Completions of μ-algebras*

Luigi Santocanale
LIF/CMI, CNRS, Université de Provence,
39, rue Joliot-Curie - F-13453 Marseille Cedex 13
E-mail: luigi.santocanale@lif.univ-mrs.fr

Abstract

We define the class of algebraic models of μ-calculi and study whether every such model can be embedded into a model which is a complete lattice. We show that this is false in the general case and focus then on free modal μ-algebras, i.e. Lindenbaum algebras of the propositional modal μ-calculus. We prove the following fact: the MacNeille-Dedekind completion of a free modal μ-algebra is a complete modal algebra, hence a modal μ-algebra (i.e. an algebraic model of the propositional modal μ-calculus). The canonical embedding of the free modal μ-algebra into its Dedekind-MacNeille completion preserves the interpretation of all the terms in the class $Comp(\Sigma_1, \Pi_1)$ of the alternation-depth hierarchy.

The proof uses algebraic techniques only and does not directly rely on previous work on the completeness of the modal μ-calculus.

1 Overview

When L is a complete lattice, the least fixed-point $\mu_x.f$ of a monotone function $f : L \longrightarrow L$ enjoys a remarkable property. We like to say that the least fixed-point is *constructive*: the equality

$$\mu_x.f = \bigvee_{\alpha \in Ord} f^\alpha(\bot) \tag{1}$$

holds and provides a method to construct $\mu_x.f$ from the bottom of the lattice. The expressions $f^\alpha(\bot)$, indexed by ordinals, are commonly called the *approximants* of $\mu_x.f$. They are defined by transfinite induction as expected: $f^0(\bot) = \bot$, $f^{\alpha+1} = f(f^\alpha(\bot))$, and $f^\alpha(\bot) = \bigvee_{\beta<\alpha} f^\beta(\bot)$.

A careful reading of Tarski's fixpoint theorem [24] reveals that the completeness assumption is not needed for f to have a fixed-point: if L is merely a poset, existence of a least prefixed-point implies the existence of a least fixed-point. A least prefixed-point of f is an element $\mu_x.f \in L$ satisfying

$$f(\,\mu_x.f\,) \leq \mu_x.f \tag{2}$$

$$f(y) \leq y \qquad \Rightarrow \qquad \mu_x.f \leq y \tag{3}$$

These two properties provide a natural way to axiomatize by equations and equational implications least fixed-points whenever a definable order relation is given. The equational implications (2) and (3) – the latter known as the Park induction rule [18, 6] – have often been used to axiomatize concrete objects where implicit or explicit fixed-points are at work: relational algebras with transitive closures [16], regular languages [11], powersets of Kripke frames [22, 10]. Model theoretic considerations also induce to prefer models of theories axiomatized by equational implications to classes of models that are complete lattices: the formers build up a quasivariety, free models exist, etc. A primary goal of this paper is to clarify the relationships between standard models, the complete lattices, and algebraic models. Rephrasing our goal, we want to compare the constructive presentation of the least fixed-point (1) and its logical presentation by means of (2) and (3).

1.1 μ-theories and μ-algebras

The following is a generic logical framework within which to develop a theory of ordered algebras with least fixed-point operators. Analogous frameworks [3, 1] can be coded within this framework.

Definition 1.1 A *μ-theory* is a first order theory with the following properties:

- it is an extension of the theory of bounded lattices,
- it comes with fixed-point pairs, that is, pairs of terms $(f, \mu_x.f)$ axiomatized by (2) and (3) so that (the interpretation of) f is an order preserving operation in the variable x, and (the interpretation of) $\mu_x.f$ is a least prefixed-point of f,

*Research developed with the support of the ACI Géométrie du Calcul (GEOCAL).

- its axioms are either equations or equational implications.

A μ-*algebra* is model of a fixed μ-theory. A μ-algebra is *complete* if its lattice reduct is a complete lattice.

We shall investigate relations between μ-algebras and complete μ-algebras, with the particular goal of understanding when a μ-algebra embeds into a complete one. Contrarily to what happens for several algebraic structures related to logic (Boolean algebras, modal algebras K, Heyting algebras, quantales), we show next that this is not always possible for μ-algebras.

1.2 μ-algebras are not completable

Let us introduce some notation first. For a fixed-point pair $(f, \mu_x.f)$ and a μ-algebra $\mathcal{A}$, we consider the interpretation of f on $\mathcal{A}$, $f : \mathcal{A}^{\{x\}\cup Y} \longrightarrow \mathcal{A}$. For $v \in \mathcal{A}^Y$ fixed, we use the notation $f_v : \mathcal{A} \longrightarrow \mathcal{A}$ for the monotone function that arises if we evaluate all the variables in Y to v, and consistently we denote its least prefixed-point by $\mu_x.f_v$. We say then that $f_v : \mathcal{A} \longrightarrow \mathcal{A}$ is a fixed-point polynomial. To simplify the notation, we shall also omit the subscript v and say that $f : \mathcal{A} \longrightarrow \mathcal{A}$ is a fixed-point polynomial.

Example 1.2 Choose a μ-algebra $\mathcal{A}$ and a fixed-point polynomial $f : \mathcal{A} \longrightarrow \mathcal{A}$ for which the chain of finite approximants

$$\perp < f(\perp) < f^2(\perp) < \ldots < f^n(\perp) < \ldots$$

is infinite. Define the infinite sequences ϕ_n by

$$\phi_n = (\underbrace{\perp, \ldots, \perp}_{n-\text{times}}, \perp, f(\perp), f^2(\perp), \ldots\,)\,, \quad n \geq 0\,,$$

and consider them as elements of the product algebra $\mathcal{A}^\omega$. Since f is computed pointwise, observe that $f(\phi_n)$ is equal to ϕ_{n-1} for all but a finite number of coordinates.

Define the equivalence relation $\sim$ on $\mathcal{A}$ by saying that two infinite sequences are equivalent if they coincide in all but a finite number of coordinates. The quotient $\mathcal{A}^\omega/\sim$ is a reduced product of $\mathcal{A}$ and all the equations and equational implications that hold in $\mathcal{A}$ hold in $\mathcal{A}^\omega/\sim$ as well, cf. [4, chapter 6]. In particular, $\mathcal{A}^\omega/\sim$ is also a μ-algebra.

Denote by $\bar{\phi}_n$ the equivalence class of ϕ_n and recall that in $\mathcal{A}^\omega/\sim$ the least fixed-point $\mu_x.f$ is the equivalence class of the infinite sequence with constant value $\mu_x.f$. The relations

$$f(\bar{\phi}_n) \leq \bar{\phi}_{n-1}, \quad n \geq 1 \qquad\qquad \mu_x.f \not\leq \bar{\phi}_0 \qquad (4)$$

hold[1] in $\mathcal{A}^\omega/\sim$ and we claim that such configuration is not compatible with $\mathcal{A}^\omega/\sim$ being complete. If $\bigwedge_{n\geq 0} \bar{\phi}_n$ exists then

$$f(\bigwedge_{n\geq 0} \bar{\phi}_n) \;\leq\; f(\bar{\phi}_{n+1}) \;\leq \bar{\phi}_n\,,$$

for all $n \geq 0$, and therefore $f(\bigwedge_{n\geq 0} \bar{\phi}_n) \leq \bigwedge_{n\geq 0} \bar{\phi}_n$. Then $\mu_x.f \leq \bigwedge_{n\geq 0} \bar{\phi}_n \leq \bar{\phi}_0$ gives a contradiction.

Finally observe that such a configuration is preserved by any extension of $\mathcal{A}^\omega/\sim$, and therefore this μ-algebra has no complete extension. □

[1] Actually, the stronger relations $f(\bar{\phi}_n) = \bar{\phi}_{n-1}$ hold in $\mathcal{A}^\omega/\sim$.

We say that a μ-theory (or the quasivariety of the μ-algebras) is non-trivial if we can find a μ-algebra $\mathcal{A}$ and a fixed point polynomial f for which its finite approximants are all distinct. If this is not possible, then for each fixed-point pair $(f, \mu_x.f)$ some equation of the form $\mu_x.f = f^n(\perp)$ holds, showing that all the least fixed-point are superflous. We collect these observations in a theorem.

Theorem 1.3 Any non-trivial quasivariety of μ-algebras contains a μ-algebra which does not admit an embedding into a complete μ-algebra.

It is possible to show that, for some simple μ-theory, if a μ-algebra has no configuration such as (4), then the principal filter embedding is a morphism of μ-algebras. Unfortunately, the principal filter embedding becomes soon useless, for example for μ-theories where greatest fixed-points are also considered. Therefore, we shall look at other conditions ensuring that a μ-algebra has an embedding into a complete μ-algebra.

1.3 The propositional modal μ-calculus: free modal μ-algebras

Recall that a free μ-algebra embeds into a complete one if and only if the class of complete μ-algebras generates the class of all μ-algebras. If we adopt the perspective of algebraic logic, the statement that free μ-algebras embed into complete ones amounts to a completeness theorem for the logic with respect to the semantics of all complete models.

It is often the case that free μ-algebras embed into complete ones, for example this is the case for free μ-lattices [19] and free modal μ-algebras (i.e. Lindenbaum algebras for the propositional modal μ-calculus [10]). The rest of this paper will be concerned with studying *free modal μ-algebras*. Their μ-theory, i.e. the theory of modal μ-algebras, essentially is nothing else but the propositional modal μ-calculus. The terms of the theory are generated according to the grammar:

$$t = p \,|\, x \,|\, \top \,|\, t_1 \wedge t_2 \,|\, \neg t \,|\, \langle\sigma\rangle t \,|\, \mu_x.t\,,$$

where σ ranges on a finite set of actions Act and the fixed-point generation rule only applies when the variable x occurs under an even number of negations. The reader has surely recognized the framework of multimodal algebras, in addition to which, we have least fixed-points. Accordingly, the axioms of the theory are those of multimodal algebras K as well as (2) and (3) for the fixed point pairs $(t, \mu_x.t)$. In the grammar we have distinguished a generator p from a variable x. This will be useful when considering the interpretation of terms as operations on free modal μ-algebras, where the generators become operations. This kind of term generation is standard from fixed-point theory [17], but it is also possible to code these terms as terms generated from an infinite signature using substitution only [15]. Finally, it can be shown that modal μ-algebras form a variety of algebras [21].

The completeness results for the propositional modal μ-calculus [10, 26] paired with the small Kripke model property [23] imply that *a free modal μ-algebra has an embedding into* an infinite product of finite modal μ-algebras. This infinite product is of course *a complete lattice*. In the rest of the paper we shall prove a weaker embedding result concerning Σ_1-terms, defined by the grammar

$$t \;=\; x \,|\, p \,|\, \neg p \,|\, \top \,|\, t \wedge t \,|\, \bot \,|\, t \vee t \,|\, \langle \sigma \rangle t \,|\, [\sigma] t \,|\, \mu_x.t \,,$$[2]

and Σ_1-operations – we say that $f : \mathcal{A}^X \longrightarrow \mathcal{A}$ is a Σ_1-operation if it is the interpretation of a Σ_1-term. The class of Π_1-terms is defined as above with the exception that least fixed-point formation is replaced by greatest fixed-point formation.[3] The class of $Comp(\Sigma_1, \Pi_1)$-operations is obtained by closing under substitution the union of Σ_1 and Π_1. Our result can be stated as follows:

Theorem 1.4 Let $\mathcal{F}$ be a free modal μ-algebra. There exists a complete modal algebra $\overline{\mathcal{F}}$ and an injective morphism of Boolean modal algebras $i : \mathcal{F} \longrightarrow \overline{\mathcal{F}}$ which preserves all the $Comp(\Sigma_1, \Pi_1)$-operations of the algebra $\mathcal{F}$.

With respect to [26], where algorithmic and game-theoretic ideas as well as tableaux manipulations are the main tools, we shall use use purely algebraic and order theoretic tools. Under some respect, our work can be understood as an effort to translate ideas from [10, 26] into an algebraic and order theoretic framework. Our work, in the spirit of algebraic logic, has also been motivated by our insuccess to give a proper account of tableaux for the propositional modal μ-calculus from a proof-theoretic perspective. For example, we have not been able to adapt interesting proof-theoretic ideas [13] to the alternation-free fragment of the modal μ-calculus.

We sketch in the rest of the section the strategy followed to prove Theorem 1.4. The algebra $\overline{\mathcal{F}}$ is the MacNeille-Dedekind completion[4]of $\mathcal{F}$. For our goals, we recall that if L is a Boolean algebra, then $\overline{L}$ is a Boolean algebra as well, cf. [2]. Also, we need the following statement:

Lemma 1.5 Let L be a lattice and $\overline{L}$ be its MacNeille-Dedekind completion. A left adjoint[5] $f : L \longrightarrow L$ has an extension – necessarily unique – to a left adjoint $f^\vee : \overline{L} \longrightarrow \overline{L}$.

Using the notation of [8], if g is right adjoint to f, then $g^\wedge$ is right adjoint to $f^\vee$. A first step towards our main result will be to prove:

Claim 1.6 The modal operators $\langle \sigma \rangle$ of free modal μ-algebra are left adjoints.

Using Lemma 1.5 and the Claim we can state:

Proposition 1.7 The MacNeille-Dedekind completion of a free modal μ-algebra is a multi-modal algebra K and the principal ideal embedding $i : \mathcal{F} \longrightarrow \overline{\mathcal{F}}$ is a morphism of multi-modal algebras.[6]

Since $\overline{\mathcal{F}}$ is a complete lattice, it is a complete modal μ-algebra, and therefore we are also interested in preservation of fixed-points. To this goal we shall use the following lemma, which is a standard consequence of the embedding i being continuous (and cocontinuous).

Lemma 1.8 Let $\mathcal{A}$ be a μ-algebra, $i : \mathcal{A} \longrightarrow \overline{\mathcal{A}}$ its Dedekind-MacNeille completion, and f_v a fixed-point polynomial. If f_v is preserved by i and $\mu_x.f_v$ is constructive, then the least fixed-point $\mu_x.f_v$ is preserved by i.

We shall prove that all the Σ_1-operations are preserved by showing that all these functions are constructive:

Claim 1.9 Any fixed-point Σ_1-polynomial $f_v : \mathcal{F} \longrightarrow \mathcal{F}$ over a free modal μ-algebra satisfies the relation

$$\mu_x.f_v = \bigvee_{n \geq 0} f_v^n(\bot)\,. \tag{5}$$

The Claim and Lemma 1.8 imply that each Σ_1-operation on a free modal μ-algebra is preserved. By duality, the same holds for Π_1-operations and, consequently, all the operations in the class $Comp(\Sigma_1, \Pi_1)$ are preserved.

2 A last rule for free modal μ-algebras

In this section we prove that free modal μ-algebras enjoy a property that could be called a last rule in the language of proof theory. Usually, such a property is a consequence

[2]The fixed-point formation rule is no longer constrained here.

[3]By duality, greatest fixed-points are definable in the given signature.

[4]An elegant and simple introduction to the MacNeille-Dedekind completion of a lattice can be found in [14].

[5]Recall that an order preserving $f : L \longrightarrow M$ is a left adjoint if there exists $g : M \longrightarrow L$ (the right adjoint) such that $f(x) \leq y$ if and only if $x \leq g(y)$, for all $x \in L$ and $y \in M$.

[6]The same statement holds if we replace "free modal μ-algebra" with "free multi-modal algebra K".

of a cut-elimination theorem. The property is analogous to Whitman's condition for free lattices, cf. [27, 7]. Since the property is the starting point of the path that leads to prove Claims 1.6 and 1.9, we shall detail its proof, which is it analogous to Day's proof of Whitman conditions and Freyd's covering of a topos.

We briefly recall the universal property of a modal μ-algebra $\mathcal{F}_P$, freely generated by a set P. Such an algebra comes with a function $j : P \longrightarrow \mathcal{F}_P$ such that for each $(f, \mathcal{A})$ where $\mathcal{A}$ is a modal μ-algebra and $f : P \longrightarrow \mathcal{A}$ there exists a unique μ-algebra morphism $\tilde{f} : \mathcal{F}_P \longrightarrow \mathcal{A}$ such that $f = \tilde{f} \circ j$. A generator in $\mathcal{F}_P$ is of the form $j(p)$ for some $p \in P$. It can be easily shown that j is injective (cf. the end of this section) and in the following we shall abuse notation and identify P with $j(P)$.

Theorem 2.1 Let $\mathcal{F}$ be a free modal μ-algebra and Λ be a finite set of literals (generators or negated generators). The following implication holds in $\mathcal{F}$: if

$$\bigwedge \Lambda \wedge \bigwedge_{\sigma \in \Sigma} (\, [\sigma] x_\sigma \wedge \bigwedge_{y \in Y_\sigma} \langle \sigma \rangle y \,) \leq \bot \,,$$

then either $p, \neg p \in \Lambda$ for some generator p, or $x_\sigma \wedge y \leq \bot$ for some $\sigma \in \Sigma$ and $y \in Y_\sigma$.

We prove first that:

Proposition 2.2 Let $\mathcal{F}$ be a free modal algebra. The implication

$$\begin{gathered} \bigwedge \Lambda \;\wedge\; \bigwedge_{\sigma \in \Sigma} (\, [\sigma] \bigvee Y_\sigma \;\wedge\; \bigwedge_{y \in Y_\sigma} \langle \sigma \rangle y \,) \leq \bot \\ \text{implies} \\ \bigwedge \Lambda \leq \bot \text{ or } \exists \sigma \in \Sigma, y \in Y_\sigma \text{ s.t. } y \leq \bot \end{gathered} \tag{6}$$

holds in $\mathcal{F}$, where Λ is a finite set of literals, $\Sigma \subseteq Act$, and, for each $\sigma \in \Sigma$, Y_σ is a finite possibly empty set of elements of $\mathcal{F}$.

Proof. Let $\mathcal{A}$ be any modal algebra and suppose that for each $\sigma \in \Sigma$ we are given a set Y_σ such that $y \not\leq \bot$ for each $y \in Y_\sigma$. For each $\sigma \in \Sigma$ and $y \in Y_\sigma$ let $\chi_\sigma^y : \mathcal{A} \longrightarrow 2$ be morphism of Boolean algebras such that $\chi_\sigma^y(y) = \top$. Define

$$\chi_\sigma(z) = \begin{cases} \bigvee_{y \in Y_\sigma} \chi_\sigma^y(z) \,, & \sigma \in \Sigma \,, \\ \bot \,, & \sigma \notin \Sigma \,. \end{cases}$$

For $\sigma \in \Sigma$, observe that $\chi_\sigma(z) = \bot$ if Y_σ is empty and otherwise that $\chi_\sigma(z) = \top$ if and only if $\chi_\sigma^y(z) = \top$ for some $y \in Y_\sigma$.

We define a modal algebra structure on the product Boolean algebra $\mathcal{A} \times 2$. The modal operators are defined as:

$$\langle \sigma \rangle (z, w) = (\langle \sigma \rangle z, \chi_\sigma(z)) \,,$$

and they are easily seen to be normal, since the χ_σ preserve joins. Also, observe that the first projection $\mathtt{pr}_{\mathcal{A}} : \mathcal{A} \times 2 \longrightarrow \mathcal{A}$ is a morphism of modal algebras.

Suppose now that $\mathcal{A}$ is freely generated by a set P, $\mathcal{A} = \mathcal{F}_P$, and let Λ be a set of literals such that $\bigwedge \Lambda \not\leq \bot$. Choose a function $f : P \longrightarrow \mathcal{A} \times 2$ with these properties: (i) $f(p) \in \{\, (p, \bot), (p, \top) \,\}$ for each $p \in P$, (ii) $f(p) = (p, \top)$ if $p \in \Lambda$ and $f(p) = (p, \bot)$ if $\neg p \in \Lambda$ (clearly such a function exists).

Let $\tilde{f} : \mathcal{F}_P \longrightarrow \mathcal{F}_P \times 2$ be the extension of f to a modal-algebra homomorphism, observe that $\mathtt{pr}_{\mathcal{A}} \circ \tilde{f} = \mathtt{id}_{\mathcal{F}_P}$, since this relation holds on generators, and that $\tilde{f}(l) = (l, \top)$ for $l \in \Lambda$. Suppose that

$$\bigwedge \Lambda \wedge \bigwedge_{\sigma \in \Sigma} ([\sigma] \bigvee Y_\sigma \wedge \bigwedge_{y \in Y_\sigma} \langle \sigma \rangle y) \leq \bot \,.$$

Apply the morphism $\tilde{f}$ to the above expression to obtain

$$\begin{aligned} (\bigwedge \Lambda \wedge \bigwedge_{\sigma \in \Sigma} ([\sigma] \bigvee Y_\sigma \wedge \bigwedge_{y \in Y_\sigma} \langle \sigma \rangle y) \,, \, a \wedge \bigwedge_{\sigma \in \Sigma} (b^\sigma \wedge c^\sigma) \,) \\ \leq (\bot, \bot) \,, \end{aligned}$$

where

$$\begin{aligned} a &= \bigwedge_{l \in \Lambda} \mathtt{pr}_2(\tilde{f}(l)) = \bigwedge_{l \in \Lambda} \top \; = \top \,, \quad \text{since } \tilde{f}(l) = (l, \top) \\ b^\sigma &= \neg \chi_\sigma(\neg \bigvee Y_\sigma) = \top \,, \end{aligned}$$

– this relation is trivial if Y_σ is empty, and otherwise note that $\chi_\sigma(\neg \bigvee Y_\sigma) = \top$ iff $\chi_\sigma^y(\neg \bigvee Y_\sigma) = \top$ for some $y \in Y_\sigma$, which cannot be because of $\bot = \chi_\sigma^y(y \wedge \neg \bigvee Y_\sigma) = \chi_\sigma^y(y) \wedge \chi_\sigma^y(\neg \bigvee Y_\sigma) = \top$ – and finally

$$c^\sigma = \bigwedge_{y \in Y_\sigma} \chi_\sigma(y) = \bigwedge_{y \in Y_\sigma} \top = \top \,.$$

We obtain $a \wedge \bigwedge_{\sigma \in \Sigma} b^\sigma \wedge c^\sigma = \top$ which contradicts $a \wedge \bigwedge_{\sigma \in \Sigma} b^\sigma \wedge c^\sigma \leq \bot$. ☐

We extend now the previous result from modal algebras to modal μ-algebras.

Proposition 2.3 The implication (6) holds in a free modal μ-algebra.

Proof. The proposition follows since in the previous proof, if $\mathcal{A}$ is a modal μ-algebra, then $\mathcal{A} \times 2$ is a modal μ-algebra and the first projection is a morphism of modal μ-algebras. This can be seen as follows: suppose that we have defined the interpretation of a term f in the algebra $\mathcal{A} \times 2$ so that the first projection preserves the interpretation. This is equivalent to saying that the following diagram commutes:

$$\begin{tikzcd} \mathcal{A} \times 2 \arrow[r, "f_v"] \arrow[d, "\mathtt{pr}"] & \mathcal{A} \times 2 \arrow[d, "\mathtt{pr}"] \\ \mathcal{A} \arrow[r, "f_{\mathtt{pr}(v)}"] & \mathcal{A} \end{tikzcd}$$

Then $f_v = \langle \mathtt{pr} \circ f_v, \psi \rangle = \langle f_{\mathtt{pr}(v)} \circ \mathtt{pr}, \psi \rangle$ for some $\psi : \mathcal{A} \times 2 \longrightarrow 2$. Considering that for each $x \in \mathcal{A}$ $\mu_y.\psi(x, y)$ exists – 2 is a complete lattice – we can use the Bekic property to argue that the least fixed-point of f_v exists and is equal to the pair $(\mu_x.f_{\mathtt{pr}(v)}, \mu_y.\psi(\mu_x.f_{\mathtt{pr}(v)}, y))$. Therefore we can interpret the term $\mu_x.f$ as expected in the algebra $\mathcal{A} \times 2$ and the first projection $\mathtt{pr}_{\mathcal{A}}$ preserves this interpretation.

Since all the terms of the theory of modal μ-algebras are generated either by substitution or by formation of fixed-points from the terms of the theory of multi-modal algebras, a straightforward induction can be used to argue that $\mathcal{A} \times 2$ is a modal μ-algebra. $\square$

Lemma 2.4 On any modal algebra $\mathcal{A}$ condition (6) is equivalent to

$$\bigwedge \Lambda \wedge \bigwedge_{\sigma \in Act} ([\sigma] x_\sigma \wedge \bigwedge_{y \in Y_\sigma} \langle \sigma \rangle y) \leq \bot$$
$$\text{implies} \qquad (7)$$
$$\bigwedge \Lambda \leq \bot \text{ or } \exists \sigma \in Act, y \in Y_\sigma \text{ s.t. } x_\sigma \wedge y \leq \bot\,.$$

The Lemma – whose proof we skip for lack of space – has almost lead us to prove Theorem 2.1: we are left to argue that if $\bigwedge \Lambda \leq \bot$ in a free modal μ-algebra, then $p, \neg p \in \Lambda$ for some generator p. To this goal, it is enough to observe that the property holds in a Boolean algebra $\mathcal{B}_P$ freely generated by the set P, and that this algebra embeds in the free modal μ-algebra generated by P. The last statement is justified as follows: the free Boolean algebra can be given a trivial structure of a modal algebra, and being finite and complete it is therefore a modal μ-algebra; the canonical map $\mathcal{B}_P \longrightarrow \mathcal{F}_P$ is therefore split by a map in the other direction.

2.1 Modal operators are adjoints

We begin exemplifying the consequences of Theorem 2.1 by proving Claim 1.6. The Claim can also be understood by saying that reverse or backward modalities are definable in a free modal μ-algebra. More importantly, the property stated in the Claim is analogous to definability of Brzozowski's derivatives in free Kleene-algebras [12]. Much of the theory presented here has its origins in this paper.

Proposition 2.5 [i.e. Claim 1.6] On a free modal μ-algebra each operator $\langle \sigma \rangle$ is a left adjoint.

Proof. Observe that each element of the free modal μ-algebra is a conjunction of special elements b of the form

$$b = \bigvee \Lambda \vee \bigvee_{\tau \in Act} (\, \langle \tau \rangle x_\tau \vee \bigvee_{y \in Y_\tau} [\tau] y \,)\,,$$

where Λ is a set of literals. This can be explained as follows. Recall that every term of the modal μ-calculus is provably equivalent to terms with negation only in front of generators and every bound fixed-point variable is in the scope of some modal operator, see [10]. Using fixed-point equalities it is possible to extract from a term its first modal level and the statement follows from distributivity.

We begin defining the right adjoint for these special elements: $\mathrm{r}_\sigma(b) = \top$ if $b = \top$ and $\mathrm{r}_\sigma(b) = x_\sigma$ otherwise. Suppose that $\langle \sigma \rangle x \leq b$. If $b = \top$ then clearly $x \leq \top = \mathrm{r}_\sigma(x)$. If $b \neq \top$, then we deduce $x \leq x_\sigma = \mathrm{r}_\sigma(b)$: indeed this is a consequence of duality and Theorem 2.1, taking into account that all other disjuncts of the consequent of 2.1 imply $b = \top$. Conversely, $\langle \sigma \rangle \mathrm{r}_\sigma(b) \leq b$ implies that $x \leq \mathrm{r}_\sigma(b)$ implies $\langle \sigma \rangle (x) \leq b$. Note also that $\mathrm{r}_\sigma(b)$ does not depend on the representation of b, as it is uniquely determined by the property $x \leq \mathrm{r}_\sigma(b)$ iff $\langle \sigma \rangle x \leq b$.

Finally, if $x = \bigwedge_{j \in J} b_j$, then we define $\mathrm{r}_\sigma(x) = \bigwedge_{j \in J} \mathrm{r}_\sigma(b_j)$. $\square$

2.2 The Kleene star is constructive

An important property of $\mathrm{r}_\sigma(z)$ – defined in the proof of Proposition 2.5 – is that it is computed out of the syntax of z. More precisely, $\mathrm{r}_\sigma(z)$ is computed as a meet of terms belonging to the Fisher-Ladner closure, see [10], of a term representing z. The Fisher-Ladner closure has to be thought as the space of subterms of z, in particular it is finite. Consequently, the set $\{\, \mathrm{r}^n_\sigma(z) \mid n \geq 0 \,\}$ is finite and $\bigwedge_{n \geq 0} \mathrm{r}^n_\sigma(z)$ exists in a free modal μ-algebra. We exemplify how to exploit this fact by proving that $\mu_y.(x \vee \langle \sigma \rangle y)$ is the supremum over the chain of its finite approximants. We shall use the standard Propositional Dynamic Logic notation and let $\langle \sigma^* \rangle x = \mu_y.(x \vee \langle \sigma \rangle y)$.

Lemma 2.6 The relation $\langle \sigma^* \rangle a = \bigvee_{n \geq 0} \langle \sigma \rangle^n a$ holds in a free modal μ-algebra.

Proof. Let b such that for each $n \geq 0$ we have $\langle \sigma \rangle^n a \leq b$. Transpose this relation to obtain $a \leq \mathrm{r}^n_\sigma(b)$ for each $n \geq 0$, hence $a \leq \bigwedge_{n \geq 0} \mathrm{r}^n_\sigma(b)$. We claim that $\bigwedge_{n \geq 0} \mathrm{r}^n_\sigma(b)$ is a $\langle \sigma \rangle$-

prefixed-point. Indeed:

$$\langle\sigma\rangle \bigwedge_{n\geq 0} \mathrm{r}^n_\sigma(b) \leq \bigwedge_{n\geq 0} \langle\sigma\rangle \mathrm{r}^n_\sigma(b) \qquad \langle\sigma\rangle \text{ is order preserving}$$

$$\leq \bigwedge_{n\geq 0} \langle\sigma\rangle \mathrm{r}^{n+1}_\sigma(b) \leq \bigwedge_{n\geq 0} \mathrm{r}^n_\sigma(b)$$

by the counit relation $\langle\sigma\rangle\mathrm{r}_\sigma x \leq x$.

Thus $\bigwedge_{n\geq 0} \mathrm{r}^n_\sigma(b)$ is a $\langle\sigma\rangle$-prefixed-point above a and therefore $\langle\sigma^*\rangle a \leq \bigwedge_{n\geq 0} \mathrm{r}^n_\sigma(b) \leq b$. □

3 On $\mathcal{O}_f$-adjoints

The proof that the Kleene star modality is constructive relies on the modality $\langle\sigma\rangle$ being a left adjoint. A similar idea cannot work in the general case: for example $[\sigma]$ is not a left adjoint since it doesn't preserve joins. To deal with the general case, adjoints must be generalized as follows.

Definition 3.1 An order preserving function $f : L \longrightarrow M$ (L and M being posets) is a left $\mathcal{O}_f$*-adjoint* if for each $m \in M$ there exists a finite set $\mathcal{C}(f;m)$ such that for all $x \in L$ $f(x) \leq m$ if and only if $x \leq c$ for some $c \in \mathcal{C}(f;m)$. That is: f is a $\mathcal{O}_f$-adjoint if for each $m \in M$ the set

$$\{\, x \mid f(x) \leq m \,\}$$

is a finitely generated lower set.

We shall say that $\mathcal{C}(f;m)$ is the set of f-covers of m or the covering set of f and m. It is easily seen that f is a left adjoint if and only if $\{\, x \mid f(x) \leq m \,\}$ is a principal ideal: thus every left adjoint is a left $\mathcal{O}_f$-adjoint. Also, it is easy to see that f is a left $\mathcal{O}_f$-adjoint if and only if

$$\mathcal{O}_f(f) : \mathcal{O}_f(L) \longrightarrow \mathcal{O}_f(M)$$

is a left adjoint. Here $\mathcal{O}_f(P)$ is the set of finitely generated lower sets of the poset P and $\mathcal{O}_f(f)$ is the obvious map induced by this functorial construction. A left $\mathcal{O}_f$-open adjoint is a $Pro(\mathcal{D})$-adjoint as defined in [25] with $\mathcal{D}$ the class of all finite discrete categories. Similar but slightly different is the notion of a multiadjoint [5].

We begin presenting an interesting order theoretic property of $\mathcal{O}_f$-adjoints:

Lemma 3.2 A $\mathcal{O}_f$-adjoint f is continuous: if I is a directed set and $\bigvee I$ exists, then $\bigvee_{i\in I} f(i)$ exists and is equal to $f(\bigvee I)$.

Proof. Suppose that for all $i \in I$ $f(i) \leq m$. We can find $c_i \in \mathcal{C}(f;m)$ such that $i \leq c_i$. Since I is directed and the c_i are finite, we can find i_0 such that $i \leq c_{i_0}$ for all $i \in I$ and consequently $\bigvee I \leq c_{i_0}$. It follows that $f(\bigvee I) \leq f(c_{i_0}) \leq m$. □

We can argue that being a left $\mathcal{O}_f$-adjoint is a stronger property than merely being continuous by considering the binary meet $\wedge : \mathcal{B} \times \mathcal{B} \longrightarrow \mathcal{B}$ on an infinite Boolean algebra $\mathcal{B}$. The binary meet is continuous – an order preserving function $f : L \times M \longrightarrow N$ is continuous as a function from the product if and only if it is continuous in each variable – but it is not a left $\mathcal{O}_f$-adjoint. This statement can be verified by computing a candidate covering set $\mathcal{C}(\wedge;\bot)$. Since $x \wedge \neg x \leq \bot$, then we should be able to find $(\alpha_x, \beta_x) \in \mathcal{C}(f;\bot)$ such that $x \leq \alpha_x$, $\neg x \leq \beta_x$, and moreover $\alpha_x \wedge \beta_x \leq \bot$. It follows that $\alpha_x \leq \neg\beta_x \leq x$, and $\alpha_x = x$. Thus, for an infinite Boolean algebra the covering set $\mathcal{C}(\wedge;\bot)$ has to be infinite.

We list next some well known properties of (left) $\mathcal{O}_f$-adjoints:

Proposition 3.3

1. An order preserving function $f : L \longrightarrow M$ is a left adjoint if and only if it is a $\mathcal{O}_f$-adjoint and preserves finite joins.
2. If a lattice M is finitely meet-generated by a subset $B \subseteq M$, then $f : L \longrightarrow M$ is a $\mathcal{O}_f$-adjoint if and only if the covering set $\mathcal{C}(f;b)$ exists for each $b \in B$.
3. The identity is a $\mathcal{O}_f$-adjoint, and $\mathcal{O}_f$-adjoints are closed under composition.
4. If the domain posets are meet semilattices, then the projections $\mathtt{pr}_i : L_1 \times L_2 \longrightarrow L_i$, $i = 1, 2$, are $\mathcal{O}_f$-adjoints. Moreover $\langle f_1, f_2\rangle : L \longrightarrow M_1 \times M_2$ is a $\mathcal{O}_f$-adjoint provided that $f_i : L \longrightarrow M_i$, $i = 1, 2$, are $\mathcal{O}_f$-adjoints.
5. Constant functions are $\mathcal{O}_f$-adjoints.
6. Finite joins are $\mathcal{O}_f$-adjoints. If L is an Heyting algebra (or a Browerian semilattice), then $f(x) = k \wedge x : L \longrightarrow L$ is a $\mathcal{O}_f$-adjoint, where k is a constant.

3.1 $\mathcal{O}_f$-adjoints and fixed points

We analyze next $\mathcal{O}_f$-adjoints for which it makes sense to consider least fixed points. These have the form $f : L^x \times M^y \longrightarrow L$. For such an f we define the graph $\mathcal{G}(f, L)$: its vertices are elements of L and we declare that $l \to l'$ if $(l', y) \in \mathcal{C}(f;l)$ for some $y \in M$. We write $\mathcal{G}(f, l)$ for the full subgraph of $\mathcal{G}(f, L)$ of vertices that are reachable from l: $l' \in L$ is a vertex of $\mathcal{G}(f, l)$ if and only if there is a path from l to l' in $\mathcal{G}(f, L)$.

Definition 3.4 We say that a $\mathcal{O}_f$-adjoint of the form $f : L^x \times M^y \longrightarrow L$ is *finitary* in the variable x if for each $l \in L$ the graph $\mathcal{G}(f, l)$ is finite.

Lemma 3.5 If $f : L^x \times M^y \longrightarrow L$ is a finitary $\mathcal{O}_f$-adjoint and $\mu_x.f(x, y)$ exists for each $y \in M$, then $\mu_x.f : M^y \longrightarrow L$ is again a $\mathcal{O}_f$-adjoint.

Proof. By an x-path from l we mean a sequence $(l_i, r_i)_{i=1,\dots,n}$ such that $(l_i, r_i) \in \mathcal{C}(f; l_{i-1})$; by convention we let $l_0 = l$ and say that n is the length the x-path. We also consider infinite x-paths, i.e. infinite sequences of this kind: since the graph $\mathcal{G}(f, l)$ is finite there is only a finite number of r_i's and therefore the meet $\bigwedge_{i \geq 1} r_i$ exists in M. Infinite x-paths allow us to define $r \in \mathcal{C}(\mu_x.f; l)$ iff $r = \bigwedge_{i \geq 1} r_i$ for some infinite x-path $(l_i, r_i)_{i \geq 0}$ from l. The set $\mathcal{C}(\mu_x.f; l)$ is finite since $\mathcal{G}(f, l)$ is finite. We begin verifying that $\mu_x.f(r) \leq l$ if $r \in \mathcal{C}(\mu_x.f; l)$. By monotonicity, $f(l_{i+1}, r) \leq f(l_{i+1}, r_{i+1}) \leq l_i$ for all $i \geq 0$. Choose $i < j$ such that $(l_i, r_i) = (l_j, r_j)$ and let $k = j - i$, then $f^k(l_i, r) \leq f^k(l_j, r_j) \leq l_i$ hence $\mu_x.f(x, r) \leq l_i$ and moreover $\mu_x.f(x, r) = f_r^i(\mu_x.f(x, r)) \leq f_r^i(l_i) \leq l_0$.

Conversely, assume that $\mu_x.f(x, y) \leq l_0$: we can use the fixed point equation to deduce $f(\mu_x.f(x, y), y) \leq l_0$ which in turn implies $(\mu_x.f(x, y), y) \leq (l_1, r_1)$ for some pair $(l_1, r_1) \in \mathcal{C}(f; l_0)$. By iterating the procedure, we can construct an infinite x-path (l_i, r_i) from l such that for all $i \geq 1$ we have $(\mu_x.f(x, y), y) \leq (l_i, r_i)$ and therefore $y \leq \bigwedge_{i \geq 1} r_i$. □

It is a natural operation to prune the covering sets $\mathcal{C}(f; m)$ and extract the antichain of maximal elements. A maximal element in $\mathcal{C}(\mu_x.f; l)$ is a meet indexed by some pan, by which we mean a finite path in $\mathcal{G}(f, l)$ that can be split into a simple path followed by a simple cycle.

Lemma 3.6 If $f : L^x \times M^y \longrightarrow L$ is a finitary $\mathcal{O}_f$-adjoint and $\mu_x.f$ exists then $\mu_x.f(x, y) = \bigvee_{n \geq 0} f_y^n(\bot)$.

Proof. Assume l is such that $f_y^n(\bot) \leq l$ for each $n \geq 0$. Let k be the number of vertices in the graph $\mathcal{G}(f, l)$ and observe that the relation $f_y^{k+1}(\bot) \leq l$ implies that we can find an x-path $(l_i, r_i)_{i=1,\dots,k+1}$ from l with the property that $y \leq r_i$ for $y = 1, \dots, k + 1$. By the pigeonhole principle, we can assume this is an infinite x-path $(l_i, r_i)_{\geq 1}$ from l such that $y \leq r_i$ for $i \geq 1$. Thus $\bigwedge r_i \in \mathcal{C}(\mu_x.f; l)$ and therefore $\mu_x.f(x, y) \leq \mu_x.f(x, \bigwedge r_i) \leq l$. □

3.2 $\mathcal{O}_f$-adjoints on free modal μ-algebras

We have seen that meets provide a counter-example for $\mathcal{O}_f$-adjointness. In [9] the authors define what turns out to be a best approximation of meets as $\mathcal{O}_f$-adjoints. They first define the arrow term by:

$$\xrightarrow{\sigma} X = [\sigma] \bigvee X \wedge \bigwedge_{x \in X} \langle \sigma \rangle x \,, \tag{8}$$

and, for a set of literals Λ, for a subset $\Sigma \subseteq Act$, and for disjoint sets of variables $\{ X_\sigma \}_{\sigma \in \Sigma}$, they also define the special conjunction term by:

$$\bigwedge\nolimits_{\Lambda, \Sigma} \{ X_\sigma \} = \Lambda \wedge \bigwedge_{\sigma \in \Sigma} \xrightarrow{\sigma} X_\sigma \,. \tag{9}$$

By 2.4, in a free modal μ-algebra a special conjunction is inconsistent, $\bigwedge_{\Lambda, \Sigma} \{ X_\sigma \} = \bot$ if either the literals in Λ are inconsistent or there is some $\sigma \in \Sigma$ and $x \in X_\sigma$ such that $x \leq \bot$. Recall also the relations $\langle \sigma \rangle x = \xrightarrow{\sigma} \{ x, \top \}$ and $[\sigma] x = \xrightarrow{\sigma} \{ x \} \vee \xrightarrow{\sigma} \emptyset$ showing that special conjunctions can be taken as primitive function symbols of the signature of the theory of modal μ-algebras.

Lemma 3.7 Special conjunctions are finitary $\mathcal{O}_f$-adjoints in free μ-algebras.

Proof. Recall from 2.5 that the free modal μ-algebra is finitely meet-generated by elements of the form

$$b = \bigvee \Gamma \vee \bigvee_{\tau \in Act} (\langle \tau \rangle z_\tau \vee \bigvee_{y \in Y_\tau} [\tau] y) \,, \tag{10}$$

where Γ is a set of literals. By Proposition 3.3.2, it is enough to define the set of covers for these elements, thus let b be as above. Observe that for $b = \top$ we always have $\mathcal{C}(f; \top) = \{ \top \}$. Hence, let us suppose that $b \neq \top$ and that $\bigwedge_{\Lambda, \Sigma}(v) \leq b$, where $v \in \mathcal{F}^X$ and X is the disjoint union of the X_σ, $\sigma \in \Sigma$. We want to apply Theorem 2.1 to the relation $\bigwedge_{\Lambda, \Sigma}(v) \leq b$. It is useful to expand this relation to:

$$\begin{aligned} \bigwedge \Lambda \ \wedge \ \bigwedge_{\sigma \in \Sigma} ([\sigma] \bigvee_{x \in X_\sigma} v(x) \wedge \bigwedge_{x \in X_\sigma} \langle \sigma \rangle v(x)) \\ \leq \bigvee \Gamma \vee \bigvee_{\tau \in Act} (\langle \tau \rangle z_\tau \vee \bigvee_{y \in Y_\tau} [\tau] y) \,. \end{aligned}$$

Since $b \neq \top$, one of the following cases arises:

1. $\bigwedge \Lambda \leq \bigvee \Gamma$,
2. there exists $\sigma \in \Sigma$ and $x \in X_\sigma$ such that $v(x) \leq z_\sigma$,
3. there exists $\sigma \in \Sigma$ and $y \in Y_\sigma$ such that $v(x) \leq x_\sigma \vee y$ for each $x \in X_\sigma$.

Suppose the first case holds: let $w_\top \in \mathcal{F}^X$ be the vector with constant value $\top$, clearly $v \leq w_\top$, and moreover $\bigwedge_{\Lambda, \Sigma}(w_\top) \leq b$.
If the second case holds, then we can define

$$w_{z_\sigma}(x) = \begin{cases} \top, & x \in X_\tau, \tau \neq \sigma \\ z_\sigma, & x \in X_\sigma \,. \end{cases}$$

Clearly $v \leq w_{z_\sigma}$ and conversely

$$\bigwedge\nolimits_{\Lambda, \Sigma}(w_\sigma) \ \leq \xrightarrow{\sigma} \{ z_\sigma \} \ \leq \ \langle \sigma \rangle z_\sigma \leq b \,.$$

Finally, if the third case holds, then we can define

$$w_{z_\sigma, y}(x) = \begin{cases} \top, & x \in X_\tau, \tau \neq \sigma \\ z_\sigma \vee y, & x \in X_\sigma \,. \end{cases}$$

We have $v \leq w_{z_\sigma}$ and conversely

$$\bigwedge\nolimits_{\Lambda,\Sigma}(w_\sigma) \;\leq \xrightarrow{\sigma} \{\, z_\sigma \vee y \,\} \;\leq\; [\sigma](z_\sigma \vee y)$$
$$\leq\; \langle\sigma\rangle x_\sigma \vee [\sigma]y \;\leq\; b\,.$$

Thus we have shown that we can define the covering set $\mathcal{C}(\bigwedge_{\Lambda,\Sigma}; b)$ as

$$\{\, w_\top \,\} \cup \{\, w_{z_\sigma} \,|\, \sigma \in \Sigma \,\} \cup \{\, w_{z_\sigma,y} \,|\, \sigma \in \Sigma, y \in Y_\sigma \,\}\,.$$

To end the proof, we remark that covers of an element $c \in \mathcal{F}$ are meets of subterms of a term representing c, showing that special conjunctions are finitary. □

It is now easy to argue that $[\sigma]$ is a $\mathcal{O}_f$-adjoint on a free modal μ-algebra. By Proposition 3.3, this is a consequence of $[\sigma]$ belonging to the cone generated by joins and special conjunctions.

4 Some constructive systems of equations

An order preserving $F : \mathcal{F}^X \times \mathcal{F}^Y \longrightarrow \mathcal{F}^X$ can also be thought to be a system of equations whose least solution is given by the least fixed-point. The set X is the set of bound variables of the system, and Y is the set of free variables. The Bekic property[7] ensures that such a system of equations has a least solution if F is built up from operations of the theory of modal μ-algebras. In this section we shall prove that, for many such F on a free modal μ-algebra, the least prefixed-point is the supremum over the chain of its finite approximants. The results of the previous sections allow us to easily derive this property for a restricted set of systems called here disjunctive-simple. Then, we freely use ideas and tools from [1, §9] to enlarge the class of systems that can be proved to be constructive. A technical improvement consists in showing that the methods exposed in [1] can be adapted in order to argue about existence of infinite suprema and approximants.

Definition 4.1 Let X and Y be two sets of variables. With respect to X and Y, we say that a term of the theory of modal μ-algebras

- is *simple* if it is a distributive combination of terms of the form $\bigwedge Y' \wedge \bigwedge_{\emptyset,\Sigma}\{\, X_\sigma \,\}$, where $Y' \subseteq Y$ and $X_\sigma \subseteq X \cup Y$,
- is *disjunctive-simple* if it is a join of terms of the form $\bigwedge Y' \wedge \bigwedge_{\emptyset,\Sigma}\{\, D_\sigma \,\}$, where $Y' \subseteq Y$ and each $d \in D_\sigma$ is a join of a set of variables: $d = \bigvee X'$ with $X' \subseteq X \cup Y$.

For a μ-algebra $\mathcal{A}$ we say that a map $F = \langle F_x \rangle_{x \in X} : \mathcal{A}^X \times \mathcal{A}^Y \longrightarrow \mathcal{A}^X$

- is *simple* if each component $F_x : \mathcal{A}^X \times \mathcal{A}^Y \longrightarrow \mathcal{A}$ is the interpretation of a term simple w.r.t. X and Y,
- is *disjunctive-simple* if each component $F_x : \mathcal{A}^X \times \mathcal{A}^Y \longrightarrow \mathcal{A}$ is the interpretation of a term that is disjunctive-simple w.r.t. X and Y.

Proposition 4.2 Let $\mathcal{F}$ be a free μ-algebra, $G : \mathcal{F}^X \times \mathcal{F}^Y \longrightarrow \mathcal{F}^X$ be a disjunctive-simple map, and let $k \in \mathcal{F}^Y$. Then $G_k : \mathcal{F}^X \longrightarrow \mathcal{F}^X$ is a finitary $\mathcal{O}_f$-adjoint.

For lack of space, we only present a skecth of the proof. *Proof.* Proposition 3.3 and Lemma 3.7 imply that for each $z \in X$ the z component of G_k – which we shall denote G_z abusing notation – is a $\mathcal{O}_f$-adjoint. Item 4 of Proposition 3.3 implies then that G_k is a $\mathcal{O}_f$-adjoint. Thus we are mainly concerned with arguing that G_k is finitary.

To this goal, let ${}^\wedge\langle l_1, \ldots, l_n \rangle$ denote the meet-semilattice generated by $l_1, \ldots, l_n$. We claim that G_k is finitary if the following condition holds: for each $u \in \mathcal{F}^X$ we can find a set $\{\, l_1, \ldots, l_n \,\}$ such that (i) $u(x) \in \{\, l_1, \ldots, l_n \,\}$ for each $x \in X$, and (ii) for $j = 1, \ldots, n$, $z \in X$, and $c \in \mathcal{C}(G_z; l_j)$, $c(x) \in {}^\wedge\langle l_1, \ldots, l_n \rangle$ for each $x \in X$.

Each G_z has the form $\bigvee_{i \in I_z} y_{z,i} \wedge \bigwedge_{\emptyset,\Sigma_i}\{\, D_\sigma \,\}$. For each $z \in Z$ and $i \in I_z$ let $t_{z,i}$ be a term representing the constant $\neg y_{x,i}$. Chose $u \in \mathcal{F}^X$ and, for each $x \in X$, let s_x be a term representing $u(x)$. Let S be the set of interpretation of subterms of the $t_{z,i}$ and s_x. We let $\{l_1, \ldots, l_n\} = {}^\vee\langle S \rangle$ be the join closure of S. □

Lemma 3.6 and the previous Lemma imply that a disjunctive-simple $G : \mathcal{F}^X \times \mathcal{F}^Y \longrightarrow \mathcal{F}^X$ satisfies

$$\mu_X.G = \bigvee_{n \geq 0} G_v^n(\bot)\,,$$

for each $v \in \mathcal{F}^Y$. Our next goal is to transfer this property from a disjunctive-simple G to a simple F. The main tool is the following Lemma:

Proposition 4.3 Consider a commuting diagram of posets with bottom

$$\begin{tikzcd}
L \arrow[r, "f"] \arrow[d, "i"] & L \arrow[d, "i"'] \\
M \arrow[r, "g"] & M \arrow[u, bend right, "\pi"']
\end{tikzcd}$$

where i is split by an order preserving π, $\pi \circ i = \mathtt{id}_L$. Let α be a limit ordinal and suppose that the approximant $g^\alpha(\bot)$ exists. If $i(f^\beta(\bot)) = g^\beta(\bot)$ for $\beta < \alpha$, then the approximant $f^\alpha(\bot)$ exists and is equal to $\pi(g^\alpha(\bot))$. If moreover i is continuous, then $i(f^\alpha(\bot)) = g^\alpha(\bot)$.

[7] See for example Proposition 2.1 and 2.2 in [20].

Proof. We want to prove that $\pi(f^\alpha(\bot)) = \bigvee_{\beta<\alpha} f^\beta(\bot)$. Let us begin supposing that, for some $l \in L$ and every $\beta < \alpha$, $f^\beta(\bot) \leq l$. Apply i to these relations and deduce that $g^\beta(\bot) \leq i(l)$ for $\beta < \alpha$, hence $g^\alpha(\bot) \leq i(l)$. Apply then π to deduce $\pi(g^\alpha(\bot)) \leq l$. Conversely, $f^\beta(\bot) \leq \pi(g^\alpha(\bot))$ if $\beta < \alpha$: these relations are obtained from $i(f^\beta(\bot)) = g^\beta(\bot) \leq g^\alpha(\bot)$ by applying π. If moreover i is continuous, then:

$$\begin{aligned} i(f^\alpha(\bot)) &= i(\bigvee_{\beta<\alpha} f^\beta(\bot)) = \bigvee_{\beta<\alpha} i(f^\beta(\bot)) \\ &= \bigvee_{\beta<\alpha} g^\beta(\bot) = g^\alpha(\bot). \quad \square \end{aligned}$$

We transfer the constructiveness property from a disjunctive-simple G to simple F as follows. For a finite set of variables X, let $\mathcal{P}_+(X)$ be the set of nonempty subsets of X. For each $S \in \mathcal{P}_+(X)$, the map $i_S : L^X \longrightarrow L$, defined by $i_S(x) = \bigwedge_{j\in S} x_j$, is continuous. These maps are collected together to produce a continuous map $i : L^X \longrightarrow L^{\mathcal{P}_+(X)}$. For each $x \in X$ there is a projection $\mathtt{pr}_{\{x\}} : L^{\mathcal{P}_+(X)} \longrightarrow L$ and these projections, collected into a common projection $\mathtt{pr} : L^{\mathcal{P}_+(X)} \longrightarrow L^X$, split i: $\mathtt{pr} \circ i = \mathtt{id}_{\mathcal{F}^X}$.

Proposition 4.4 For each simple $F : \mathcal{F}^X \times \mathcal{F}^Y \longrightarrow \mathcal{F}^X$ there is a $G : \mathcal{F}^{\mathcal{P}_+(X)} \times \mathcal{F}^Y \longrightarrow \mathcal{F}^{\mathcal{P}_+(X)}$ which is disjunctive-simple and such that the diagram

$$\begin{tikzcd} \mathcal{F}^X \times \mathcal{F}^Y \arrow[r, "F"] \arrow[d, "i\times \mathtt{id}_{\mathcal{F}^Y}"] & \mathcal{F}^X \arrow[d, "i"] \\ \mathcal{F}^{\mathcal{P}_+(X)} \times \mathcal{F}^Y \arrow[r, "G"] & \mathcal{F}^{\mathcal{P}_+(X)} \end{tikzcd}$$

commutes.

Proposition 4.4 means that for each nonempty subset $S \subseteq X$ we can find a disjunctive-simple $G_S : \mathcal{F}^{\mathcal{P}_+(X)} \longrightarrow \mathcal{F}$ such that

$$\bigwedge_{j\in S} F_j(x) = G_S(\bigwedge_{j\in S_1} x_j, \ldots, \bigwedge_{j\in S_k} x_j),$$

where $S_1, \ldots, S_k$ is the list of nonempty subsets of X. Its proof strictly follows [1, §9.4] and therefore we omit it.

Together with Lemma 4.3 and the properties of disjunctive-simple systems, the Proposition implies that a simple $F : \mathcal{F}^X \times \mathcal{F}^Y \longrightarrow \mathcal{F}^X$ satisfies

$$\mu_X.F_v = \bigvee_{n\geq 0} F_v^n(\bot),$$

for each $v \in \mathcal{F}^Y$.

For our goals, the following concept turns out to be useful:

Definition 4.5 We say that an order preserving map $f : L \times M \longrightarrow L$ is *regular* if it is continuous and constructive.

Recall that being constructive means that the approximant $f_v^\alpha(\bot)$ exists for each $v \in M$ and each ordinal α. It is easily seen that for a continuous f existence of the supremum $\bigvee_{n\geq 0} f_v^n(\bot)$ suffices for existence of all approximants. As an example, we have seen that all the simple $G : \mathcal{F}^X \times \mathcal{F}^Y \longrightarrow \mathcal{F}^X$ are constructive. Since each $G_x : \mathcal{F}^X \times \mathcal{F}^Y \longrightarrow \mathcal{F}$ is also continuous, G is continuous as well: hence *each simple G is regular*. A system is *elementary* if each F_x is among $\top, \wedge, \bot, \vee, \xrightarrow{\sigma}$.

Proposition 4.6 Each elementary system is regular.

To prove the Proposition, we say that two systems $F : \mathcal{F}^X \times \mathcal{F}^Y \longrightarrow \mathcal{F}^X$ and $G : \mathcal{F}^X \times \mathcal{F}^Y \longrightarrow \mathcal{F}^X$ are *equivalent* if for each $v \in \mathcal{F}^Y$, the two chains of finite approximants $\{ F_v^n(\bot) \}_{n\geq 0}$ and $\{ G_v^n(\bot) \}_{n\geq 0}$ are cofinal into each other: for each $n \geq 0$ there exists $k \geq 0$ such that $F^n(\bot) \leq G^k(\bot)$, and vice-versa. The equivalence of systems can be generalized to systems $F : \mathcal{F}^X \times \mathcal{F}^Y \longrightarrow \mathcal{F}^X$ and $G : \mathcal{F}^{X'} \times \mathcal{F}^Y \longrightarrow \mathcal{F}^{X'}$ with different sets of bound variables. In this case two systems are equivalent iff the partial chains of finite approximants $\{ \mathtt{pr}_{X\cap X'}(F_v^n(\bot)) \}_{n\geq 0}$ and $\{ \mathtt{pr}_{X\cap X'}(G_v^n(\bot)) \}_{n\geq 0}$ are cofinal into each other and the supremum of one chain of approximant exists if and only if the supremum of the other chain of approximant exists. If F and G are continuous and equivalent, then F is regular if and only if G is regular. Also, if the regular F and G are equivalent, then $\mathtt{pr}_{X\cap X'} \circ \mu_X.F = \mathtt{pr}_{X\cap X'} \circ \mu_{X'}.G$.

A variable x is guarded in a term t if either it does not occur in t or it occurs in t within the scope of a modal operator. A system $F : \mathcal{F}^X \times \mathcal{F}^Y \longrightarrow \mathcal{F}^X$ is guarded if for each each $x, x' \in X$, x' is *guarded* in each F_x. The following Lemma is partly analogous to the well known fact that every formula of the modal μ-calculus is equivalent to a guarded one [10].

Lemma 4.7 For each elementary system F there exists a guarded system G which is equivalent to F. For each guarded system G there exists a simple system H which is equivalent to G.

Proposition 4.6 follows.

5 Σ_1-operations are constructive

Since we can define all terms of the theory of modal μ-algebras using the arrow terms (8) in place of standard modal operators, cf. §3.2, we redefine Σ_1-terms accordingly:

$$t \;=\; x \,|\, \top \,|\, t \wedge t \,|\, \bot \,|\, t \vee t \,|\, \xrightarrow{\sigma} T \,|\, \mu_x.t\,.$$

Here x is a variable and T is a set of previously defined terms. Our goal is now to prove:

Theorem 5.1 If $\mathcal{F}$ is a free modal μ-algebra, then every Σ_1-operation $f : \mathcal{F}^y \times \mathcal{F}^{Y \setminus \{y\}} \longrightarrow \mathcal{F}$ is regular (for each $y \in Y$).

In this way we will have accomplished a proof of Claim 1.9. Our goal is achieved by considering Σ_1-operations as solutions of systems of equations, as stated in the next Lemma.

Lemma 5.2 For each Σ_1-operation $f : L^Y \longrightarrow L$ there exists an elementary system $F : L^X \times L^Y \longrightarrow L^X$ and $x \in X$ such that $f = \mathtt{pr}_x \circ \mu_X.F$.

To prove Theorem 5.1 we need a Bekic-like property for regular functions:

Lemma 5.3 Suppose that $F : L^x \times M^y \times N^z \longrightarrow L$ is regular in x and $G : L^x \times M^y \times N^z \longrightarrow M$ is continuous. Then $\langle F, G \rangle : L \times M \times N \longrightarrow L \times M$ is regular if and only if $G(\mu_x.F(x,y,z), y, z) : M \times N \longrightarrow M$ is regular in y.

Proof of Theorem 5.1. By Lemma 5.2 $f = \mathtt{pr}_x \circ \mu_X.F$ for some elementary $F : \mathcal{F}^X \times \mathcal{F}^Z \longrightarrow \mathcal{F}^X$ and some $x \in X$. Choose $y \in Z$ and observe that the system

$$\langle F, \mathtt{pr}_x \rangle : \mathcal{F}^X \times \mathcal{F}^y \times \mathcal{F}^{Z \setminus \{y\}} \longrightarrow \mathcal{F}^X \times \mathcal{F}^y$$

is elementary and therefore it is regular. Since $\mathtt{pr}_x$ is continuous, we can use the previous Lemma with $G = \mathtt{pr}_x$ and deduce that $f = \mathtt{pr}_x \circ \mu_X.F : \mathcal{F}^Z \longrightarrow \mathcal{F}^y$ is regular for $y \in Z$. $\square$

References

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

[2] B. Banaschewski and G. Bruns. Categorical characterization of the MacNeille completion. *Arch. Math. (Basel)*, 18:369–377, 1967.

[3] S. L. Bloom and Z. Ésik. *Iteration theories*. EATCS Monographs on Theoretical Computer Science. Springer-Verlag, Berlin, 1993. The equational logic of iterative processes.

[4] C. C. Chang and H. J. Keisler. *Model theory*, volume 73 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Publishing Co., Amsterdam, third edition, 1990.

[5] Y. Diers. Some spectra relative to functors. *J. Pure Appl. Algebra*, 22:57–74, 1981.

[6] Z. Ésik. Completeness of Park induction. *Theoret. Comput. Sci.*, 177(1):217–283, 1997. Mathematical foundations of programming semantics (Manhattan, KS, 1994).

[7] R. Freese, J. Ježek, and J. B. Nation. *Free lattices*. American Mathematical Society, Providence, RI, 1995.

[8] M. Gehrke and J. Harding. Bounded lattice expansions. *J. of Algebra*, (231):345–371, 2001.

[9] D. Janin and I. Walukiewicz. Automata for the modal μ-calculus and related results. In *Mathematical foundations of computer science 1995 (Prague)*, volume 969 of *Lecture Notes in Comput. Sci.*, pages 552–562. Springer, Berlin, 1995.

[10] D. Kozen. Results on the propositional μ-calculus. *Theoret. Comput. Sci.*, 27(3):333–354, 1983.

[11] D. Kozen. A completeness theorem for Kleene algebras and the algebra of regular events. *Inf. Comput.*, 110(2):366–390, 1994.

[12] D. Kozen. Myhill-Nerode relations on automatic systems and the completeness of Kleene algebra. In *STACS 2001 (Dresden)*, volume 2010 of *Lecture Notes in Comput. Sci.*, pages 27–38. Springer, Berlin, 2001.

[13] M. Lange and C. Stirling. Model checking games for branching time logics. *J. Logic Comput.*, 12(4):623–639, 2002. International Conference on Temporal Logic, 2000 (Leipzig).

[14] R. N. McKenzie, G. F. McNulty, and W. F. Taylor. *Algebras, lattices, varieties. Vol. I.* The Wadsworth & Brooks/Cole Mathematics Series. Wadsworth & Brooks/Cole Advanced Books & Software, Monterey, CA, 1987.

[15] E. Nelson. Iterative algebras. *Theoret. Comput. Sci.*, 25(1):67–94, 1983.

[16] K. Ng and A. Tarski. Relational algebras with transitive closure. *Notices Amer. Math. Soc*, 24:A29-A30, 1977. Abstract 742-02-09.

[17] D. Niwinski. On fixed-point clones. In *Automata, languages and programming, Proc. 13th Int. Colloq., Rennes/France 1986, Lect. Notes Comput. Sci. 226, 464-473* . 1986.

[18] D. Park. Fixpoint induction and proofs of program properties. In *Machine intelligence, 5*, pages 59–78. American Elsevier, New York, 1970.

[19] L. Santocanale. Free μ-lattices. *J. Pure Appl. Algebra*, 168(2-3):227–264, 2002. Category theory 1999 (Coimbra).

[20] L. Santocanale. μ-bicomplete categories and parity games. *Theor. Inform. Appl.*, 36(2):195–227, 2002. Fixed points in computer science (Florence, 2001).

[21] L. Santocanale. On the equational definition of the least prefixed point. *Theoret. Comput. Sci.*, 295(1-3):341–370, 2003. Mathematical foundations of computer science (Mariánské Lázně, 2001).

[22] K. Segerberg. A completeness theorem in the modal logic of programs. *Notices Amer. Math. Soc*, 24:A552, 1977. Abstract 77T-E69.

[23] R. S. Streett and E. A. Emerson. An automata theoretic decision procedure for the propositional mu-calculus. *Inform. and Comput.*, 81(3):249–264, 1989.

[24] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific J. Math.*, 5:285–309, 1955.

[25] W. Tholen. Pro-categories and multiadjoint functors. *Can. J. Math.*, 36:144–155, 1984.

[26] I. Walukiewicz. Completeness of Kozen's axiomatisation of the propositional μ-calculus. *Inform. and Comput.*, 157(1-2):142–182, 2000. LICS 1995 (San Diego, CA).

[27] P. M. Whitman. Free lattices. *Ann. of Math. (2)*, 42:325–330, 1941.