr}(\phi) \in \mathsf{L}_\mu$, containing a new propositional variable p, with the following property: for each model $\mathcal{M}$, each subset $S \subseteq |\mathcal{M}|$, and each $s \in |\mathcal{M}|$,*

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

Moreover, for each ordinal α, $\mathtt{tr}(\phi)_{\mathcal{M}[S/p]}^\alpha(\emptyset) = \phi_{\mathcal{M}_{\upharpoonright S}}^\alpha(\emptyset)$.

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

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

We introduce in this section the fragment $\mathcal{C}_{\aleph_1}(x)$ of the modal μ-calculus whose formulas, when interpreted as monotone functions of the variable x, are all $\aleph_1$-continuous. We show

how to construct a formula $\phi' \in \mathcal{C}_{\aleph_1}(x)$ from a given formula ϕ such that ϕ is κ-continuous, for some κ, if and only if ϕ and ϕ' are equivalent formulas. We argue therefore that the problem whether a formula is κ-continuous for some κ is decidable and, moreover, that there are no interesting notions of κ-continuity, for the modal μ-calculus, besides those for the cardinals $\aleph_0$ and $\aleph_1$.

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

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

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

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 [10], 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 achievement of [10] 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 $\mathcal{C}_{\aleph_0}(x)$.

Observe that the set of κ-continuous functions from $P(|\mathcal{M}|)^n$ to $P(|\mathcal{M}|)$, with $n \geqslant 1$, contains constants, projections, intersections and unions, as well as the unary functions $\phi_{\mathcal{M}}$ with $\phi = \langle a \rangle x$ for some $a \in \textit{Act}$. Moreover, this set is closed under composition and diagonalisation, and so Propositions 6 and 7 immediately yield the following result:

▶ **Proposition 11.** *Every formula in the fragment $\mathcal{C}_{\aleph_1}(X)$ is $\aleph_1$-continuous in X.*

4.1 Syntactic considerations

▶ **Definition 12.** 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 13.** A path in $G(\phi)$ is *bad* if one of its nodes corresponds to a subformula occurrence of the form $[a]\psi$. A bad cycle in $G(\phi)$ is a bad path starting and ending at the same vertex.

Recall that a path in a digraph is simple if it does not visit twice the same vertex. The rooted digraph $G(\phi)$ is a tree with back-edges; in particular, it has this property: for every node, there exists a unique simple path from the root to this node.
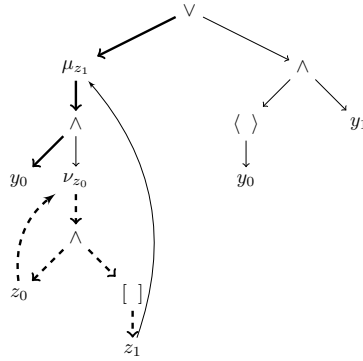
▶ **Definition 14.** We say that an occurrence of a free variable x of ϕ is
1. *bad* if there is a bad path in $G(\phi)$ from the root to it;
2. *not-so-bad* (or *boxed*) if the unique simple path in $G(\phi)$ from the root to it is bad;
3. *very bad* if it is bad and not boxed.

▶ **Example 15.** 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:

Figure 1 The digraph of a formula in L_μ.

- The free occurrence of z_1 in the digraph of $\nu_{z_0}.(z_0 \wedge [\ ]z_1)$ (in dashed) is bad but not-so-bad.
- The free occurrence of y_0 in the left branch of the digraph (in bold) is very bad. The other occurrence of y_0 is not bad.
- The unique free occurrence of y_1 in ϕ is not bad.

▶ **Lemma 16.** *For every set X of variables and every $\phi \in \mathsf{L}_\mu$, 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.*

4.2 The $\mathcal{C}_{\aleph_1}(x)$-flattening of formulas

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

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

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

▶ **Lemma 17.** *Let $\phi \in \mathsf{L}_\mu$. We have $\phi^{\sharp x} \cdot [x/\overline{x}] = \phi$.*

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

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

and henceforward we shorten it up to $\phi^\flat$.

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

▶ **Definition 18.** A formula ϕ is *almost-good* w.r.t. a set X of variables if no occurrence of a variable $x \in X$ is very bad. A formula ϕ is *almost-good* if it is almost-good w.r.t. $\{\, x \,\}$.

▶ **Lemma 19.** *If ϕ is an almost-good formula, then both $\phi^{\sharp x}$ and $\phi^{\flat}$ belong to $\mathcal{C}_{\aleph_1}(x)$.*

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

$$(\langle a\rangle\psi)^{\Box X} := \langle a\rangle(\psi)^{\Box X}, \qquad (\psi_1 @\, \psi_2)^{\Box X} := (\psi_1)^{\Box X} @\, (\psi_2)^{\Box X}, \quad \text{with } @ \in \{\, \wedge, \vee \,\},$$
$$(Q_z.\psi)^{\Box X} := \psi_0[\psi_1/\overline{z}]\,, \text{ where } \psi_0 := Q_z.\psi_2,\ \psi_1 := Q_{\overline{z}}.\psi_0\,,\ \psi_2 := (\psi^{\Box X \cup \{z\}})^{\sharp z},$$

with $Q \in \{\, \mu, \nu \,\}$. That is, in the last clause, ψ_2 is obtained from $\psi^{\Box X \cup \{z\}}$ by renaming all the boxed occurrences of z to $\overline{z}$. Observe that the first defining clause implies that

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

▶ **Proposition 20.** *The formula $\phi^{\Box X}$ is almost-good w.r.t. X and it is equivalent to ϕ.*

We can finally state the main result up to now.

▶ **Theorem 21.** *Every formula ϕ is equivalent to a formula ψ with $\psi^{\sharp x}$ and $\psi^{\flat}$ in $\mathcal{C}_{\aleph_1}(x)$.*

4.3   Comparing the closures of ϕ and $\phi^{\flat}$

We develop here some syntactic considerations allowing us to relate the closures of ϕ and $\phi^{\flat}$. In turn, this will make it possible to relate the positions of the games $\mathcal{G}(\mathcal{M}, \phi)$ and $\mathcal{G}(\mathcal{M}, \phi^{\flat})$, so to construct, in the proof of Proposition 24, a winning strategy in the latter game from a winning strategy in the former.

▶ **Lemma 22.** *If $\overline{x}$ is a free variable of ϕ and κ is either a variable not bound in ϕ or a constant, then*

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

In particular, we have

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

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

4.4   The continuous fragments

Our next goal is to prove some sort of converse to Proposition 11.

A pointed model $\langle \mathcal{M}, s \rangle$ is a *tree model* if the rooted digraph $\langle |\mathcal{M}|, \bigcup_{a \in Act} R_a, s \rangle$ is a tree. Let κ be a cardinal. A tree model $\langle \mathcal{M}, s \rangle$ is *κ-expanded* if, for each $a \in Act$, whenever xR_ax', there are at least κ a-successors of x that are bisimilar to x'. The following lemma is a straightforward generalization of [10, Proposition 1].

▶ **Lemma 23.** *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 24.** *If $\mathcal{M}, s \Vdash \phi$ and ϕ is κ-continuous in x, then $\mathcal{M}, s \Vdash \phi^\flat$.*

Proof. Suppose that $\mathcal{M} = (|\mathcal{M}|, \{\, R_a \mid a \in A \,\}, v)$ is a model and that $s_0 \Vdash \phi$. We want to prove that $s_0 \Vdash \phi^\flat$. Notice first that, by Lemma 23, we can assume that $\langle \mathcal{M}, s_0 \rangle$ is κ-expanded tree model.

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

In the following let $\mathcal{N} = \mathcal{M}[U/x]$ (notice that $\mathcal{N}$ is not anymore κ-expanded). Since $\mathcal{N}, s_0 \Vdash \phi$, let us fix a winning strategy for Eva in the game $\mathcal{G}(\mathcal{N}, \phi)$ from position (s_0, ϕ). We define next a strategy for Eva in the game $\mathcal{G}(\mathcal{N}, \phi^\flat)$ from position $(s_0, \phi^\flat)$. Observe first that, by Lemma 22, positions in $\mathcal{G}(\mathcal{N}, \phi)$ (respectively, $\mathcal{G}(\mathcal{N}, \phi^\flat)$) are of the form $(s, \psi[x/\overline{x}])$ (resp., $(s, \psi[\bot/\overline{x}])$) for a formula $\psi \in \mathrm{CL}(\phi^{\sharp x})$. Therefore, at the beginning of the play, Eva plays in $\mathcal{G}(\mathcal{N}, \phi^\flat)$ simulating the moves of the given winning strategy for the game $\mathcal{G}(\mathcal{N}, \phi)$. The simulation goes on until the play reaches a pair of positions $p = (s, [a]\chi \sigma^{\phi^{\sharp x}}_{[a]\chi} \cdot [x/\overline{x}])$ and $p' = (s, [a]\chi \sigma^{\phi^{\sharp x}}_{[a]\chi} \cdot [\bot/\overline{x}])$, for some subformula $[a]\chi$ of $\phi^{\sharp x}$, where $\chi = \chi'[\overline{x}/x]$ for some subformula χ' of ϕ.

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

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

▶ **Claim.** $s \in [\![[a]\psi]\!]_{\mathcal{N}}$ *implies* $s \in [\![[a]\psi[\bot/x]]\!]_{\mathcal{N}}$.

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 $\mathcal{N}, s' \Vdash \psi[\bot/x]$. To this goal, recalling that $\psi[\bot/x] \in \mathsf{L}_\mu[Prop \setminus \{x\}, Act]$ and using Proposition 9, it is enough to prove that $\langle \mathcal{N}, s' \rangle$ is $(Prop \setminus \{x\}, Act)$-bisimilar to some $\langle \mathcal{N}, s'' \rangle$ such that $\mathcal{N}, s'' \Vdash \psi[\bot/x]$.

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 $\mathcal{N}, s'' \Vdash \psi[\bot/x]$. Yet $\langle \mathcal{M}, s'' \rangle$ and $\langle \mathcal{M}, s' \rangle$ are bisimilar and since $\mathcal{N}$ is obtained from $\mathcal{M}$ just by modifying the value of the variable x, $\langle \mathcal{N}, s'' \rangle$ and $\langle \mathcal{N}, s' \rangle$ are $(Prop \setminus \{x\}, Act)$-bisimilar. As stated before, this together with $\mathcal{N}, s'' \Vdash \psi[\bot/x]$ imply that $\mathcal{N}, s' \Vdash \psi[\bot/x]$. ◀

To complete the proof of Proposition 24 we need to argue that the strategy so defined for Eva to play in the game $\mathcal{G}(\mathcal{M}, \phi^\flat)$ is winning. The only difficulty in asserting this is to exclude

the case where the initial simulation leads to a pair of positions of the form $(s, \overline{x}[x/\overline{x}])$ and $(s, \overline{x}[\perp/\overline{x}])$. This is however excluded since in $\phi^{\sharp x}$ all the occurrences of $\overline{x}$ are boxed, so we are enforced to go through the second step of the strategy. ◄

▶ **Proposition 25.** *If, for some regular cardinal κ, $\phi \in \mathsf{L}_\mu$ is κ-continuous, then ϕ is equivalent to $\phi^\flat$.*

Proof. Notice that, by monotonicity in the variable x, $\phi^\flat \to \phi$ is a tautology. Proposition 24 exhibits the converse implication as another tautology. ◄

▶ **Theorem 26.** *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 21, ϕ is equivalent to a formula ψ with $\psi^\flat \in \mathcal{C}_{\aleph_1}(x)$. Clearly, ψ is κ-continuous as well, so it is equivalent to $\psi^\flat$ by Proposition 25. It follows that ϕ is equivalent to $\psi^\flat \in \mathcal{C}_{\aleph_1}(x)$. ◄

As a consequence of the previous Theorem 26, we obtain the following result.

▶ **Theorem 27.** *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 28.** *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 $\phi' = (\phi^{\square x})^\flat$. It is then enough to observe that there are effective processes to construct the formula ϕ' and to check whether ϕ is equivalent to ϕ'. ◄

5   Large closure ordinals

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

▶ **Definition 29.** 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 $\mathsf{cl}(\phi) = \alpha$) if, for each model $\mathcal{M}$, the function $\phi_\mathcal{M}$ converges to its least fixed-point in at most α steps, and there exists a model $\mathcal{M}$ in which $\phi_\mathcal{M}$ converges to its least fixed-point in exactly α steps.

▶ **Lemma 30.** *If α is a closure ordinal, then there is a formula $\phi(x)$ such that $\mathsf{cl}(\phi(x)) = \alpha$ and that is* total, *meaning that $[\![\mu_x.\phi(x)]\!]_\mathcal{M} = |\mathcal{M}|$, for each model $\mathcal{M}$.*

▶ **Proposition 31.** *If a formula $\phi(x)$ belongs to the syntactic fragment $\mathcal{C}_{\aleph_1}(x)$, then it has a closure ordinal $\mathsf{cl}(\phi(x))$ and ω_1 is an upper bound for $\mathsf{cl}(\phi(x))$.*

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 5 that $\phi_\mathcal{M}$ converges to its least fixed-point in at most ω_1 steps. ◄

5.1 ω_1 is a closure ordinal

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

$$\Phi(x) := \nu_z.(\langle v \rangle x \wedge \langle h \rangle z) \vee [v]\bot \,. \tag{1}$$

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

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

By Proposition 31, the formula $\Phi(x)$ has a closure ordinal and $\mathsf{cl}(\Phi(x)) \leqslant \omega_1$. In order to prove that $\mathsf{cl}(\Phi(x)) = \omega_1$, we are going to construct a model $\mathcal{M}_{\omega_1}$ where $\Phi^{\omega_1}_{\mathcal{M}_{\omega_1}}(\emptyset) \not\subseteq \Phi^{\alpha}_{\mathcal{M}_{\omega_1}}(\emptyset)$ for each $\alpha < \omega_1$.

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

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

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

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

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

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

▶ **Theorem 32.** *The closure ordinal of $\Phi(x)$ is ω_1.*

5.2 From a bimodal language to a monomodal language

The following statement generalizes to the modal μ-calculus a well known coding of poly-modal logic to monomodal logic, see [29] and [18, Section 4].

▶ **Proposition 33.** *For each bimodal formula ϕ of the modal μ-calculus, we construct a monomodal formula $\phi^{\mathtt{sim}}$; if ϕ belongs to $\mathcal{C}_{\aleph_1}(x)$, then so does $\phi^{\mathtt{sim}}$. Moreover, for each bimodal model $\mathcal{M}$ we can also construct a monomodal model $\mathcal{M}^{sim}$, together with an injective function $(-)^\circ : |\mathcal{M}| \to |\mathcal{M}^{sim}|$ such that, for each $s \in |\mathcal{M}|$, $\mathcal{M}, s \Vdash \phi$ if and only if $M^{sim}, s^\circ \Vdash \phi^{\mathtt{sim}}$.*

▶ **Theorem 34.** *The monomodal formula $\Phi^{\mathtt{sim}}$ has closure ordinal ω_1.*

Proof. Since the translation $\phi \mapsto \phi^{\mathtt{sim}}$ sends formulas in $\mathcal{C}_{\aleph_1}(x)$ to formulas in $\mathcal{C}_{\aleph_1}(x)$, $\Phi^{\mathtt{sim}}$ is $\aleph_1$-continuous and therefore it has a closure ordinal bounded by ω_1. To argue that the closure ordinal of $\Phi^{\mathtt{sim}}$ is equal to ω_1 it is enough to consider the model $\mathcal{M}^{sim}_{\omega_1}$ and rely on Proposition 33. ◀

5.3   Closure under ordinal sum

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

▶ **Theorem 35.** *Suppose $\phi_0(x)$ and $\phi_1(x)$ are monomodal formulas that have, respectively, α and β as closure ordinals. For a variable p occurring neither in ϕ_0 nor in ϕ_1, for $\chi := \chi_0 \wedge \chi_1$ with $\chi_0 = \neg p \to ([\;\;]\neg p \wedge (\neg p \wedge \mu_z.\phi_0(z)))$ and $\chi_1 := p \to ([\;\;](\neg p \to \mu_z.\phi_0(z)) \wedge \mu_z.\mathtt{tr}(\phi_1(z)))$, and for*

$$\psi(x) := (\neg p \wedge \phi_0(x)) \vee (\mathtt{tr}(\phi_1)(x) \wedge [\;](\neg p \to x)),$$

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

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

We continue by understanding how $\psi_{\mathcal{N}}$ acts on an acceptable model $\mathcal{N}$. To this goal, let $\mathcal{N}_0$ and $\mathcal{N}_1$ be the submodels of $\mathcal{N}$ induced by $v(\neg p)$ and $v(p)$, respectively. To ease the reading, let also $N_0 := v(\neg p)$, $N_1 := v(p)$, $\phi_{\mathcal{N}_0} := \phi_{0\mathcal{N}_0}$ and $\phi_{\mathcal{N}_1} := \phi_{1\mathcal{N}_1}$, so $\phi_{\mathcal{N}_0} : P(N_0) \to P(N_0)$ and $\phi_{\mathcal{N}_1} : P(N_1) \to P(N_1)$. Observe that since $\mathcal{N}, s \models \neg p \to [\;]\neg p$ for every $s \in |\mathcal{N}|$, N_0 is a closed subset of $|\mathcal{N}|$. Then, by Proposition 10, we have $\psi_{\mathcal{N}}(X) \cap N_0 = \phi_{\mathcal{N}_0}(X \cap N_0)$ and $\mathtt{tr}(\phi_1)_{\mathcal{N}_1}(X) = \phi_{\mathcal{N}_1}(X \cap N_1)$ for each $X \subseteq |\mathcal{N}|$. Now let $\nabla(X) := N_1 \cap [\;]_{\mathcal{N}}(N_0 \to X)$. We consider that the domain of ∇ is $P(N_0)$ while its codomain is $P(N_1)$. Therefore, $\psi_{\mathcal{N}}$ is of the form

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

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

$$\nabla(X) = N_1, \quad \text{whenever } X \supseteq \phi^{\alpha}_{\mathcal{N}_0}(\emptyset)(= N_0). \tag{3}$$

▶ **Proposition 36.** *On every acceptable model $\mathcal{N}$ the equality $\psi^{\alpha+\beta}_{\mathcal{N}}(\emptyset) = |\mathcal{N}|$ holds and, consequently, the formula $\psi(x)$ converges before $\alpha + \beta$ steps.*

Proof. Since N_0 is a closed subset of $|\mathcal{N}|$, by Proposition 10, we have

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

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

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

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

◄

▶ **Proposition 37.** *There exists an acceptable model $\mathcal{N}$ on which $\psi(x)$ converges exactly after $\alpha + \beta$ steps.*

Proof. Since the formulas $\phi_0(x)$ and $\phi_1(x)$ have, respectively, α and β as closure ordinals, by Proposition 30 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 (5) is actually an equality. But then we apply equations (4) and (5) to obtain $\psi_{\mathcal{N}}^{\alpha}(\emptyset) = \phi_{\mathcal{N}_0}^{\alpha}(\emptyset) \neq \phi_{\mathcal{N}_0}^{\delta}(\emptyset) = \psi_{\mathcal{N}}^{\delta}(\emptyset)$ and $\psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset) = N_0 \cup \phi_{\mathcal{N}_1}^{\gamma}(\emptyset)$, for ordinals $\delta < \alpha$ and γ. Finally, $\psi_{\mathcal{N}}^{\alpha+\beta}(\emptyset) = |\mathcal{N}| = N_0 \cup \phi_{\mathcal{N}_1}^{\beta}(\emptyset) \neq N_0 \cup \phi_{\mathcal{N}_1}^{\gamma}(\emptyset) = \psi_{\mathcal{N}}^{\alpha+\gamma}(\emptyset)$, for $\gamma < \beta$. ◄

──── **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 Ronchi Della Rocca [23], pages 30–44. URL: `http://www.dagstuhl.de/dagpub/978-3-939897-60-6`, `doi:10.4230/LIPIcs.CSL.2013.30`.
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, December 2009. URL: `https://hal.archives-ouvertes.fr/hal-00396431`.
4   D. Berwanger, E. Grädel, and G. Lenzi. On the variable hierarchy of the modal μ-calculus. In J. Bradfield, editor, *CSL 2002, Proceedings*, pages 352–366, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg. `doi:10.1007/3-540-45793-3_24`.
5   M. Bojanczyk, C. Dittmann, and S. Kreutzer. Decomposition theorems and model-checking for the modal μ-calculus. In *CSL-LICS'14*, pages 17:1–17:10, New York, NY, USA, 2014. ACM. `doi:10.1145/2603088.2603144`.
6   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.
7   J. C. Bradfield. The modal mu-calculus alternation hierarchy is strict. *Theoretical Computer Science*, 195(2):133–153, 1998. `doi:10.1016/S0304-3975(97)00217-X`.
8   R. Cousot and P. Cousot. Constructive versions of Tarski's fixed point theorems. *Pacific Journal of Mathematics*, 82(1):43–57, 1979.
9   M. Czarnecki. How fast can the fixpoints in modal μ-calculus be reached? In L. Santocanale, editor, *7th Workshop on Fixed Points in Computer Science, FICS 2010*, page 89, Brno, Czech Republic, August 2010. Available from `Hal`: `https://hal.archives-ouvertes.fr/hal-00512377`.
10 G. Fontaine. Continuous fragment of the mu-calculus. In M. Kaminski and S. Martini, editors, *CSL 2008. Proceedings*, volume 5213 of *Lecture Notes in Computer Science*, pages 139–153. Springer, 2008. `doi:10.1007/978-3-540-87531-4_12`.

11 J. Fortier and L. Santocanale. Cuts for circular proofs: semantics and cut-elimination. In Ronchi Della Rocca [23], pages 248–262. URL: `http://www.dagstuhl.de/dagpub/978-3-939897-60-6`, `doi:10.4230/LIPIcs.CSL.2013.248`.

12 V. Vaccaro G. Sambin. A new proof of Sahlqvist's theorem on modal definability and completeness. *Journal of Symbolic Logic*, 54:992–999, 1989.

13 M. J. Gouveia and L. Santocanale. ℵ₁ and the modal μ-calculus. Preprint, available from `Hal`: `https://hal.archives-ouvertes.fr/hal-01503091`, March 2017.

14 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*, pages 263–277, Berlin, Heidelberg, 1996. Springer Berlin Heidelberg. `doi:10.1007/3-540-61604-7_60`.

15 M. Jurdziński. Deciding the winner in parity games is in UP ∩ co-UP. *Information Processing Letters*, 68(3):119–124, 1998. `doi:10.1016/S0020-0190(98)00150-1`.

16 B. Jónsson. On the canonicity of sahlqvist identities. *Studia Logica*, 53(4):473–491, 1994. URL: `http://www.jstor.org/stable/20015747`.

17 D. Kozen. Results on the propositional mu-calculus. *Theoretical Computer Science*, 27:333–354, 1983. `doi:10.1016/0304-3975(82)90125-6`.

18 M. Kracht and F. Wolter. Simulation and transfer results in modal logic – A survey. *Studia Logica*, 59(1):149–177, 1997. `doi:10.1023/A:1004900300438`.

19 J.-L. Lassez, V. L. Nguyen, and L. Sonenberg. Fixed point theorems and semantics: A folk tale. *Inf. Process. Lett.*, 14(3):112–116, 1982. `doi:10.1016/0020-0190(82)90065-5`.

20 G. Lenzi. The modal μ-calculus: a survey. *Task Quarterly*, 9(3):29–316, 2005.

21 D. Niwiński and I. Walukiewicz. Games for the μ-calculus. *Theoretical Computer Science*, 163(1):99–116, 1996. `doi:10.1016/0304-3975(95)00136-0`.

22 J. Obdržálek. Fast mu-calculus model checking when tree-width is bounded. In W. A. Hunt and F. Somenzi, editors, *CAV 2003, Proceedings*, pages 80–92, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. `doi:10.1007/978-3-540-45069-6_7`.

23 S. Ronchi Della Rocca, editor. *Computer Science Logic 2013 (CSL 2013), CSL 2013, September 2-5, 2013, Torino, Italy*, volume 23 of *LIPIcs*. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2013. URL: `http://www.dagstuhl.de/dagpub/978-3-939897-60-6`.

24 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. `doi:10.1016/S0049-237X(08)70728-6`.

25 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`. `doi:10.1051/ita:2002010`.

26 L. Santocanale. Completions of μ-algebras. *Ann. Pure Appl. Logic*, 154(1):27–50, 2008. `doi:10.1016/j.apal.2007.11.001`.

27 T. Studer. On the proof theory of the modal mu-calculus. *Studia Logica*, 89(3):343–363, 2008.

28 A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5:285–309, 1955.

29 S. K. Thomason. Independent propositional modal logics. *Studia Logica*, 39(2):143–144, 1980. `doi:10.1007/BF00370317`.

30 I. Walukiewicz. Completeness of Kozen's axiomatisation of the propositional μ-calculus. *Information and Computation*, 157(1):142–182, 2000. `doi:10.1006/inco.1999.2836`.