

The Topological Mu-Calculus: completeness and decidability

Alexandru Baltag
University of Amsterdam

Email: thealexandrubaltag@gmail.com

Nick Bezhanishvili
University of Amsterdam

Email: N.Bezhanishvili@uva.nl

David Fernández-Duque
Ghent University

Email: David.FernandezDuque@UGent.be

Abstract—We study the topological μ -calculus, based on both Cantor derivative and closure modalities, proving completeness, decidability and FMP over general topological spaces, as well as over T_0 and T_D spaces. We also investigate relational μ -calculus, providing general completeness results for all natural fragments of μ -calculus over many different classes of relational frames. Unlike most other such proofs for μ -calculus, ours is model-theoretic, making an innovative use of a known Modal Logic method (–the ‘final’ submodel of the canonical model), that has the twin advantages of great generality and essential simplicity.

I. INTRODUCTION

The modal μ -calculus is one of the most powerful extensions of modal logic, of great use in computer science applications. It is decidable, but very expressive, embedding many modal/temporal logics, such as PDL, CTL and CTL*, that are widely applied in program specification and verification.

The completeness of the modal μ -calculus was a difficult problem and remained open for many years. Even today, there are very few completeness results for axiomatic systems for μ -calculus with respect to standard classes of Kripke models (e.g., [1], [2], [3], see also a more recent proof theoretic approach [4]). Prior to our work, there seemed to be no general model-theoretic method to establish completeness for various natural fragments of μ -calculus over various classes of models.

An alternative interpretation of modal logic is based not on Kripke frames, but on topological spaces. This semantics is in fact older and can be traced back to McKinsey and Tarski [5]. When the modal \diamond is interpreted as topological closure and the modal \square as topological interior, one obtains a semantics for the modal logic S4 and its extensions, generalizing Kripke semantics over transitive, reflexive frames. The logic of all topological spaces in this semantics is S4. We refer to [6] for an overview of the rich landscape of results on topological completeness of modal logics above S4.

McKinsey and Tarski also suggested a second topological semantics, obtained by interpreting the modal \diamond as Cantor derivative.¹ Esakia [7], [8] showed that the derivative logic of all topological spaces is the modal logic $wK4 = K + (\diamond\diamond p \rightarrow \diamond p \vee p)$. This is also the modal logic of all *weakly transitive* frames, i.e. those for which the reflexive closure of the accessibility relation is transitive. It is well-known that

the modal logic of transitive frames is K4 [9], [10], which moreover corresponds to a natural class of topological spaces denoted T_D . Another natural class are T_0 spaces, whose modal logic is also finitely axiomatizable; we discuss T_0 spaces and T_D spaces in the context of modal logic in Section III. Modal logics axiomatizing well-known classes of spaces also include the Gödel-Löb logic GL, which is complete with respect to the class of scattered spaces [11], [12].

Topological structures are of great interest to Computer Science. As noticed by Vickers [13] and Abramsky [14], the notion of observability and its logic require a topological setting. Abstract notions of computability also involve topological structures, and a famous example is Scott topology. More recently, developments in Formal Learning Theory [15], [16], Distributed Computing [17] and Epistemic Logic in Multi-Agent Systems [18], [19], [20], have taken a topological turn. In particular, recent epistemic work [20], [19] on modelling and reasoning about evidence and knowability uses topological structures. Research on spatial reasoning, in both topological and metric incarnations, is also of significant interest for AI. The addition of fixed point operators allows us to reason about non-trivial properties of topological spaces: for example, the well-known Cantor-Bendixon theorem states that any topological space has a *perfect core*, i.e. a maximal subset equal to its own derivative. The perfect core is not modally definable (in terms of derivative or closure modalities), but it is definable in the μ -calculus with the derivational semantics. Parikh [21] showed the relevance of Cantor derivative and the perfect core for multi-agent epistemic puzzles and applications.²

Our main aim in this paper is to investigate the *topological* μ -calculus based on the Cantor derivative modality, as well as its weaker version based on the closure modality. As a secondary aim, we explore (various fragments of) the *relational* μ -calculus, on (various classes of) weakly transitive frames. As such, our contribution in this paper is two-fold. First, we develop a new model-theoretic method of proving completeness for systems of μ -calculus over weakly transitive frames. This method applies to a wide range of logics, including many well-known ones. Concretely, we show that if a modal logic Λ is a canonical cofinal subframe logic, then its modal μ -variant,

¹Recall that the derivative $d(A)$ of a set A consists of all limit points of A .

²In on-going work, we show that the perfect core and its logic have deep connections with the topic of learnability from observations, as well as with epistemic paradoxes, such as the Surprise Examination.

obtained by adding the fixed-point axiom and induction rule, is Kripke complete and enjoys the finite model property. This implies that the modal μ -variants of the well known modal logics $wK4$, wKT_0 , $K4$, $K4D$, $K4.1$, $K4.2$, $K4.3$, $S4$, $S4.1$, $S4.2$, and $S4.3$ have the FMP³ and are decidable. Second, we show that the derivational μ -calculus is completely axiomatized on all topological spaces, all T_0 spaces, and all T_D spaces, by the μ -variants of the logics $wK4$, $wK4T_0$, and $K4$, respectively. We also give a new proof of the known fact that the weaker μ -calculus based on topological closure is completely axiomatized by the μ -variant of the modal logic $S4$.

Our model-theoretic proof is based on restricting the canonical model to the set of *final theories*, i.e. theories which satisfy a natural maximality condition. A similar construction has been employed by Fine [22] to prove FMP for subframe logics over $K4$ (without fixed point operators). Zakharyashev [23] generalized this to show FMP for cofinal subframe logics over $K4$, and [24] extended this result to cofinal subframe logics over $wK4$. Our Kripke-completeness results apply to the μ -variants of the same class of logics. The crucial new insight is that the truth lemma extends to the full μ -calculus over the set of final theories, despite not doing so for the full canonical model. Topological completeness follows then from more or less standard constructions and the observation that the logics of the classes of all topological spaces, all T_0 spaces, and all T_D spaces are all subframe logics.

These results are new, with one proviso concerning T_D spaces. First, note that the transitive closure of a binary relation is definable in standard relational μ -calculus (on arbitrary frames). Thus, FMP for μ -calculus over transitive frames follows immediately from Kozen's proof of FMP for general μ -calculus [25]. Second, Goldblatt and Hodkinson [26] have completely axiomatized a modal logic (with the so-called *tangled derivative* modality), that is co-expressive with derivational μ -calculus over T_D spaces, by a result of Dawar and Otto [27]. But, as explained in Section IX, even in the T_D case, our work has the added benefit of providing a direct proof of completeness/decidability of full μ -calculus over these spaces, without relying on the corresponding results for standard μ -calculus. Moreover, dropping the T_D assumptions drastically changes the behavior of the μ -calculus in at least two important ways. Weakly transitive closure does *not* seem to be definable in μ -calculus, and so decidability over arbitrary (as well as T_0) spaces does *not* follow from any known results. Finally, as shown in Section IV, the above-mentioned co-expressivity result fails on arbitrary (or T_0) spaces: μ -calculus on such spaces does *not* collapse to its tangled fragment. For this reason, we work here with the full language of μ -calculus.

The structure of this paper In Section II we introduce derivative spaces, a general notion subsuming both topological spaces and weakly-transitive frames. Section III defines μ -calculus over such spaces and states our main completeness

³In fact, there are continuum-many such logics [10], so our results apply to uncountably many classes of frames.

result. In Section IV, we show that the tangled fragment is not expressively complete in this setting. Section V investigates truth-preserving maps and relations between derivative spaces. Section VI presents the stepping stones of the main completeness proof. Section VII generalizes this to an infinite class of fixed-point logics, while Section VIII extends it to T_0 and T_D spaces. We end in Section IX with some concluding remarks and a comparison with related work. All the proof details are in the Appendix.

II. DERIVATIVE SPACES

Although our primary focus in this paper is the derivational μ -calculus on topological spaces, for technical reasons it is useful to consider a slightly more general class of structures.

Definition II.1. A derivative space is a pair (\mathcal{X}, d) , where \mathcal{X} is a set of 'points', and $d : \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{X})$ is an operator on subsets of \mathcal{X} , satisfying the following properties, for all $X, Y \subseteq \mathcal{X}$:

- $d(\emptyset) = \emptyset$;
- $d(X \cup Y) = d(X) \cup d(Y)$;
- $d(d(X)) \subseteq X \cup d(X)$.

The conjunction of the first two conditions above is known as normality, while the third condition is known as weak idempotence.

The notion of derivative space is the concrete set-theoretic instantiation of the more abstract concept of *derivative algebra*, introduced by Esakia [8] (as a generalization of a notion with the same name introduced by McKinsey and Tarski [5]).

Example II.2 (topological closure spaces). A special case of derivative spaces is given by closure spaces: these are derivative spaces (\mathcal{X}, c) that additionally satisfy $X \subseteq c(X)$ (and, a fortiori, $c(c(X)) \subseteq c(X)$). These strengthened conditions are known as the Kuratowski axioms, that define topological spaces in terms of their closure operator.⁴ When considered as a special case of derivative spaces, with $d(X) := c(X)$ given by topological closure, topological spaces will be called topological closure spaces.

Example II.3 (topological derivative spaces). Our main example of derivative spaces in this paper are structures (\mathcal{X}, d) , based on an underlying topological (closure) space (\mathcal{X}, c) (satisfying the Kuratowski axioms), but with the derivative operator given by the so-called Cantor derivative, i.e. by taking $d(X)$ to be the set of limit points of X :

$$\begin{aligned} d(X) &:= \{y \in \mathcal{X} : y \in c(X - \{y\})\} \\ &= \{y \in \mathcal{X} : \forall U \in \mathcal{N}(y) X \cap (U - \{y\}) \neq \emptyset\}, \end{aligned}$$

where $\mathcal{N}(y)$ is the family of (open) neighborhoods of y in the space (\mathcal{X}, c) . It is easy to see that (\mathcal{X}, d) is a derivative space, which we'll refer to as a topological derivative space. The closure operator can be recovered as $c(X) = X \cup d(X)$.

⁴Given a closure space, let $X \subseteq \mathcal{X}$ be closed whenever $X = c(X)$, and open whenever its complement is closed. This gives us the more common definition of topology as a family of open or closed sets. So closure spaces are exactly the same notion as topological spaces.

So, every topological space gives rise to a derivative space in at least two different ways (as a closure space, and as a topological derivative space), though we are mostly interested in the second one. The converse is also true:

Closure and interior in derivative spaces Given a derivative space (X, d) , we define the *closure* and *interior* operators $c, i : \mathcal{P}(X) \rightarrow \mathcal{P}(X)$, by putting

$$c(X) := X \cup d(X), \quad i(X) := X - c(X - X).$$

It is easy to see that these satisfy all the Kuratowski axioms.

This means that every derivative space induces a topological space. Moreover, in a *topological* derivative space (with Cantor derivative over some topological space), the induced closure operator (as defined above) coincides with the underlying topological closure. But in general, this matching does not work the other way around: given an arbitrary derivative space, its derivative does *not* necessarily coincide with the Cantor derivative in the induced topology (given by the above-defined closure operator). It follows that *not every derivative space is a topological derivative space*. A counterexample is given by the next special case.

Example II.4 (weakly transitive Kripke frames). A weakly transitive frame (or wK4 frame) is a Kripke structure (W, \rightarrow) (also known as a ‘transition system’), consisting of a set of ‘states’ (or ‘possible worlds’) W , together with a binary relation $\rightarrow \subseteq W \times W$ (known as an ‘accessibility’ or ‘transition’ relation), assumed to be weakly transitive: i.e., for all states $w, s, t \in W$, if $w \rightarrow s \rightarrow t$ then either $w = t$ or $w \rightarrow t$. We denote by \rightarrow^* the reflexive closure $\text{Id} \cup \rightarrow$ of \rightarrow , which (due to weak transitivity) coincides with its transitive-reflexive closure $\text{Id} \cup \bigcup_{n \geq 1} \rightarrow^n$.

We also denote by $\overleftrightarrow{\rightarrow}$ the strict part of \rightarrow , i.e. $w \overleftrightarrow{\rightarrow} v$ if $w \rightarrow v \not\rightarrow w$; and write $w \longleftrightarrow v$ if $w \rightarrow v \rightarrow w$ and $w \longleftrightarrow^* v$ if $w \longleftrightarrow v$ or $w = v$. For any state $w \in W$, we put $w \uparrow := \{s \in W : w \rightarrow s\}$ for the set of its successors, and also put $w \uparrow^* := \{s \in S : w \rightarrow^* s\} = \{w\} \cup w \uparrow$; more generally, for any set $X \subseteq W$, we put $X \uparrow := \{s \in W : x \rightarrow s \text{ for some } x \in X\} = \bigcup_{x \in X} x \uparrow$, and similarly put $X \uparrow^* := \{s \in W : x \rightarrow^* s \text{ for some } x \in X\} = X \cup X \uparrow$. By applying the same definitions to the converse \leftarrow , we obtain the corresponding notions of down-closure $w \downarrow, w \downarrow^*, X \downarrow, X \downarrow^*$.

It is easy to see that every weakly transitive frame gives rise to a derivative space (X, d_{\rightarrow}) , obtained by taking $X := W$, and taking the derivative d_{\rightarrow} to be usual modal ‘Diamond’ operator:

$$d_{\rightarrow}(X) := X \downarrow = \{w \in W : X \cap w \uparrow \neq \emptyset\} \\ = \{w \in W : \exists s w \rightarrow s \in X\}.$$

Moreover, the induced closure $c_{\rightarrow}(X)$ (as defined above in arbitrary derivative spaces) is given by $c_{\rightarrow}(X) = X \downarrow^*$.

In general, weakly transitive frames are *not* topological derivative spaces. But the intersection of the two classes is of independent interest, as shown by the next two examples:

Example II.5 (Alexandroff closure spaces as S4 Kripke frames). A topological space (X, c) is Alexandroff if its closure operator distributes over arbitrary unions: $c(\bigcup_i X_i) = \bigcup_i c(X_i)$. Given $x, y \in X$, define $x \rightarrow y$ if $x \in c\{y\}$. Then, it is not hard to check that if X is Alexandroff, then \rightarrow is a reflexive-transitive relation, i.e. (X, \rightarrow) is an S4 Kripke frame, and moreover the relational derivative coincides in this case with the topological closure: $d_{\rightarrow} = c$. As it is well-known, the converse also holds: every S4 frame (X, \rightarrow) gives rise to an Alexandroff closure space, by putting $c_{\rightarrow}(X) := X \downarrow = X \downarrow^*$ for the closure/derivative operator. This time, the equivalence is complete: starting from either side, and applying successively these two transformations, we obtain the original structure. So Alexandroff topological closure spaces are essentially the same as S4 Kripke frames.

Example II.6 (Alexandroff derivative spaces as irreflexive wK4 frames). Another way to convert an Alexandroff space (X, c) into a relational structure is to define $x \rightarrow y$ if $x \in d\{y\} = c\{y\} - \{y\}$, for all $x, y \in X$. Then \rightarrow is weakly transitive and irreflexive, and the relational derivative d_{\rightarrow} coincides in this case with the Cantor derivative induced by c . Conversely, every irreflexive wK4 frame (X, \rightarrow) gives rise to an Alexandroff derivative space (X, d) , by putting $c_{\rightarrow}(X) := X \downarrow^*$ for the topological closure, and taking d to be induced Cantor derivative in the resulting topology (for which one can check that $d(X) = X \downarrow$). Once again, the equivalence is complete: by applying successively these transformations, we obtain the original structure. So Alexandroff topological derivative spaces are essentially the same as irreflexive wK4 frames.

D-neighborhoods For every point $x \in X$ in a derivative space (X, d) , we can define the family of d -neighborhoods of x :

$$\mathcal{N}_d(x) := \{X \subseteq X : x \notin d(X - X)\}$$

Note that, in general, d -neighborhoods are *not* neighborhoods of x in the topology given by the closure $c(X)$ induced by d . In fact, in a topological derivative space (where derivative means Cantor derivative), a d -neighborhood $X \in \mathcal{N}_d(x)$ is just a ‘punctured neighborhood’ of x , i.e. a set with the property that $U - \{x\} \subseteq X$ for some open neighborhood $U \ni x$. On the other hand, in a topological closure space (where the ‘derivative’ is just the topological closure), d -neighborhoods coincide with standard topological neighborhoods. Finally, in a weakly transitive frame (W, \rightarrow) , a set $X \subseteq W$ is a d -neighborhood of a state $x \in W$ iff $x \uparrow \subseteq X$.

We can now characterize the derivative in terms of d -neighborhoods, in a way that generalizes the definition of Cantor derivative in topological spaces:

Lemma II.7. For every set $X \subseteq X$ in a derivative space (X, d) , we have

$$d(X) = \{y \in X : \forall U \in \mathcal{N}_d(y) U \cap X \neq \emptyset\}.$$

This leads to an equivalent presentation of derivative spaces as a special case of monotonic neighborhood structures [28]: a

neighborhood derivative space is a pair $(\mathcal{X}, \mathcal{N})$, where \mathcal{X} is a set of points, and $\mathcal{N} : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{P}(\mathcal{X}))$ is a map that assigns to each point $x \in \mathcal{X}$ a family $\mathcal{N}(x) \subseteq \mathcal{P}(\mathcal{X})$ of ‘neighborhoods’ of x , satisfying the following conditions

- 1) $\mathcal{X} \in \mathcal{N}(x)$;
- 2) if $X \in \mathcal{N}(x)$ and $X \subseteq Y$, then $Y \in \mathcal{N}(x)$;
- 3) if $X, Y \in \mathcal{N}(x)$, then $X \cap Y \in \mathcal{N}(x)$;
- 4) if $x \in X \in \mathcal{N}(x)$, then $\{y \in \mathcal{X} : X \in \mathcal{N}(y)\} \in \mathcal{N}(x)$.

Each derivative space (\mathcal{X}, d) gives rise to a neighborhood derivative space by taking

$$\mathcal{N}(x) := \mathcal{N}_d(x) = \{X \subseteq \mathcal{X} : x \notin d(\mathcal{X} - X)\}$$

to be the set of all d -neighborhoods. Conversely, every neighborhood derivative space $(\mathcal{X}, \mathcal{N})$ gives rise to a derivative space, via the following generalization of Cantor derivative:

$$d(X) := \{y \in \mathcal{X} : \forall U \in \mathcal{N}(y) U \cap X \neq \emptyset\}.$$

This is a full equivalence between derivative spaces and neighborhood derivative spaces: starting from either side, and applying the above two transformations, we obtain the original structure.

III. MU CALCULUS ON DERIVATIVE SPACES: MAIN RESULTS

For reasons having to do with our intended applications, as well as to simplify some proof details, in this paper we take the *greatest fixed point operator* $\nu x.\varphi$ as primitive, and define the least fixed point $\mu x.\varphi$ as an abbreviation.⁵

Syntax: Let \mathbb{P} be a set of *propositional variables*. We recursively define the set \mathcal{L}_μ of *formulas*, together with a map $\text{Free} : \mathcal{L}_\mu \rightarrow \mathcal{P}(\mathbb{P})$, associating to each formula $\varphi \in \mathcal{L}_\mu$ its set of *free variables* $\text{Free}(\varphi) \subseteq \mathbb{P}$. The definition is by simultaneous recursion, with formulas $\varphi \in \mathcal{L}_\mu$ given by

$$\varphi ::= \top \mid x \mid \neg\varphi \mid \varphi \wedge \varphi' \mid \diamond\varphi \mid \nu x.\varphi$$

where: $x \in \mathbb{P}$; in the construct $\varphi \wedge \varphi'$, no variables occur free in φ and bound in φ' , or vice versa; and in the construct $\nu x.\varphi$, formula φ is *positive in x* (i.e. whenever x occurs in φ , we have that $x \in \text{Free}(\varphi)$ and x occurs only in the scope of an even number of negations). The set $\text{Free}(\varphi)$ of free variables of a formula φ is simultaneously defined by recursion:

$$\begin{aligned} \text{Free}(\top) &:= \emptyset, & \text{Free}(x) &:= \{x\}, \\ \text{Free}(\varphi \wedge \varphi') &:= \text{Free}(\varphi) \cup \text{Free}(\varphi'), \\ \text{Free}(\neg\varphi) &= \text{Free}(\diamond\varphi) := \text{Free}(\varphi), \\ \text{Free}(\nu x.\varphi) &:= \text{Free}(\varphi) - \{x\}. \end{aligned}$$

A variable is *bound* in φ if it occurs in φ but is not in $\text{Free}(\varphi)$. For any set of variables $P \subseteq \mathbb{P}$, we denote by \mathcal{L}_μ^P the set of all formulas $\varphi \in \mathcal{L}_\mu$ having $\text{Free}(\varphi) \subseteq P$. Note in particular that $\mathcal{L}_\mu = \mathcal{L}_\mu^{\mathbb{P}}$.

We use the notation $\bar{x} = (x_1, \dots, x_n)$ to denote finite strings of variables $x_1, \dots, x_n \in \mathbb{P}$, and denote by λ the empty string.

⁵This setting is of course equivalent to the more standard presentation, that takes $\mu x.\varphi$ as primitive.

When we want to make explicit some of the free variables, we write $\varphi(\bar{x})$ for a formula in which all variables in the string \bar{x} are free (if occurring at all).

Subformulas The *subformula relation* \sqsubset is the smallest transitive relation on formulas satisfying the following properties: $\varphi \sqsubset (\neg\varphi)$, $(\diamond\varphi)$, $(\nu x.\varphi)$, and $\varphi \sqsubset (\varphi \wedge \varphi')$, $(\varphi' \wedge \varphi)$. The set $\text{Sub}(\varphi)$ of all (improper) subformulas of φ is defined as $\text{Sub}(\varphi) := \{\varphi' : \varphi' \sqsubset \varphi\} \cup \{\varphi\}$.

Semantics. An *atomic valuation* on a derivative space (\mathcal{X}, d) is a map $\|\cdot\| : \mathbb{P} \rightarrow \mathcal{P}(\mathcal{X})$ associating to each propositional atom $x \in \mathbb{P}$ some set of states $\|x\| \subseteq \mathcal{X}$. For each atomic valuation $\|\cdot\| : \mathbb{P} \rightarrow \mathcal{P}(\mathcal{X})$, tuple $\bar{x} = (x_1, \dots, x_n)$ of variables and corresponding tuple $\bar{X} = (X_1, \dots, X_n)$ of sets of points $X_i \subseteq \mathcal{X}$, we denote by $\|\cdot\|_{\bar{x}:=\bar{X}}$ the valuation that assigns to each variable x_i the set X_i and agrees with the original valuation $\|\cdot\|$ on all the other atoms.

A *derivative model* $\mathcal{M} = (\mathcal{X}, d, \|\cdot\|)$ consists of a derivative space (\mathcal{X}, d) , together with an atomic valuation $\|\cdot\| : \mathbb{P} \rightarrow \mathcal{P}(\mathcal{X})$. The semantics is given by extending the atomic valuation to a map $\|\cdot\| : \mathcal{L}_\mu \rightarrow \mathcal{P}(\mathcal{X})$, which we call *extended valuation* (and for which we use the same notation $\|\cdot\|$ as for the corresponding atomic valuation). The definition of the extended valuation is by recursion on subformulas: for propositional variables this is already given by the atomic valuation map of the model \mathcal{M} , while in the rest we put

$$\begin{aligned} \|\top\| &= \mathcal{X}, & \|\neg\varphi\| &= \mathcal{X} - \|\varphi\|, & \|\varphi \wedge \varphi'\| &= \|\varphi\| \cap \|\varphi'\|, \\ \|\diamond\varphi\| &= d(\|\varphi\|), & \|\nu x.\varphi\| &= \bigcup \{X \subseteq \mathcal{X} : X \subseteq \|\varphi\|_{x:=X}\}. \end{aligned}$$

For formulas $\varphi = \varphi(\bar{x})$ and corresponding tuples of sets \bar{X} , we will sometimes write $\|\varphi(\bar{X})\|$ instead of $\|\varphi\|_{\bar{x}:=\bar{X}}$; in order to avoid subscript overload. With this notation, e.g., the clause for νx becomes: $\|\nu x.\varphi\| = \bigcup \{X \subseteq \mathcal{X} : X \subseteq \|\varphi(X)\|\}$.

Whenever $x \in \|\varphi\|$ for some point $x \in \mathcal{X}$, we also write $x \models_{\mathcal{M}} \varphi$, and say that φ is *true* (or *satisfied*) at point x in the model \mathcal{M} . As usual, when the model is understood, we skip the subscript, writing $x \models \varphi$. Conversely, we may write $\|\cdot\|_{\mathcal{M}}$ instead of $\|\cdot\|$ when we wish to specify the relevant model. We say that φ is *valid on the model* \mathcal{M} if $\|\varphi\|_{\mathcal{M}} = \mathcal{X}$, i.e. φ is true at all points of \mathcal{M} ; similarly, φ is *satisfied on the model* \mathcal{M} if $\|\varphi\|_{\mathcal{M}} \neq \emptyset$. By abstracting away from the specific valuation, we say that φ is *valid on the space* (\mathcal{X}, d) if for every valuation $\|\cdot\|$ on \mathcal{X} , φ is valid on the model $(\mathcal{X}, d, \|\cdot\|)$; and φ is *satisfied on the space* (\mathcal{X}, d) if there exists a valuation $\|\cdot\|$ on \mathcal{X} , s.t. φ is satisfied on the model $(\mathcal{X}, d, \|\cdot\|)$. Finally, φ is *valid (on a class \mathcal{C} of derivative models, or of derivative spaces)* if it is valid on all models/spaces (in the class \mathcal{C}).

Note that in the special case of wK4 relational models (W, \rightarrow) , the above semantics of \diamond coincides with the standard Kripke semantics. As a consequence, on relational frames our semantics for μ -calculus coincides with the standard one.

Abbreviations: We have the usual abbreviations \perp , $\varphi \vee \psi$, $\varphi \Rightarrow \psi$, $\varphi \Leftrightarrow \psi$, $\Box\varphi$. The *least fixed-point formula* $\mu x.\varphi(x, \bar{y})$ can be defined as $\neg\nu x\neg\varphi(\neg x, \bar{y})$. Finally, we define *closure*

and *interior* modalities, as well as *tangled derivative* $\diamond^\infty\Gamma$ and *tangled closure* $\langle *\rangle^\infty\Gamma$ (for finite sets of formulas Γ), with the *perfect core* modality $\diamond^\infty\varphi$ as a special case:

$$\begin{aligned} \langle *\rangle\varphi &:= \varphi \vee \diamond\varphi, & [*]\varphi &:= \varphi \wedge \square\varphi, \\ \diamond^\infty\Gamma &:= \nu x. \bigwedge_{\gamma \in \Gamma} \diamond(x \wedge \gamma), & \langle *\rangle^\infty\Gamma &:= \nu x. \bigwedge_{\gamma \in \Gamma} \langle *\rangle(x \wedge \gamma), \\ & & \diamond^\infty\varphi &:= \diamond^\infty\{\varphi\}. \end{aligned}$$

Note that the definitions of $\langle *\rangle\varphi$ and $[*]\varphi$ do not use any fixed points. But, to justify these notations, one can easily check that, in the special case of weakly transitive frames, $[*]$ and $\langle *\rangle$ are the standard Kripke modalities for the reflexive-transitive closure \longrightarrow^* of the accessibility relation (which, as already mentioned, coincides on these frames with its reflexive closure). More generally, in derivative spaces, $[*]\varphi$ is equivalent to $\nu x.(\varphi \wedge \square x)$, while $\langle *\rangle\varphi$ is equivalent to $\neg[*]\neg\varphi$ and thus to $\mu x.(\varphi \vee \diamond x)$. In fact, $\|[*]\varphi\|$ and $\|\langle *\rangle\varphi\|$ coincide with the *interior* $i(\|\varphi\|)$ and respectively the *closure* $c(\|\varphi\|)$, as defined in derivative spaces. In particular, in the case of *topological* derivative spaces (where d is Cantor derivative), these coincide with the underlying topological interior and closure operators. As for $\diamond^\infty\Gamma$ and $\langle *\rangle^\infty\Gamma$, they are variants of the *tangle modality* introduced in a relational setting by Dawar and Otto [27], who showed that μ -calculus over transitive frames collapses to tangle logic based on $\diamond^\infty\Gamma$. Their topological interpretations were developed by Fernandez-Duque [29], who distinguished between the *tangled derivative* $\diamond^\infty\Gamma$ and *tangled closure* $\langle *\rangle^\infty\Gamma$, and axiomatized the logic of tangled closure. More recently, Goldblatt and Hodkinson [26] axiomatized the logic of tangled derivative $\diamond^\infty\Gamma$ over transitive frames, and showed that it is equivalent to the logic over T_D spaces. Finally, the perfect core modality $\diamond^\infty\varphi$ is a special case of tangle, that captures Cantor's *perfect core*: the largest subset of the state space that is equal to its own Cantor derivative.

Substitution and natural sublanguages Given a formula $\varphi = \varphi(\bar{x})$ and a tuple of formulas $\bar{\theta} = (\theta_1, \dots, \theta_n)$, we denote by $\varphi(\bar{\theta})$ the result of substituting every variable in \bar{x} by the corresponding formula in $\bar{\theta}$. Note that we have

$$\|\varphi(\bar{\theta})\| = \|\varphi(\|\bar{\theta}\|)\|$$

(where on the right hand we used an instance of the above-mentioned simplified notation $\|\varphi(\bar{X})\|$ for $\|\varphi\|_{\bar{x}:=\bar{X}}$). A *natural sublanguage* of \mathcal{L}_μ is any set $\mathcal{L} \subseteq \mathcal{L}_\mu$ which contains \top , is closed under substitution, and such that if $\varphi, \psi \in \mathcal{L}_\mu$ then also $\neg\varphi, \varphi \wedge \psi, \diamond\psi \in \mathcal{L}$. The basic modal language is a natural sublanguage and will be denoted \mathcal{L}_\diamond .

The following characterization of $\|\nu y.\varphi\|$ is also well-known in the literature:

Proposition III.1. *Let $(\mathcal{X}, d, \|\cdot\|)$ be any derivative model and $\varphi = \varphi(y, \bar{x})$ be a formula that is positive in y . Then we have the following:*

- 1) *the unary operator $Y \mapsto \|\varphi(Y, \bar{X})\|$ is monotonic: if $Y \subseteq Y'$ then $\|\varphi(Y, \bar{X})\| \subseteq \|\varphi(Y', \bar{X})\|$;*

- 2) *$\|\nu y.\varphi(y, \bar{X})\|$ is the greatest fixed point of the operator $Y \mapsto \|\varphi(Y, \bar{X})\|$, i.e. the largest set $Y \subseteq \mathcal{X}$ s.t. $Y = \|\varphi(Y, \bar{X})\|$;*
- 3) *$\|\nu y.\varphi(y, \bar{X})\| = \bigcap_{\alpha \in \text{On}} \varphi_y^\alpha(\bar{X})$, where On is the class of ordinals and the transfinite sequence of sets $\varphi_y^\alpha(\bar{X}) \subseteq \mathcal{X}$ is defined by ordinal recursion: $\varphi_y^\alpha(\bar{X}) = \bigcap_{\beta < \alpha} \|\varphi(\varphi_y^\beta(\bar{X}), \bar{X})\|$ (and so in particular $\varphi_y^0(\bar{X}) = \mathcal{X}$).*

Proof. Well known (and easy to check). \square

Definition III.2 (μ -wK4). *We define the logic μ -wK4 to be the least set of formulas of \mathcal{L}_μ containing the following axioms and closed under the following rules (for all formulas φ, ψ , and formulas $\theta = \theta(x)$ that are positive in x):*

- *All the instances of the Axioms and Rules of Propositional Logic.*
- **Necessitation Rule:** *From φ , infer $\square\varphi$.*
- **Distribution Axiom** (=Kripke's Axiom K): $\square(\varphi \Rightarrow \psi) \Rightarrow (\square\varphi \Rightarrow \square\psi)$.
- **Weak Transitivity:** $\diamond\diamond\varphi \Rightarrow (\varphi \vee \diamond\varphi)$.
- **Fixed Point Axiom:** $\nu x.\theta \Rightarrow \theta(\nu x.\theta)$.
- **Induction Rule:** *From $\varphi \Rightarrow \theta(\varphi)$, infer $\varphi \Rightarrow \nu x.\theta$.*

We will also be interested in variants of μ -wK4. If Λ is any normal modal logic over \mathcal{L}_\diamond (in the sense of [9]) that extends wK4, then μ - Λ is the extension of μ -wK4 with all axioms of Λ , closed under uniform substitution with arbitrary formulas in \mathcal{L}_μ . If \mathcal{L} is a natural sublanguage of \mathcal{L}_μ , then μ - $\Lambda^\mathcal{L}$ is the restriction of μ - Λ to \mathcal{L} , in the sense that all axioms and rules may only be applied when all formulas belong to \mathcal{L} .

Proposition III.3 (Soundness). *The logic μ -wK4 is sound for the class of derivative spaces, and so in particular for the class of weakly transitive frames. If Λ is any extension of wK4, then μ - Λ is sound for the class of Λ -spaces, i.e. the class of spaces validating all theorems of Λ .*

Proof. The Necessitation Rule and the Distribution Axiom are sound because of the normality conditions imposed on the derivative operator d , while Weak Transitivity is sound due to the weak idempotence of d . The soundness of the Fixed Point Axiom and of the Induction Rule follows in the usual way from our (standard) semantics for fixed-point formulas. The same argument applies to any extension of wK4 and its class of derivative spaces. \square

Our goal is to show that this system is also (weakly) complete, and that the logic is decidable. But for this, we need first look at some theorems of the above axiomatic system.

Proposition III.4. *The following schemas are provable in the logic μ -wK4 (for all formulas φ, ψ , and formulas $\theta = \theta(x)$ that are positive in x):*

- 1) $\nu x.\theta \Leftrightarrow \theta(\nu x.\theta)$
- 2) $([*]\varphi \wedge \theta(\psi)) \Rightarrow \theta([*]\varphi \wedge \psi)$
- 3) $[*](\varphi \Rightarrow \psi) \Rightarrow (\theta(\varphi) \Rightarrow \theta(\psi))$
- 4) $[*](\varphi \Rightarrow \theta(\varphi)) \Rightarrow (\varphi \Rightarrow \nu x.\theta)$

Proof. Claim 1) is an well-known, easy consequence of the Fixed Point Axiom and the Induction Rule.

For claim 2), it is useful to check first the following special cases:

- 2a** $([*]\varphi \wedge \Diamond\psi) \Rightarrow \Diamond([*]\varphi \wedge \psi)$
- 2b** $([*]\varphi \wedge \Box\psi) \Rightarrow \Box([*]\varphi \wedge \psi)$
- 2c** $([*]\varphi \wedge \nu x.\theta) \Rightarrow \nu x.([*]\varphi \wedge \theta)$
- 2d** $([*]\varphi \wedge \mu x.\theta) \Rightarrow \mu x.([*]\varphi \wedge \theta)$

Checking that these special instances of 2) follow from the axioms is an easy verification. Given them, one can prove 2) by induction on the complexity of $\theta(x)$, written in positive form (i.e. with negations only in front of propositional variables, other than x , and using in rest only conjunctions, disjunctions, \Diamond , \Box and the fixed-point operators νx and μx). The atomic cases are immediate, and the inductive steps for conjunction and disjunction follow trivially by propositional logic, while the other inductive steps are taken care by the instance 2)a-2)d above.

To prove claim 3), first note that $([*](\varphi \Rightarrow \psi) \wedge \varphi) \Rightarrow \psi$ is a theorem in our axiom system. By using the monotonicity of the positive formula $\theta(x)$ (itself provable in the system), we can derive the theorem $\theta([*](\varphi \Rightarrow \psi) \wedge \varphi) \Rightarrow \theta(\psi)$. Putting this together with $([*](\varphi \Rightarrow \psi) \wedge \theta(\varphi)) \Rightarrow \theta([*](\varphi \Rightarrow \psi) \wedge \varphi)$ (which is just a special instance of claim 2)), we obtain $([*](\varphi \Rightarrow \psi) \wedge \theta(\varphi)) \Rightarrow \theta(\psi)$, from which the desired conclusion follows by propositional reasoning.

Finally, to prove claim 4), we start with the obvious theorem $[*](\varphi \Rightarrow \theta(\varphi)) \Rightarrow (\varphi \Rightarrow \theta(\varphi))$, from which we get $([*](\varphi \Rightarrow \theta(\varphi)) \wedge \varphi) \Rightarrow \theta(\varphi)$, and thus also $([*](\varphi \Rightarrow \theta(\varphi)) \wedge \varphi) \Rightarrow ([*](\varphi \Rightarrow \theta(\varphi)) \wedge \theta(\varphi))$. Putting this together with $[*](\varphi \Rightarrow \theta(\varphi)) \wedge \theta(\varphi) \Rightarrow \theta([*](\varphi \Rightarrow \theta(\varphi)) \wedge \varphi)$ (itself an instance of claim 2)), we obtain $([*](\varphi \Rightarrow \theta(\varphi)) \wedge \varphi) \Rightarrow \theta([*](\varphi \Rightarrow \theta(\varphi)) \wedge \varphi)$. Applying then the Induction Rule, we derive $([*](\varphi \Rightarrow \theta(\varphi)) \wedge \varphi) \Rightarrow \nu x.\theta$, from which the desired conclusion follows by propositional reasoning. \square

We are now ready to state the first of our main results. Below, recall that Λ has the *finite model property* if for any formula φ , φ is a theorem of Λ iff φ is valid over the class of finite Λ -models. The logic Λ has the *strong finite model property* if the size of a finite countermodel for φ can be bounded by a function computable from the length of φ .

Theorem III.5 (Completeness, FMP and Decidability). *Let \mathcal{L} be a natural sublanguage of \mathcal{L}_μ . The logic $\mu\text{-wK4}^{\mathcal{L}}$ is (weakly) complete for the class of all weakly transitive frames, as well as for the class of Alexandroff spaces (irreflexive weakly transitive frames). Hence, $\mu\text{-wK4}^{\mathcal{L}}$ is complete for the class of all topological spaces, and thus also for the larger class of all derivative spaces. The logic has the strong finite model property (wrt all the above classes), and so its satisfiability problem is decidable.*

This will be proven in Section VI, while in Section VII we generalize this result to many other classes of relational structures and the corresponding logics.

We conclude this section by discussing two extensions of wK4 that are of interest in the context of topological semantics. Recall that a topological space (\mathcal{X}, c) is T_0 if given $x, y \in \mathcal{X}$ with $x \neq y$, either $x \notin c\{y\}$ or $y \notin c\{x\}$ (i.e., the two do not have the same set of neighborhoods). It is known (see [30]) that the derivational modal logic of T_0 spaces is the system

$$\text{wK4T}_0 := \text{wK4} + p \wedge \Diamond(q \wedge \Diamond p) \rightarrow \Diamond p \vee \Diamond(q \wedge \Diamond q).$$

Moreover, wK4T_0 corresponds to the class of wK4 frames (W, \rightarrow) so that $w \rightarrow v \rightarrow w$ implies that $w \rightarrow w$ or $v \rightarrow v$. Frames satisfying this property are *weakly reflexive frames*. If we define the *cluster* of $w \in W$ to be the set of points v so that $v \rightarrow^* w \rightarrow^* v$ (equivalently: s.t. either $w \leftrightarrow v$ or $w = v$), then a weakly transitive frame (W, \rightarrow) is weakly reflexive iff every cluster has at most one irreflexive point.

The second extension we are interested in is K4 , given by $\text{wK4} + \Diamond\Diamond p \rightarrow \Diamond p$. It is well known (see, e.g., [10], [8], [30]) that this is the logic of all transitive frames, and that it is also the logic of all T_D spaces. These are topological spaces (\mathcal{X}, c) such that every point is isolated in its own closure; i.e., if $x \in \mathcal{X}$, there is an open set U with $\{x\} = U \cap c\{x\}$. These results readily extend to the derivative μ -calculus.

- Theorem III.6.** 1) *The logic $\mu\text{-wK4T}_0$, obtained by adding to wK4T_0 the above Fixed Point Axiom and Induction Rule, is sound and (weakly) complete for the class of all T_0 topological spaces.*
- 2) *The logic $\mu\text{-K4}$, obtained by adding to K4 the above Fixed Point Axiom and Induction Rule, is sound and (weakly) complete for the class of all T_D topological spaces.*

We will prove this result in Section VIII. A related completeness result for the T_D case has already been proven in [26]. But that result concerns only the (semantically equivalent) tangled modal logic, while ours is about the full language of μ -calculus.

Before proving Theorem III.5, we will make a detour to discuss the tangled derivative in the context of wK4T_0 models. The results in the following section are not needed to establish our main theorems, but they motivate our choice to work over the full μ -calculus rather than focusing on tangled fragments.

IV. EXPRESSIVE INCOMPLETENESS OF TANGLE LOGIC

A natural question is whether topological μ -calculus collapses to a simpler modal fragment; if so, then a complete axiomatization of the simpler fragment would in principle suffice, and might be easier to prove than for the full calculus. As mentioned in the Introduction, this is exactly what happened on T_D spaces. Dawar and Otto [27] showed that the full μ -calculus is expressively equivalent to the so-called *tangled derivative logic* $\mathcal{L}_{\Diamond\infty}$ over the class of (finite) K4 frames, and

thus also over T_D spaces; while Goldblatt and Hodkinson [26] completely axiomatized $\mathcal{L}_{\diamond\infty}$ over these classes.⁶

In this section, we show that the Dawar-Otto result does *not* hold for general spaces, and in fact *not even* for T_0 spaces: the tangled derivative logic is no longer expressively equivalent to the μ -calculus over the class of $wK4T_0$ frames, and hence over the class of all T_0 spaces.

For each finite set of formulas Γ , consider the *tangled derivative* $\diamond^\infty\Gamma$ and *tangled closure* $\langle *\rangle^\infty\Gamma$ of Γ , as defined in Section III. Let $\mathcal{L}_{\diamond\infty}$ and $\mathcal{L}_{\langle *\rangle^\infty}$ be the natural sublanguages of the μ -calculus whose only fixed points are of the respective forms above. To prove that $\mathcal{L}_{\diamond\infty}$ is not expressively complete for μ -calculus over $wK4T_0$ frames, we will show that $\langle *\rangle^\infty$ is not definable in $\mathcal{L}_{\diamond\infty}$.

For this, we define a ‘spine’ model \mathcal{S} based on the ordinal $\omega+3$. We briefly recall that ω denotes the first infinite ordinal, and follow the set-theoretic convention that each ordinal is identified with its set of predecessors. We moreover use interval notation on the ordinals: (α, β) is the set of ordinals ξ with $\alpha < \xi < \beta$.

With this in mind, we set $\mathcal{S} = (\omega+3, \longrightarrow, \|\cdot\|)$, where

- 1) $\alpha \longrightarrow \beta$ if one of the following occurs:
 - a) $\alpha > \beta$;
 - b) $\alpha = \beta$ and α is odd (including $\omega+1$), or
 - c) $\alpha = \omega+1$ and $\beta = \omega+2$.
- 2) $\alpha \in \|\!|p\|\!$ iff α is odd, $\|\!|q\|\! = \emptyset$ for all $q \neq p$.

Lemma IV.1. \mathcal{S} is a $wK4T_0$ model.

Proof. Weak transitivity is easily checked using a case distinction. The T_0 condition is checked by noting that all clusters are singletons, except for $\{\omega+1, \omega+2\}$. But only $\omega+2$ is irreflexive, as needed. \square

The goal is to show that over \mathcal{S} , no $\mathcal{L}_{\diamond\infty}$ formula is equivalent to $\langle *\rangle^\infty\{p, \neg p\}$. First, we evaluate the latter.

Lemma IV.2. Over \mathcal{S} , $\|\langle *\rangle^\infty\{p, \neg p\}\| = \{\omega+1, \omega+2\}$.

Proof. We have that $\alpha \in \|\langle *\rangle^\infty\{p, \neg p\}\|$ if and only if there is an infinite chain

$$\alpha \longrightarrow^* \beta_0 \longrightarrow^* \beta_1 \longrightarrow^* \beta_2 \longrightarrow^* \dots$$

such that $\beta_i \notin V(p)$ when i is even, $\beta_i \in V(p)$ when i is odd. From the latter it follows that $\beta_i \neq \beta_{i+1}$. Since $\omega+3$ with the usual ordering is well-founded, such a chain can only occur in the ‘ill-founded’ part of our model, namely $\{\omega+1, \omega+2\}$. However, the infinite chain

$$\omega+2 \longrightarrow^* \omega+1 \longrightarrow^* \omega+2 \longrightarrow^* \dots$$

witnesses that $\{\omega+1, \omega+2\} \subseteq \|\langle *\rangle^\infty\{p, \neg p\}\|$. \square

⁶On the other hand, Goldblatt and Hodkinson [26] showed that \diamond^∞ is not definable in $\mathcal{L}_{\langle *\rangle^\infty}$ over the class of $K4$ frames, and hence over the class of T_D spaces. It follows that $\mathcal{L}_{\langle *\rangle^\infty}$ is not expressively complete, even over the class of $K4$ frames.

Lemma IV.3. If φ is any formula of $\mathcal{L}_{\diamond\infty}$ then there exists $n_\varphi < \omega$ such that for every $\alpha, \beta \in (n_\varphi, \omega+3)$ which are either both even or both odd, $\alpha \in \|\varphi\|$ iff $\beta \in \|\varphi\|$.

Proof. By induction on the complexity of φ . The base case follows from the definition of $\|\cdot\|$ and the cases for Booleans are straightforward. Consider the case where $\varphi = \diamond\psi$. By the induction hypothesis, n_ψ is well-defined and finite, and we can take $n_\varphi = n_\psi + 2$. Then, if $\alpha, \beta > n_\varphi$ and $\alpha \in \|\diamond\psi\|$, there is α' such that $\alpha \longrightarrow \alpha'$ and $\alpha' \in \|\psi\|$. If $\alpha' < n_\psi$, set $\beta' = \alpha'$; otherwise, take $\beta' \in \{n_\psi+1, n_\psi+2\}$ of the same parity as α' . We then see that $\beta \longrightarrow \beta'$, so that $\beta \in \|\diamond\psi\|$.

Finally, consider $\varphi = \diamond^\infty\Gamma$. We may assume that $\Gamma \neq \emptyset$, since $\diamond^\infty\emptyset$ is tautologically true. Let $n_\varphi = \max_{\gamma \in \Gamma} n_\gamma + 2$. Suppose that $\alpha, \beta > n_\varphi$, and that $\alpha \in \|\diamond^\infty\Gamma\|$. Let α_* be the least element of $\|\diamond^\infty\Gamma\|$ with respect to the usual ordinal ordering. First assume that $\alpha_* < \omega+1$. Then, for all $\gamma \in \Gamma$, there is $\xi_\gamma \in \|\gamma \wedge \diamond^\infty\Gamma\|$ such that $\alpha_* \longrightarrow \xi_\gamma$. But then, by the definition of \longrightarrow and the minimality of α_* we must have that $\xi_\gamma = \alpha_*$, and thus α_* satisfies every element of Γ . Note also that α_* is reflexive, so α_* must be odd.

It follows that $\alpha_* \leq n_\varphi$, since otherwise the odd element of $\{n_\varphi-1, n_\varphi\}$ satisfies all formulas of Γ and hence $\diamond^\infty\Gamma$, contradicting the minimality of α_* . But then, from $\beta > n_\varphi$ we see that $\beta \longrightarrow \alpha_*$, so $\beta \in \|\diamond^\infty\Gamma\|$.

Finally, we consider the case where $\alpha_* \geq \omega+1$. In fact, we will show that this case is impossible. Note that in this case $\omega+2 \in \|\diamond^\infty\Gamma\|$. As before, for all $\gamma \in \Gamma$, there is $\xi_\gamma \in \|\gamma \wedge \diamond^\infty\Gamma\|$ such that $\omega+2 \longrightarrow \xi_\gamma$. But by minimality of α_* , the only option is to have $\xi_\gamma = \omega+1$, so in fact $\omega+1$ satisfies all elements of Γ . Reasoning as above, letting $\alpha' \in \{n_\varphi-1, n_\varphi\}$ be odd, we see that α' satisfies all formulas of Γ . But then $\alpha' \in \|\diamond^\infty\Gamma\|$, contradicting the choice of α_* . \square

Remark IV.4. From the proof it can be estimated that it suffices to take $n_\varphi = 2|\varphi|$.

Since $n < \omega, \omega+2$ for all $n < \omega$ and the two are even, we obtain the following special case.

Corollary IV.5. In \mathcal{S} , ω and $\omega+2$ satisfy the same formulas of $\mathcal{L}_{\diamond\infty}$.

However, we have seen that $\omega+2 \in \|\langle *\rangle^\infty\{p, \neg p\}\|$, but $\omega \notin \|\langle *\rangle^\infty\{p, \neg p\}\|$. We may thus conclude that $\langle *\rangle^\infty\{p, \neg p\}$ is not definable.

Theorem IV.6. The formula $\varphi = \langle *\rangle^\infty\{p, \neg p\}$ is not definable in $\mathcal{L}_{\diamond\infty}$, even by an infinite set of formulas.

Given that $\langle *\rangle^\infty\{p, \neg p\}$ is definable in the μ -calculus but not in $\mathcal{L}_{\diamond\infty}$, we obtain the following result.

Corollary IV.7. Not every formula of the μ -calculus is definable in $\mathcal{L}_{\diamond\infty}$ over the class of $wK4T_0$ frames.

For this reason, in this paper we will work over the full μ -calculus, rather than the tangled derivative fragment.

V. TRUTH-PRESERVING MAPS

In this section, we focus on the relational semantics, and review and generalize some well-known properties of μ -calculus [31]: locality and invariance under bounded morphisms.

Definition V.1 (D-morphisms and P -morphisms). A d-morphism between derivative spaces (\mathcal{X}, d) and (\mathcal{X}', d') is a function $\pi: \mathcal{X} \rightarrow \mathcal{X}'$ such that $\pi^{-1}d'(X') = d\pi^{-1}(X')$ for all sets $X' \subseteq \mathcal{X}'$.

If π is surjective, we say that the space \mathcal{X}' is a d-morphic image of the space \mathcal{X} .

For any set $P \subseteq \mathbb{P}$, a P -morphism between derivative models $\mathcal{M} = (\mathcal{X}, d, \|\cdot\|)$ and $\mathcal{M}' = (\mathcal{X}', d', \|\cdot\|')$ is a d-morphism $\pi: \mathcal{X} \rightarrow \mathcal{X}'$ s.t. $\pi^{-1}\|x\|' = \|x\|$ for all atoms $x \in P$.

If π is surjective, we say that the model \mathcal{M}' is a P -morphic image of the model \mathcal{M} .

Remark V.2. The notion of P -morphism is a generalization to derivative spaces of the well-known concept of \mathfrak{p} -morphism [9], albeit relativized to a set of variables P . The restriction to such a set of variables (particularly, when P is finite) will be essential in many of our proofs.

Lemma V.3. If $\pi: \mathcal{X} \rightarrow \mathcal{X}'$ is a P -morphism between derivative models $\mathcal{M} = (\mathcal{X}, d, \|\cdot\|)$ and $\mathcal{M}' = (\mathcal{X}', d', \|\cdot\|')$, then for all μ -calculus formulas $\varphi = \varphi(\bar{x}) \in \mathcal{L}_\mu^P$ and tuples of sets \bar{X}, \bar{X}' s.t. $X_i = \pi^{-1}X'_i$ for all i , we have:

$$\|\varphi(\bar{X})\| = \pi^{-1}\|\varphi(\bar{X}')\|'.$$

Proof. Induction on the complexity of the formula $\varphi = \varphi(\bar{x})$. The atomic case $\varphi = x_i$ follows immediately from the assumption that $X_i = \pi^{-1}X'_i$, while the atomic case $\varphi = y$ for $y \in P$ not occurring in \bar{x} follows from atomic requirement on P -morphisms. The Boolean cases follow from two well-known properties of the inverse map: $\pi^{-1}(\mathcal{X}' - X') = \mathcal{X} - \pi^{-1}X'$ and $\pi^{-1}(X' \cap Y') = \pi^{-1}X' \cap \pi^{-1}Y'$. The derivative case $\diamond\varphi$ is an immediate consequence of the definition of d-morphism.

The case $\nu y.\varphi(y, \bar{x})$. We assume the induction hypothesis for φ , and we first prove the following

Claim: If $\varphi_y^\alpha(\bar{X})$ is the transfinite sequence of sets in Proposition III.1(3), then for all ordinals α we have

$$\pi^{-1}\varphi_y^\alpha(\bar{X}') = \varphi_y^\alpha(\bar{X}).$$

We prove this Claim by subinduction on α :
 $\pi^{-1}\varphi_y^\alpha(\bar{X}') = \pi^{-1}(\bigcap_{\beta < \alpha} \|\varphi(\varphi_y^\beta(\bar{X}'), \bar{X}')\|') = \bigcap_{\beta < \alpha} \pi^{-1}\|\varphi(\varphi_y^\beta(\bar{X}'), \bar{X}')\|' = \bigcap_{\beta < \alpha} \|\varphi(\varphi_y^\beta(\bar{X}), \bar{X})\| = \varphi_y^\alpha(\bar{X})$ (where at the third step we used both the induction hypothesis for φ and the subinduction hypothesis for β , as well as the background assumption about \bar{X}, \bar{X}').

Given the Claim, we can now prove the inductive case for $\nu y.\varphi(y, \bar{x})$:
 $\pi^{-1}\|\nu y.\varphi(y, \bar{X}')\|' = \pi^{-1}(\bigcap_\alpha \varphi_y^\alpha(\bar{X}')) = \bigcap_\alpha \pi^{-1}\varphi_y^\alpha(\bar{X}) = \|\nu y.\varphi(y, \bar{X})\|$. \square

It is useful to keep in mind the special case where the tuple of substitution variables is empty.

Corollary V.4 (Invariance under P -morphisms). If $\pi: \mathcal{X} \rightarrow \mathcal{X}'$ is a P -morphism between derivative models $\mathcal{M} = (\mathcal{X}, d, \|\cdot\|)$ and $\mathcal{M}' = (\mathcal{X}', d', \|\cdot\|')$, then for all $\varphi \in \mathcal{L}_\mu^P$ we have:

$$\|\varphi\| = \pi^{-1}\|\varphi\|'.$$

Proof. Apply Lemma V.3 to the empty tuple of variables, and corresponding empty tuples of sets. (Alternatively: let $\bar{x} = (x_1, \dots, x_n)$ be a tuple enumerating all the free variables in φ . Then apply Lemma V.3 to this tuple, and to tuples of sets $\bar{X} := (\|x_1\|, \dots, \|x_n\|)$ and $\bar{X}' := (\|x_1\|', \dots, \|x_n\|')$.) \square

In practice, P -morphisms are most useful when they are surjective, as they then preserve validity of formulas.

Corollary V.5.

- 1) If a derivative model $\mathcal{M}' = (\mathcal{X}', d', \|\cdot\|')$ is a P -morphic image of a model $\mathcal{M} = (\mathcal{X}, d, \|\cdot\|)$, then the two models validate (satisfy) the same formulas of \mathcal{L}_μ^P .
- 2) If a derivative space (\mathcal{X}', d') is a d-morphic image of space (\mathcal{X}, d) , then every formula that is satisfiable on (\mathcal{X}', d') is also satisfiable on (\mathcal{X}, d) ; equivalently: every formula that is valid on (\mathcal{X}, d) is also valid on (\mathcal{X}', d') .

Proof. To check part 1, we prove the satisfiability version. Let $\varphi \in \mathcal{L}_\mu^P$ and let $\pi: \mathcal{X} \rightarrow \mathcal{X}'$ be a surjective P -morphism. By Corollary V.4, we have $\|\varphi\| = \pi^{-1}\|\varphi\|'$. Combining this with the functionality and surjectivity of π , we obtain the equivalence: $\|\varphi\| \neq \emptyset$ iff $\|\varphi\|' \neq \emptyset$.

For part 2: we again check the satisfiability version. Let $\pi: \mathcal{X} \rightarrow \mathcal{X}'$ be a surjective d-morphism, and let φ be satisfiable on (\mathcal{X}', d') , i.e. there exists some valuation $\|\cdot\|'$ satisfying φ at some point of \mathcal{X}' . Take the map $\|\cdot\| := \pi^{-1}\|\cdot\|'$ defined on \mathcal{X} . Then $\|\cdot\|$ is a valuation on \mathcal{X} that makes π a surjective \mathbb{P} -morphism, hence by part 1, $\|\cdot\|$ satisfies φ at some point of \mathcal{X} . \square

It is useful to have a more ‘bisimulation-like’ characterization of d-morphisms. Using the equivalence between derivative spaces and neighborhood derivative spaces, we can characterize d-morphisms in terms of d-neighborhoods:

Lemma V.6. Let $\pi: \mathcal{X} \rightarrow \mathcal{X}'$ be a map between derivative spaces (\mathcal{X}, d) and (\mathcal{X}', d') . Then the following are equivalent:

- 1) π is a d-morphism;
- 2) the conjunction of the following back-and-forth conditions holds for all points $x \in \mathcal{X}$ and all sets $X \subseteq \mathcal{X}$ and $X' \subseteq \mathcal{X}'$:
 - (back) $X' \in \mathcal{N}_{d'}(\pi(x))$ implies $\pi^{-1}(X') \in \mathcal{N}_d(x)$, and
 - (forth) $X \in \mathcal{N}_d(x)$ implies $\pi(X) \in \mathcal{N}_{d'}(\pi(x))$;
- 3) $\pi^{-1}(X') \in \mathcal{N}_d(x)$ iff $X' \in \mathcal{N}_{d'}(\pi(x))$, for all $x \in \mathcal{X}$ and $X' \subseteq \mathcal{X}'$.

This follows from the general theory of bounded morphisms in monotonic neighborhood models [32], [28]: indeed, the third equivalent statement is exactly the definition of a

bounded morphism in monotonic neighborhood semantics. When both spaces are topological derivative spaces, the back-and-forth conditions refer to punctured neighborhoods. When both are topological closure spaces, we obtain the usual notion of *interior map*. The case where \mathcal{X} is a topological space and \mathcal{X}' a Kripke frame is of a special interest:

Corollary V.7. *Let $\pi : \mathcal{X} \rightarrow \mathcal{X}'$ be a map between a topological derivative space (\mathcal{X}, d) and a weakly transitive frame $(\mathcal{X}', \rightarrow)$. Then the following are equivalent:*

- 1) π is a d -morphism;
- 2) the conjunction of the following back-and-forth conditions holds for all points $x \in \mathcal{X}$:
 - $\pi(U - \{x\}) \subseteq \pi(x)\uparrow$, for some open neighborhood U of x , and
 - $\pi(x)\uparrow \subseteq \pi(U - \{x\})$, for all open neighborhoods U of x .

Finally, when both spaces are weakly transitive frames, we recover the standard notion of bounded frame morphism:

P -morphisms and P -bisimulations between relational models. When both \mathcal{X} and \mathcal{X}' are weakly transitive frames, it is easy to see that our notion of P -morphism matches the standard modal notion of p -morphisms (also known as “bounded P -morphisms”), i.e. *functional P -bisimulations*.

Definition V.8. *Let $\mathcal{M}_1 = (W_1, \rightarrow_1, \|\cdot\|_1)$ and $\mathcal{M}_2 = (W_2, \rightarrow_2, \|\cdot\|_2)$ be relational models. A relation $B \subseteq W_1 \times W_2$ is a P -bisimulation if, for all states $w_1 \in W_1, w_2 \in W_2$, $(w_1, w_2) \in B$ implies three conditions: (a) $w_1 \in \|p\|_1$ iff $w_2 \in \|p\|_2$ (Atomic Preservation); (b) if $w_1 \rightarrow_1 s_1$ then there exists some $s_2 \in W_2$ with $w_2 \rightarrow_2 s_2$ and $(s_1, s_2) \in B$ (Forth Condition); (c) if $w_2 \rightarrow_2 s_2$ then there exists some $s_1 \in W_1$ with $w_1 \rightarrow_1 s_1$ and $(s_1, s_2) \in B$ (Back Condition).*

Then, a bounded P -morphism is just a functional P -bisimulation. It is well known that relational P -bisimulations between weakly transitive relational models $\mathcal{M}_1 = (W_1, \rightarrow_1, \|\cdot\|_1)$ and $\mathcal{M}_2 = (W_2, \rightarrow_2, \|\cdot\|_2)$ are exactly the relations of the form $\pi_1^{-1}; \pi_2$, where $\pi_1 : \mathcal{M} \rightarrow \mathcal{M}_1$ and $\pi_2 : \mathcal{M} \rightarrow \mathcal{M}_2$ are P -morphisms from some other weakly transitive model \mathcal{M} into the two models, and $;$ is relational composition.⁷

Invariance under bisimilarity The relation of P -bisimilarity \simeq_P on a given model $\mathcal{M} = (W, \rightarrow, \|\cdot\|)$ is the largest P -bisimulation relation $\simeq_P \subseteq W \times W$. When $P = \mathbb{P}$, we drop the subscript, writing e.g. $s \simeq w$ and talking simply of ‘bisimulation’ and ‘bisimilarity’. It is easy to see that P -bisimilarity is an equivalence relation on W . The following fact is a widely known feature of μ -calculus:

Proposition V.9 (Invariance under Bisimilarity). *The valuation $\|\varphi\|$ of every formula $\varphi \in \mathcal{L}_\mu^P$ is closed under P -bisimilarity: for all $s, w \in W$, if $s \simeq_P w$ and $s \in \|\varphi\|$, then $w \in \|\varphi\|$.*

⁷This relationship between P -bisimulations and spans of bounded P -morphisms is well-known in modal logic, and has led to the general definition of coalgebraic bisimulation, as a span of coalgebraic morphisms.

Proof. This is well-known (and easy to verify directly). \square

Locality Another known fact is that μ -calculus is “local”: the truth value of a formula φ at a state depends only on the accessible part of the model (i.e., the so-called generated submodel). This can be generalized as follows:

Lemma V.10. *Let $\varphi = \varphi(\bar{x}, \bar{y})$ be a formula. Then we have the following:*

- 1) $\|\varphi(\bar{X}, \bar{Y})\| \cap w\uparrow^* = \|\varphi(\bar{X}, \overline{Y \cap w\uparrow^*})\| \cap w\uparrow^*$, for all states $w \in W$ and tuples of sets of states \bar{X}, \bar{Y} ;
- 2) If $\bar{y} = (\bar{z}, y)$ and $\varphi = \varphi(\bar{x}, \bar{y}) = \varphi(\bar{x}, \bar{z}, y)$ is positive in y , then

$$\varphi_y^\alpha(\bar{X}, \bar{Z}) \cap w\uparrow^* = \varphi_y^\alpha(\bar{X}, \overline{Z \cap w\uparrow^*}) \cap w\uparrow^*$$

for all states $w \in W$, ordinals $\alpha \in \text{On}$ and tuples of sets of states \bar{X}, \bar{Z} . (Here, φ_y^α is the sequence introduced in Proposition III.1(3).)

Proof. We show the two claims by simultaneous induction on the subformula-complexity of φ . For claim (1), the base case $\varphi = x$, as well as the inductive case for Boolean operators, are trivial.

The case of $\diamond\varphi$ for (1): $\|\diamond\varphi(\bar{X}, \bar{Y})\| \cap w\uparrow^* = \{w \in W : \exists s \in w\uparrow \text{ s.t. } s \in \|\varphi(\bar{X}, \bar{Y})\| \} \cap w\uparrow^*$. By the induction hypothesis (for φ and $s \in w\uparrow \subseteq w\uparrow^*$), this is equal to $\{w \in W : \exists s \in w\uparrow \text{ s.t. } s \in \|\varphi(\bar{X}, \overline{Y \cap w\uparrow^*})\| \} \cap w\uparrow^*$, i.e. to $\|\diamond\varphi(\bar{X}, \overline{Y \cap w\uparrow^*})\| \cap w\uparrow^*$, as desired.

The case of $\nu z.\varphi$ for (1): Using Proposition III.1(3) and the inductive hypothesis (2) for φ , we have: $\|\nu z.\varphi(\bar{X}, \bar{Y})\| \cap w\uparrow^* = \bigcap_{\alpha \in \text{On}} \varphi_z^\alpha(\bar{X}, \bar{Y}) \cap w\uparrow^* = \bigcap_{\alpha \in \text{On}} \varphi_z^\alpha(\bar{X}, \overline{Y \cap w\uparrow^*}) \cap w\uparrow^* = \|\nu z.\varphi(\bar{X}, \overline{Y \cap w\uparrow^*})\| \cap w\uparrow^*$.

To prove claim (2) for φ , assume claim (1) for $\varphi = \varphi(\bar{x}, \bar{z}, y)$ (for all set tuples), and prove (2) by subinduction on the ordinal α :

$$\begin{aligned} \varphi_y^\alpha(\bar{X}, \bar{Z}) \cap w\uparrow^* &= \bigcap_{\beta < \alpha} \|\varphi(\bar{X}, \bar{Z}, \varphi_y^\beta(\bar{X}, \bar{Z}))\| \cap w\uparrow^* \\ &= \bigcap_{\beta < \alpha} \|\varphi(\bar{X}, \overline{\bar{Z} \cap w\uparrow^*}, \varphi_y^\beta(\bar{X}, \bar{Z}))\| \cap w\uparrow^* \\ &= \bigcap_{\beta < \alpha} \|\varphi(\bar{X}, \overline{\bar{Z} \cap w\uparrow^*}, \varphi_y^\beta(\bar{X}, \overline{\bar{Z} \cap w\uparrow^*}))\| \cap w\uparrow^* \\ &= \varphi_y^\alpha(\bar{X}, \overline{\bar{Z} \cap w\uparrow^*}) \cap w\uparrow^*, \end{aligned}$$

where we used first the induction hypothesis for φ , then the subinduction hypothesis for $\beta < \alpha$. \square

Asserting properties locally above a point Given a point $w \in W$, and given a property $P(X_1, \dots, X_n)$ involving sets $X_1, \dots, X_n \subseteq W$, we say that $P(X_1, \dots, X_n)$ holds above w if we have $P(X_1 \cap w\uparrow^*, \dots, X_n \cap w\uparrow^*)$. In particular, for two sets $X, Y \subseteq W$, we say that $X = Y$ holds above w iff $X \cap w\uparrow^* = Y \cap w\uparrow^*$.

Depth of a point in a model Recall that $\overrightarrow{\neq}$ is the strict preorder induced by \rightarrow . Given a weakly transitive model

$\mathcal{M} = (W, \longrightarrow, \|\cdot\|)$, and a point $w \in W$, a *strict (finite) w-chain* is a finite sequence of points of the form $w = w_0 \xrightarrow{\neq} w_1 \xrightarrow{\neq} \dots \xrightarrow{\neq} w_n$. The number n is called the *length* of our finite chain. The *depth* $\text{dpt}(w)$ of the point $w \in W$ is the supremum of the lengths of all strict w -chains. In general, we have $\text{dpt}(w) \geq 0$, with $\text{dpt}(w) = 0$ iff for every $s \in W$, $w \longrightarrow s$ implies $s \longrightarrow w$; and $\text{dpt}(w) = \omega$ iff there exist w -chains of every length $n \in \mathbb{N}$. The *depth* $\text{dpt}(\mathcal{M})$ of the model \mathcal{M} is the supremum of the depths of all points of the model:

$$\text{dpt}(\mathcal{M}) := \sup\{\text{dpt}(w) : w \in W\}.$$

Lemma V.11. *Let $\mathcal{M} = (W, \longrightarrow, \|\cdot\|)$ be a weakly transitive model, and $w, s \in W$ be two points. Then we have the following:*

- 1) if $w \longrightarrow^* s$, then $\text{dpt}(w) \geq \text{dpt}(s)$;
- 2) if $w \longleftrightarrow s$, then $\text{dpt}(w) = \text{dpt}(s)$;
- 3) if $w \longrightarrow s$ and $\text{dpt}(w) = \text{dpt}(s) < \omega$, then $w \longleftrightarrow s$;
- 4) if $w \xrightarrow{\neq} s$ and $\text{dpt}(s)$ is finite, then $\text{dpt}(w) > \text{dpt}(s)$.

Proof. Easy verification. \square

Our goal in the next section is to prove Proposition III.5, in particular the completeness of our axiomatization with respect to irreflexive, weakly transitive frames. But for this, recall first that modal logic cannot express irreflexivity. The following result allows us to drop the irreflexivity condition:

Lemma V.12. *For every weakly-transitive model \mathcal{M} , there exists some irreflexive weakly-transitive model $\widetilde{\mathcal{M}}$ that validates/satisfies the same μ -calculus formulas as \mathcal{M} . Moreover, if \mathcal{M} is finite, then $\widetilde{\mathcal{M}}$ can be taken to be finite as well.*

Proof. Given any weakly transitive model $\mathcal{M} = (W, \longrightarrow, \|\cdot\|)$, we associate to it an irreflexive and weakly transitive model $\widetilde{\mathcal{M}} = (\widetilde{W}, \widetilde{\longrightarrow}, \|\widetilde{\cdot}\|)$, by first taking

$$\begin{aligned} \widetilde{W} := & \{x \in W : x \text{ is irreflexive}\} \\ & \cup \{(x, i) \in W \times \{0, 1\} : x \longrightarrow x\}. \end{aligned}$$

It is useful to consider a map $\pi : \widetilde{W} \rightarrow W$, given by $\pi(x, i) := x$ (for reflexive points $x \in W$) and $\pi(x) := x$ (for irreflexive points $x \in \widetilde{W}$). Using this, we can define the accessibility relation on \widetilde{W} by putting

$$\widetilde{x} \widetilde{\longrightarrow} \widetilde{y} \text{ if } \pi(\widetilde{x}) \longrightarrow \pi(\widetilde{y}) \text{ and } \widetilde{x} \neq \widetilde{y},$$

for all $\widetilde{x}, \widetilde{y} \in \widetilde{W}$; and we define the valuation on \widetilde{W} by

$$\|\widetilde{p}\| := \{\widetilde{x} \in \widetilde{W} : \pi(\widetilde{x}) \in \|p\|\}.$$

It is easy to see that $\widetilde{\mathcal{M}}$ is an irreflexive and weakly transitive relational model, and that the map $\pi : \widetilde{W} \rightarrow W$ is a \mathbb{P} -morphism. Since by Proposition V.9, all formulas of μ -calculus are invariant under bisimulation, the two models are equivalent with respect to our syntax. \square

So, to prove Proposition III.5, it is enough to show *completeness and FMP for weakly transitive frames*. This is topic of the next section.

VI. PROOF OF THE MAIN COMPLETENESS/FMP RESULT

In this section, we prove our main completeness result (Theorem III.5). Throughout the section, we fix a consistent formula φ_0 , and let $P_0 = \text{Free}(\varphi_0)$. We also fix some *finite* set $\Sigma \subseteq \mathcal{L}_\mu$, with the following properties: $\varphi_0 \in \Sigma$; Σ is closed under subformulas; Σ is closed up to logical equivalence (in our axiomatic system) under negation $\neg\varphi$ and under $\langle * \rangle\varphi$ operators. The existence of such a finite set Σ (for every formula φ_0) follows from the fact that $\langle * \rangle$ is provably an S4-type modality, together with the well-known fact that there are only finitely many non-equivalent modalities in the modal system S4 [10, Ch. 3]. Note that $\varphi_0 \in \Sigma$, and $P_0 \subseteq \Sigma$ is finite.

Plan of the Proof We start with the canonical model Ω (comprising all maximally consistent theories), a standard construction in modal logic. But we should stress that the *canonical model is not our intended model*. Indeed, the usual Truth Lemma *fails* for the μ -calculus in the canonical model: consistent μ -calculus formulas are *not* necessarily satisfied in the canonical model by the theories that contain them.⁸ In fact, the notion of truth in the canonical model will play *no role* in this paper: we never evaluate our formulas in Ω . Instead, we only use a few basic syntactic properties of this model.

Next, we select a special submodel of the canonical model Ω^Σ (called the Σ -*final model*). Essentially, this consists of the theories whose cluster is locally definable by some formula in Σ . Our goal will be to show that the Truth Lemma *does* hold in Ω^Σ for Σ -formulas. It is easy to show that Ω^Σ satisfies the usual \diamond -Witness Lemma for formulas in Σ , but extending this to fixed points requires some work.

An important role will be played by the notion of Σ -bisimilarity, a strengthening of the standard notion of bisimilarity, in which the Atomic Permanence clause is replaced by the requirement that Σ -bisimilar theories agree on Σ -formulas. Since it is stronger than usual P_0 -bisimilarity, Σ -bisimilarity still preserves the truth values of μ -calculus formulas, as long as their free variables belong to Σ .

Another key ingredient in our proof is the fact that Ω^Σ is “essentially” a finite object: though possibly infinite in size, it has finite ‘depth’, and moreover it contains only finitely many Σ -bisimilarity classes. As a consequence, all relevant fixed points are attained at some fixed finite stage of the iterative process from Proposition III.1(3).

We will then use these ingredients to prove our Truth Lemma for the final model Ω^Σ . The inductive step for the fixed-point formulas uses the fact that the valuation of these formulas is locally definable by some Σ -formula.

Once completeness is obtained in this way, we will prove the finite model property by taking the quotient of the final model Ω^Σ modulo Σ -bisimilarity.

⁸To see this, consider atoms $(p_n)_{n < \omega}$ and check that for every n , the set $\Phi_n := \{p_n, \neg \diamond^\infty \top\} \cup \{[*](p_i \Rightarrow \diamond p_{i+1}) : i < \omega\}$ is consistent (since all finite subsets are satisfiable). Use the Canonical Truth Lemma for Basic Modal Logic (and the fact that $[*]$ is definable in it) to construct $(T_n)_{n < \omega}$ with $\Phi_n \subseteq T_n$ and $T_0 \rightarrow T_1 \rightarrow \dots \rightarrow T_n \rightarrow \dots$. Thus, $T_0 \models \diamond^\infty \top$ although $(\neg \diamond^\infty \top) \in T_0$.

Canonical Model The standard ‘canonical model’ construction provides an (infinite) weakly transitive model. A *theory* is a maximally consistent set of formulas in \mathcal{L}_μ (i.e. a set $T \subseteq \mathcal{L}_\mu$ that is consistent and has no proper consistent extension). We denote by Ω the family of all theories. The *canonical accessibility relation* \longrightarrow between two such theories $T, T' \in \Omega$ is given as usual by putting

$$T \longrightarrow T' \text{ iff } \forall \varphi \text{ (if } \Box\varphi \in T \text{ then } \varphi \in T'),$$

and the canonical valuation is given by

$$\|x\| := \{T \in \Omega : x \in T\}.$$

The *canonical model* is the structure $(\Omega, \longrightarrow, \|\cdot\|)$. Since the weak-transitivity condition is Sahlqvist, it immediately follows that *the canonical model is weakly transitive* (though not irreflexive); see [9], [10] for details on Sahlqvist formulas and their properties. As a consequence, *the reflexive closure*, which we denote \longrightarrow^* , of the canonical relation coincides with its reflexive-transitive closure.

We will make use of a few well-known properties of the canonical model, given by the next four results (see, e.g., [9]).

Lemma VI.1 (Lindenbaum Lemma). *Every consistent set Φ of formulas can be extended to a maximal consistent set $T \in \Omega$ s.t. $\Phi \subseteq T$.*

Lemma VI.2 (Canonical \diamond -Witness Lemma). *For every theory $T \in \Omega$ and formula $\varphi \in \mathcal{L}_\mu$, we have that $\diamond\varphi \in T$ iff there exists some theory $T' \in \Omega$ s.t. $T \longrightarrow T' \ni \varphi$.*

We also have an equivalent statement in \Box -form:

$$\Box\varphi \in T \text{ iff } \forall T' \in \Omega \text{ (if } T \longrightarrow T' \text{ then } \varphi \in T').$$

The left-to-right implication in the first statement above is known as the (Canonical) \diamond -Existence Lemma. The proofs are well-known (see, e.g., [9, Ch. 4]), and these results imply that the so-called Truth Lemma holds in the canonical model for the \diamond -fragment of our logic.

In fact, we can extend this to a Canonical $\langle * \rangle$ -Witness Lemma, using the following result

Lemma VI.3. *For theories $T, T' \in \Omega$, we have:*

$$T \longrightarrow^* T' \text{ iff } \forall \varphi \text{ (if } [*]\varphi \in T \text{ then } \varphi \in T').$$

Proof. The *left-to-right implication*: Assume that $T \longrightarrow^* T'$. If $T = T'$, then $[*]\varphi \in T$ implies by definition that $\varphi \in T = T'$, as desired. If $T \neq T'$, then we must have $T \longrightarrow T'$, and then $[*]\varphi \in T$ implies by definition that $\Box\varphi \in T$, which implies that $\varphi \in T'$ (by the Canonical \diamond -Witness Lemma), as desired.

The *right-to-left implication*: Assume that we have $\forall \varphi ([*]\varphi \in T \implies \varphi \in T')$. To show that $T \longrightarrow^* T'$, we assume that $T \neq T'$, and we need to prove that $T \longrightarrow T'$. Since $T \neq T'$, there exists some formula $\theta \in T$ with $\theta \notin T'$. To show the desired conclusion, let ϕ be any arbitrary formula s.t. $\Box\phi \in T$, and we need to prove that $\varphi \in T'$. From $\theta \in T$, we infer $(\varphi \vee \theta) \in T$; similarly, from $\Box\phi \in T$, we infer $\Box(\varphi \vee \theta) \in T$. Putting these together, we obtain $[*](\varphi \vee \theta) \in T$.

By our assumption, this implies that $(\varphi \vee \theta) \in T'$, and since $\theta \notin T'$, we conclude that $\varphi \in T'$, as desired. \square

As a consequence of Lemma VI.3, we immediately get:

Lemma VI.4 (Canonical $\langle * \rangle$ -Witness Lemma). *For every formula φ and theory $T \in \Omega$, we have that $\langle * \rangle\varphi \in T$ iff there exists some theory $T' \in \Omega$ s.t. $T \longrightarrow^* T' \ni \varphi$.*

Final Theories Given a formula θ , a theory $T \in \Omega$ is *θ -final* if we have: $\theta \in T$, and for all theories $S \in \Omega$, if $T \longrightarrow S$ and $\theta \in S$ then $S \longrightarrow T$ (hence $T \longleftrightarrow S$). Given a set Σ of formulas, a theory $T \in \Omega$ is *Σ -final* (or ‘final’, for short) if it is θ -final for some formula $\theta \in \Sigma$.

Final Model Let Σ be any set of formulas. The *final model* is the canonical submodel⁹ determined by the set $\Omega^\Sigma := \{T \in \Omega : T \text{ is } \Sigma\text{-final}\}$ of all final theories.

The final model may be infinite, but we can show that it has finite depth:

Lemma VI.5 (Finite Depth Lemma). *The final model Ω^Σ has depth bounded by $|\Sigma| - 1$. In other words: for every chain of Σ -final theories $T_0 \xrightarrow{\not\rightarrow} T_1 \xrightarrow{\not\rightarrow} \dots T_n$, we have that $n \leq |\Sigma| - 1$.*

Proof. Suppose, towards a contradiction, that $T_0 \longrightarrow T_1 \longrightarrow \dots T_n$ is a strict chain of Σ -final theories of length $n \geq |\Sigma|$. Since all T_i are Σ -final, there exist formulas $\theta_0, \dots, \theta_n \in \Sigma$ s.t. T_i is θ_i -final (and hence $\theta_i \in T_i$) for all $i \leq n$. But this is a sequence of $n + 1 \geq |\Sigma| + 1 > |\Sigma|$ formulas in Σ , so some formula θ must be repeated. Let θ be such a repeating formula in the enumeration, and let i and j be indices such that $i < j$ and $\theta_i = \theta_j = \theta$.

So we have $T_i \longrightarrow T_{i+1} \longrightarrow^* T_j$, with both T_i and T_j being θ -final, and so also $T_i \longrightarrow^* T_j$. We have two cases: either $T_i \longrightarrow T_j$ or $T_i = T_j$. We claim that in both cases we have $T_{i+1} \longrightarrow^* T_i$. To show this, consider first the case $T_i \longrightarrow T_j$. By θ -finality we get $T_i \longleftrightarrow T_j$, hence $T_i \longrightarrow T_{i+1} \longrightarrow^* T_j \longleftrightarrow T_i$, and thus $T_i \longrightarrow T_{i+1} \longrightarrow^* T_i$, as desired. In the second case, we assume $T_i = T_j$, so we immediately obtain $T_{i+1} \longrightarrow^* T_j = T_i$, as desired.

So we showed that we have $T_i \longrightarrow T_{i+1} \longrightarrow^* T_i$. There are again two cases: either $T_i \longrightarrow T_{i+1} \longrightarrow T_i$, or $T_i \longrightarrow T_{i+1} = T_i$. In the first case, we immediately conclude that $T_i \longleftrightarrow T_{i+1}$, which contradicts the ‘strictness’ of our chain. In the second case, we have $T_i \longrightarrow T_{i+1} = T_i \longrightarrow T_{i+1} = T_i$, so we again conclude that $T_i \longleftrightarrow T_{i+1}$, in contradiction with our ‘strictness’ assumption. \square

In order to prove completeness with respect to the final model, we first need to show that every consistent formula belongs to some final theory. This is achieved by combining the Lindenbaum Lemma with the following.

⁹Any subset $X' \subseteq X$ of the set of worlds of a relational model $M = (X, \longrightarrow, \|\bullet\|)$ determines a unique *submodel*, obtained by taking: X' as its set of worlds; the restriction of \longrightarrow to X' as its accessibility relation; and the valuation given by $\|p\| \cap X'$.

Lemma VI.6 (Final Lemma). *If $\varphi \in T \in \Omega$, then there exists some φ -final theory $T^* \in \Omega$ such that $T \longrightarrow^* T^*$ (and obviously, $\varphi \in T^*$, by finality).*

Proof. We will use a well-known variant of Zorn's Lemma, stated for preorders: a preordered set (\mathcal{S}, \leq) has a maximal element if every chain has an upper bound. (Here, being maximal in a preordered set means that there is no strictly larger element.)

Let $\varphi \in T \in \Omega$. Take $\mathcal{S} := \{T' \in \Omega : T \longrightarrow^* T' \ni \varphi\}$, with the relation \longrightarrow^* as its preorder. Let $\mathcal{S}' \subseteq \mathcal{S}$ be a chain of theories in \mathcal{S} . To show that it has an upper bound, take the set

$$\Phi := \{\varphi\} \cup \{[*]\theta : [*]\theta \in T' \text{ for some } T' \in \mathcal{S}'\}$$

We show that Φ is consistent: suppose this is not the case. Then there exists some finite such inconsistent subset $\Phi' = \{\varphi\} \cup \{[*]\theta_1, \dots, [*]\theta_n\}$, with $[*]\theta_1 \in T_1, \dots, [*]\theta_n \in T_n$ for some theories $T_1, T_2, \dots, T_n \in \mathcal{S}'$. Since \mathcal{S}' is a chain, we can assume that $T_1, T_2, \dots, T_{n-1} \longrightarrow^* T_n$, and thus $[*]\theta_1, \dots, [*]\theta_n \in T_n$. Since $T_n \in \mathcal{S}$, we also have $\varphi \in T_n$, so $\Phi' \subseteq T_n$, which contradicts the consistency of T_n .

Applying now Lindenbaum's Lemma, there exists some maximally consistent extension $S \in \Omega$ with $\Phi \subseteq S$. By construction (and using Lemma VI.3), we have $T' \longrightarrow^* S$ for all $T' \in \mathcal{S}'$, so S is an upper bound for the chain \mathcal{S} . Applying Zorn's lemma, we obtain a \longrightarrow^* -maximal element $T^* \in \mathcal{S}$. In particular, this means that $\varphi \in T^*$ and $T \longrightarrow^* T^*$, as desired. To prove that T^* is φ -final, suppose that $T^* \longrightarrow S \ni \varphi$; we have to show that $S \longrightarrow T^*$. By the \longrightarrow^* -maximality of T^* , we must have $S \longrightarrow^* T^*$, i.e. either $S \longrightarrow T^*$ or $S = T^*$. If $S \longrightarrow T^*$, then we are done. If $S = T^*$, then $S = T^* \longrightarrow S = T^*$, so we get again $S \longrightarrow T^*$, as desired. \square

Using similar reasoning, we may establish an analogue of the \diamond -Witness Lemma for final theories:

Lemma VI.7 (Final \diamond -Witness Lemma). *For any theory $T \in \Omega$ and formula φ , we have that $\diamond\varphi \in T$ iff there exists some φ -final theory T' such that $T \longrightarrow T'$. (Obviously, we have $\varphi \in T'$ in this case, by finality.)*

Proof. The left-to-right implication: by the Canonical \diamond -Witness Lemma VI.2, $\diamond\varphi \in T$ implies the existence of some theory S with $T \longrightarrow S$ and $\varphi \in S$. By the Final Lemma VI.6, there exists some φ -final theory S^* with $S \longrightarrow S^*$ and $\varphi \in S^*$. If $T \longrightarrow S^*$, then we can take $T' := S^*$ and we are done (since S^* is φ -final and $T \longrightarrow S \ni \varphi$, as desired). If $T \not\longrightarrow S^*$, then from this and $T \longrightarrow S \longrightarrow S^*$ we get by weak transitivity that $T = S^*$, and so $T \longrightarrow S \longrightarrow S^* = T$. In this case, we can take $T' := S$. Indeed, since we already know that $T \longrightarrow S \ni \varphi$, to finish the proof we only need to check that S is φ -final. For this, let $U \in \Omega$ be any theory with $S \longrightarrow U \ni \varphi$; we need to show that $U \longrightarrow S$. From $S^* = T \longrightarrow S \longrightarrow U$, we obtain by weak transitivity that either $U = S^* = T \longrightarrow S$ (and we are done), or $S^* \longrightarrow U \ni \varphi$. In the second case, by the φ -finality of S^* , we

have $U \longrightarrow S^* = T \longrightarrow S$; by weak transitivity, we obtain either $U \longrightarrow S$ (and we are done) or $U = S \longrightarrow U = S$. So, in all cases, we concluded that $U \longrightarrow S$, as desired.

The converse follows directly from the Canonical \diamond -Witness Lemma VI.2, as a special case. \square

It will be useful to observe that θ -final theories are closely related to $\langle *\rangle\theta$ -final theories.

Lemma VI.8. *Let $T \in \Omega$ be a θ -final theory. Then:*

- 1) *T is also $\langle *\rangle\theta$ -final.*
- 2) *For every $S \in \Omega$ s.t. $T \longrightarrow^* S$, we have $\langle *\rangle\theta \in S$ iff either $T = S$ or $T \longleftrightarrow S$.*
- 3) *All theories $S \in \Omega$ satisfying $T \longleftrightarrow S$ are $\langle *\rangle\theta$ -final.*

Proof. Assume T is θ -final. To show that it is also $\langle *\rangle\theta$ -final, observe that we have $\langle *\rangle\theta \in T$ (since $\theta \implies \langle *\rangle\theta$ is a theorem in our logic). Second, let $S \in \Omega$ be s.t. $T \longrightarrow S$ and $\langle *\rangle\theta \in S$, and we need to prove that $S \longrightarrow T$. Since $\langle *\rangle\theta \in S$, we have either $\theta \in S$ or $\diamond\theta \in S$. In the first case, from $T \longrightarrow S \ni \theta$ and the fact that T is θ -final, we conclude that $S \longrightarrow T$, as desired. In the second case, from $\diamond\theta \in S$ we infer (by the Canonical \diamond -Witness Lemma) that there exists $S' \in \Omega$, with $S \longrightarrow S' \ni \theta$. Since $T \longrightarrow S \longrightarrow S'$, by weak transitivity we have either $T = S'$ or $T \longrightarrow S'$. If $T = S'$, then we conclude $S \longrightarrow S' = T$, and we are done. If $T \longrightarrow S'$, then since T is θ -final and $\theta \in S'$, we get $S' \longrightarrow T$. Thus we have $S \longrightarrow S' \longrightarrow T$, hence by weak transitivity we get that either $S \longrightarrow T$ (and we are done) or $S = T$ (in which case $S = T \longrightarrow S = T$, so we again obtain $S \longrightarrow T$, as desired).

For the second claim of the Lemma: assuming $S \in \Omega$ s.t. $T \longrightarrow^* S$ (i.e. $T = S$ or $T \longrightarrow S$), we need to show that $\langle *\rangle\theta \in S$ holds iff either $T = S$ or $T \longleftrightarrow S$. The case $T = S$ is obvious. In the case $T \longrightarrow S$, the left-to-right implication follows from the fact that T is $\langle *\rangle\theta$ -final. As for the converse: assuming $T \longleftrightarrow S$, and using the fact that $\diamond\theta \in T$, we obtain $\diamond\langle *\rangle\theta \in S$ (by the Diamond Existence Lemma and $S \longrightarrow T$), and so $\langle *\rangle\theta \in S$ (because $\diamond\langle *\rangle\theta \implies \langle *\rangle\theta$ is a theorem in our axiomatic system).

For the third claim of the Lemma: assume that $S \in \Omega$ is s.t. $T \longleftrightarrow S$. Since this implies that $T \longrightarrow^* S$, we are in the conditions of the second claim, and hence we can apply it to derive from $T \longleftrightarrow S$ that $\langle *\rangle\theta \in S$. To show finality, let $S' \in \Omega$ be s.t. $S \longrightarrow S'$ and $\langle *\rangle\theta \in S'$; we need to prove that $S' \longrightarrow S$. From $T \longrightarrow^* S$ and $S \longrightarrow S'$, we get $T \longrightarrow^* S'$. From this and $\langle *\rangle\theta \in S'$, we obtain by the second claim that we have either $S' \longleftrightarrow T$ or $S' = T$. Both cases, combined with the fact that $T \longrightarrow^* S$, give us $S' \longrightarrow^* S$. This means that we either have $S' \longrightarrow S$ (as desired) or else $S' = S$ (in which case we have $S' = S \longrightarrow S' = S$, hence we again get $S' \longrightarrow S$, as desired). \square

Locality in the final model For the rest of this section, whenever we talk about 'locality', we refer to the final model Ω^Σ . In particular, for $T \in \Omega^\Sigma$, we use the notations $T \uparrow := \{S \in \Omega^\Sigma : T \longrightarrow S\}$, $T \uparrow^* := \{S \in \Omega^\Sigma : T \longrightarrow^* S\}$, and whenever we write that a property holds locally "above T ", we mean that it holds above T in Ω^Σ .

Notation. It is useful to introduce the notation

$$\widehat{\varphi} := \{T \in \Omega^\Sigma : \varphi \in T\}$$

for all formulas $\varphi \in \mathcal{L}_\mu$. From the definition of the canonical valuation on (the canonical model, and hence on) the final submodel, it is obvious that we have $\|x\|_{\Omega^\Sigma} = \widehat{x}$, for all atoms $p \in P$. Our goal is to extend this observation to all sentences in Σ .

Σ -Bisimilarity in the Final Model We can apply the concept P -bisimilarity \simeq_P to the final model Ω^Σ (for any set of variables $P \subseteq \mathbb{P}$), and in fact the special case of P_0 -bisimilarity \simeq_{P_0} will be relevant for our proof. But it is convenient to introduce a stronger notion: a relation $B \subseteq \Omega^\Sigma \times \Omega^\Sigma$ is a Σ -bisimulation if it satisfies the same back-and-forth clauses as a usual P -bisimulation, but the Atomic Preservation clause is replaced by the requirement that: $(T, T') \in B$ implies $T \cap \Sigma = T' \cap \Sigma$. The relation $\simeq_\Sigma \subseteq \Omega^\Sigma \times \Omega^\Sigma$ of Σ -bisimilarity is defined as the largest Σ -bisimulation relation on Ω^Σ .

It is easy to see that \simeq_Σ is an equivalence relation on Ω^Σ , and that it is stronger than P_0 -bisimilarity: if $T \simeq_\Sigma T'$ then $T \simeq_{P_0} T'$. Using this and the above-mentioned well-known fact about invariance of μ -calculus under standard bisimilarity, we immediately obtain the following:

Lemma VI.9. *All the formulas $\varphi \in \mathcal{L}_\mu^{P_0}$ are invariant under Σ -bisimilarity, i.e. if $T, T' \in \Omega^\Sigma$ satisfy $T \simeq_\Sigma T'$, then for all $\varphi \in \mathcal{L}_\mu^{P_0}$ we have $T \in \|\varphi\|$ iff $T' \in \|\varphi\|$.*

It is useful to introduce a more “local” version of closure under bisimilarity.

Closure under Σ bisimilarity above a point This is just a special case of the general notion of asserting a property locally: a set $X \subseteq \Omega^\Sigma$ is closed under Σ -bisimilarity above a theory $T \in \Omega^\Sigma$ if $X \cap T \uparrow^*$ is closed under Σ -bisimilarity.

Of course, global closure implies local closure: if a set $X \subseteq \Omega^\Sigma$ is closed under Σ -bisimilarity, then it is also closed under Σ -bisimilarity above every $T \in \Omega^\Sigma$. Note also that: X is closed under Σ -bisimilarity above T iff $X \cap T \uparrow^*$ is.

Convention on global/local versions Sometimes we want to assert that both the global and the local version of a statement hold in the final model Ω^Σ : e.g. if a certain premise holds, either globally or locally, then a certain conclusion holds, either globally or locally. To do this in a compact manner, we will state the global version, but adding in brackets the words “above T ”, to include the local version as well. An example is the following result:

Proposition VI.10. *If $\overline{X} = (X_1, \dots, X_n)$ is a tuple of sets $X_i \subseteq \Omega^\Sigma$ that are closed under Σ -bisimilarity (above some $T \in \Omega^\Sigma$), and $\varphi = \varphi(\overline{x}) \in \Sigma$ is a Σ -formula, then $\|\varphi(\overline{X})\|$ is also closed under Σ -bisimilarity (above T).*

Proof. We prove only the local version (since the proof of the global statement is just a simplification the local proof, obtained by omitting every mention of T). Let $Q = \{q_1, \dots, q_n\}$ be a set of n ‘fresh’ propositional atoms (with $\mathbb{P} \cap Q = \emptyset$). We

extend the valuation of the final model Ω^Σ to all the atoms in $\mathbb{P} \cup Q$, by putting $\|q_i\| = X_i \cap T \uparrow^*$, for all $i \leq n$. Then, using the fact that all X_i are invariant under Σ -bisimilarity above T (together with the fact that $P_0 \subseteq \Sigma$), it is easy to see that Σ -bisimilarity above T implies $P_0 \cup Q$ -bisimilarity above T , i.e.: if $T', T'' \in T \uparrow^*$ are s.t. $T' \simeq_\Sigma T''$, then $T' \simeq_{P_0 \cup Q} T''$. Putting this together with the fact that $\varphi(\overline{q}) \in \mathcal{L}_\mu^{P_0 \cup Q}$ is closed under $P_0 \cup Q$ -bisimilarity and using Lemma V.10, we conclude that $\|\varphi(\overline{X})\| \cap T \uparrow^* = \|\varphi(\overline{X \cap T \uparrow^*})\| \cap T \uparrow^* = \|\varphi(\|\overline{q}\|\|) = \|\varphi(\overline{q})\|$ is closed under Σ -bisimilarity above T , and hence $\|\varphi(\overline{X})\|$ is also closed under Σ -bisimilarity above T . \square

Next, we will use the following easy observation:

Lemma VI.11. *If $T, T' \in \Omega^\Sigma$ are such that $T \longleftrightarrow T'$ and $T \cap \Sigma = T' \cap \Sigma$, then $T \simeq_\Sigma T'$.*

Proof. Take

$$B := \{(T, T') \in \Omega^\Sigma \times \Omega^\Sigma : T \cap \Sigma = T' \cap \Sigma, \text{ and } T \longleftrightarrow^* T'\}$$

Clearly, to prove our lemma it is enough to show that B is a Σ -bisimulation.

For this, assume $(T, T') \in B$, and we have to check that (T, T') satisfy the three clauses in the definition of a P -bisimulation:

“Atomic” preservation is automatically ensured by the fact that $T \cap \Sigma = T' \cap \Sigma$.

For the forth condition, let $S \in \Omega^\Sigma$ such that $T \longrightarrow S$. We need to show that this, together with $(T, T') \in B$, implies the existence of some $S' \in \Omega^\Sigma$ with $T' \longrightarrow S'$ and $(S, S') \in B$:

If $T = T'$, we can take $S' := S$, and we are done, since $(S, S') = (S, S) \in B$. Otherwise, we have $T \longleftrightarrow T'$ and $T \longrightarrow S$. From these, we infer that we have either $T' \longrightarrow S$ or $T' = S$. In the first case, we can take again $S' := S$, and we are done. In the second case, $T' = S$ yields $T \longleftrightarrow T' = S$, hence we can take $S' := T$: we then have $T' \longrightarrow T = S'$ and $(S, S') = (T', T) \in B$ (by the symmetry of B and the fact that $(T, T') \in B$), as desired.

The back condition follows from the satisfaction of the forth condition and the symmetry of B . \square

Notations: (sets of) Σ -bisimilarity classes. It is convenient to introduce a notation for Σ -bisimilarity classes over the final model: for every final theory $T \in \Omega^\Sigma$, we put

$$T_\Sigma := \{S \in \Omega^\Sigma : T \simeq_\Sigma S\}$$

for the Σ -bisimilarity class of T . For every set $\mathcal{S} \subseteq \Omega^\Sigma$ of final theories, we put

$$\mathcal{S}_\Sigma := \{S_\Sigma : S \in \mathcal{S}\}$$

for the set of Σ -bisimilarity classes of theories in \mathcal{S} . In particular, for the case of the set Ω^Σ of all final theories, we simplify the notation, writing

$$\Omega_\Sigma := (\Omega^\Sigma)_\Sigma = \{S_\Sigma : S \in \Omega^\Sigma\}$$

for the set of Σ -bisimilarity classes of all Σ -final theories. Similarly, for each number n , we put

$$\begin{aligned}\Omega_\Sigma^n &:= \{T \in \Omega^\Sigma : \text{dpt}(T) \leq n\}_\Sigma \\ &= \{T_\Sigma : T \in \Omega^\Sigma \text{ with } \text{dpt}(T) \leq n\}\end{aligned}$$

for the set of all Σ -bisimilarity classes of theories of depth no larger than n . By the Finite Depth Lemma VI.5, we have $\Omega_\Sigma = \Omega_\Sigma^N$ for some natural number N .

Proposition VI.12. *There are only finitely many distinct bisimilarity classes in the final model Ω^Σ .*

Proof. It is enough to show that, for each natural number n , the set Ω_Σ^n is finite (since the desired conclusion will obviously follow from the above observation that the set of all final Σ -bisimilarity classes Ω_Σ coincides with Ω_Σ^N for some number N).

The finiteness of $|\Omega_\Sigma^n|$ for all n follows immediately by induction from the following two claims:

- 1) $|\Omega_\Sigma^0| \leq 2^{|\Sigma|} \cdot 2^{2^{|\Sigma|}}$;
- 2) $|\Omega_\Sigma^n| \leq 2^{|\Sigma|} \cdot 2^{2^{|\Sigma|}} \cdot 2^{|\Omega_\Sigma^{n-1}|}$ for all $n > 0$.

To prove these two claims, note first that, for every final theory $T \in \Omega^\Sigma$, its bisimilarity class T_Σ is uniquely determined by the pair $(T \cap \Sigma, T \uparrow_\Sigma)$, where $T \uparrow_\Sigma = \{S_\Sigma : T \longleftrightarrow S\}$ is the set of Σ -bisimilarity classes of T 's successors. We can split further this second component into two parts, depending on whether these bisimilarity classes are of the same depth as T or of lower depth. In other words: for a final theory T of depth n , its Σ -bisimilarity class T_Σ is uniquely determined by the triplet

$$(T \cap \Sigma, T \uparrow_\Sigma - \Omega_\Sigma^{n-1}, T \uparrow_\Sigma \cap \Omega_\Sigma^{n-1}),$$

where the third component is empty when $n = 0$.

To count these triplets, note that the number of distinct possibilities for the first component of the triple $T \cap \Sigma \subseteq \Sigma$ is at most $2^{|\Sigma|}$. Further, since $T \longleftrightarrow S$ and $\text{dpt}(S) \geq n = \text{dpt}(T)$ implies $T \longleftrightarrow S$, we have that $T \uparrow_\Sigma - \Omega_\Sigma^{n-1} \subseteq \{S_\Sigma : T \longleftrightarrow S\}$. But, by Lemma VI.11 (and the fact that $T \longleftrightarrow S, S'$ implies by weak transitivity that we have either $S = S'$ or $S \longleftrightarrow S'$), distinct elements $S_\Sigma \neq S'_\Sigma$ of this last set must have $S \cap \Sigma \neq S' \cap \Sigma$. So the set $\{S_\Sigma : T \longleftrightarrow S\}$ has at most $2^{|\Sigma|}$ elements, and thus the number of distinct possibilities for the second component of the triple is at most $2^{2^{|\Sigma|}}$. Finally, for the third component, we have $T \uparrow_\Sigma \cap \Omega_\Sigma^{n-1} \subseteq \Omega_\Sigma^{n-1}$ for $n > 0$ (and is empty for $n = 0$), so the number of distinct possibilities for the third component is at most $2^{|\Omega_\Sigma^{n-1}|}$ for $n > 0$ (and is $= 0$ for $n = 0$). The above two claims immediately follow. \square

Corollary VI.13. *For every fixed-point formula $\nu y.\varphi(y, \bar{x})$ where the values of \bar{x} are closed under Σ -bisimilarity, the iterative process in Proposition III.1(3) reaches its fixed point on the final model (above some $T \in \Omega^\Sigma$) at some finite stage. More precisely: for all tuples \bar{X} of subsets of Ω^Σ that are closed under Σ -bisimilarity (above some $T \in \Omega^\Sigma$), there exists some N s.t. we have that*

$$\|\nu y.\varphi(y, \bar{X})\| = \bigcap_n \varphi_y^n(\bar{X}) = \varphi_y^N(\bar{X}) \text{ holds (above } T),$$

where $\varphi_y^0(\bar{X}) := \Omega^\Sigma$, $\varphi_y^{n+1}(\bar{X}) = \|\varphi(\varphi_y^n(\bar{X}), \bar{X})\|$ (and all the formulas are interpreted in the final model Ω^Σ).

Proof. It is obvious that the sequence

$$\varphi_y^0(\bar{X}) \supseteq \varphi_y^1(\bar{X}) \supseteq \dots \varphi_y^n(\bar{X}) \supseteq \dots$$

stabilizes, reaching the fixed point (above T) at the first stage N s.t. $\varphi_y^N(\bar{X}) = \varphi_y^{N+1}(\bar{X})$ holds (above T), provided that such a finite number N exists. To show that such an N exists, suppose towards a contradiction that all the sets $\varphi_y^n(\bar{X}) \setminus \varphi_y^{n+1}(\bar{X})$ are non-empty (above T). For every n , let $T_n \in \varphi_y^n(\bar{X}) \setminus \varphi_y^{n+1}(\bar{X})$. From the definition of the sequence $\varphi_y^n(\bar{X})$ and Proposition VI.10, it follows (by an easy induction) that all the sets $\varphi_y^n(\bar{X})$ are closed under Σ -bisimilarity (above T). So, when T_n (in $T \uparrow^*$) is eliminated in the move from stage n to stage $n+1$, the whole Σ -bisimilarity class of T_n (above T) is also eliminated. But since there are only finitely many Σ -bisimilarity classes (above T) in Ω^Σ , this elimination process cannot go forever. In fact, an upper bound for the stabilizing stage N is given by the number of bisimilarity classes. \square

Lemma VI.14 (Functional Truth Lemma). *For every formula $\varphi = \varphi(\bar{y}) \in \Sigma$ in which the variables in the string $\bar{y} = (y_1, \dots, y_n)$ are free (or do not occur), every Σ -final theory $T \in \Omega^\Sigma$, and every tuple $\bar{\theta} = (\theta_1, \dots, \theta_n)$ of formulas $\theta_i \in \mathcal{L}_\mu^{P_0}$ s.t. $\widehat{\theta}_i$ is closed under Σ -bisimilarity above T , we have:*

- (1) $T \in \|\varphi(\bar{\theta})\|$ iff $T \in \widehat{\varphi(\bar{\theta})}$;
- (2) if $\varphi = \varphi(\bar{y}) = \varphi(z, \bar{y})$ is positive in z , then for all natural numbers $n \in \mathbb{N}$, we have:
 - $\varphi_z^n(\bar{\theta}) = \widehat{\varphi_z^n(\bar{\theta})}$ holds above T ;
 - moreover, $\varphi_z^n(\bar{\theta})$ is closed under Σ -bisimilarity above T ;

where: $\|\varphi\|$ is the interpretation of φ in the final model Ω^Σ ; $\varphi_z^n(\bar{\theta})$ is an instance of the sequence of sets in Corollary VI.13 (i.e. it is recursively defined by putting $\varphi_z^0(\bar{\theta}) := \Omega^\Sigma$, $\varphi_z^{n+1}(\bar{\theta}) := \|\varphi(\varphi_z^n(\bar{\theta}), \bar{\theta})\|$); and $\widehat{\varphi_z^n(\bar{\theta})}$ is a sequence of formulas, recursively defined by putting $\varphi_z^0(\bar{\theta}) := \top$, $\varphi_z^{n+1}(\bar{\theta}) := \varphi(\varphi_z^n(\bar{\theta}), \bar{\theta})$.

Proof. We prove both assertions (1) and (2) by double induction on the depth $\text{dpt}(T)$ of $T \in \Omega^\Sigma$ and on the subformula-complexity of φ .

Proof of assertion (1):

The base cases $\varphi := y_i$ and $\varphi := x \in P_0$, as well as the Boolean cases $\neg\varphi$ and $\varphi \wedge \varphi'$, are trivial.

Case $\diamond\varphi$. We have the sequence of equivalencies: $T \in \|\diamond\varphi(\bar{\theta})\|$ iff $\exists S \in \|\varphi(\bar{\theta})\|$ s.t. $T \longrightarrow S$ (by the semantic clause for \diamond in the final model) iff $\exists S \in \widehat{\varphi(\bar{\theta})}$ s.t. $T \longrightarrow S$ (by the induction hypothesis for φ) iff $T \in \widehat{\diamond\varphi(\bar{\theta})}$ (by the Final \diamond -Witness Lemma VI.7).

Case $\nu x.\varphi$ with $\varphi = \varphi(x, \bar{y})$. Since T is Σ -final, there exists some $\rho \in \Sigma$ s.t. T is ρ -final, and so (by Lemma VI.8)

T is also $\langle * \rangle \rho$ -final, and moreover $\langle * \rangle \rho$ locally defines T 's cluster above T . Also by Lemma VI.8, all theories $S \in \Omega$ s.t. $T \longleftrightarrow S$ are also $\langle * \rangle \rho$ -final, hence they all belong to Ω^Σ (since $\langle * \rangle \rho$ is provably equivalent to some Σ -formula)¹⁰. For each theory S in the cluster of T (i.e. s.t. either $T \longleftrightarrow S$ or $T = S$), we put

$$\chi_S := \bigwedge \{ \psi : \psi \in S \cap \Sigma \}.$$

Note that, for any theory $T' \in \Omega$, we have $\chi_S \in T'$ iff $T' \cap \Sigma = S \cap \Sigma$. Put

$$\chi := \bigvee \{ \chi_S : S \in \|\nu x.\varphi(\bar{\theta})\|, \text{ and } T \longleftrightarrow^* S \}.$$

Take now the sentence

$$\eta := ((*)\rho \wedge \chi) \vee ([*]\neg\rho \wedge \nu x.\varphi(x, \bar{\theta}))$$

Claim 1: $\|\nu x.\varphi(\bar{\theta})\| = \widehat{\eta}$ holds above T .

Proof of Claim 1: Let $S \in T\uparrow^*$, i.e. s.t. $T \longrightarrow S$. We need to show that: $S \in \|\nu x.\varphi(\bar{\theta})\|$ iff $\eta \in S$. For this, we distinguish two cases.

Case 1: $T \longleftrightarrow S$ or $T = S$. By Lemma VI.8, we have $\langle * \rangle \rho \in S$. Then the desired conclusion follows from the following sequence of equivalencies: $\eta \in S$ iff $\chi \in S$ iff $\exists S' \in \|\nu x.\varphi(\bar{\theta})\|$ s.t. ($S' \longleftrightarrow T$ or $S' = T$) & $S' \cap \Sigma = S \cap \Sigma$ iff $S \in \|\nu x.\varphi(\bar{\theta})\|$ (by Lemma VI.11, Proposition VI.10 and the assumption that all θ_i are closed under Σ -bisimilarity above T).

Case 2: $T \not\longleftrightarrow S$ and $T \neq S$, hence $T \not\longleftrightarrow^* S$, and thus $\text{dpt}(S) < \text{dpt}(T)$. By Lemma VI.8, we have $\langle * \rangle \rho \notin S$, so $[*]\neg\rho \in S$. Once again, the desired conclusion follows from the sequence of equivalencies: $\eta \in S$ iff $\nu x.\varphi(\bar{\theta}) \in S$ iff $S \in \nu x.\varphi(\bar{\theta})$ iff $S \in \|\nu x.\varphi(\bar{\theta})\|$ (by the induction hypothesis for theories $S \in \Omega^\Sigma$ with $\text{dpt}(S) < \text{dpt}(T)$).

Given Claim 1, we can now prove:

Claim 2. $\widehat{\eta}$ is closed under Σ -bisimilarity above T .

Proof of Claim 2: By Claim 1, we have $\widehat{\eta} \cap T\uparrow^* = \|\nu x.\varphi(\bar{\theta})\| \cap T\uparrow^*$, and the right-hand side can be easily seen to be closed under Σ -bisimilarity above T (using Proposition VI.10 and the assumption that all θ_i are closed under Σ -bisimilarity above T). Hence, $\widehat{\eta}$ is also closed under Σ -bisimilarity above T .

Claim 3. $(\eta \Rightarrow \nu x.\varphi(x, \bar{\theta})) \in T$.

Proof of Claim 3: Suppose not. Then by Proposition III.4(4), we must have $[*](\eta \Rightarrow \varphi(\eta, \bar{\theta})) \notin T$. By the Canonical $\langle * \rangle$ -Witness Lemma VI.4, there exists $T' \in \Omega$ (not necessarily final!) such that $T \longrightarrow^* T'$, $\eta \in T'$ and $\varphi(\eta, \bar{\theta}) \notin T'$. Once again, we distinguish two cases.

Case 1: $\langle * \rangle \rho \in T'$. From $T \longrightarrow^* T'$ and the $\langle * \rangle \rho$ -finality of T , we obtain that T' is also $\langle * \rangle \rho$ -final, hence $T' \in \Omega^\Sigma$ and thus $T' \in \widehat{\eta} \cap T\uparrow^*$ (since $\eta \in T' \in \Omega^\Sigma$ and $T \longrightarrow^* T'$). We have the following sequence of equalities:

$$\begin{aligned} \widehat{\eta} \cap T\uparrow^* &= \|\nu x.\varphi(\bar{\theta})\| \cap T\uparrow^* = \|\varphi(\|\nu x.\varphi(\bar{\theta})\|, \bar{\theta})\| \cap T\uparrow^* = \\ &= \|\varphi(\|\nu x.\varphi(\bar{\theta})\| \cap T\uparrow^*, \bar{\theta})\| \cap T\uparrow^* = \|\varphi(\widehat{\eta} \cap T\uparrow^*, \bar{\theta})\| \cap T\uparrow^* = \\ &= \|\varphi(\widehat{\eta}, \bar{\theta})\| \cap T\uparrow^* = \widehat{\varphi(\eta, \bar{\theta})} \cap T\uparrow^* \end{aligned}$$

(where we used repeatedly Claim 1, Lemma V.10, the fact that $\|\nu x.\varphi(\bar{\theta})\|$ is a fixed point of $X \mapsto \|\varphi(X, \bar{\theta})\|$, as well as the induction hypothesis for φ , combined with the fact that all θ_i 's and $\widehat{\eta}$ are closed under Σ -bisimilarity above T). Using the above equalities, we get from $T' \in \widehat{\eta} \cap T\uparrow^*$ to $T' \in \widehat{\varphi(\eta, \bar{\theta})}$, which contradicts the above assumption that $\varphi(\eta, \bar{\theta}) \notin T'$.

Case 2: $[*]\neg\rho \in T'$. From this, together with $\eta \in T'$, we obtain (using reasoning in the axiomatic system) that $([*]\neg\rho \wedge \nu x.\varphi(x, \bar{\theta})) \in T'$. Using the Fixed Point Axiom, we get $([*]\neg\rho \wedge \varphi(\nu x.\varphi(x, \bar{\theta}), \bar{\theta})) \in T'$. Applying Proposition III.4(2) (and the fact that φ is positive in x), we infer that $\varphi([*]\neg\rho \wedge \nu x.\varphi(x, \bar{\theta})) \in T'$, then applying Proposition III.4(3) (as well as the fact that $([*]\neg\rho \wedge \nu x.\varphi(x, \bar{\theta})) \Rightarrow \eta$ is provable in propositional logic, hence by Necessitation $[*](([*]\neg\rho \wedge \nu x.\varphi(x, \bar{\theta})) \Rightarrow \eta)$ is a theorem in our system), we obtain that $\varphi(\eta, \bar{\theta}) \in T'$, which again contradicts the above assumption that $\varphi(\eta, \bar{\theta}) \notin T'$.

Given the above three Claims, let us prove the case $\nu x.\varphi$. For the *left-to-right direction*: assume that $T \in \|\nu x.\varphi(\bar{\theta})\|$. Then by Claim 1, we have $T \in \widehat{\eta}$, hence $\eta \in T$, and thus by Claim 3, we also have $\nu x.\varphi(x, \bar{\theta}) \in T$, i.e. $T \in \widehat{\nu x.\varphi(x, \bar{\theta})}$, as desired.

For the *converse*: assume that $T \in \widehat{\nu x.\varphi(x, \bar{\theta})}$. Using reasoning in the axiomatic system (making essential use of the Fixed Point Axiom), we see that $\nu x.\varphi(x, \bar{\theta}) \Rightarrow \varphi_x^n(\bar{\theta})$ for all n , so we get $T \in \widehat{\varphi_x^n(\bar{\theta})}$ for all n . By the inductive assertion (2) of our Lemma (for φ), we obtain that $T \in \varphi_x^n(\bar{\theta})$ for all n . So $T \in \bigcap_n \varphi_x^n(\bar{\theta})$. But, by Corollary VI.13 (and the fact that all θ_i are closed under Σ -bisimilarity above T), this last set is equal above T with the greatest fixed point of the operator $X \mapsto \|\varphi(X, \bar{\theta})\|$, i.e. with $\|\nu x.\varphi(\bar{\theta})\|$, and so we obtain the desired conclusion.

Proof of assertion (2): We prove assertion (2) of our Lemma for φ , using the fact that we proved assertion (1) of the Lemma for φ . The proof is by induction on n . For $n = 0$: $\varphi_z^0(\bar{\theta}) = W = \varphi_z^0(\bar{\theta})$, and W is obviously closed under Σ -bisimilarity. For the inductive step $n + 1$, assume the assertion is true for n . Then, for all theories $S \in T\uparrow^*$, we have the following sequence of equivalencies:

$$\begin{aligned} S \in \varphi_z^{n+1}(\bar{\theta}) &\text{ iff } S \in \|\varphi(\varphi_z^n(\bar{\theta}), \bar{\theta})\| \text{ iff } S \in \|\varphi(\widehat{\varphi_z^n(\bar{\theta})}, \bar{\theta})\| \\ &\text{ (by the inductive hypothesis (2) for } n) \text{ iff } S \in \varphi(\varphi_z^n(\bar{\theta}), \bar{\theta}) \\ &\text{ (by the induction hypothesis (1) for } \varphi \text{ and } S, \text{ and using the} \\ &\text{ closure of } \widehat{\varphi_z^n(\bar{\theta})} \text{ and of } \varphi_z^n(\bar{\theta}) \text{ under } \Sigma\text{-bisimilarity above } T, \\ &\text{ by the inductive hypothesis (2) for } n, \text{ and as a consequence their} \\ &\text{ closure under } \Sigma\text{-bisimilarity above } S \in T\uparrow^*) \text{ iff } S \in \widehat{\varphi_z^{n+1}(\bar{\theta})}. \end{aligned}$$

That takes care of the first item in assertion (2) of our Lemma. As for the *second item* of this assertion (closure of $\widehat{\varphi_z^n(\bar{\theta})}$ under Σ -bisimilarity above T): first, using Proposition

¹⁰This follows from $\rho \in \Sigma$, together with the fact that Σ is closed under the $\langle * \rangle$ operator up to logical equivalence.

VI.10 and the recursive definition of $\varphi_z^n(\bar{\theta})$, an easy induction on n shows that all $\varphi_z^n(\bar{\theta})$ are closed under Σ -bisimilarity above T ; from this, together with the already proven first item of assertion (2), we conclude that all $\varphi_z^n(\bar{\theta})$ are also closed under Σ -bisimilarity above T . \square

Lemma VI.15 (Truth Lemma). *For every formula $\varphi \in \Sigma$, we have:*

$$\|\varphi\|_{\Omega^\Sigma} = \widehat{\varphi}.$$

Proof. For $\varphi = \varphi(y_1, \dots, y_n) \in \Sigma$, apply the Functional Truth Lemma VI.14 to formula $\theta_i := y_i$. \square

Weak completeness for wK4 frames follows immediately from Lemma VI.15 (cf. Appendix). By Lemma V.12, this also applies to irreflexive wK4 frames, hence to topological derivative spaces, and thus to arbitrary derivative spaces.

Proof of Completeness for wK4 frames, topological derivative spaces, and general derivative spaces:

Recall that we started with a consistent formula φ_0 , and a set Σ s.t. $\varphi_0 \in \Sigma$ and Σ is closed under negations and $\langle * \rangle$ up to logical equivalence. Take some Σ -final theory $T_0 \in \Omega^\Sigma$ with $\varphi_0 \in T_0$ (-such a theory exists by the Lindenbaum Lemma combined with the Final Lemma VI.6). Since $\varphi_0 \in T_0 \in \Omega^\Sigma$, the above Truth Lemma VI.15 shows that $T_0 \models \varphi_0$ holds in Ω^Σ . Hence, our axiomatic system is complete for the class of weakly transitive relational models. By Lemma V.12, we can add irreflexivity: the logic is the same, so the system is also complete for the class of irreflexive and weakly transitive models. But, as already mentioned in Example II.6, this class coincides with the class of Alexandroff topological derivative models. So the system is also complete for topological derivative models (and thus also for general derivative models).

As for *finite model property*, this can be shown by taking the quotient of Ω^Σ modulo Σ -bisimilarity:

Final Quotient The *final quotient* $(\Omega_\Sigma, \longrightarrow_\Sigma, \|\cdot\|_\Sigma)$ is defined as the “strongly extensional Σ -quotient” of the final model Ω^Σ ; i.e. the set of worlds Ω_Σ consists of all equivalence classes $T_\Sigma = \{S \in \Omega^\Sigma : T \simeq_\Sigma S\}$, the accessibility relation is given by putting $T_\Sigma \longrightarrow_\Sigma S_\Sigma$ if there are $T' \in T_\Sigma, S' \in S_\Sigma$ s.t. $T \longrightarrow S$, and the valuation is given by putting, for each $p \in \mathbb{P}$, $T_\Sigma \in \|p\|_\Sigma$ iff there is $T' \in T_\Sigma$ s.t. $T' \in \|p\|$.

Proposition VI.16 (Finite Model Property).

- 1) *The final quotient Ω_Σ is finite (with an upper bound given by a computable function of $|\Sigma|$);*
- 2) *for every formula $\varphi \in \mathcal{L}_\mu^{P_0}$, we have that: φ is true in the final model at some final theory $T \in \Omega^\Sigma$ iff φ is true in the final quotient at the Σ -bisimilarity class T_Σ ;*
- 3) *μ -calculus has FMP (wrt relational, topological and derivative-space semantics).*

Proof. The finite bound follows immediately from Proposition VI.12.

The second part follows from the easily checked fact that the map $T \mapsto T_\Sigma$ is a functional P_0 -bisimulation between the two models, and that all μ -calculus sentences $\varphi \in \mathcal{L}_\mu^{P_0}$ are invariant under P_0 -bisimulations.

To check part 3 (FMP) for weakly transitive models, it is enough to check that the relation \longrightarrow_Σ is weakly transitive in the final quotient Ω_Σ . For this, suppose that $T_\Sigma, S_\Sigma, U_\Sigma \in \Omega_\Sigma$ are s.t. $T_\Sigma \longrightarrow_\Sigma S_\Sigma \longrightarrow_\Sigma U_\Sigma$. By the definition of \longrightarrow_Σ , this means that we can assume $T \longrightarrow S$ and $S' \longrightarrow U$, for some theory $S' \simeq_\Sigma S$. Since Σ -bisimilarity is a bisimulation relation, this implies that there exists some $U' \simeq_\Sigma U$ s.t. $S \longrightarrow U'$. But from $T \longrightarrow S \longrightarrow U'$, we obtained (by weak transitivity) that we have either $T = U'$, in which case $T_\Sigma = U'_\Sigma = U_\Sigma$, or else $T \longrightarrow U'$, in which case $T_\Sigma \longrightarrow_\Sigma U'_\Sigma = U_\Sigma$, as desired.

Using again Lemma V.12 (and the fact that the copying construction in its proof preserves a model’s finiteness), we get FMP for irreflexive and weakly transitive models, i.e. (by the equivalence in Example II.6) for Alexandroff topological derivative models (and hence for arbitrary topological derivative models, as well as general derivative models). \square

Axiomatization of natural sublanguages The proofs in this section can all be carried out within any natural sublanguage \mathcal{L} of \mathcal{L}_μ . In particular, natural sublanguages are closed under all operations used to define Σ , and moreover the formulas considered in the proof are all built from elements of Σ using substitution, Booleans or applications of modalities. This finishes the proof of Theorem III.5, establishing completeness and FMP for all logics μ -wK4 $^\mathcal{L}$.

Next, we get similar results for many logics above μ -wK4.

VII. GENERALIZATION TO COFINAL SUBFRAME LOGICS

Our completeness and finite model property uses only a handful of properties of the logic wK4, and can readily be extended to a wide class of related logics. To be precise, we will now show that FMP holds for any canonical cofinal subframe logic above wK4 enriched with fixed-points.

Definition VII.1. ([10, Ch. 9]) *Let (W, \longrightarrow) be a weakly transitive frame. A subset $X \subseteq W$ is called cofinal if $X \uparrow \subseteq X \downarrow^*$. That is, for every $x \in W$, if there is $y \in X$ such that $y \longrightarrow x$, then there is $z \in X$ with $x \longrightarrow^* z$.*

Let Λ be any normal modal logic over \mathcal{L}_\diamond that extends wK4. Recall that a Kripke frame (W, \longrightarrow) is called a Λ -frame if it validates all the formulas in Λ , and that a modal logic Λ is *canonical* if the underlying frame of the canonical model for Λ is a Λ -frame. Every logic axiomatized by Sahlqvist formulas is canonical [9]. Recall also that a canonical logic Λ is *cofinal subframe* if and only if for every Λ -frame $\mathcal{F} = (W, \longrightarrow)$ and every cofinal $U \subseteq X$, the restriction of \mathcal{F} to U is also a Λ -frame [10].

Examples of canonical cofinal subframe logics above wK4 are wKT $_0$, K4, K4D = K4 + $\diamond \top$, K4.1 = K4 + $\square \diamond p \rightarrow \diamond \square p$, K4.2 = K4 + $\diamond \square p \rightarrow \square \diamond p$, K4.3 = K4 + $\square(\square^+ p \rightarrow q) \vee$

$\Box(\Box^+q \rightarrow p)$, $S4 = K4 + \Box p \rightarrow p$, $S4.1 = S4 + \Box\Diamond p \rightarrow \Diamond\Box p$, $S4.2 = S4 + \Diamond\Box p \rightarrow \Box\Diamond p$, $S4.3 = S4 + \Box(\Box p \rightarrow q) \vee \Box(\Box q \rightarrow p)$, $S5 = S4 + p \rightarrow \Box\Diamond p$, etc. (see [10, Chapter 9]).

We are ready to state Theorem III.5 in full generality.

Theorem VII.2. *Let Λ be a canonical cofinal subframe logic over $wK4$, and \mathcal{L} be a natural sublanguage of \mathcal{L}_μ . Then, $\mu\text{-}\Lambda^\mathcal{L}$ is sound for the class of Λ -frames, and complete for the class of finite Λ -frames.*

Proof. We just follow the proof of the previous section. First, we note that by canonicity, the Kripke frame underlying the canonical model of $\mu\text{-}\Lambda$ is a Λ -frame. This does require some checking, as canonicity only tells us that the canonical model for Λ is a Λ -frame, but this can be done by observing that the canonical model of $\mu\text{-}\Lambda$ is a generated submodel of the canonical model of Λ . We proceed as in the proof of weak completeness in the previous section with a small modification that if \top does not belong to Σ , then we add it to Σ . The fact that $\top \in \Sigma$ ensures that the final model Ω^Σ contains all final point of the canonical frame. Therefore, Ω^Σ is based on a cofinal subframe of the canonical frame. Hence, the underlying frame of Ω^Σ is a Λ -frame. Finally, as p-morphic images preserve the validity of modal μ -formulas, the finite p-morphic image of Ω^Σ is a finite Λ -frame. Thus, every consistent \mathcal{L}_μ -formula is satisfied in a model based on a Λ -frame, and hence, on a $\mu\text{-}\Lambda$ -frame, implying the FMP of $\mu\text{-}\Lambda$. \square

There exist continuum many canonical cofinal subframe logics above $K4$ ([10, Theorem 11.28 and Exercise 11.14]). Hence the above theorem covers continuum many logics.¹¹ Of course, only countable many of them have a recursively enumerable set of axioms: for those logics, decidability follows from Theorem VII.2. Next, we single out some important ones.¹²

Corollary VII.3. *The logics $\mu\text{-}wKT_0$, $\mu\text{-}K4$, $\mu\text{-}K4D$, $\mu\text{-}K4.1$, $\mu\text{-}K4.2$, $\mu\text{-}K4.3$, $\mu\text{-}S4$, $\mu\text{-}S4.1$, $\mu\text{-}S4.2$, $\mu\text{-}S4.3$ have the FMP and are decidable.*

VIII. COMPLETENESS FOR T_0 AND T_D SPACES

The simple world-duplication construction underlying the last step of the topological completeness proof in Section III.5 does not work in the case of T_0 and T_D spaces. So we will use d-morphisms to prove topological completeness for these cases. In the process we give an alternative to the above-mentioned proof of completeness for arbitrary spaces, although this new proof has the disadvantage that it does not yield the finite model property in this setting.

For this it suffices, given a $wK4$ frame $\mathcal{F} = (W, \rightarrow)$, to construct a topological space (X, τ) and a d-morphism $\pi: X \rightarrow W$ as characterized by Lemma V.7, in such a way

¹¹As far as we are aware, this is a first non-trivial example of a completeness result for modal fixed-point logics that covers so many logical systems.

¹²Topological completeness of $wK4$, wKT_0 and $K4$ has already been discussed in Section 3. We also recall that $S4.1$ is the logic of spaces whose dense sets form a filter, that $S4.2$ is the logic of extremally disconnected spaces [6, Sec. 2.6] and that $S4.3$ is the logic of hereditarily extremally disconnected spaces [33].

that if \mathcal{F} is a $wK4T_0$ frame then X will be T_0 , and if \mathcal{F} is a $K4$ frame, then X will be T_D .

Definition VIII.1. *Let $\mathcal{F} = (W, \rightarrow)$ be a $wK4$ frame. We build a topological space $(X, \tau) = (X_\mathcal{F}, \tau_\mathcal{F})$ and a map $\pi: X \rightarrow W$ as follows. Let W^r be the set of reflexive points of W and W^i be the set of irreflexive points. Then, set*

$$X = (W^r \times \mathbb{N}) \cup (W^i \times \{\omega\}),$$

and say that $U \subseteq X$ is open if whenever $(w, \alpha) \in U$, the following two properties are satisfied:

- 1) There is $n \in \mathbb{N}$ such that for all $(v, \beta) \in X$, $v \longleftarrow w$ and $\beta \geq n$ implies that $(v, \beta) \in U$.
- 2) If $(v, \beta) \in X$ and $w \not\rightarrow^* v$ then $(v, \beta) \in U$.

Finally, set $\pi(w, \alpha) = w$.

In other words, if an open set contains (w, α) then it contains all copies of v whenever $w \rightarrow^* v$, except possibly for cofinitely many in the case that $v \longleftarrow w$.

Lemma VIII.2. *If $\mathcal{F} = (W, \rightarrow)$ is any $wK4$ frame then $\tau_\mathcal{F}$ is a topology on $X_\mathcal{F}$ and $\pi: X_\mathcal{F} \rightarrow W$ is a d-morphism.*

Proof. Let $(X, \tau) = (X_\mathcal{F}, \tau_\mathcal{F})$. We omit the proof that τ is a topology, which proceeds by routine verification. To see that π is a d-morphism, we appeal to Lemma V.7. Let $(w, \alpha) \in X$ and note that $w = \pi(w, \alpha)$. Define

$$O = \{(v, \beta) : w \rightarrow^* v \text{ and } \beta \geq 0\}.$$

It should be clear that O is a neighborhood of (w, α) . Moreover, if $(v, \beta) \in O \setminus \{(w, \alpha)\}$, then $w \rightarrow^* v$. It follows that $w \rightarrow v$, except in the case where $w = v$ and w is irreflexive. But then $\alpha = \beta = \omega$, so that $(v, \beta) = (w, \alpha)$, contradicting $(v, \beta) \in O \setminus \{(w, \alpha)\}$. Hence $O \setminus \{(w, \alpha)\} \subseteq \pi^{-1}(w\uparrow)$, as needed.

Now let U be any neighborhood of (w, α) and $v \in w\uparrow$; we must show that $v \in \pi(U \setminus \{(w, \alpha)\})$. Then, there is $n \in \mathbb{N}$ such that $\beta > n$ and $(v, \beta) \in X$ implies that $(v, \beta) \in U$. If v is reflexive, choose any $\beta \neq \alpha$ such that $n < \beta < \omega$. Then, $(v, \beta) \in U \setminus \{(w, \alpha)\}$ and $\pi(v, \beta) = v$, as required. If v is irreflexive, then from $w \rightarrow v$ we obtain $v \neq w$. Thus $(v, \omega) \in U \setminus \{(w, \alpha)\}$ and $\pi(v, \omega) = v$.

Finally, π is surjective since $\pi(w, 0) = w$ if w is reflexive and $\pi(w, \omega) = w$ if w is irreflexive. \square

Lemma VIII.3. *If $\mathcal{F} = (W, \rightarrow)$ is any $wK4$ frame then:*

- 1) If \mathcal{F} is a $wK4T_0$ frame then $X_\mathcal{F}$ is T_0 .
- 2) If \mathcal{F} is a $K4$ frame then $X_\mathcal{F}$ is T_D .

Proof. Let $X = X_\mathcal{F}$. First assume that \mathcal{F} is a $wK4T_0$ frame and let $(w, \alpha) \neq (v, \beta) \in X$. If $w \not\rightarrow^* v$ then $\{(u, \gamma) \in X : w \rightarrow^* u\}$ is a neighborhood of (w, α) which is not a neighborhood of (v, β) . The case where $v \not\rightarrow^* w$ is symmetric, so we may assume $w \longleftrightarrow^* v$. If $\beta < \omega$ then

$$U = \{(u, \gamma) \in X : w \rightarrow^* u\} \setminus \{(v, \beta)\}$$

is a neighborhood of (w, α) not containing (v, β) . The case where $\alpha < \omega$ is symmetric. So we are left with the case where

$\alpha = \beta = \omega$. Since $(w, \alpha) \neq (v, \beta)$, it follows that $w \neq v$. Since $w \longleftrightarrow v$ and \mathcal{F} is a wK4T_0 frame, we cannot have that both w, v are irreflexive; but if w is reflexive then $\alpha < \omega$, contrary to our assumption, and similarly if v is reflexive then $\beta < \omega$. We conclude that the case $\alpha = \beta = \omega$ is impossible.

Now assume that \mathcal{F} is a K4 frame and let $(w, \alpha) \in X$; we must find open U and closed F such that $U \cap F = \{(w, \alpha)\}$. Let

$$U = \{(v, \beta) \in X : w \longrightarrow^* v\}$$

$$F = \{(v, \beta) \in X : v \longrightarrow w\} \cup \{(w, \alpha)\}.$$

It should be clear that $U \cap F = \{(w, \alpha)\}$ and that U is open, so we check only that F is closed; that is, that the complement of F is open. So, define $O := X \setminus F$. Let $(v, \beta) \in O$. If $(u, \gamma) \in X$ is such that $v \longrightarrow u$, then we cannot have $u \longrightarrow^* w$ by transitivity, since this would lead to $v \longrightarrow w$ and $(v, \beta) \in F$. Thus $u \not\longrightarrow w$ and $u \neq w$, yielding $(u, \gamma) \notin F$ regardless of γ . Next we check that there is n so that if $\gamma > n$, $u \longleftarrow^* v$ and $(u, \gamma) \in X$, it follows that $(u, \gamma) \in O$. If $v \not\longleftarrow^* w$, then we may set $n = 0$. For then, $u \longleftarrow^* v$ and $v \not\longrightarrow w$ yield $u \not\longrightarrow w$ and $u \neq w$, so that $(u, \gamma) \in O$ regardless of γ . If $v \longleftarrow^* w$, we claim that w is reflexive. If not, then $v \longleftarrow^* w$ yields $v = w$ since \mathcal{F} is a K4 frame, and the definition of X yields $\alpha = \beta = \omega$, so that $(v, \beta) = (w, \alpha)$ and $(v, \beta) \in O$ is impossible. Thus w is reflexive, so that $\alpha < \omega$. But then, $u \longleftarrow^* v$ and $\gamma > \alpha$ yield $(u, \gamma) \in O$, so we may set $n = \alpha$. \square

We can now proceed to prove topological completeness for T_0 and T_D spaces. In fact, the proof also works for the wK4 case of arbitrary spaces (but unlike the proof in the previous section it does *not* give us finite model property)¹³:

- Theorem VIII.4.** 1) $\mu\text{-wK4}$ is sound and complete for the class of all topological spaces.
 2) $\mu\text{-wK4T}_0$ is sound and complete for the class of all T_0 topological spaces.
 3) $\mu\text{-K4}$ is sound and complete for the class of all T_D topological spaces.

Proof. Let Λ be any of the logics $\mu\text{-wK4}$, $\mu\text{-wK4T}_0$ or $\mu\text{-K4}$. Soundness of Λ for its class of spaces follows from Lemma III.3 and the fact that each fixed point-free fragment is sound for the respective class of spaces (see Section III). Since Λ is a canonical subframe logic, by Kripke completeness, if φ is not derivable then it is falsifiable on some Λ -frame $\mathcal{F} = (W, \longrightarrow)$. Then, by Lemma VIII.3, $(X_{\mathcal{F}}, \tau_{\mathcal{F}})$ is a Λ -space and \mathcal{F} is a d-morphic image of $X_{\mathcal{F}}$, so φ is also falsifiable on $(X_{\mathcal{F}}, \tau_{\mathcal{F}})$, as needed. \square

¹³On the other hand, neither wK4T_0 nor K4 have the finite topological model property, so the next result cannot be improved upon.

IX. CONCLUSION AND COMPARISON WITH OTHER WORK

In this paper, we have studied the μ -calculus over arbitrary topological spaces, as well as some natural subclasses, and obtained a general soundness and completeness result for the standard axiomatization.

Our results are novel for several reasons. First, in the setting of Kripke semantics, neither completeness nor the FMP for weakly transitive frames were known, nor do they follow immediately from known results. Moreover, our completeness proof appears to be the first such result for a variant of μ -calculus that simultaneously applies to infinitely many logics and their respective classes of frames.

From the topological perspective, neither completeness nor decidability for non- T_D spaces were known, nor they follow from known results. Unlike the transitive/ T_D case, our logics do not embed into standard μ -calculus, or any of its known decidable extensions. This is in sharp contrast to the T_D /transitive case, where FMP and decidability follow via a simple encoding into standard μ -calculus.¹⁴ Moreover, we showed that the tangled derivative is not expressively complete over the class of all topological (or even T_0) spaces, so we had to give a completeness proof that applies to the full language of μ -calculus.

But note that even on T_D spaces, our proof is the first to directly establish completeness over such spaces of a Kozen-type axiomatization for full μ -calculus (rather than for some semantically equivalent modal logic). Prior work on T_D spaces, mainly by Goldblatt and Hodkinson [26], had focused only on the tangled fragment. Though this fragment is known to be co-expressive with μ -calculus over T_D spaces (and transitive frames), completeness for the full μ -calculus over these spaces only follows if we combine the results in [26] with Walukiewicz's proof of Kripke completeness for μ -calculus. In contrast, our proof of completeness is self-contained (for both the T_D and the non- T_D case), taking advantage of the weak transitivity to give a streamlined proof tailored for the topological setting.

Furthermore, our results are based on an innovative use of the proof techniques using final submodels (due to Fine and Zakharyashev). This method has not been applied previously in a setting with fixed points, and provides a novel, general and relatively simple approach to dealing with fixed point logics over wK4 frames (for which the filtration method, used in [26] and elsewhere in the study of fixed point logics, does not seem to work). In fact, even for the much easier case of topological *closure spaces*, our method provides a simpler and more uniform way to reprove existing results: while Goldblatt and

¹⁴As already mentioned, the transitive closure of a relation can be encoded in μ -calculus (and thus the decidability of μ -calculus over transitive frames follows immediately from Kozen's result on the decidability of μ -calculus over arbitrary frames). In contrast, the weakly-transitive closure of a relation does *not* seem to be definable in μ -calculus, and not even in its recent hybrid extension [34]. Weakly-transitive closure is definable only if one adds the binding operator from hybrid logic. But this increases the complexity of hybrid μ -calculus, and the resulting logic is no longer known to have FMP (or to even be decidable).

Hodkinson [35] had to do a lot of work to show the FMP for S4-tangle logic (and thus also for the semantically equivalent μ -S4), in Corollary VII.3 we get this result essentially for free from our general methods.

There are many open questions left within the context of topological fixed point logics. The problem of finding a simple but expressively complete fragment of the μ -calculus over wK4, in analogy to the tangled fragment for logics over K4, remains open. But we conjecture that topological μ -calculus does indeed collapse to a simpler natural fragment, possibly the alternation-free fragment, with a proof along the lines of the similar argument for transitive frames in [36]. Anticipating such a development, we have set up our main completeness result in a modular fashion so that, if such a fragment \mathcal{L} is ever found, the completeness for its natural axiomatization μ -wK4 $^{\mathcal{L}}$ will follow immediately from Theorem III.5.

Another line of inquiry that we leave open here is the problem of extending our methods to classes of spaces which enjoy topologically natural properties that do not correspond to any cofinal subframe logic. The prime example here is that of *connected spaces*, whose modal logic is well understood in presence of the universal modality [37]. We believe that our methods can be extended to such settings, but some non-trivial modifications would be required.

REFERENCES

- [1] I. Walukiewicz, “Completeness of Kozen’s axiomatisation of the propositional μ -calculus,” *Inf. Comput.*, vol. 157, no. 1-2, pp. 142–182, 2000.
- [2] L. Santocanale, “Completions of μ -algebras,” *Ann. Pure Appl. Log.*, vol. 154, no. 1, pp. 27–50, 2008.
- [3] L. Santocanale and Y. Venema, “Completeness for flat modal fixpoint logics,” *Ann. Pure Appl. Log.*, vol. 162, no. 1, pp. 55–82, 2010.
- [4] B. Afshari and G. Leigh, “Cut-free completeness for modal μ -calculus,” in *32nd Annual ACM/IEEE Symposium on Logic in Computer Science LICS*. IEEE Press, 2017, pp. 1–12.
- [5] J. C. C. McKinsey and A. Tarski, “The algebra of topology,” *Ann. of Math.*, vol. 45, pp. 141–191, 1944.
- [6] J. v. Benthem and G. Bezhanishvili, “Modal logics of space,” in *Handbook of spatial logics*. Dordrecht: Springer, 2007, pp. 217–298.
- [7] L. Esakia, “Weak transitivity—a restitution,” in *Logical investigations, No. 8 (Russian) (Moscow, 2001)*. Moscow: “Nauka”, 2001, pp. 244–255.
- [8] —, “Intuitionistic logic and modality via topology,” *Annals of Pure and Applied Logic*, vol. 127, no. 1-3, pp. 155–170, 2004, provinces of logic determined.
- [9] P. Blackburn, M. de Rijke, and Y. Venema, *Modal Logic*. Cambridge University Press, 2001.
- [10] A. Chagrov and M. Zakharyashev, *Modal Logic*. New York: The Clarendon Press, 1997.
- [11] L. Esakia, “Diagonal constructions, the Löb formula and rarefied Cantor’s scattered spaces,” in *Studies in logic and semantics*. Tbilisi: “Metsniereba”, 1981, pp. 128–143.
- [12] L. Beklemishev and D. Gabelaia, “Topological interpretations of provability logic,” in *Leo Esakia on duality in modal and intuitionistic logics*. Springer, 2014, pp. 257–290.
- [13] S. Vickers, *Topology via Logic*. Cambridge: Cambridge University Press, 1989.
- [14] S. Abramsky, “Domain theory in logical form,” *Ann. Pure Appl. Logic*, vol. 51, no. 1-2, pp. 1–77, 1991.
- [15] M. Brecht and A. Yamamoto, “Topological properties of concept spaces,” *Information and Computation*, vol. 208, no. 4, pp. 327–340, 2010.
- [16] A. Baltag, N. Gierasimczuk, and S. Smets, “On the solvability of inductive problems: A study in epistemic topology,” in *TARK 2015, Proceedings of the 15th Conference on Theoretical Aspects of Rationality and Knowledge*. ENTCS, 2015, pp. 81–98.
- [17] D. K. M. Herlihy and S. Rajsbaum, *Distributed Computing Through Combinatorial Topology*. Elsevier-Morgan Kaufman, 2013.
- [18] J. L. E. Goubault and S. Rausbaum, “A simplicial complex model for dynamic epistemic logic to study distributed task computability,” *Information and Computation*, 2020.
- [19] A. Özgün, “Evidence in epistemic logic : A topological perspective,” Ph.D. dissertation, ILLC, Univ. of Amsterdam and Univ. of Lorraine, 2017.
- [20] A. Baltag, N. Bezhanishvili, A. Özgün, and S. Smets, “Justified belief and the topology of evidence,” in *Logic, Language, Information, and Computation - 23rd International Workshop, WoLLIC 2016*, ser. Lecture Notes in Computer Science, J. A. Väänänen, Å. Hirvonen, and R. J. G. B. de Queiroz, Eds., vol. 9803. Springer, 2016, pp. 83–103.
- [21] R. Parikh, “Finite and infinite dialogues,” in *Logic from Computer Science*. Springer, 1992, pp. 481–497.
- [22] K. Fine, “Logics containing K4. II,” *J. Symb. Log.*, vol. 50, no. 3, pp. 619–651, 1985.
- [23] M. Zakharyashev, “Canonical formulas for K4. II. Cofinal subframe logics,” *J. Symb. Log.*, vol. 61, no. 2, pp. 421–449, 1996.
- [24] G. Bezhanishvili, S. Ghilardi, and M. Jibladze, “An algebraic approach to subframe logics. Modal case,” *Notre Dame J. Formal Log.*, vol. 52, no. 2, pp. 187–202, 2011.
- [25] D. Kozen, “A finite model theorem for the propositional μ -calculus,” *Studia Logica*, vol. 47, pp. 233–241, 1988.
- [26] R. Goldblatt and I. Hodkinson, “The finite model property for logics with the tangle modality,” *Stud Logica*, vol. 106, no. 1, pp. 131–166, 2018.
- [27] A. Dawar and M. Otto, “Modal characterisation theorems over special classes of frames,” *Ann. Pure Appl. Log.*, vol. 161, no. 1, pp. 1–42, 2009.
- [28] E. Pacuit, *Neighborhood semantics for modal logic*. Springer, 2017.
- [29] D. Fernández-Duque, “Tangled modal logic for spatial reasoning,” in *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, T. Walsh, Ed. IJCAI/AAAI, 2011, pp. 857–862.
- [30] G. Bezhanishvili, L. Esakia, and D. Gabelaia, “Spectral and T_0 -spaces in d-semantics,” in *8th International Tbilisi Symposium on Logic, Language, and Computation. Selected papers.*, ser. Lecture Notes in Artificial Intelligence, N. Bezhanishvili, S. Löbner, K. Schwabe, and L. Spada, Eds. Springer, 2011, pp. 16–29.
- [31] Y. Venema, “Lectures on the modal μ -calculus,” 2008.
- [32] H. H. Hansen, “Monotonic modal logics,” 2003, master’s Thesis, Available as: ILLC report: PP-2003-24.
- [33] G. Bezhanishvili, N. Bezhanishvili, J. Lucero-Bryan, and J. van Mill, “S4.3 and hereditarily extremally disconnected spaces,” *Georgian Mathematical Journal*, vol. 22, no. 4, pp. 469–475, 2015.
- [34] U. Sattler and M. Vardi, “The hybrid μ -calculus,” in *Automated Reasoning, First International Joint Conference, IJCAR 2001, Siena, Italy, June 18-23, 2001, Proceedings*, ser. Lecture Notes in Computer Science, R. Goré, A. Leitsch, and T. Nipkow, Eds., vol. 2083. Springer, 2001, pp. 76–91.
- [35] R. Goldblatt and I. Hodkinson, “Spatial logic of tangled closure operators and modal μ -calculus,” *Ann. Pure Appl. Log.*, vol. 168, no. 5, pp. 1032–1090, 2017.
- [36] G. d’Agostino and G. Lenzi, “On the μ -calculus over transitive and finite transitive frames,” *Theoretical Computer Science*, vol. 411, pp. 4273–4290, 2010.
- [37] V. Shehtman, ““Everywhere” and “here,”” *J. Appl. Non-Classical Logics*, vol. 9, no. 2-3, pp. 369–379, 1999, issue in memory of George Gargov.