

Filtration and canonical completeness for continuous modal μ -calculi

Jan Rooduijn*

ILLC
University of Amsterdam
j.m.w.rooduijn@uva.nl

Yde Venema

ILLC
University of Amsterdam
y.venema@uva.nl

The continuous modal μ -calculus is a fragment of the modal μ -calculus, where the application of fixpoint operators is restricted to formulas whose functional interpretation is Scott-continuous, rather than merely monotone. By game-theoretic means, we show that this relatively expressive fragment still allows two important techniques of basic modal logic, which notoriously fail for the full modal μ -calculus: filtration and canonical models. In particular, we show that the Filtration Theorem holds for formulas in the language of the continuous modal μ -calculus. As a consequence we obtain the finite model property over a wide range of model classes. Moreover, we show that if a basic modal logic L is canonical and the class of L -frames admits filtration, then the logic obtained by adding continuous fixpoint operators to L is sound and complete with respect to the class of L -frames. This generalises recent results on a strictly weaker fragment of the modal μ -calculus, *viz.* PDL.

1 Introduction

Filtration and canonical models This paper concerns two key methods in the theory of modal logic, both of which were introduced in their modern forms by Lemmon & Scott in [11]. First, *filtration*, which allows one to shrink a Kripke model into a finite one, by identifying states that agree on the truth of some given finite set of formulas. The *Filtration Theorem* then states that the equivalence classes in the finite model satisfy the same formulas as their members do in the original model. Filtration is the most important tool for proving the finite model property and the decidability of modal logics. For an overview of recent developments in the theory of filtration, see [2].

The other method central to this paper is that of *canonical models*. This well-known technique for proving the completeness of modal logics is related to Henkin's method for first-order logic. Given a modal logic L , it allows one to construct the *canonical model* \mathbb{S}^L of L with the powerful property that a formula φ is consistent in the logic L if and only if it is satisfiable in \mathbb{S}^L . It follows that L is complete with respect to any class of frames containing the *canonical frame*, *i.e.* the frame underlying \mathbb{S}^L . Thus, when then the canonical frame is a frame for L - in this case L is said to be *canonical* - the logic L is complete with respect to the class of frames for L .

Modal fixpoint logics Modal fixpoint logics are extensions of basic modal logic by operators capable of expressing certain kinds of recursive statements. They are of particular interest for computer science, where they are used to express important properties of processes. Examples of modal fixpoint logics are common knowledge logic (CKL), provability logic (GL), propositional dynamic logic (PDL) and computation tree logic (CTL). The central modal fixpoint logic, in which each of the aforementioned logics can be interpreted, is the modal μ -calculus (μ ML), introduced by Kozen in [9]. It extends basic modal logic with *explicit* least and greatest fixed point operators, resulting in a large gain of expressive

*The research of this author has been made possible by a grant from the Dutch Research Council NWO, project nr. 617.001.857.

power. Although many desirable properties, such as decidability and bisimulation invariance, withstand this gain in expressive power, the methods of filtration and canonical models do not.

In fact, the method of canonical models breaks down already in the case of relatively simple modal fixpoint logics. The reason is that these logics generally lack the compactness property, preventing the use of infinite maximally consistent sets. If, however, the method of filtration *does* work for such a logic L , then the canonical model method can often be salvaged. This roughly works as follows. One begins by taking the canonical model \mathbb{S}^L . Due to the compactness failure, this model is *non-standard*, meaning that the frame underlying \mathbb{S}^L fails to satisfy some desired properties. However, by applying filtration to \mathbb{S}^L we obtain a finite model (a *finitary canonical model*), whose underlying frame often does satisfy these desired properties. This procedure for instance underlies the completeness proof for PDL by Kozen & Parikh in [10]. In the book [6], Goldblatt applies the same procedure to several modal fixpoint logics, including CTL.

In the recent paper [8], Kikot, Shapirovsky & Zolin, prove a result of this kind that is relatively wide in scope. They show that if a basic modal logic L allows the method of filtration, then so does its expansion with the transitive closure modality. By iterating this procedure they show the same for the expansion of L by all modalities of PDL. Subsequently, if the original basic modal logic L moreover is canonical, the completeness of this PDL-expansion of L can be obtained by applying filtration to its canonical model.

The continuous modal μ -calculus In this paper we consider the methods of filtration and canonical models for a specific fragment of μML , which is called the *continuous modal μ -calculus* and is denoted $\mu_c\text{ML}$. In the paper [5], Fontaine shows that there are two equivalent ways to define $\mu_c\text{ML}$. First semantically, as the fragment of the modal μ -calculus where the application of fixpoint operators is restricted to formulas whose functional interpretation is Scott-continuous, rather than merely monotone. And second syntactically, as the fragment where the modal operator \Box and the fixpoint operator ν are not allowed to occur in the scope of a μ -operator (and dually for the ν -operator). To the best of our knowledge, the logic $\mu_c\text{ML}$ was mentioned first in van Benthem [1] under the name ‘ ω - μ -calculus’. It is related, and perhaps equivalent in expressive power, to the logic CPDL of concurrent propositional dynamic logic, cf. Carreiro [3, section 3.2] for more information.

There are at least two reasons why the continuous μ -calculus is an interesting logic; first, the continuity condition that is imposed on the formation of fixpoint formulas ensures that the construction of a definable fixpoint using its ordinal approximations will always be finished after ω many steps. And second, in the same manner that the full μ -calculus is the bisimulation-invariant fragment of monadic second-order logic [7], $\mu_c\text{ML}$ has the same expressive power as *weak* monadic second-order logic, when it comes to bisimulation-invariant properties [4].

The goal of the present paper is to show that we can add two more desirable properties to this list: (i) the Filtration Theorem holds for $\mu_c\text{ML}$ and (ii) completeness for sufficiently nice logics in the language of $\mu_c\text{ML}$ can be proven using finitary canonical models.

Since $\mu_c\text{ML}$ is strictly more expressive than PDL [5, 3], this is a proper generalisation of the aforementioned results from the paper [8]. On the other hand, because the failure of filtration for μML is witnessed by the formula $\mu x.\Box x$, the syntactic restrictions characterising $\mu_c\text{ML}$ seem to be not only sufficient, but also necessary for filtration. This indicates that $\mu_c\text{ML}$ might be positioned as a maximal filtration-allowing language between the basic modal language and the full language of the modal μ -calculus. We leave it for future work to make this statement mathematically precise and to investigate its correctness.

Overview of the paper In Section 2 we define the syntax of the continuous modal μ -calculus, the game semantics and other basic notions. In Section 3 we treat filtration. After giving the necessary definitions, we will use game-theoretic arguments to prove the Filtration Theorem for the language $\mu_c\text{ML}$. As a corollary, we obtain the finite model property for this language interpreted over a wide range of model classes. In Section 4 we prove our completeness result, again using game-theoretic methods.

Unlike for PDL, there is no obvious way to construct a non-standard canonical model for $\mu_c\text{ML}$. Because of this, we define the finitary canonical model used in our completeness proof directly, instead of as some filtration of a non-standard canonical model. This causes Section 3 and Section 4 to contain some rather similar constructions and proofs. We leave it for future work to unify these two.

2 The continuous modal μ -calculus

Syntax The continuous modal μ -calculus will be defined using the syntactic characterisation given by Fontaine in [5]. We fix a countably infinite set P of propositional variables.

Definition 2.1. By simultaneous induction we define the following three languages.

- (i) The syntax $\mu_c\text{ML}$ of the *continuous modal μ -calculus*:

$$\varphi ::= p \mid \neg p \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond \varphi \mid \square \varphi \mid \mu x. \varphi' \mid \nu x. \varphi''$$

where $p, x \in P$ and $\varphi' \in \text{Con}_{\{x\}}(\mu_c\text{ML})$, and $\varphi'' \in \text{Cocon}_{\{x\}}(\mu_c\text{ML})$.

- (ii) For $X \subseteq P$, the fragment $\text{Con}_X(\mu_c\text{ML})$ of $\mu_c\text{ML}$ -formulas that are *continuous* in X :

$$\varphi ::= x \mid \alpha \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond \varphi \mid \mu y. \varphi'$$

where $x \in X, y \in P, \alpha \in \mu_c\text{ML}$ X -free, and $\varphi' \in \text{Con}_{X \cup \{y\}}(\mu_c\text{ML})$.

- (iii) For $X \subseteq P$, the fragment $\text{Cocon}_X(\mu_c\text{ML})$ of $\mu_c\text{ML}$ -formulas that are *cocontinuous* in X :

$$\varphi ::= x \mid \alpha \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \square \varphi \mid \nu y. \varphi'$$

where $x \in X, y \in P, \alpha \in \mu_c\text{ML}$ X -free, and $\varphi' \in \text{Cocon}_{X \cup \{y\}}(\mu_c\text{ML})$. △

If one of the above fragments is subscripted by a singleton $\{x\}$, we will simply write x instead. We will use *formula* to refer to a $\mu_c\text{ML}$ -formula. We define the subformula relation \trianglelefteq and, for a given formula ξ , the sets $\text{Sf}(\xi)$ of subformulas, $\text{FV}(\xi)$ of free variables and $\text{BV}(\xi)$ of bound variables of ξ in the usual way. Given two formulas φ, ψ and a propositional variable x , we define $\varphi[\psi/x]$ to be the result of replacing each free occurrence of x in φ by ψ . We will assume an implicit mechanism of α -conversion in order to avoid the capture of free variables of ψ by binders in φ in the substitution $\varphi[\psi/x]$.

We say that a formula is *tidy* if the sets of its free and its bound variables are disjoint. A formula φ is called *clean* if, in addition, we can associate with each bound variable x , a unique fixpoint binder η_x and a unique formula δ_x such that $\eta_x. \delta_x$ is a subformula of φ . In this case, if $\eta_x = \mu$ ($\eta_x = \nu$), the variable x is said to be a μ -variable (ν -variable). We will sometimes denote by $\bar{\eta}$ the dual of η . Note that every subformula of a clean formula is itself clean. For convenience we will assume that every formula is tidy. Finally, we will use ML to refer to the basic modal language (over the set P of propositional variables).

Definition 2.2. The *FL-closure* of a set Φ of μ_c ML-formulas is the least $\Psi \supseteq \Phi$ such that:

- (i) If $\neg p \in \Psi$, then $p \in \Psi$;
- (ii) If $\varphi \circ \psi \in \Psi$ for $\circ \in \{\vee, \wedge\}$, then $\varphi, \psi \in \Psi$;
- (iii) If $\heartsuit \varphi \in \Psi$ for $\heartsuit \in \{\diamond, \square\}$, then $\varphi \in \Psi$;
- (iv) If $\eta x. \varphi \in \Psi$ for $\eta \in \{\mu, \nu\}$, then $\varphi[\eta x. \varphi/x] \in \Psi$.

We write $Cl(\Phi)$ for the FL-closure of Φ and say that Φ is *FL-closed* if $Cl(\Phi) = \Phi$. If $\Phi = \{\varphi\}$ is a singleton, we simply write $Cl(\varphi)$. \triangle

It is a well-known fact that the closure of a finite set of formulas is finite. Note, moreover, that in the FL-closure of a set of tidy formulas, every formula is tidy.

We say of a subformula $\varphi \trianglelefteq \xi$ that it is a *free* subformula of ξ , and write $\varphi \trianglelefteq_f \xi$, if $\varphi \in Cl(\xi)$. Equivalently, a subformula $\varphi \trianglelefteq \xi$ is a free subformula of ξ whenever every free variable of φ is a free variable of ξ .

Algebraic semantics As usual, formulas will be interpreted in Kripke models.

Definition 2.3. A *Kripke frame* is a pair (S, R) consisting of a set S of *states* together with an *accessibility relation* $R \subseteq S \times S$. A *Kripke model* is a triple (S, R, V) , where (S, R) is a Kripke frame and $V : P \rightarrow \mathcal{P}(S)$ a *valuation function*. \triangle

Given some accessibility relation R , we often write sRt instead of $(s, t) \in R$. The algebraic semantics of the continuous μ -calculus extends that of the basic modal language. Given a valuation function $V : P \rightarrow \mathcal{P}(S)$, we write $V[x \mapsto X]$ for the function given by $V[x \mapsto X](x) = X$ and $V[x \mapsto X](y) = V(y)$ for $y \neq x$.

Definition 2.4. We define for every formula φ its *meaning* $\llbracket \varphi \rrbracket^{\mathbb{S}} \subseteq S$ in any model $\mathbb{S} = (S, R, V)$ by the following induction on formulas:

$$\begin{aligned} \llbracket \mu x. \varphi \rrbracket^{\mathbb{S}} &:= \bigcap \{X \subseteq S : \llbracket \varphi \rrbracket^{\mathbb{S}[x \mapsto X]} \subseteq X\} \\ \llbracket \nu x. \varphi \rrbracket^{\mathbb{S}} &:= \bigcup \{X \subseteq S : X \subseteq \llbracket \varphi \rrbracket^{\mathbb{S}[x \mapsto X]}\} \end{aligned}$$

and the propositional and modal cases are as usual. \triangle

We say that ξ is *satisfied* at a state s of the model \mathbb{S} , and write $\mathbb{S}, s \Vdash \xi$ whenever $s \in \llbracket \xi \rrbracket^{\mathbb{S}}$. As usual, we say that ξ is *valid* in \mathbb{S} , written $\mathbb{S} \models \xi$, whenever ξ is satisfied at every state s of \mathbb{S} , and *valid* in the frame (S, R) , written $(S, R) \models \xi$, whenever $(S, R, V) \models \xi$ for every valuation function $V : P \rightarrow \mathcal{P}(S)$.

Two formulas are called *equivalent* whenever they have the same meaning in every Kripke model. It is easy to see that every formula has an equivalent alphabetic variant which is clean.

Game semantics A well-known equivalent characterisation of the meaning of a formula uses the formalism of infinite games. We assume familiarity with this kind of games.

Definition 2.5. Given a clean formula ξ , we define the *dependency order* $<_{\xi}$ on $BV(\xi)$ as the least strict partial order such that $x <_{\xi} y$ whenever $\delta_x \triangleleft \delta_y$ and $y \triangleleft \delta_x$. \triangle

Note that for formulas of the continuous μ -calculus $x <_{\xi} y$ implies that x is a μ -variable if and only if y is a μ -variable. In other words, the continuous modal μ -calculus is *alternation free*.

Definition 2.6. Let ξ be a clean formula and let $\mathbb{S} = (S, R, V)$ be a Kripke model. The *evaluation game* $\mathcal{E}(\xi, \mathbb{S})$ takes positions in $Sf(\xi) \times S$ and has the following ownership function and admissible moves.

Position	Player	Admissible moves
$(\varphi_1 \vee \varphi_2, s)$	\exists	$\{(\varphi_1, s), (\varphi_2, s)\}$
$(\varphi_1 \wedge \varphi_2, s)$	\forall	$\{(\varphi_1, s), (\varphi_2, s)\}$
$(\diamond \varphi, s)$	\exists	$\{(\varphi, t) : sRt\}$
$(\square \varphi, s)$	\forall	$\{(\varphi, t) : sRt\}$
$(\eta x. \delta_x, s)$	-	$\{(\delta_x, s)\}$
(x, s) with $x \in \text{BV}(\xi)$	-	$\{(\delta_x, s)\}$
(p, s) with $p \in \text{FV}(\xi)$ and $s \in V(p)$	\forall	\emptyset
$(\neg p, s)$ with $p \in \text{FV}(\xi)$ and $s \in V(p)$	\exists	\emptyset
(p, s) with $p \in \text{FV}(\xi)$ and $s \notin V(p)$	\exists	\emptyset
$(\neg p, s)$ with $p \in \text{FV}(\xi)$ and $s \notin V(p)$	\forall	\emptyset

For γ a match in $\mathcal{E}(\xi, \mathbb{S})$, we denote the first position of γ by $\text{first}(\gamma)$ and, if γ is finite, the last position by $\text{last}(\gamma)$. A finite match γ is won by one of the players whenever $\text{last}(\gamma)$ is owned by its opponent and this opponent's set of admissible moves is empty (in this case the opponent is said to have gotten *stuck*). An infinite match is won by \exists (\forall) if the $<_{\xi}$ -highest variable that is unfolded infinitely often is a v -variable (a μ -variable). We write $(\varphi, s) \in \text{Win}_{\exists}(\mathcal{E}(\xi, \mathbb{S}))$ to denote that \exists has a winning strategy in the game $\mathcal{E}(\xi, \mathbb{S})$ initialised at position (φ, s) . \triangle

The following lemma contains some basic facts about the course of play in evaluation games for $\mu_c\text{ML}$. Items (1) and (2) hold because the continuous μ -calculus is alternation free. Item (3) is specific to the continuous modal μ -calculus, in the sense that it does not hold for the more expressive *alternation free* μ -calculus (see e.g. [12] for a formal definition of this language).

Lemma 2.7. *Let \mathbb{S} be a model and let ξ be a clean formula.*

1. *In any infinite match of the game $\mathcal{E}(\xi, \mathbb{S})$, either all variables that are unfolded infinitely often are μ -variables, or all are v -variables.*
2. *If a match of the game $\mathcal{E}(\xi, \mathbb{S})$ progresses from a position $(s, \eta x. \delta)$ to a position $(t, \bar{\eta} y. \theta)$, then in between it must pass a position (r, φ) with $\varphi \triangleleft_f \xi$.*
3. *If a match of the game $\mathcal{E}(\xi, \mathbb{S})$ progresses from a position $(s, \mu x. \delta)$ to a position $(t, \square \psi)$, then in between it must pass a position (r, φ) with $\varphi \triangleleft_f \xi$.*

We say of an enumeration $\{x_1, \dots, x_n\}$ of $\text{BV}(\xi)$ that it *respects* the dependency order if $x_i <_{\xi} x_j$ implies $i < j$. Since any partial order can be extended to a linear order, every formula ξ admits an enumeration of its bound variables that respects the dependency order. For the rest of this paper we fix such an enumeration of $\text{BV}(\xi)$ for every clean formula ξ .

Definition 2.8. Let ξ be a clean formula with $\text{BV}(\xi) = \{x_1, \dots, x_n\}$. For any subformula $\varphi \trianglelefteq \xi$, we define its *expansion* with respect to ξ as:

$$\text{exp}_{\xi}(\varphi) := \varphi[\eta x_1 \delta_{x_1}/x_1] \cdots [\eta x_n \delta_{x_n}/x_n]. \quad \triangle$$

Note that when $\varphi \triangleleft_f \xi$, it holds that $\text{exp}_{\xi}(\varphi) = \varphi$. The following well-known theorem provides the central link between the algebraic and the game semantics.

Theorem 2.9. *For any clean formula ξ and subformula $\varphi \trianglelefteq \xi$ it holds that:*

$$(\varphi, s) \in \text{Win}_{\exists}(\mathcal{E}(\xi, \mathbb{S})) \Leftrightarrow \mathbb{S}, s \Vdash \text{exp}_{\xi}(\varphi).$$

for any model \mathbb{S} and state s of \mathbb{S} .

In particular, for any clean formula ξ and $\varphi \sqsubseteq_f \xi$ we have $\mathbb{S}, s \Vdash \varphi$ if and only if \exists has a winning strategy in the game $\mathcal{E}(\xi, \mathbb{S})$ initialised at the position (φ, s) .

Another useful fact, originally provided by Dexter Kozen in [9], is the following.

Proposition 2.10. *For any clean formula ξ :*

$$Cl(\xi) = \{\text{exp}_\xi(\varphi) : \varphi \sqsubseteq \xi\}.$$

Axiomatisation We give an axiomatisation of the continuous modal μ -calculus based on an axiomatisation introduced by Dexter Kozen for the full modal μ -calculus in [9].

Definition 2.11. The logic $\mu_c\text{K}$ is the least logic containing the following axioms and closed under the following rules.¹

Axioms.

1. A complete set of axioms for classical propositional logic.
2. Normality: $\neg \diamond \perp$.
3. Additivity: $\diamond(p \vee q) \leftrightarrow (\diamond p \vee \diamond q)$.
4. For every $\varphi \in \text{Con}_x(\mu_c\text{ML})$, the prefixpoint axiom:

$$\varphi[\mu x. \varphi/x] \rightarrow \mu x. \varphi.$$

Rules.

1. Modus Ponens: from $\varphi \rightarrow \psi$ and φ , derive ψ .
2. Monotonicity: from $\varphi \rightarrow \psi$, derive $\diamond \varphi \rightarrow \diamond \psi$.
3. Uniform Substitution: from φ , derive $\varphi[\psi/x]$.
4. The least prefixpoint rule: from $\varphi[\gamma/x] \rightarrow \gamma$ with $\varphi \in \text{Con}_x(\mu_c\text{ML})$, derive $\mu x. \varphi \rightarrow \gamma$. △

We will consider axiomatic extensions of $\mu_c\text{K}$ that are closed under the rules above. We will use μ_c -logic to refer to such an extension. The term *logic* will be used to refer to any normal modal logic. If L is a logic in the basic modal language, we use μ_c - L to denote the least μ_c -logic containing L . Moreover, we will use $\text{Mod}(L)$ ($\text{Fr}(L)$) to denote the class of models (frames) on which every formula in L is valid. If (S, R, V) belongs to $\text{Mod}(L)$ ((S, R) belongs to $\text{Fr}(L)$) we say that (S, R, V) is an L -model ((S, R) is an L -frame) and write $(S, R, V) \models L$ ($(S, R) \models L$).

3 Filtration

Filtration is a well-known method in the theory of basic modal logic. In this section we define filtration and related notions for the continuous modal μ -calculus and show that some of their most important properties transfer to this more expressive language.

Filtration

Definition 3.1. Let $\mathbb{S} = (S, R, V)$ be a Kripke model and let Σ be a finite and FL-closed set of formulas. Let $\sim_\Sigma^{\mathbb{S}}$ be the equivalence relation given by:

$$s \sim_\Sigma^{\mathbb{S}} s' \text{ if and only if } \mathbb{S}, s \Vdash \varphi \Leftrightarrow \mathbb{S}, s' \Vdash \varphi \text{ for all } \varphi \in \Sigma.$$

A Σ -filtration of \mathbb{S} through Σ is a model $\mathbb{S}^\Sigma = (S^\Sigma, R^\Sigma, V^\Sigma)$ such that:

¹Because we have defined $\mu_c\text{ML}$ in negation normal form, we formally also need to add the dual version of each axiom and rule. Moreover, we should have rules expressing that \diamond and \square and, respectively, μ and ν are duals. For reasons of space and clarity we omit these technical details.

- (i) $S^\Sigma = S / \sim_\Sigma^{\mathbb{S}}$
- (ii) $R^{\min} \subseteq R^\Sigma \subseteq R^{\max}$;
- (iii) $V^\Sigma(p) = \{\bar{s} : s \Vdash p\}$ for every propositional variable $p \in \Sigma$.

where:

$$R^{\min} := \{(\bar{s}, \bar{t}) : \text{there are } s' \sim_\Sigma^{\mathbb{S}} s \text{ and } t' \sim_\Sigma^{\mathbb{S}} t \text{ such that } Rs't'\},$$

$$R^{\max} := \{(\bar{s}, \bar{t}) : \text{for all } \Box\varphi \in \Sigma; \text{ if } s \Vdash \Box\varphi, \text{ then } t \Vdash \varphi\}.$$

where \bar{s} denotes the equivalence class with representative s . △

The relation R^{\min} will be called the *finest filtration* and the relation R^{\max} the *coarsest*.

Filtration Theorem for the continuous modal μ -calculus If f is a strategy for the player \exists (\forall) in a game \mathcal{G} , we say of a (possibly infinite) \mathcal{G} -match γ that it is *f-guided* whenever every choice made by \exists (\forall) in the match γ is the choice dictated by the strategy f .

Theorem 3.2 (Filtration Theorem). *Let Σ be a finite and FL-closed set of formulas and let $\mathbb{S} = (S, R, V)$ be a Kripke model. For every filtration $\bar{\mathbb{S}} = (\bar{S}, \bar{R}, \bar{V})$ of \mathbb{S} through Σ it holds that*

$$\mathbb{S}, s \Vdash \xi \Leftrightarrow \bar{\mathbb{S}}, \bar{s} \Vdash \xi,$$

for every clean formula $\xi \in \Sigma$.

Proof. Because negation is definable in our language, it suffices to prove the implication in just one direction, which in our case will be the direction \Rightarrow . Throughout this proof we will write \mathcal{G} for the game $\mathcal{E}(\xi, \mathbb{S})$ and $\bar{\mathcal{G}}$ for the game $\mathcal{E}(\xi, \bar{\mathbb{S}})$. As hypothesis we assume that \exists has a winning strategy f in the game \mathcal{G} initialised at position (ξ, s) ; we wish to show that $(\xi, \bar{s}) \in \text{Win}_\exists(\bar{\mathcal{G}})$.

The main idea of the proof is to obtain a winning strategy for \exists in $\bar{\mathcal{G}}$ by playing a ‘shadow match’ in \mathcal{G} . That is, we will simulate in \mathcal{G} every move played by \forall in our $\bar{\mathcal{G}}$ -match, and, to determine a move for \exists in $\bar{\mathcal{G}}$, we copy the move dictated in \mathcal{G} by the strategy f . If we manage to do this, then whenever the match in $\bar{\mathcal{G}}$ is at some position (φ, \bar{s}) , the shadow match in \mathcal{G} will be at a position (φ, s) (note that this is indeed the case for the initial positions). It turns out that this works well for all positions, except those of the form $(\Box\varphi, \bar{s})$. At those positions, a problem arises when \forall chooses a position (φ, \bar{t}) such that $\bar{s}\bar{R}\bar{t}$, but not sRt . This move by \forall in $\bar{\mathcal{G}}$ can then not be simulated in the shadow match, because (φ, t) is not an admissible move for \forall in \mathcal{G} . However, using the fact that $\bar{R} \subseteq R^{\max}$, we will be able to show that if $\bar{s}\bar{R}\bar{t}$ and $(\Box\varphi, s) \in \text{Win}_\exists(\mathcal{G})$, then $(\varphi, t) \in \text{Win}_\exists(\mathcal{G})$. We will use this to initiate a new shadow match in \mathcal{G} whenever encounter a position of the form $(\Box\varphi, \bar{s})$. A key observation will be that we only need to initiate a new shadow match at most finitely many times, because formulas of the form $\Box\varphi$ do not occur in the scope of least fixed point operators in the language $\mu_c\text{ML}$.

More formally, we say that for $I \in \omega \cup \{\omega\}$, a $\bar{\mathcal{G}}$ -match $\bar{\gamma} = (\varphi_i, \bar{t}_i)_{i \in I}$ is *linked* to some \mathcal{G} -match $\gamma = (\psi_i, s_i)_{i \in I}$ whenever for every $i \in I$ it holds that $\varphi_i = \psi_i$ and $\bar{s}_i = \bar{t}_i$. Moreover, we say that $\bar{\gamma}$ *follows* γ whenever some final segment of $\bar{\gamma}$ is linked to γ .

Claim. Let $\bar{\gamma}$ be a finite $\bar{\mathcal{G}}$ -match that follows some f -guided \mathcal{G} -match γ , where f is a winning strategy for \mathcal{G} initialised at $\text{first}(\gamma)$. Then:

- If the formula in $\text{last}(\gamma)$ is not of the form $\Box\theta$, then \exists can ensure that after the next round in $\bar{\mathcal{G}}$, there is some admissible move (ψ_{n+1}, t_{n+1}) in \mathcal{G} such that the resulting $\bar{\mathcal{G}}$ -match follows the \mathcal{G} -match $\gamma \cdot (\psi_{n+1}, t_{n+1})$ and the latter remains f -guided.

- If the formula in $\text{last}(\gamma)$ is of the form $\Box\theta$, then \exists can at least ensure that after the next round, the resulting $\overline{\mathcal{G}}$ -match follows a new match (θ, t_0) for which \exists has a winning strategy.

The above claim is proven by a case distinction on the main connective of the formula in $\text{last}(\gamma)$. We treat the most difficult cases of \Diamond and \Box , leaving the rest to the reader.

Suppose $\text{last}(\gamma)$ is of the form $(\Diamond\theta, t_n)$. Let (θ, t_{n+1}) be the next move instructed by the assumed winning strategy f . Then $t_n R t_{n+1}$ and thus, because $\overline{R} \subseteq R^{\min}$ and $s_n \sim t_n$, we have $\overline{s_n} R \overline{t_{n+1}}$. Therefore \exists can simply choose the position $(\theta, \overline{t_{n+1}})$.

If $\text{last}(\gamma)$ is of the form $(\Box\theta, t_n)$, consider the move $(\theta, \overline{s_{n+1}})$ chosen by \forall in $\overline{\mathcal{G}}$. We have,

$$\begin{aligned}
(\Box\theta, t_n) \in \text{Win}_{\exists}(\mathcal{G}) &\Rightarrow \mathbb{S}, t_n \Vdash \text{exp}_{\xi}(\Box\theta) && \text{(Theorem 2.9)} \\
&\Rightarrow \mathbb{S}, s_n \Vdash \text{exp}_{\xi}(\Box\theta) && (\text{exp}_{\xi}(\Box\theta) \in \Sigma \text{ and } s_n \sim t_n) \\
&\Rightarrow \mathbb{S}, s_n \Vdash \Box \text{exp}_{\xi}(\theta) && \text{(Definition of exp)} \\
&\Rightarrow \mathbb{S}, s_{n+1} \Vdash \text{exp}_{\xi}(\theta) && (\Box \text{exp}_{\xi}(\theta) \in \Sigma \text{ and } \overline{s_n} R^{\max} \overline{s_{n+1}}) \\
&\Rightarrow (\theta, s_{n+1}) \in \text{Win}_{\exists}(\mathcal{G}). && \text{(Theorem 2.9)}
\end{aligned}$$

Thus we may choose (θ, s_{n+1}) as the new match that is followed by $\overline{\gamma} \cdot (\theta, \overline{s_{n+1}})$.

Using the fact that (ξ, s) is linked to (ξ, \overline{s}) as induction base, and the above claim as induction step, we obtain a strategy g for \exists in $\overline{\mathcal{G}}$ initialised at (ξ, \overline{s}) . We claim that g is a winning strategy. Indeed, if a g -guided match $\overline{\gamma}$ ends in finitely many steps, then either \forall got stuck on a formula of the form $\Box\theta$, or the final position is of the form (π, \overline{t}) for some $\pi = p, \neg p \in \text{FV}(\xi)$. Without loss of generality, suppose $\pi = p$. By construction $\overline{\gamma}$ follows a \mathcal{G} -match γ such that $\text{last}(\gamma) = (\pi, r) \in \text{Win}_{\exists}(\mathcal{G})$ for some state $r \in S$ with $\overline{r} = \overline{t}$. But this means that $r \in V(p)$ and thus, since $p \in \Sigma$, also $\overline{t} \in \overline{V}(p)$. Hence \exists indeed wins the match $\overline{\gamma}$.

If a g -guided match $\overline{\gamma}$ lasts infinitely long, then by item (1) of Lemma 2.7, there must be some point after which either only μ -variables, or only ν -variables, are unfolded. In the latter case the match is indeed winning for \exists . We will now argue that this is the only possibility, because the former case cannot occur. The reason is that if from some point on in $\overline{\gamma}$ only μ -variables are unfolded, then the syntax of μ_c ML dictates that from some point on in $\overline{\gamma}$ no formula of the form $\Box\theta$ will occur. By construction, this means that the infinite $\overline{\mathcal{G}}$ -match $\overline{\gamma}$ follows an infinite \mathcal{G} -match γ which is guided by a winning strategy for \exists . But this is a contradiction, because the match γ , by the fact that it is linked to an infinite final segment of $\overline{\gamma}$, contains infinitely many μ -unfoldings. \square

Note that the above argument would not go through for the alternation free μ -calculus, since we would no longer be able to guarantee that we create at most finitely many shadow matches in the case of infinitely many μ -unfoldings. A well-known counterexample to the Filtration Theorem for the alternation free μ -calculus is the formula $\mu x. \Box x$.

Admissibility of filtration Having established that filtrations preserve satisfaction of μ_c ML-formulas, we will now investigate to which classes of models filtration can be applied.

Definition 3.3. A class of models \mathcal{M} is said to *admit filtration* with respect to a language D if for every model \mathbb{S} in \mathcal{M} and every finite FL-closed set of D -formulas Σ , the class \mathcal{M} contains a filtration of \mathbb{S} through Σ . A class of frames \mathcal{F} is said to *admit filtration* if the class of models $\{(S, R, V) : (S, R) \in \mathcal{F}\}$ does. \triangle

One might expect that admitting filtration with respect to the basic modal language is a weaker property than admitting filtration with respect to a proper extension of the language. However, for the language $\mu_c\text{ML}$ it turns out that this is not the case, at least for those classes of models that are determined by some logic.

We will show this by making use of the following technical sufficient condition.

Lemma 3.4. *Let \mathcal{M} be a class of models that admits filtration wrt ML. Suppose that for every model $\mathbb{S} := (S, R, V) \in \mathcal{M}$ and finite FL-closed set $\Sigma \subset \mu_c\text{ML}$, there is a valuation $V' : P \rightarrow \mathcal{P}(S)$ and a translation $\tau : \Sigma \rightarrow \text{ML}$ such that:*

1. $V'(p) = p$ for all $p \in P$ occurring in Σ ;
2. The model $\mathbb{S}' := (S, R, V')$ belongs to \mathcal{M} .
3. The set $\tau[\Sigma] \subset \text{ML}$ is FL-closed.
4. The translation τ commutes with \Box , i.e. $\tau(\Box\varphi) = \Box\tau(\varphi)$ for all $\Box\varphi \in \Sigma$.
5. For every $\xi \in \Sigma$ and $s \in S$ it holds that: $\mathbb{S}, s \Vdash \xi \Leftrightarrow \mathbb{S}', s \Vdash \tau(\xi)$.

Then \mathcal{M} admits filtration wrt $\mu_c\text{ML}$.

Proof. Using conditions (2) and (3) and the assumption that \mathcal{M} admits filtration with respect to ML, there is a filtration $\mathbb{S}^{\tau[\Sigma]} \in \mathcal{M}$ of \mathbb{S}' through $\tau[\Sigma]$. We claim that $\mathbb{S}^{\tau[\Sigma]}$ simultaneously is a filtration of \mathbb{S} through Σ .

By assumption (5), the equivalence relations $\sim_{\Sigma}^{\mathbb{S}}$ and $\sim_{\tau[\Sigma]}^{\mathbb{S}'}$ on S coincide. From this we obtain condition (i) of Definition 3.1, as well as the first inclusion of condition (ii). For the second inclusion, suppose that $\bar{\exists}R^{\tau[\Sigma]}\bar{t}$ and $\mathbb{S}, s \Vdash \Box\varphi$ for some $\Box\varphi \in \Sigma$. We must show that $\mathbb{S}, t \Vdash \varphi$. By assumption (5), we have $\mathbb{S}', s \Vdash \tau(\Box\varphi)$ and thus, by assumption (4), also $\mathbb{S}', s \Vdash \Box\tau(\varphi)$. Since $R^{\tau[\Sigma]}$ is contained in the coarsest filtration of \mathbb{S}' through $\tau[\Sigma]$ and $\Box\tau(\varphi) = \tau(\Box\varphi) \in \tau[\Sigma]$, we obtain $\mathbb{S}', t \Vdash \tau(\varphi)$. Applying the other direction of assumption (5), we obtain $\mathbb{S}, t \Vdash \varphi$, as required. Finally, condition (iii) follows directly from assumption (1). \square

The proof of the following lemma resembles that of Theorem 3.8 in [8].

Lemma 3.5. *For any logic L, the class $\text{Mod}(L)$ admits filtration wrt ML iff it admits filtration wrt $\mu_c\text{ML}$.*

Proof. The implication from right to left is trivial. For the other direction we will use Lemma 3.4. Let $\mathbb{S} = (S, R, V)$ be a model such that $\mathbb{S} \models L$ and let Σ be a finite FL-closed set of $\mu_c\text{ML}$ -formulas. Let $\varphi_1, \dots, \varphi_n$ be an enumeration of formulas of the form $\eta x.\psi$ in Σ . For every such formula φ_i , we pick a unique propositional variable p_i not occurring in Σ .

We define the following alternative valuation $V' : P \rightarrow \mathcal{P}(S)$.

$$V'(p) := \begin{cases} \llbracket \varphi_i \rrbracket^{\mathbb{S}} & \text{if } p = p_i \text{ for some } \varphi_i \in \Sigma; \\ V(p) & \text{otherwise,} \end{cases}$$

and define $\mathbb{S}' := (S, R, V')$. A straightforward induction on formulas now shows that for every formula $\xi \in \mu_c\text{ML}$ and state $s \in S$:

$$\mathbb{S}', s \Vdash \xi \Leftrightarrow \mathbb{S}', s \Vdash \xi[\varphi_1/p_1] \cdots [\varphi_n/p_n]. \quad (1)$$

We claim that $\mathbb{S}' \in \text{Mod}(\text{L})$. Indeed, we have

$$\begin{aligned}
\xi \in \text{L} &\Rightarrow \xi[\varphi_1/p_1] \cdots [\varphi_n/p_n] \in \text{L} && (\text{L is closed under uniform substitution}) \\
&\Rightarrow \mathbb{S} \models \xi[\varphi_1/p_1] \cdots [\varphi_n/p_n] && (\mathbb{S} \models \text{L}) \\
&\Rightarrow \mathbb{S}' \models \xi[\varphi_1/p_1] \cdots [\varphi_n/p_n] && (V \text{ and } V' \text{ agree on all relevant propositional variables}) \\
&\Rightarrow \mathbb{S}' \models \xi && (\text{by (1) from right to left})
\end{aligned}$$

Now let the translation $\tau : \Sigma \rightarrow \text{ML}$ be the translation that commutes with all propositional and modal symbols, and acts on fixpoint operators in the following way:

$$\tau(\eta x. \psi) := p_i \text{ where } \eta x. \psi = \varphi_i.$$

We leave it to the reader to verify that $\tau[\Sigma]$ is FL-closed. Finally, another straightforward induction shows that for every formula $\xi \in \Sigma$ and state $s \in \mathbb{S}$:

$$\mathbb{S}, s \Vdash \xi \Leftrightarrow \mathbb{S}', s \Vdash \tau(\xi).$$

This finishes the proof, for all conditions of Lemma 3.4 are met. \square

Note that the above proof does not rely on any specific properties of the language $\mu_c\text{ML}$. In fact, it could also have been carried out for the full language μML of the modal μ -calculus. As a corollary, we obtain the finite model property.

Corollary 3.6 (Finite Model Property). *Let L be a logic such that $\text{Mod}(\text{L})$ admits filtration with respect to ML, and let ϕ be a formula of the continuous μ -calculus. Then ϕ is valid in every L-model if and only if ϕ is valid in every finite L-model.*

Proof. Let ϕ be a formula such that $\mathbb{S} \not\models \phi$ for some $\mathbb{S} \models \text{L}$. Without loss of generality, we may assume that ϕ is clean. Letting $\Sigma = \text{Cl}(\phi)$, there is, by Lemma 3.5 and the fact that $\text{Mod}(\text{L})$ admits filtration, a filtration \mathbb{S}^Σ of \mathbb{S} through Σ such that $\mathbb{S}^\Sigma \models \text{L}$. Observe that number of states of \mathbb{S}^Σ is at most $2^{|\text{Cl}(\phi)|}$ and thus finite. By Theorem 3.2, it holds that $\mathbb{S}^\Sigma \not\models \phi$, as required. \square

For instance, since the class of symmetric models is the class of KB-models, the continuous modal μ -calculus has the finite model property over this class.

4 Canonical completeness

In this section we prove our completeness result. In the first paragraph, we will define the finitary canonical models of an arbitrary μ_c -logic L and prove the Truth Lemma. In the second paragraph we will show that a finitary canonical model can be obtained for the logic $\mu_c\text{-L}$, where L is any canonical basic modal logic such that $\text{Fr}(\text{L})$ admits filtration. As a direct consequence we obtain that $\mu_c\text{-L}$ is sound and complete with respect to $\text{Fr}(\text{L})$.

Finitary canonical models For the entirety of this paragraph we fix an arbitrary μ_c -logic L. We define the negation operator $\sim : \mu_c\text{ML} \rightarrow \mu_c\text{ML}$ in the usual way. In particular, that means that we define $\sim \eta x. \phi := \bar{\eta} x. \sim \phi[\neg x/x]$. We leave it to the reader to verify that $\text{L} \vdash (\sim \phi \wedge \phi) \leftrightarrow \perp$ and $\text{L} \vdash (\sim \phi \vee \phi) \leftrightarrow \top$.

Definition 4.1. Let Σ be a set of formulas. If for all $\phi \in \Sigma$ it holds that $\sim \phi \in \Sigma$, then Σ is said to be \sim -closed. \triangle

We say that Σ is \sim FL-closed if it is both FL-closed and \sim -closed. Note that for every finite set of μ_c ML-formulas, the \sim -closure of its FL-closure is a finite \sim FL-closed extension.

A set Γ of formulas is said to be *L-inconsistent* if $\mathsf{L} \vdash (\gamma_1 \wedge \dots \wedge \gamma_n) \rightarrow \perp$ for some $\gamma_1, \dots, \gamma_n \in \Gamma$. We say of a formula φ that it is L-inconsistent whenever $\{\varphi\}$ is.

Definition 4.2. A set of formulas Γ is called *maximally L-consistent* if it is consistent and maximal in that respect, *i.e.* for every other set of formulas Γ' :

$$\text{If } \Gamma \subset \Gamma', \text{ then } \Gamma' \text{ is L-inconsistent.} \quad \triangle$$

By a standard argument it can be shown that every L-consistent set of formulas has a maximally L-consistent extension. The proof of the following lemma is also standard and left to the reader.

Lemma 4.3. *Let Γ be a maximally L-consistent set. Then:*

- (i) *If $\mathsf{L} \vdash \varphi$, then $\varphi \in \Gamma$;*
- (ii) *$\sim\varphi \in \Gamma$ if and only if $\varphi \notin \Gamma$;*
- (iii) *$\varphi \vee \psi \in \Gamma$ if and only if $\varphi \in \Gamma$ or $\psi \in \Gamma$;*
- (iv) *$\mu x.\varphi \in \Gamma$ if and only if $\varphi[\mu x.\varphi/x] \in \Gamma$.*

Definition 4.4. Let Σ be a finite \sim FL-closed set of formulas. A *model over Σ with respect to L* is any model (S, R, V) such that:

- $S = \{\Gamma \cap \Sigma : \Gamma \text{ is maximally L-consistent}\}$.
- $R^{\min} \subseteq R \subseteq R^{\max}$, where:

$$\begin{aligned} AR^{\min}B &: \Leftrightarrow \bigwedge A \wedge \diamond B \text{ is L-consistent} \\ AR^{\max}B &: \Leftrightarrow \text{for all } \Box\varphi \in \Sigma : \Box\varphi \in A \Rightarrow \varphi \in B. \end{aligned}$$

- $V(p) = \{s \in S : p \in s\}$ for all $p \in \Sigma$. \triangle

For A some finite set of formulas, we will usually write ψ_A for the conjunction $\bigwedge A$. In the following we will assume a fixed model over some finite and \sim FL-closed set Σ with respect to L, which will be denoted by $\mathbb{S}^\Sigma = (S^\Sigma, R^\Sigma, V^\Sigma)$. We will often drop the superscript Σ 's and \mathbb{S} 's. Moreover, if in the following we refer to provability or consistency, this will be tacitly assumed to be in the logic L.

The following existence lemma is standard in the context of (finitary) canonical models for modal logics.

Lemma 4.5. *For any formula $\varphi \in \mu_c$ ML and state $A \in S$:*

$$\psi_A \wedge \diamond\varphi \text{ is consistent if and only if } \psi_B \wedge \varphi \text{ is consistent for some } ARB.$$

In particular, it follows that for all $\diamond\varphi \in \Sigma$ we have $\diamond\varphi \in A$ if and only if $\varphi \in B$ for some ARB . The following lemma follows from the fact that Σ is \sim -closed.

Lemma 4.6. *For every $A, B \in S$ it holds that $\psi_A \wedge \psi_B$ is consistent iff $A = B$.*

Given a finite collection U of finite sets of formulas, we write ψ_U for the disjunction of all ψ_X for $X \in U$, *i.e.*

$$\psi_U = \bigvee_{X \in U} \psi_X.$$

Note that by the previous lemma, for any $U \subseteq S$ and $A \in S$, the formula $\psi_U \wedge \psi_A$ is consistent if and only if $A \in U$.

We wish to prove the following lemma.

Lemma 4.7. (*Truth Lemma*) *If $A \in \mathcal{S}$ and $\xi \in \Sigma$ is clean, then*

$$\xi \in A \Rightarrow A \in \llbracket \xi \rrbracket. \quad (\text{T})$$

We shall prove this by a double induction on formulas, of which the inner induction is captured by Lemma 4.10.

Definition 4.8. Let ξ a formula with $BV(\xi) = \{x_1, \dots, x_n\}$. We define the *name-expansion* $n\text{-exp}_\xi^{\mathbb{S}}(\varphi)$ of a subformula φ of ξ in \mathbb{S} as follows:

$$n\text{-exp}_\xi^{\mathbb{S}}(\varphi) := \varphi[\psi_{U_1}/x_1] \cdots [\psi_{U_n}/x_n],$$

where $U_i := \llbracket \exp_\xi(\delta_{x_i}) \rrbracket^{\mathbb{S}}$ for every $1 \leq i \leq n$. △

Whenever clear from context, we drop the subscript and superscript from $n\text{-exp}_\xi^{\mathbb{S}}$. The main property of name-expansions that we will use is the following.

Lemma 4.9. *For any clean formula ξ and bound μ -variable $x_i \in BV(\xi)$:*

$$\text{If } \mathbb{L} \vdash n\text{-exp}(\delta_{x_i}) \rightarrow \psi_{U_i}, \text{ then } \mathbb{L} \vdash n\text{-exp}(\mu x_i. \delta_{x_i}) \rightarrow \psi_{U_i}.$$

Proof. Let χ be the formula $n\text{-exp}_\xi(\delta_{x_i})$, but without the substitution $[\psi_{U_i}/x_i]$. Then the to-be-proven implication becomes:

$$\text{If } \mathbb{L} \vdash \chi[\psi_{U_i}/x_i] \rightarrow \psi_{U_i}, \text{ then } \mathbb{L} \vdash \mu x_i. \chi \rightarrow \psi_{U_i},$$

but this is simply an application of the least prefixpoint rule. □

Lemma 4.10. *Let ξ be a clean formula in Σ such that for every free strict subformula of ξ the implication (T) holds. Then for every subformula of ξ of the form $\mu x_i. \delta_{x_i}$ it holds that:*

$$\mathbb{L} \vdash n\text{-exp}(\mu x_i. \delta_{x_i}) \rightarrow \psi_{U_i}.$$

Proof. We proceed by induction on the complexity of subformulas of ξ . Let $\mu x_i. \delta_{x_i}$ be a subformula of ξ and suppose, as inductive hypothesis, that the thesis holds for every strict subformula of $\mu x_i. \delta_{x_i}$ (of the form $\mu x_j. \delta_{x_j}$). By Lemma 4.9, it suffices to show that $\mathbb{L} \vdash n\text{-exp}(\delta_{x_i}) \rightarrow \psi_{U_i}$. For this, in turn, it is enough to show that:

$$\text{For any } A \in \mathcal{S}^\Sigma \text{ such that } \psi_A \wedge n\text{-exp}(\delta_{x_i}) \text{ is consistent, it holds that } A \Vdash \exp(\delta_{x_i}). \quad (2)$$

This is because if $\mathbb{L} \vdash n\text{-exp}(\delta_{x_i}) \rightarrow \psi_{U_i}$ were not the case, then $n\text{-exp}(\delta_{x_i}) \wedge \neg \psi_{U_i}$ would be consistent. It follows that there is a maximally \mathbb{L} -consistent set Γ extending this formula. Letting $A := \Gamma \cap \Sigma$, we obtain that $\psi_A \wedge n\text{-exp}(\delta_{x_i})$ is consistent, but $A \not\Vdash \exp(\delta_{x_i})$, for $A \wedge \psi_{U_i}$ is inconsistent.

Applying Theorem 2.9, we will show (2) by constructing a winning strategy for \exists in $\mathcal{G} := \mathcal{E}(\xi, \mathbb{S})$ initialised at (δ_{x_i}, A) . The idea is to show that \exists has a strategy f ensuring for some initial segment of the match that at each position (θ, B) reached, it holds that $\theta \trianglelefteq \delta_{x_i}$ and the conjunction $\psi_B \wedge n\text{-exp}(\theta)$ is consistent. For the rest of this proof we shall call such a position *good*. We then show that this sequence of good positions eventually leads to a good position (θ, B) such that one of the following holds:

- (i) θ is a free subformula of ξ .
- (ii) θ is of the form $\mu x_j. \delta_{x_j}$.
- (iii) $\theta = x_j$ for some bound variable x_j of ξ .

Good positions of this form will be called *perfect*. The proof rests on the following three claims:

Claim 1. If (θ, B) is the last position of some finite \mathcal{G} -match γ of which every position is good, but not perfect, then \exists can ensure that the next position will also be good.

Claim 2. There can be no infinite match of which every position is good, but not perfect.

Claim 3. Any perfect position is winning for \exists in \mathcal{G} .

Suppose we have established these three claims. Then, since the initial position of \mathcal{G} is good by assumption, it follows from Claim 1 that \exists can maintain this property until a perfect position is reached. By Claim 2, such a position must be reached after finitely many steps, from which, by Claim 3, there must be some strategy that \exists can take on in order to win the match.

Proof of Claim 1. By a case distinction on the shape of θ , we will show that \exists can ensure that the next position will also be good. First note that θ cannot have a main connective in $\{\square, \vee\}$, because then, by items (2) and (3) Lemma 2.7, the match γ must have passed through some formula $\alpha \triangleleft_f \xi$. This is impossible since every position in γ is assumed not to be perfect. Moreover, the formula θ can neither be a bound nor a free variable of ξ , for in both cases the position (θ, B) would be perfect. Finally, by the same reason it cannot be the case that θ is of the form $\mu x_j. \delta_{x_j}$. This leaves the following three cases:

- θ is of the form $\theta_1 \vee \theta_2$. Then $\psi_B \wedge n\text{-exp}(\theta_1 \vee \theta_2) = \psi_B \wedge (n\text{-exp}(\theta_1) \vee n\text{-exp}(\theta_2))$ is consistent, so for some $k \in \{1, 2\}$ it must hold that $\psi_B \wedge n\text{-exp}(\theta_k)$ is consistent. We let \exists choose accordingly.
- θ is of the form $\theta_1 \wedge \theta_2$. Then $\psi_B \wedge n\text{-exp}(\theta_1 \wedge \theta_2)$ is consistent. It follows that both the formulas $\psi_B \wedge n\text{-exp}(\theta_1)$ and $\psi_B \wedge n\text{-exp}(\theta_2)$ are consistent. Thus both moves available to \forall result in good positions.
- θ is of the form $\diamond \delta$. Then $n\text{-exp}(\theta) = \diamond n\text{-exp}(\delta)$, so by the existence lemma there is some *BRC* such that $(C, n\text{-exp}(\delta))$ is good, which we let \exists choose.

Proof of Claim 2. This follows from the fact that any infinite \mathcal{G} -match must pass through some bound variable of ξ .

Proof of Claim 3. Let (θ, B) be a perfect position in \mathcal{G} . We consider the three different types of perfect positions one-by-one.

- (i) In this case we have $\theta \triangleleft_f \xi$, which means that $\theta \in Cl(\xi) \subseteq \Sigma$. Moreover, since (θ, B) is good, it holds that $\psi_B \wedge \theta$ is consistent, whence $\theta \in B$. The lemma's hypothesis gives $B \Vdash \theta$, supplying \exists with the required strategy.
- (ii) In this case θ is of the form $\mu x_j. \delta_{x_j}$. By the fact that (θ, B) is good, we have that the formula $\psi_B \wedge n\text{-exp}(\mu x_j. \delta_{x_j})$ is consistent and $\theta \triangleleft \delta_{x_j}$, hence $\theta \triangleleft \mu x_j. \delta_{x_j}$. Therefore, we can apply the induction hypothesis to conclude that $\psi_B \wedge \psi_{U_j}$ is consistent. It follows that $B \in \llbracket \text{exp}(\delta_{x_j}) \rrbracket$, so an application of Theorem 2.9 gives the required strategy for \exists .
- (iii) θ is a bound variable x_j of ξ . Then the fact that $\psi_B \wedge n\text{-exp}(x_j)$ is consistent implies that $B \in U_{x_j}$, from which we can obtain the required strategy for \exists in the same way as in the previous case.

This finishes the proof of Lemma 4.10 □

Proof of Lemma 4.7. We proceed by induction on ξ . Suppose that the thesis holds for all subformulas of ξ . We will show that \exists has a winning strategy in the game $\mathcal{E}(\xi, \mathbb{S})$ initialised at (ξ, A) .

The point is that \exists can initially ensure that at each position (θ, B) reached, the formula $\psi_B \wedge \text{exp}_\xi(\theta)$ is consistent (note that by hypothesis this is the case for the initial position). Let γ be a match where \exists employs this strategy. If at some point in γ a μ -formula is reached, let $(\mu x_i. \delta_{x_i}, B)$ be the first such position. The syntactic restrictions on μ_c ML ensure that $\mu x_i. \delta_{x_i}$ will be a free subformula of ξ , whence $n\text{-exp}(\mu x_i. \delta_{x_i}) = \mu x_i. \delta_{x_i}$. Therefore we can invoke Lemma 4.10 to obtain $A \in \llbracket \text{exp}(\mu x_i. \delta_{x_i}) \rrbracket$. Theorem 2.9 supplies \exists with a strategy to follow from here on out.

Now suppose that no μ -formula is reached in some complete match γ . If γ is infinite, it must be winning for \exists . Finally, if γ is finite, the player \forall must have gotten stuck, or at some point a free variable of ξ is reached. The latter is also winning for \exists because of the assumption that for every position (θ, B) reached, it holds that $\psi_B \wedge \text{exp}(\theta)$ is consistent. \square

Completeness The goal of this paragraph is to prove completeness for certain well-behaved μ_c -logics. Given a logic L , we define its canonical model as usual.

Definition 4.11. The *canonical model* $\mathbb{S}^L := (S^L, R^L, V^L)$ of a logic L is given by:

- $S^L := \{\Gamma : \Gamma \text{ is maximally } L\text{-consistent}\}.$
- $V^L(p) := \{\Gamma : p \in \Gamma\}.$
- $\Gamma R^L \Delta :\Leftrightarrow (\Box \varphi \in \Gamma \Rightarrow \varphi \in \Delta).$ \triangle

For (infinitary) canonical models there is also a standard existence lemma:

Lemma 4.12. *For any state Γ of a canonical model \mathbb{S}^L :*

If $\diamond \varphi \in \Gamma$, then there is a state Δ such that $\Gamma R^L \Delta$ and $\varphi \in \Delta$.

Generally, a μ_c -logic L will lack the compactness property. It is well-known that this prevents one to prove a Truth Lemma for the (standard) canonical model of L . Indeed, if there are unsatisfiable sets of formulas which are finitely satisfiable, then, because derivations are finite objects, there will be maximally consistent sets which are unsatisfiable.

The following lemma is analogous to Lemma 3.5.

Lemma 4.13. *Let L be a logic and let \mathcal{F} be a class of frames that admits filtration and contains the canonical frame (S^L, R^L) . For any finite and \sim FL-closed set Σ , the class \mathcal{F} contains a frame underlying some model \mathbb{S} over Σ with respect to L .*

Proof. We will apply a form of filtration to the canonical model \mathbb{S}^L . As in the proof of Lemma 3.5, we let $\varphi_1, \dots, \varphi_n$ be an enumeration of the formulas of the form $\eta x. \psi$ in Σ . For every such formula φ_i , we pick a unique propositional variable p_i not occurring in Σ .

We will define an alternative valuation function $V' : P \rightarrow \mathcal{P}(S^L)$. In contrast to the proof of Lemma 3.5, we will not let the valuation of p_i be the *meaning* of φ_i in \mathbb{S}^L , but rather we let p_i be true at those Γ for which $\varphi_i \in \Gamma$. Note that if a Truth Lemma would hold for \mathbb{S}^L , these two options would be equivalent.

$$V'(p) := \begin{cases} \{\Gamma : \varphi_i \in \Gamma\} & \text{if } p = p_i \text{ for some } \varphi_i \in \Sigma; \\ V(p) & \text{otherwise.} \end{cases}$$

Let \mathbb{S}' be the model \mathbb{S}^L , but with V' as valuation function. We define the translation $\tau : \Sigma \rightarrow \text{ML}$ in the same way as we did in the proof of Lemma 3.5. The set $\tau[\Sigma]$ is again \sim FL-closed. A straightforward induction shows that for every $\xi \in \Sigma$:

$$\mathbb{S}', \Gamma \Vdash \xi \Leftrightarrow \tau(\xi) \in \Gamma. \quad (3)$$

Since the frame underlying \mathbb{S}' is in \mathcal{F} , we can apply the assumed admissibility of filtration to obtain a filtration $\mathbb{S}^{\tau[\Sigma]} = (\mathcal{S}^{\tau[\Sigma]}, \mathcal{R}^{\tau[\Sigma]}, \mathcal{V}^{\tau[\Sigma]})$ of \mathbb{S}' through $\tau[\Sigma]$ such that the frame $(\mathcal{S}^{\tau[\Sigma]}, \mathcal{R}^{\tau[\Sigma]})$ belongs to \mathcal{F} .

We will finish the proof by showing that $\mathbb{S}^{\tau[\Sigma]}$ is isomorphic to a model over Σ . We define the set of states $S^\Sigma := \{\Gamma \cap \Sigma : \Gamma \text{ is maximally L-consistent}\}$ and claim that the map

$$h : [\Gamma] \mapsto \Gamma \cap \Sigma$$

is a well-defined bijection from $S^{\mathbb{L}} / \sim_{\tau[\Sigma]}^{\mathbb{S}'}$ to S^Σ . For well-definedness, suppose $\Gamma \sim_{\tau[\Sigma]}^{\mathbb{S}'} \Gamma'$ and let $\varphi \in \Sigma$. Using the equivalence (3), we have

$$\varphi \in \Gamma \Leftrightarrow \mathbb{S}', \Gamma \Vdash \tau(\varphi) \Leftrightarrow \mathbb{S}', \Gamma' \Vdash \tau(\varphi) \Leftrightarrow \varphi \in \Gamma',$$

as required.

Injectivity is similar: if $\Gamma \cap \Sigma = \Gamma' \cap \Sigma$, then for all $\tau(\varphi) \in \tau[\Sigma]$, we have:

$$\Gamma \Vdash \tau(\varphi) \Leftrightarrow \varphi \in \Gamma \Leftrightarrow \varphi \in \Gamma' \Leftrightarrow \Gamma' \Vdash \tau(\varphi).$$

For surjectivity, take $\Gamma \cap \Sigma$ for some any $\Gamma \in S^{\mathbb{L}}$. Then $h([\Gamma]) = \Gamma \cap \Sigma$, as required.

Now let the relation $R^\Sigma \subseteq S^\Sigma \times S^\Sigma$ and the valuation $V^\Sigma : P \rightarrow \mathcal{P}(S^\Sigma)$ be given by transporting the structure of $\mathbb{S}^{\tau[\Sigma]}$ along h . More precise, we let

$$AR^\Sigma B := h^{-1}(A)R^{\tau[\Sigma]}h^{-1}(B).$$

We claim that $R^{\min} \subseteq R^\Sigma \subseteq R^{\max}$.

First, suppose that $AR^{\min}B$. Then $\psi_A \wedge \diamond \psi_B$ is L-consistent. Pick some $\Gamma \in S^{\mathbb{L}}$ containing both ψ_A and $\diamond \psi_B$. By Lemma 4.12, there is a $\Delta \in S^{\mathbb{L}}$ such that $\Gamma R^{\mathbb{L}} \Delta$ and $\psi_B \in \Delta$. Since $R^{\tau[\Sigma]}$ contains the finest filtration, we have $[\Gamma]R^{\tau[\Sigma]}[\Delta]$ and thus $h([\Gamma])R^\Sigma h([\Delta])$. The required result follows from the fact that $h([\Gamma]) = A$ and $h([\Delta]) = B$.

Now suppose that $AR^\Sigma B$. We will show that $AR^{\max}B$. To that end, let $\Box \varphi \in A$. Pick $\Gamma \supset A$ and $\Delta \supset B$ from $S^{\mathbb{L}}$. Since $[\Gamma] = h^{-1}(A)$ and $[\Delta] = h^{-1}(B)$, we have $[\Gamma]R^{\tau[\Sigma]}[\Delta]$. We now use the fact that $R^{\tau[\Sigma]}$ is contained in the coarsest filtration. This means that for all $\Box \psi \in \tau[\Sigma]$ such that $\mathbb{S}', \Gamma \Vdash \Box \psi$, we have $\mathbb{S}', \Delta \Vdash \psi$. By assumption we have $\Box \varphi \in \Gamma$, whence the equivalence (3) gives $\mathbb{S}', \Gamma \Vdash \tau(\Box \varphi)$, i.e. $\mathbb{S}', \Gamma \Vdash \Box \tau(\varphi)$. It follows that $\mathbb{S}', \Delta \Vdash \tau(\varphi)$. Finally, another application of the equivalence (3) yields $\varphi \in \Delta$, hence $\varphi \in B$, as required.

Lastly, for any $p \in \Sigma$, we define

$$V^\Sigma(p) := \{A \in S^\Sigma : h^{-1}(A) \in V^{\tau[\Sigma]}(p)\} = \{A \in S^\Sigma : p \in A\},$$

which suffices. □

Theorem 4.14. *Let \mathbb{L} be a canonical logic in the basic modal language such $\text{Fr}(\mathbb{L})$ admits filtration. Then μ_c - \mathbb{L} is sound and complete with respect to $\text{Fr}(\mathbb{L})$.*

Proof. Soundness follows from the fact the fixpoint axioms and rules are sound on the class of all frames. For completeness, let $\varphi \in \mu_c \text{ML}$ be L-consistent; we will show that φ is satisfiable in a model based on a L-frame. Without loss of generality we may assume that φ is clean. Let Σ be the \sim -FL-closure of $\{\varphi\}$. Note that by canonicity the canonical frame $(S^{\mathbb{L}}, R^{\mathbb{L}})$ is contained in $\text{Fr}(\mathbb{L})$. Therefore, we can use Lemma 4.13 to obtain a model \mathbb{S}^Σ over Σ with respect to \mathbb{L} which is based on an L-frame. By the L-consistency of φ , there is a state $A \in S^\Sigma$ such that $\varphi \in A$. Finally, Lemma 4.7 gives $\mathbb{S}^\Sigma, A \Vdash \varphi$, as required. □

For instance, the logic μ_c -KB is sound and complete with respect to the class of symmetric frames. Some other examples of basic modal logics that satisfy the hypotheses of the above theorem are: K, T, K4, S4 and S5.

References

- [1] Johan van Benthem (2006): *Modal frame correspondences and fixed-points*. *Studia Logica* 83, pp. 133–155, doi:10.1007/s11225-006-8301-9.
- [2] Johan van Benthem & Nick Bezhanishvili: *Modern faces of filtration*. ILLC Prepublication PP-2019-13.
- [3] Facundo Carreiro (2015): *Fragments of fixpoint logics*. Ph.D. thesis, University of Amsterdam.
- [4] Facundo Carreiro, Alessandro Facchini, Yde Venema & Fabio Zanasi (2020): *The Power of the Weak*. *ACM Transactions on Computational Logic* 21(2), pp. 15:1–15:47, doi:10.1016/S0304-3975(01)00185-2.
- [5] Gaëlle Fontaine (2008): *Continuous fragment of the mu-calculus*. In: *International Workshop on Computer Science Logic*, Springer, pp. 139–153, doi:10.1007/3-540-49116-3_50.
- [6] Robert Goldblatt (1987): *Logics of time and computation*. Center for the Study of Language and Information.
- [7] David Janin & Igor Walukiewicz (1996): *On the Expressive Completeness of the Propositional μ -Calculus w.r.t. Monadic Second-Order Logic*. In: *Proceedings of the Seventh International Conference on Concurrency Theory, CONCUR '96, LNCS 1119*, pp. 263–277, doi:10.1007/3-540-61604-7_60.
- [8] Stanislav Kikot, Ilya Shapirovsky & Evgeny Zolin (2020): *Modal Logics with Transitive Closure: Completeness, Decidability, Filtration*. In Nicola Olivetti, Rineke Verbrugge, Sara Negri & Gabriel Sandu, editors: *13th Conference on Advances in Modal Logic, AiML 2020, Helsinki, Finland, August 24-28, 2020*, College Publications, pp. 369–388.
- [9] Dexter Kozen (1983): *Results on the propositional μ -calculus*. *Theoretical computer science* 27(3), pp. 333–354, doi:10.1016/0304-3975(82)90125-6.
- [10] Dexter Kozen & Rohit Parikh (1981): *An elementary proof of the completeness of PDL*. *Theoretical Computer Science* 14(1), pp. 113–118, doi:10.1016/0304-3975(81)90019-0.
- [11] John Lemmon & Dana Scott (1977): *An introduction to modal logic*. Blackwell.
- [12] Johannes Marti & Yde Venema (2021): *Focus-style proof systems and interpolation for the alternation-free μ -calculus*.