

Uniform Interpolation in Coalgebraic Modal Logic*

Fatemeh Seifan¹, Lutz Schröder¹, and Dirk Pattinson²

1 Friedrich-Alexander-Universität Erlangen-Nürnberg

2 Australian National University

Abstract

A logic has *uniform interpolation* if its formulas can be projected down to given subsignatures, preserving all logical consequences that do not mention the removed symbols; the weaker property of (*Craig*) *interpolation* allows the projected formula – the *interpolant* – to be different for each logical consequence of the original formula. These properties are of importance, e.g., in the modularization of logical theories. We study interpolation in the context of coalgebraic modal logics, i.e. modal logics axiomatized in rank 1, restricting for clarity to the case with finitely many modalities. Examples of such logics include the modal logics K and KD , neighbourhood logic and its monotone variant, finite-monoid-weighted logics, and coalition logic. We introduce a notion of one-step (uniform) interpolation, which refers only to a restricted logic without nesting of modalities, and show that a coalgebraic modal logic has uniform interpolation if it has one-step interpolation. Moreover, we identify preservation of finite surjective weak pullbacks as a sufficient, and in the monotone case necessary, condition for one-step interpolation. We thus prove or reprove uniform interpolation for most of the examples listed above.

1998 ACM Subject Classification F.4.1 Mathematical Logic – Modal logic; I.2.4 Knowledge Representation Formalisms and Methods – Modal logic; I.2.3 Deduction and Theorem Proving

Keywords and phrases Coalgebraic modal logic, uniform interpolation and weak pullback.

Digital Object Identifier 10.4230/LIPIcs.CALCO.2017.21

1 Introduction

Given a logic with a notion of formula and signature (and featuring implication for simplicity), the *Craig interpolation* property requires that every valid implication $\phi \rightarrow \psi$ has an *interpolant*, i.e. a formula ρ mentioning only the signature symbols that occur in both ϕ and ψ , such that both $\phi \rightarrow \rho$ and $\rho \rightarrow \psi$ are valid. The stricter *uniform interpolation* property additionally demands that ρ can be made to depend only on ϕ and on the signature of ψ (or, yet stricter, on the shared symbols of ϕ and ψ), rather than also on ψ itself. Both Craig interpolation and uniform interpolation are useful in the structuring and modularization of logical theories for purposes of specification and automated deduction, e.g. in large ontologies [40, 19]. Craig interpolation was originally proved for first-order logic [5] and later extended to many other systems, notably various modal logics including the basic modal logic K [8], as well as intuitionistic logic [9] and the μ -calculus [6]. Uniform interpolation is easily seen to hold for classical propositional logic but in fact fails for first-order predicate logic [16]. Intuitionistic logic [30], the basic modal logic K [10, 39], and the modal μ -calculus [17] do have uniform interpolation, while it fails for the modal logics $S4$ [11] and $K4$ [2].

In this paper, we study interpolation and uniform interpolation in the context of predicate-lifting style coalgebraic modal logic [26, 34], equivalently, of modal logics that are axiomatized

* Work by the first and second author forms part of DFG project GenMod3 (SCHR 1118/5-3)



licensed under Creative Commons License CC-BY

7th Conference on Algebra and Coalgebra in Computer Science (CALCO 2017).

Editors: Filippo Bonchi and Barbara König; Article No. 21; pp. 21:1–21:22



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

by *rank-1* axioms [33, 37]. Coalgebraic modal logic is a generic framework for modal logics whose semantics goes beyond the standard relational world, and e.g. includes probabilistic, game-based, neighbourhood-based, or weighted behaviour. It is parametrized over the choice of a *type functor* (in our setting, on the category of sets), whose coalgebras play the role of models. The name of the game in coalgebraic logic is to reduce properties of the full modal logic to properties of the *one-step* logic, which restricts to formulas with exactly one layer of modalities and is interpreted over very simple structures that essentially capture the collection of successors of a single state in a model. Following this paradigm, we identify a notion of one-step interpolation, and then establish that for a coalgebraic modal logic \mathcal{L} with finitely many modalities, the following properties imply each other in sequence

1. the modalities are *separating*, i.e. support a Hennessy-Milner-style expressivity theorem [26, 34] (implying that the type functor preserves finite sets), and the type functor preserves surjective finite weak pullbacks (which for finitary functors just means that the functor preserves surjective weak pullbacks);
2. \mathcal{L} has one-step interpolation;
3. \mathcal{L} has uniform interpolation.

Here a pullback is called surjective if it consists of surjective maps. If the modalities of \mathcal{L} are separating and monotone, then preservation of finite surjective weak pullbacks is in fact necessary for one-step interpolation.

As applications of this result, we obtain that neighbourhood logic (i.e. classical modal logic [3]), monotone modal (neighbourhood) logic [3], the relational modal logics K and KD , coalition logic [29], and logics of monoid-weighted transition systems for finite refinable commutative monoids (in particular for finite Abelian groups, even though the latter fail to admit monotone modalities) have uniform interpolation; for neighbourhood logic, coalition logic, and monoid-weighted logics, these results appear to be new.

Proofs are often omitted or only sketched; full proofs are in the appendix.

Related Work

Craig interpolation for monotone modal logic was first proved by Hansen and Kupke [14] and later improved to uniform interpolation by Santocanale and Venema [32]. Craig interpolation (but not uniform interpolation) for coalition logic was proved by Schröder and Pattinson using coalgebraic cutfree sequent systems [28]. Hansen, Kupke, and Pacuit [15] have proved Craig interpolation (but not uniform interpolation) for neighbourhood logic, using semantic methods. Uniform interpolation for coalgebraic modal logic with a generalized Moss modality based on a *quasifunctorial lax lifting* has been shown, for functors preserving finite sets, by Marti in his MSc thesis [20] (and in fact this result has been extended to coalgebraic modal fixpoint logics [21]). Logics based on diagonal-preserving lax liftings (even without assuming quasi-functoriality) satisfy an obvious variant of separation and thus support a generalized Hennessy-Milner theorem, and moreover can be translated into the language of monotone predicate liftings [20]. We leave it as an open problem to determine the relationship between quasifunctoriality and preservation of surjective weak pullbacks in presence of a separating set of monotone predicate liftings. We emphasize that our criteria for interpolation apply also to logics that fail to be separating or admit monotone modalities, hence cannot be phrased in terms of quasifunctorial lax liftings, notable examples of this type being coalition logic, neighbourhood logic, and logics of finite-Abelian-group-weighted transition systems.

In [27], the (coalgebraic) logic of exact covers was introduced; besides a generic Hennessy-Milner theorem and results on completeness and small models, a generic uniform interpolation theorem was claimed which implies that every rank-1 modal logic with finitely many (not

necessarily monotone) modalities has uniform interpolation. We show by means of a counterexample that the latter claim is incorrect; our results help delineate in which cases it can be salvaged.

2 Preliminaries

We assume basic familiarity with category theory (see [1] for an introduction). Throughout, we work over the category \mathbf{Set} of sets and functions as the base category. Given a functor $F : \mathbf{Set} \rightarrow \mathbf{Set}$, an F -coalgebra is a pair $\mathbb{X} = (X, \xi)$ consisting of a set X (of states) and a function $\xi : X \rightarrow FX$. In the spirit of coalgebraic logic, we use such coalgebras as generic models of modal logics; e.g. Kripke frames can be seen as a coalgebras $\xi : X \rightarrow \mathcal{P}X$ for the powerset functor \mathcal{P} , as they assign to each state $x \in X$ a set $\xi(x) \in \mathcal{P}(X)$ of successor states. We will later see non-relational examples. We denote by \mathcal{Q} the contravariant powerset functor, which acts on sets by taking powersets and on maps by taking preimage maps ($\mathcal{Q}f(A) = f^{-1}[A]$).

2.1 Set Functors and Weak Pullbacks

The pullback of a cospan $(f, g) = X \xrightarrow{f} Z \xleftarrow{g} Y$ in \mathbf{Set} is described as

$$(\mathbf{pb}(f, g), \pi_1, \pi_2) \text{ where } \mathbf{pb}(f, g) = \{(x, y) \in X \times Y \mid f(x) = g(y)\}$$

and π_1, π_2 are the projections to the components.

An important property of set functors in the analysis of coalgebras is weak pullback preservation [31]. A set functor F *preserves weak pullbacks*, or *weakly preserves pullbacks* if it maps pullbacks to *weak pullbacks* (equivalently maps weak pullbacks to weak pullbacks), where a weak pullback is defined categorically like a pullback but requiring only existence, not uniqueness, of mediating morphisms. In an element-wise formulation, $X \xleftarrow{\pi_1} P \xrightarrow{\pi_2} Y$ is a weak pullback of $X \xrightarrow{f} Z \xleftarrow{g} Y$ if whenever $f(x) = g(y)$ for $x \in X, y \in Y$, then there exists $p \in P$ (not necessarily unique) such that $\pi_1(p) = x, \pi_2(p) = y$. It well-known and easy to see that the identity functor, all constant functors, and the powerset functor preserve weak pullbacks and that the class of weak-pullback preserving functors is closed under products, coproducts, exponentiation with constants, functor composition, and taking finitary parts [31, 38]. In particular, all generalized Kripke polynomial functors (built from powerset, finite powerset, constant functors, and identity by taking products, coproducts, and exponentiation with constants) preserve weak pullbacks. Some negative examples are as follows.

► **Example 2.1.** 1. The *Neighbourhood functor* or double contravariant powerset functor $\mathcal{N} = \mathcal{Q}\mathcal{Q}$ maps a set X to $\mathcal{N}X = \mathcal{Q}\mathcal{Q}X$ and a function $f : X \rightarrow Y$ to $\mathcal{N}f(\alpha) = \{A \subseteq Y \mid f^{-1}[A] \in \alpha\}$. This functor does not preserve weak pullbacks [31].

2. A variant of the neighbourhood functor is the *monotone neighbourhood functor* denoted by \mathcal{M} . Given an element $\alpha \in \mathcal{Q}\mathcal{Q}X$ we define $\mathbf{Up}(\alpha)$ to be the set

$$\mathbf{Up}(\alpha) := \{Y \subseteq X \mid Y \supseteq Z \text{ for some } Z \in \alpha\},$$

and we say that α is *upwards closed* if $\alpha = \mathbf{Up}(\alpha)$. The functor \mathcal{M} is then given on sets X by $\mathcal{M}X = \{\alpha \in \mathcal{Q}\mathcal{Q}(X) \mid \alpha \text{ upwards closed}\}$ and on maps $f : X \rightarrow Y$ by $\mathcal{M}f(\alpha) = \{A \subseteq Y \mid f^{-1}[A] \in \alpha\} = \mathcal{Q}\mathcal{Q}(f)(\alpha)$. Like \mathcal{N} , \mathcal{M} does not preserve weak pullbacks [14].

3. Another functor that does not preserve weak pullbacks is F_2^3 , defined as a subfunctor of the cubing functor $X \mapsto X^3$ by $F_2^3X = \{(x_1, x_2, x_3) \in X^3 \mid |\{x_1, x_2, x_3\}| \leq 2\}$ [12].

2.2 Coalgebraic Modal Logic

We briefly recall the syntax and semantics of coalgebraic modal logic.

We fix a countable set V of *propositional variables*. The syntax of a coalgebraic modal logic $\mathcal{L}(\Lambda)$ is then determined by the choice a *modal signature* Λ consisting of modal operators with assigned arities: the set $\mathcal{F}(\Lambda)$ of (*modal*) Λ -*formulas* is defined by the grammar

$$\mathcal{L}(\Lambda) \ni \phi, \psi ::= v \mid \perp \mid \neg\phi \mid \phi \wedge \psi \mid \heartsuit(\phi_1, \dots, \phi_n),$$

where $v \in V$ and $\heartsuit \in \Lambda$ is an n -ary modality (we deviate slightly from usual practice in coalgebraic modal logic by including propositional variables in the syntax rather than regarding them as nullary modalities; this is in order to facilitate the definition of interpolation). Other Boolean operators ($\top, \vee, \rightarrow, \leftrightarrow$) are defined in the standard way. We write $\text{rk}(\phi)$ for the *rank* of ϕ , i.e. the maximal nesting depth of modal operators in ϕ .

As indicated above, the type of systems underlying the semantics of $\mathcal{L}(\Lambda)$ is then determined by the choice of a set functor F whose coalgebras play the role of frames. The interpretation of the modal operators is then defined in terms of predicate liftings for F :

► **Definition 2.2** (Predicate liftings). An n -ary *predicate lifting* for F is a natural transformation $\lambda : \mathcal{Q}(-)^n \rightarrow \mathcal{Q} \circ F$, where $\mathcal{Q}(-)^n$ denotes the n -fold product of \mathcal{Q} with itself. We say that λ is *monotone* if $\lambda_X(Y_1, \dots, Y_n) \subseteq \lambda_X(Z_1, \dots, Z_n)$, whenever $Y_i \subseteq Z_i \subseteq X$ for each i . Equivalently, we can describe λ by its *transpose* $\lambda^b : F \rightarrow \mathcal{Q}\mathcal{Q}^n$ given by $\lambda_X^b(t) = \{(Y_1, \dots, Y_n) \in \mathcal{Q}\mathcal{Q}^n X \mid t \in \lambda_X(Y_1, \dots, Y_n)\}$.

By the Yoneda lemma, we have the following equivalent description of predicate liftings [34].

► **Fact 2.3.** *The n -ary predicate liftings for F are in one-to-one correspondence with subsets of $F(2^n)$, where $2 = \{\top, \perp\}$; for $n = 1$, such a subset U determines a predicate lifting λ by $\lambda_X(A) = \{t \in FX \mid T\chi_A(t) \in U\}$ where $\chi_A : X \rightarrow 2$ is the characteristic map of $A \subseteq X$.*

We then complete the semantic parametrization of $\mathcal{L}(\Lambda)$ by assigning to each n -ary modal operator $\heartsuit \in \Lambda$ an n -ary predicate lifting $\llbracket \heartsuit \rrbracket$ for F . For readability, we mostly restrict the technical exposition to unary modalities from now on; the extension to finitary modalities is just a matter of adding indices.

► **Definition 2.4.** An F -*model* (X, ξ, τ) consists of an F -coalgebra $\mathbb{X} = (X, \xi)$ and a valuation $\tau : V \rightarrow \mathcal{P}(X)$ of the propositional variables. We then inductively define a satisfaction relation \models between states of the model (X, ξ, τ) and formulas of $\mathcal{L}(\Lambda)$ by $x \models v$ iff $x \in \tau(v)$, standard clauses for Boolean connectives, and

$$x \models \heartsuit(\phi_1, \dots, \phi_n) \quad \text{iff} \quad \xi(x) \in \heartsuit_X(\llbracket \phi_1 \rrbracket, \dots, \llbracket \phi_n \rrbracket),$$

where $\llbracket \phi_i \rrbracket = \{t \mid t \models \phi_i\}$. As usual, we say that a formula ϕ is *satisfiable* if there exists a state x in some model such that $x \models \phi$, and *valid* if $x \models \phi$ for every state x in every F -model.

► **Example 2.5. 1.** The modal logic K is captured coalgebraically by taking the powerset functor \mathcal{P} as the type functor, $\Lambda = \{\diamond\}$, and

$$\llbracket \diamond \rrbracket_X(Y) = \{A \in \mathcal{P}X \mid A \cap Y \neq \emptyset\}.$$

2. Neighbourhood logic (or *classical modal logic* [3]) has $\Lambda = \{\square\}$, interpreted over the neighbourhood functor \mathcal{N} (Example 2.1.1) by

$$\llbracket \square \rrbracket_X(Y) := \{\alpha \in \mathcal{N}X \mid Y \in \alpha\}.$$

Monotone modal (neighbourhood) logic is captured in the same way, replacing \mathcal{N} with the monotone neighbourhood functor \mathcal{M} (Example 2.1.2).

We fix the data F , Λ , $\llbracket \heartsuit \rrbracket$ of the logic $\mathcal{L}(\Lambda)$ from now on.

The One-Step Logic Given any set Z , we denote by $\text{Prop}(Z)$ the set of propositional formulas over Z :

$$\text{Prop}(Z) \ni \phi ::= \perp \mid z \mid \neg\phi \mid \phi \wedge \phi \quad (z \in Z),$$

and write $\Lambda(Z)$ for the set of formulas $\heartsuit(z_1, \dots, z_n)$ where $\heartsuit \in \Lambda$ has arity n and $z_1, \dots, z_n \in Z$. We then define a *one-step formula over Z* to be an element of $\text{Prop}(\Lambda(\text{Prop}(Z)))$. Here, Z will often be a subset of V ; also, Z will sometimes be a subset of some powerset $\mathcal{P}X$, in which case we will understand every element of $\mathcal{P}X$ to be interpreted as itself. In general, we interpret both propositional formulas and one-step formulas over Z w.r.t. $\mathcal{P}(X)$ -valuations $\tau : Z \rightarrow \mathcal{P}(X)$ for some set X : We extend τ to propositional formulas using the Boolean algebra structure of $\mathcal{P}X$, obtaining for $\phi \in \text{Prop}(Z)$ a subset

$$\phi\tau \in \mathcal{P}X.$$

We write $X \models \phi\tau$ if $\phi\tau = X$. We then define the extension

$$\psi\tau \in \mathcal{P}(FX)$$

of a one-step formula $\psi \in \text{Prop}(\Lambda(\text{Prop}(Z)))$ recursively by the evident clauses for Boolean connectives, and

$$(\heartsuit\phi)\tau = \llbracket \heartsuit \rrbracket_X(\phi\tau).$$

When $Z \subseteq \mathcal{P}(X)$ and τ is just subset inclusion, then we omit τ from the notation, so $\psi \in \text{Prop}(\Lambda(\mathcal{P}X))$ denotes both a one-step formula and its interpretation in $\mathcal{P}(FX)$.

► **Definition 2.6.** A one-step formula $\psi \in \text{Prop}(\Lambda(\text{Prop}(Z)))$ is *(one-step) satisfiable over $\tau : Z \rightarrow \mathcal{P}(X)$* if $\psi\tau \neq \emptyset$, and *(one-step) satisfiable* if ψ is one-step satisfiable over τ for some τ . Dually, ψ is *(one-step) valid (over τ)* if $\neg\psi$ is (one-step) unsatisfiable (over τ). We write $FX, \tau \models \psi$ if ψ is one-step valid over τ , and $\models \psi$ if ψ is one-step valid.

We will need the following pieces of terminology and notation:

► **Definition 2.7.** For a map $f : X \rightarrow Y$, we write σ_f for the *substitution* mapping $A \in \mathcal{P}(X)$ to $f[A]$ (e.g. in one-step formulas of type $\text{Prop}(\Lambda(\mathcal{P}(X)))$), and σ_{f^-} for the substitution mapping $B \in \mathcal{P}(Y)$ to $f^{-1}[B]$. A set $A \in \mathcal{P}(X)$ is *f-invariant* if $f^{-1}[f[A]] = A$.

Clearly, all sets of the form $f^{-1}[B]$ are *f-invariant*, i.e. the *f-invariant* sets are precisely those of the form $f^{-1}[B]$. The *f-invariant* sets form a Boolean subalgebra of $\mathcal{P}(X)$.

► **Definition 2.8.** We denote by $\mathcal{S}(\mathfrak{A})$ the set of *atoms* of a finite Boolean algebra \mathfrak{A} , i.e. its minimal non-bottom elements, and $\text{can}_{\mathfrak{A}}$ for the canonical isomorphism $\mathfrak{A} \rightarrow \mathcal{P}(\mathcal{S}(\mathfrak{A}))$. Given a subalgebra \mathfrak{A}_0 of \mathfrak{A} , we have a *canonical projection* $\mathcal{S}(\mathfrak{A}) \rightarrow \mathcal{S}(\mathfrak{A}_0)$.

The following lemmas are straightforward consequences of naturality of predicate liftings:

► **Lemma 2.9.** *Given a finite Boolean subalgebra \mathfrak{A} of $\mathcal{P}X$ for a set X , $\phi \in \text{Prop}(\Lambda(\mathfrak{A}))$ is satisfiable (valid) iff $\phi\text{can}_{\mathfrak{A}}$ is satisfiable (valid).*

► **Lemma 2.10.** *Let $\mathfrak{A}_0 \subseteq \mathfrak{A}_1$ be finite Boolean subalgebras of $\mathcal{P}X$ for a set X , let $f : \mathcal{S}(\mathfrak{A}_1) \rightarrow \mathcal{S}(\mathfrak{A}_0)$ be the canonical projection, let $\phi \in \text{Prop}(\Lambda(\mathfrak{A}_0))$, and let $t \in F(\mathcal{S}(\mathfrak{A}_1))$. Then $t \in \phi\text{can}_{\mathfrak{A}_1}$ iff $Ff(t) \in \phi\text{can}_{\mathfrak{A}_0}$.*

Separation and Maximally Satisfiable Sets The key condition ensuring that $\mathcal{L}(\Lambda)$ satisfies the Hennessy-Milner property, i.e. distinguishes non-bisimilar states, is *separation* [26, 34]:

► **Definition 2.11** (Separation). We say that Λ is *separating* if for each set X , the family of maps $(\llbracket \heartsuit \rrbracket_X^b : FX \rightarrow \mathcal{Q}\mathcal{Q}X = \mathcal{N}X)_{\heartsuit \in \Lambda}$ is jointly injective.

We proceed to define the *MSS-functor* (for *maximally one-step satisfiable sets*) from F and Λ . (A related functor using *maximally one-step consistent* sets has been used to show that every rank-1 modal logic has a coalgebraic semantics [37].)

► **Definition 2.12.** A set $\Phi \subseteq \text{Prop}(\Lambda(\mathcal{P}(X)))$ is *one-step satisfiable* if the intersection of the interpretations of the formulas in Φ is non-empty, and *maximally one-step satisfiable* if Φ is maximal among such sets. The *MSS-functor* M_F^Λ is given by $M_F^\Lambda X$ being the set of maximally one-step satisfiable subsets of $\text{Prop}(\Lambda(\mathcal{P}X))$, and $M_F^\Lambda f(\Phi) = \{\phi \in \text{Prop}(\Lambda(\mathcal{P}(Y))) \mid \phi \sigma_f \in \Phi\}$ for $f : X \rightarrow Y$.

The following lemma allows us to identify F with its MSS-functor whenever Λ is separating.

► **Lemma 2.13.** *If Λ is separating, then F and M_F^Λ are isomorphic.*

3 Surjective Weak Pullbacks

We proceed to introduce the key semantic interpolation criterion, preservation of *surjective weak pullbacks*. We record explicitly:

► **Definition 3.1.** A pullback of a cospan (f, g) of maps (in Set) is *surjective* if both f and g are surjective, and *finite* if all involved sets are finite. A functor *preserves (finite) surjective weak pullbacks* if it maps (finite) surjective pullbacks to weak pullbacks.

Recall that under AC, every set functor preserves surjective maps. Also, surjective maps are stable under pullbacks, so all morphisms in a surjective pullback are surjective. non-empty binary Cartesian products $X \times Y$ are surjective pullbacks of $X \rightarrow 1 \leftarrow Y$. Moreover, the kernel pair of a map $f : X \rightarrow Y$ is a surjective pullback of the codomain restriction $X \rightarrow f[X]$.

For finitary functors, the finiteness restriction in the preservation condition is immaterial:

► **Lemma 3.2.** *If F is finitary, then F preserves (surjective) weak pullbacks iff F preserves finite (surjective) weak pullbacks.*

Of course, every functor that preserves weak pullbacks also preserves surjective weak pullbacks, e.g. the (finite or unrestricted) powerset functor, and more generally all Kripke polynomial functors. Two negative examples are as follows.

► **Example 3.3. 1.** The neighbourhood functor \mathcal{N} fails to preserve finite surjective weak pullbacks. To see this, consider the pullback of the following functions as in [31]. Let $X = \{a_1, a_2, a_3\}$, $Y = \{b_1, b_2, b_3\}$ and $Z = \{c_1, c_2\}$ and define surjective maps $f : X \rightarrow Z$ and $g : Y \rightarrow Z$ as follows: $f(a_1) = f(a_2) = c_1$, $f(a_3) = c_2$, $g(b_1) = c_1$ and $g(b_2) = g(b_3) = c_2$.

2. The functor F_2^3 fails to preserve finite surjective weak pullbacks. For a counterexample consider a surjective cospan (f, g) with $f = g$ being the constant map $\{a, b\} \rightarrow \{b\}$. For $u = (b, b, a)$ and $v = (a, b, b)$, it is impossible to find a $w \in F_2^3 \text{pb}(f, g)$ such that $F_2^3 \pi_1(w) = u$ and $F_2^3 \pi_2(w) = v$.

We proceed to see examples that fail to preserve weak pullbacks but do preserve surjective weak pullbacks.

The Monotone Neighbourhood Functor

The monotone neighbourhood functor \mathcal{M} does not preserve all weak pullbacks (Example 2.1.2). However:

► **Proposition 3.4.** *The monotone neighbourhood functor \mathcal{M} preserves surjective weak pullbacks.*

The proof is facilitated by the following fact:

► **Lemma and Definition 3.5** (Compatibility). *Let*

$$\begin{array}{ccc} P & \xrightarrow{\pi_1} & X \\ \pi_2 \downarrow & & \downarrow f \\ Y & \xrightarrow{g} & Z \end{array} \quad (1)$$

be a surjective pullback, and let $\alpha_1 \in \mathcal{M}X$, $\alpha_2 \in \mathcal{M}Y$. Then $\mathcal{M}f(\alpha_1) = \mathcal{M}g(\alpha_2)$ iff α_1 and α_2 are compatible, i.e. for every $U \in \alpha_1$ we have $\pi_2[\pi_1^{-1}[U]] \in \alpha_2$ and symmetrically.

Proof (Proposition 3.4, Sketch). Given a surjective pullback (1) and compatible $\alpha_1 \in \mathcal{M}X$, $\alpha_2 \in \mathcal{M}Y$, it is straightforward to show that

$$\beta = \text{Up}(\{\pi_1^{-1}[U] \mid U \in \alpha_1\} \cup \{\pi_2^{-1}[V] \mid V \in \alpha_2\}) \in \mathcal{M}P$$

satisfies $\mathcal{M}\pi_1(\beta) = \alpha_1$ and $\mathcal{M}\pi_2(\beta) = \alpha_2$. ◀

Monoid-weighted Functors

Given a commutative monoid M (which we write additively), the *monoid-weighted functor* S_M is defined by taking $S_M X$ to be the set of finitely supported functions $X \rightarrow M$ (i.e. functions that vanish almost everywhere), and $S_M f(\mu) = \lambda y. \sum_{f(x)=y} \mu(x)$ for $f : X \rightarrow Y$ and $\mu \in S_M X$. Examples of monoid-weighted functors include the free Abelian groups functor ($M = \mathbb{Z}$), the free vector space functor ($M = \mathbb{R}$), the finite multiset functor ($M = \mathbb{N}$), and the finite powerset functor ($M = 2 = \{\perp, \top\}$ with $+$ being disjunction).

► **Definition 3.6** (Refinability). [13] A commutative monoid M is *refinable* if whenever $\sum_{i=1}^n a_i = \sum_{j=1}^k b_j$ for $a_1, \dots, a_n, b_1, \dots, b_k \in M$, $n, k \geq 1$, then there exists an $n \times k$ -matrix over M with row sums a_i and column sums b_j .

As shown by Gumm and Schröder [13], S_M preserves weak kernel pairs iff M is refinable. In fact, refinability already ensures preservation of all weak surjective pullbacks:

► **Lemma 3.7.** *The functor S_M preserves weak surjective pullbacks iff M is refinable.*

Given that a) weak pullback preserving finitary functors are known to admit separating sets of monotone predicate liftings [18], and b) the monotone neighbourhood functor itself preserves surjective weak pullbacks but not all weak pullbacks, it is tempting to conjecture that preservation of surjective weak pullbacks is already sufficient for existence of a separating set of monotone predicate liftings. This is not true, however:

► **Definition 3.8.** A commutative monoid M is *positive* if for $a, b \in M$, $a + b = 0$ implies $a = b = 0$.

► **Proposition 3.9.** *Let M be refinable. Then S_M has a separating set of monotone predicate liftings iff M is positive.*

That is, every commutative monoid that is refinable but not positive gives rise to a monoid-weighted functor that preserves surjective weak pullbacks but does not admit a separating set of monotone predicate liftings. One class of such commutative monoids are the non-trivial Abelian groups: they clearly fail to be positive, and are easily seen to be refinable [13].

4 One-Step Interpolation

We proceed to develop our notion of one-step interpolation, and its relationship to preservation of surjective weak pullbacks. From here on, we assume throughout that *the modal signature Λ is finite*.

► **Definition 4.1.** Two Boolean subalgebras $\mathfrak{A}_1, \mathfrak{A}_2$ of $\mathcal{P}(X)$ for a set X are *interpolable* if whenever $A \subseteq B$ for $A \in \mathfrak{A}_1$ and $B \in \mathfrak{A}_2$, then there exists $C \in \mathfrak{A}_1 \cap \mathfrak{A}_2$ such that $A \subseteq C$ and $C \subseteq B$. We say that $\mathcal{L}(\Lambda)$ has *one-step interpolation* if given interpolable $\mathfrak{A}_1, \mathfrak{A}_2$ and $\phi \in \text{Prop}(\Lambda(\mathfrak{A}_1)), \psi \in \text{Prop}(\Lambda(\mathfrak{A}_2))$ such that $\text{FX} \models \phi \rightarrow \psi$, there is always an *interpolant* $\rho \in \text{Prop}(\Lambda(\mathfrak{A}_1 \cap \mathfrak{A}_2))$ such that $\text{FX} \models \phi \rightarrow \rho$ and $\text{FX} \models \rho \rightarrow \psi$. Moreover, $\mathcal{L}(\Lambda)$ has *uniform one-step interpolation* if the interpolant can be made to depend only on ϕ and $\mathfrak{A}_1 \cap \mathfrak{A}_2 =: \mathfrak{A}_0$; it is then called a *uniform \mathfrak{A}_0 -interpolant* of ϕ .

It is in fact not hard to see that under our running assumptions, these two notions coincide, so refer to them just as *one-step interpolation*:

► **Lemma 4.2.** *The logic $\mathcal{L}(\Lambda)$ has one-step uniform interpolation iff $\mathcal{L}(\Lambda)$ has one-step interpolation.*

Proof. ‘Only if’ is trivial. For ‘if’, the assumption implies that, given data as in Definition 4.1,

$$i(\phi) = \bigwedge \{ \rho \in \text{Prop}(\Lambda(\text{Prop}(\mathfrak{A}_0))) \mid \text{FX} \models \phi \rightarrow \rho \},$$

effectively a finite formula because Λ is finite, is a uniform \mathfrak{A}_0 -interpolant of ϕ . ◀

► **Remark 4.3.** In any logic supporting the requisite propositional connectives, the set of formulas ϕ having a uniform interpolant is easily seen to be closed under disjunction, and similarly the set of pairs of formulas ϕ, ψ such that an interpolant between ϕ and ψ exists is closed under disjunction in ϕ and under conjunction in ψ . When establishing (one-step) uniform interpolation for ϕ , or (one-step) interpolation between ϕ and ψ , we can therefore assume that ϕ is a conjunctive clause over modalized formulas and that ψ is a disjunctive clause over modalized formulas.

Our first positive example is neighbourhood logic:

► **Example 4.4.** Neighbourhood logic has one-step interpolation (and hence, by Lemma 4.2, uniform one-step interpolation). To see this, let $\mathfrak{A}_1, \mathfrak{A}_2$ be interpolable Boolean subalgebras of $\mathcal{P}(X)$, let ϕ be a conjunctive clause over $\Lambda(\mathfrak{A}_1)$, and let ψ be a disjunctive clause over $\Lambda(\mathfrak{A}_2)$ such that $\text{FX} \models \phi \rightarrow \psi$ (this case suffices by Remark 4.3). We can assume w.l.o.g. that $\text{FX} \not\models \neg\phi$ and $\text{FX} \not\models \psi$. Then $\text{FX} \models \phi \rightarrow \psi$ implies that ϕ contains a conjunct $\epsilon \square A$ and ψ a disjunct $\epsilon \square B$, with ϵ representing either nothing or negation, such that $A = B$. Then $\epsilon \square A$ interpolates between ϕ and ψ .

Preservation of surjective weak pullbacks is sufficient for uniform one-step interpolation:

► **Lemma 4.5.** *Let Λ be separating, and let F preserve finite surjective weak pullbacks. Then $\mathcal{L}(\Lambda)$ has one-step uniform interpolation.*

Proof. Let $\mathfrak{A}_0 \subseteq \mathfrak{A}_1$ be finite Boolean subalgebras of $\mathcal{P}X$, and let $\phi \in \text{Prop}(\Lambda(\mathfrak{A}_1))$. We show that

$$i(\phi) = \bigwedge \{ \rho \in \text{Prop}(\Lambda(\mathfrak{A}_0)) \mid FX \models \phi \rightarrow \rho \}$$

(effectively a finite formula) is a uniform \mathfrak{A}_0 -interpolant for ϕ . Dually, we show for $\psi \in \text{Prop}(\Lambda(\mathfrak{A}_2))$ with $\mathfrak{A}_1, \mathfrak{A}_2$ interpolable, $\mathfrak{A}_1 \cap \mathfrak{A}_2 \subseteq \mathfrak{A}_0$, and $i(\phi) \wedge \psi$ satisfiable that also $\phi \wedge \psi$ is satisfiable. By Lemma 2.9, we have $s \in F(\mathcal{S}(\mathfrak{A}_2))$ such that $s \models (i(\phi) \wedge \psi)\text{can}_{\mathfrak{A}_2}$. Let θ be the $\text{Prop}(\Lambda(\mathfrak{A}_1 \cap \mathfrak{A}_2))$ -theory $\theta = \bigwedge \{ \rho \in \text{Prop}(\Lambda(\mathfrak{A}_1 \cap \mathfrak{A}_2)) \mid s \models \rho\text{can}_{\mathfrak{A}_2} \}$ of s . Then $\phi \wedge \theta$ is satisfiable: otherwise, $FX \models \phi \rightarrow \neg\theta$, so $FX \models i(\phi) \rightarrow \neg\theta$ by definition of $i(\phi)$, which by Lemma 2.9 contradicts $s \models (i(\phi) \wedge \theta)\text{can}_{\mathfrak{A}_2}$.

Again by Lemma 2.9, we thus have $t \in F(\mathcal{S}(\mathfrak{A}_1))$ such that $t \models (\phi \wedge \theta)\text{can}_{\mathfrak{A}_1}$. Let \mathfrak{A} be the Boolean subalgebra of $\mathcal{P}X$ generated by $\mathfrak{A}_1 \cup \mathfrak{A}_2$. Then the diagram

$$\begin{array}{ccc} \mathcal{S}(\mathfrak{A}) & \xrightarrow{\pi_1} & \mathcal{S}(\mathfrak{A}_1) \\ \pi_2 \downarrow & & \downarrow f \\ \mathcal{S}(\mathfrak{A}_2) & \xrightarrow{g} & \mathcal{S}(\mathfrak{A}_1 \cap \mathfrak{A}_2), \end{array}$$

where all maps are canonical projections, is a finite surjective pullback because $\mathfrak{A}_1, \mathfrak{A}_2$ are interpolable, hence weakly preserved by F . We claim that $Ff(t) = Fg(s)$: indeed, both sides satisfy $\theta\text{can}_{\mathfrak{A}_1 \cap \mathfrak{A}_2}$ by Lemma 2.10, and since θ is a complete $\text{Prop}(\Lambda(\mathfrak{A}_1 \cap \mathfrak{A}_2))$ -theory, equality follows by separation. It follows that we have $u \in F(\mathcal{S}(\mathfrak{A}))$ such that $F\pi_1(u) = t$ and $F\pi_2(u) = s$. Again by Lemma 2.10, $u \models \phi\text{can}_{\mathfrak{A}}$ and $u \models \psi\text{can}_{\mathfrak{A}}$, so by Lemma 2.9, $\phi \wedge \psi$ is satisfiable. \blacktriangleleft

The example of neighbourhood logic (Examples 4.4 and 3.3.1) shows that the converse of Lemma 4.5 does not hold in general. It does however hold in the monotone case:

► **Lemma 4.6.** *Let $\mathcal{L}(\Lambda)$ be monotone and separating and have one-step interpolation. Then F preserves finite surjective weak pullbacks.*

The proof relies on invariant sets, and uses the following lemma:

► **Lemma 4.7.** *Let $f : X \rightarrow Y$ be surjective, let \mathfrak{A} denote the subalgebra of $\mathcal{P}(X)$ consisting of the f -invariant sets, and let $\phi \in \text{Prop}(\Lambda(\mathfrak{A}))$. Then $TX \models \phi$ iff $TY \models \phi\sigma_f$.*

Proof (Lemma 4.6, Sketch). Let $X \xleftarrow{\pi_1} P \xrightarrow{\pi_2} Y$ be a finite surjective pullback of $X \xrightarrow{f} Z \xleftarrow{g} Y$ as in Diagram (1). As indicated in Section 2, we can identify F with its MSS functor, i.e. we assume that FX consists of maximally satisfiable subsets $\Phi \subseteq \text{Prop}(\Lambda(\mathcal{P}(X)))$. In this reading, we are given $\Phi_1 \in FX$ and $\Phi_2 \in FY$ such that $Ff(\Phi_1) = Fg(\Phi_2)$, which by a straightforward generalization of Lemma 3.5 means that Φ_1 and Φ_2 are *compatible*, i.e.

$$\phi \in \Phi_2 \text{ implies } \phi\sigma_{\pi_2^-}\sigma_{\pi_1} \in \Phi_1$$

and symmetrically. We have to show that there exists $\Phi \in FR$ such that $F\pi_1(\Phi) = \Phi_1$ and $F\pi_2(\Phi) = \Phi_2$, i.e.

$$\phi \in \Phi_1 \iff \phi\sigma_{\pi_1^-} \in \Phi \tag{2}$$

and correspondingly for Φ_2 . In (2), ‘ \Rightarrow ’ is sufficient, because the logic has negation. That is, we have to show that the set $\{ \phi\sigma_{\pi_1^-} \mid \phi \in \Phi_1 \} \cup \{ \phi\sigma_{\pi_2^-} \mid \phi \in \Phi_2 \}$ is one-step satisfiable. Since

21:10 Uniform Interpolation in Monotone Coalgebraic Modal Logic

Φ_1 and Φ_2 are effectively finite and closed under conjunctions, it thus suffices to show that whenever $\phi_1 \in \Phi_1$ and $\phi_2 \in \Phi_2$, then

$$\phi = \phi_1 \sigma_{\pi_1^-} \wedge \phi_2 \sigma_{\pi_2^-}$$

is one-step satisfiable. Assume the contrary; then $\phi_1 \sigma_{\pi_1^-} \rightarrow \neg \phi_2 \sigma_{\pi_2^-}$ is one-step valid. Now let \mathfrak{A}_1 denote the Boolean subalgebra of $\mathcal{P}(R)$ consisting of the π_1 -invariant sets, correspondingly \mathfrak{A}_2 for the π_2 -invariant sets. One checks that $\mathfrak{A}_1, \mathfrak{A}_2$ are interpolable. Since $\mathcal{L}(\Lambda)$ has one-step interpolation, we therefore find $\rho \in \text{Prop}(\Lambda(\mathfrak{A}_1 \cap \mathfrak{A}_2))$ such that $R \models \phi_1 \sigma_{\pi_1^-} \rightarrow \rho$ and $R \models \rho \rightarrow \neg \phi_2 \sigma_{\pi_2^-}$. Using surjectivity of the π_i , Lemma 4.7, and compatibility, we can derive $\rho \sigma_{\pi_1} \in \Phi_1$, $\rho \sigma_{\pi_2} \in \Phi_2$, and eventually $\neg \phi_2 \in \Phi_2$, contradicting satisfiability of Φ_2 . \blacktriangleleft

5 Uniform Interpolation

We now relate one-step interpolation to interpolation for the full logic. Recall from Section 2.2 that we work in a language with propositional variables. Given a set $V_0 \subseteq V$ of propositional variables, we write $\mathcal{F}(\Lambda, V_0)$ for the set of Λ -formulas mentioning only propositional atoms from V_0 , and put

$$\mathcal{F}_n(\Lambda, V_0) = \{\phi \in \mathcal{F}(\Lambda, V_0) \mid \text{rk}(\phi) \leq n\}.$$

For a state x in some model, we put

$$\text{Th}_{V_0}^n(x) = \{\rho \in \mathcal{F}_n(\Lambda, V_0) \mid x \models \rho\}$$

(eliding the model, which will always be clear from the context). Since Λ is assumed to be finite, we have

► **Lemma 5.1.** *For finite V_0 , $\mathcal{F}_n(\Lambda, V_0)$ is finite up to logical equivalence.*

We record explicitly:

► **Definition 5.2.** We say that $\mathcal{L}(\Lambda)$ has *interpolation* if whenever $\models \phi \rightarrow \psi$ for $\phi \in \mathcal{F}(\Lambda, V_1)$ and $\psi \in \mathcal{F}(\Lambda, V_2)$, then there exists an *interpolant* $\rho \in \mathcal{F}(\Lambda, V_1 \cap V_2)$ such that $\models \phi \rightarrow \rho$ and $\models \rho \rightarrow \psi$; and $\mathcal{L}(\Lambda)$ has *uniform interpolation* if the interpolant ρ can be made to depend only on $V_0 := V_1 \cap V_2$. We then call ρ a *uniform V_0 -interpolant* of ϕ .

We do not currently know whether one-step interpolation in the strong sense of Definition 4.1 is necessary for $\mathcal{L}(\Lambda)$ to have interpolation. However, a weaker version of one-step interpolation is necessary:

► **Lemma 5.3.** *If $\mathcal{L}(\Lambda)$ has interpolation, then the one-step logic $\text{Prop}(\Lambda(\text{Prop}(V)))$ has interpolation.*

This can be used to disprove interpolation in some examples (contradicting [27] as indicated in the introduction):

► **Example 5.4.** Let \mathcal{N}_\vee be the subfunctor of the neighbourhood functor \mathcal{N} defined by

$$\mathcal{N}_\vee X = \{\alpha \in \mathcal{N}X \mid \forall A, B \subseteq X. A \cup B = X \Rightarrow (A \in \alpha \vee B \in \alpha)\},$$

and interpret the modality \Box over \mathcal{N}_\vee like over \mathcal{N} . Take $V_1 = \{p, q\}$, $V_2 = \{r, p\}$. Then the implication $\neg \Box(p \vee q) \rightarrow \Box(\neg p \vee r)$ is valid but has no interpolant in $\text{Prop}(\{\Box\}(\text{Prop}\{p\}))$.

As to sufficiency, we have

► **Theorem 5.5.** *If $\mathcal{L}(\Lambda)$ has one-step interpolation then $\mathcal{L}(\Lambda)$ has uniform interpolation.*

Proof (Sketch). Induction on the rank, proving the stronger claim that the rank of the uniform interpolant of ϕ is at most $\text{rk}(\phi)$. Let $\phi \in \mathcal{L}_n(\Lambda, V_1)$, and let $V_0 \subseteq V_1$. We claim that

$$i(\phi) = \bigwedge \{ \phi' \in \mathcal{F}_n(\Lambda, V_0) \mid \models \phi \rightarrow \phi' \}$$

(by Lemma 5.1, effectively a finite formula) is a uniform V_0 -interpolant for ϕ . The proof reduces straightforwardly to showing that, given $\psi \in \mathcal{F}(\Lambda, V_2)$ where $V_1 \cap V_2 \subseteq V_0$ and models $D = (Y, \zeta, \tau_2)$, $C = (X, \xi, \tau_1)$ and $y_0 \in Y$, $x_0 \in X$ such that $y_0 \models_D i(\phi) \wedge \psi$ and $x_0 \models_C \phi \wedge \text{Th}_{V_0}^{n-k}(y_0)$, the formula $\phi \wedge \psi$ is satisfiable.

Using a minor variation of standard model constructions in coalgebraic modal logic [33, 36, 24], we can assume that C, D are finite dags in which all states have a well-defined *height* (distance from any initial state in a supporting Kripke frame), with x_0 and y_0 being initial states whose depth (length of the longest path starting at x_0 and y_0 , respectively) equals the rank of the relevant formulas, and in which every state x of height $n - k$ in C is uniquely determined (among the states of height $n - k$) by $\text{Th}_{V_1}^k(x)$, correspondingly for $y \in D$ and $\text{Th}_{V_2}^k(x)$. Moreover, we can assume that the models are *canonical*, i.e. every maximally satisfiable subset of $\mathcal{F}_k(\Lambda, V_1)$ is indeed satisfied at a unique state of height $n - k$ of C , as by Lemma 5.1, there are only finitely many such sets; correspondingly for D and $\mathcal{F}_k(\Lambda, V_2)$.

We now construct a model $E = (Z, \gamma, \tau)$ of $\phi \wedge \psi$ as follows. We put

$$\begin{aligned} Z = & \{ (x, y) \in X \times Y \mid n \geq \text{ht}(x) = \text{ht}(y) =: k, \text{Th}_{V_0}^{n-k}(x) = \text{Th}_{V_0}^{n-k}(y) \} \\ & \cup \{ y \in Y \mid \text{ht}(y) > n \}, \end{aligned}$$

denoting the first part by Z_0 and the second by Z_1 , and their height- k levels by Z_0^k, Z_1^k, Z^k , respectively. Note that $(x_0, y_0) \in Z_0$. It is straightforward to define the valuation τ on Z . Moreover, we define a coalgebra structure $\gamma : Z \rightarrow \text{FZ}$ such that $\gamma(z) \in \text{FZ}^{k+1}$ for $z \in Z^k$. We put $\gamma(y) = \zeta(y) \in \text{FZ}_1 \subseteq \text{FZ}$ for $y \in Z_1$ (using preservation of inclusions by F), and on states $(x, y) \in Z_0$ of maximal height n by $\gamma(x, y) = \zeta(y)$. On the rest of Z_0 we define γ by a *coherence* requirement: By construction of Z , we have a well-defined a *pseudo-satisfaction* relation \models^0 on Z given by

$$\begin{aligned} (x, y) \models^0 \rho & \iff \begin{cases} x \models_C \rho & (\rho \in \mathcal{F}_{n-\text{ht}(x)}(\Lambda, V_1)) \\ y \models_D \rho & (\rho \in \mathcal{F}_{\text{rk}(\psi)-\text{ht}(y)}(\Lambda, V_2)) \end{cases} \\ y \models^0 \rho & \iff y \models_D \rho \quad (\rho \in \mathcal{F}_{\text{rk}(\psi)-\text{ht}(y)}(\Lambda, V_2)) \end{aligned}$$

For a $\rho \in \mathcal{F}_n(\Lambda, V_1) \cup \mathcal{F}_{\text{rk}(\psi)}(\Lambda, V_2)$, we then have the *pseudo-extension* $\hat{\rho} \subseteq Z$ defined by

$$\hat{\rho} = \{ z \in Z \mid \text{rk}(\rho) \leq n - \text{ht}(z), z \models^0 \rho \}.$$

Then we say that γ is *coherent* if for $\heartsuit \rho \in \mathcal{F}_{n-k}(\Lambda, V_1) \cup \mathcal{F}_{n-k}(\Lambda, V_2)$, $k \leq n$, and $(x, y) \in Z_0^k$,

$$\gamma(x, y) \models \heartsuit(\hat{\rho} \cap Z_0^{k+1}) \iff (x, y) \models^0 \heartsuit \rho.$$

For $i = 0, 1, 2$, let \mathfrak{A}_i be the Boolean subalgebra of $\mathcal{P}Z_0^{k+1}$ of sets of the form $\hat{\rho} \cap Z_0^{k+1}$ for $\rho \in \mathcal{F}_{n-k-1}(\Lambda, V_i)$. Then, of course, $\mathfrak{A}_1 \cap \mathfrak{A}_2 \subseteq \mathfrak{A}_0$, and by induction, $\mathfrak{A}_1, \mathfrak{A}_2$ are interpolable. We define $\phi_0 \in \text{Prop}(\Lambda(\mathfrak{A}_1))$ and $\psi_0 \in \text{Prop}(\Lambda(\mathfrak{A}_2))$ by

$$\begin{aligned} \phi_0 &= \bigwedge \{ \epsilon \heartsuit(\hat{\rho} \cap Z_0^{k+1}) \mid (x, y) \models^0 \epsilon \heartsuit \rho, \rho \in \mathcal{F}_{n-k-1}(\Lambda, V_1), \epsilon \in \{ \cdot, \neg \} \} \\ \psi_0 &= \bigwedge \{ \epsilon \heartsuit(\hat{\rho} \cap Z_0^{k+1}) \mid (x, y) \models^0 \epsilon \heartsuit \rho, \rho \in \mathcal{F}_{n-k-1}(\Lambda, V_2), \epsilon \in \{ \cdot, \neg \} \}. \end{aligned}$$

By applying Lemma 4.2 and forming the uniform \mathfrak{A}_0 -interpolant $i(\phi_0)$ of ϕ_0 , and showing satisfiability of $i(\phi_0) \wedge \psi_0$, one can show that a coherent γ exists. Then, we have by induction on $\rho \in \mathcal{F}_n(\Lambda, V_1) \cup \mathcal{F}_{\text{rk}(\psi)}(\Lambda, V_2)$ that $z \vDash_E \rho$ iff $z \vDash^0 \rho$ for $\text{ht}(z) = k$ and $\text{rk}(\rho) \leq n - k$; so in particular $z_0 = (x_0, y_0) \vDash \phi$ and $z_0 \vDash \psi$, as required. \blacktriangleleft

► **Remark 5.6.** Canonical models in the sense of the above proof sketch in fact bear a strong resemblance to models based on the stages of the final sequence of functors of the type $\mathcal{P}(V_i) \times \mathbf{F}$, $i = 0, 1$ (e.g. [25]). This indicates in particular that the proof may eventually be made to bear a relationship, via duality, with Ghilardi’s method of graded modal algebras [10].

As indicated in the introduction, our results can be summed up as follows:

► **Theorem 5.7.** *Let Λ be finite. Then the following properties imply each other in sequence:*

1. Λ is separating and the type functor \mathbf{F} preserves finite surjective weak pullbacks.
 2. $\mathcal{L}(\Lambda)$ has one-step interpolation
 3. $\mathcal{L}(\Lambda)$ has uniform interpolation
 4. $\mathcal{L}(\Lambda)$ has interpolation
 5. The one-step logic $\text{Prop}(\Lambda(\text{Prop}(V)))$ has interpolation.
- Moreover, if Λ is monotone and separating, then 2. implies 1.

We note that if Λ is finite and separating, then \mathbf{F} preserves finite sets. As indicated above, we suspect but cannot currently prove that 5 implies 2, which would make items 2–5 equivalent. From Theorem 5.7, we obtain uniform interpolation for the following concrete logics:

- **Example 5.8. 1.** Whenever \mathbf{F} preserves weak pullbacks and Λ is finite and separating, then $\mathcal{L}(\Lambda)$ has uniform interpolation. This case is covered already in [20], see Remark 5.9. In particular, we obtain that the modal logics K and KD have uniform interpolation, thus reproving previous results [10, 39].
2. Since the monotone neighbourhood functor preserves surjective weak pullbacks (Section 3), we obtain that monotone modal logic has uniform interpolation, again reproving a previous result [32].
 3. If M is a finite refinable monoid, then the monoid-weighted functor \mathbf{S}_M (Section 3) preserves surjective weak pullbacks, so that any rank-1 modal logic that is expressive (i.e. separating) for \mathbf{S}_M has uniform interpolation, such as the logic with modalities $[m]$ for $m \in M$, interpreted by the predicate lifting given by $[m]_X(A) = \{\mu \in \mathbf{S}_M \mid \sum_{x \in A} \mu(x) = m\}$. This holds in particular when M is a finite Abelian group, in which case \mathbf{S}_M does not have a separating set of monotone predicate liftings so that this case is not covered by existing generic results [20]. If we take $M = \mathbb{Z}/n\mathbb{Z}$, then the modalities $[m]$ described above are modulo-constraints as found in Presburger modal logic [7]: $[m]\phi$ says that the number of successors of the current state (counting multiplicities) equals m modulo n .
 4. Neighbourhood logic fails to preserve surjective weak pullbacks (Example 3.3) but does have one-step interpolation (Example 4.4), so we obtain that neighbourhood logic has uniform interpolation.
 5. One-step interpolation has been proved, in slightly different terms, for coalition logic [28], so that our results improve the known interpolation result for coalition logic [28] to uniform interpolation.

► **Remark 5.9.** We conclude with a more detailed discussion of the relationship between our results and results on the logic of quasi-functorial lax liftings. Glossing over the ramifications of the axiomatics, a *diagonal-preserving lax lifting* L for a set functor T [22] extends T to act also on relations, satisfying monotonicity w.r.t. inclusion of relations, preservation of relational

converse and diagonal relations, and lax preservation of composition ($LR \circ LS \subseteq L(R \circ S)$). The monotone neighbourhood functor and its polyadic variants have diagonal-preserving lax liftings, and diagonal-preserving lax liftings are easily seen to be inherited along products and subfunctors, so that every functor that has a separating set of monotone predicate liftings has a diagonal-preserving lax lifting. Conversely, every finitary functor that has a diagonal-preserving lax lifting has a separating set of monotone predicate liftings, the so-called Moss liftings [20]. A lax lifting induces a modal logic with a slightly non-standard modality ∇ that generalizes Moss' modality for weak-pullback-preserving functors [23]; for functors that preserve finite sets, the ∇ -modality and the predicate-lifting based modalities are however mutually intertranslatable [20], essentially by dint of the fact that both are separating. Summing up, for a functor that preserves finite sets, a diagonal-preserving lax lifting exists iff a separating (finite) set of monotone predicate lifting exists, and the induced logics are essentially the same.

Marti [20] shows that the logic of a diagonal-preserving lax lifting L for T has uniform interpolation if T preserves finite sets and L is *quasifunctorial*, i.e. satisfies $LS \circ LR = L(S \circ R) \cap (\text{dom}(LR) \times \text{rg}(LS))$ where $\text{dom}(LR) = \{t \mid \exists t. (s, t) \in LR\}$ and $\text{rg}(LS) = \{t \mid \exists s. (s, t) \in LS\}$. We recall again that our reduction of uniform interpolation to one-step interpolation holds also in cases where either separation or monotonicity fails, such as coalition logic / alternating-time logic and neighbourhood logic, respectively. Also, we have seen examples (Abelian-group-weighted functors) where there is no monotone separating set of predicate liftings but we nevertheless obtain uniform interpolation from preservation of surjective weak pullbacks.

6 Conclusions

We have given sufficient criteria for a rank-1 modal logic (with finitely many modalities), i.e. a coalgebraic modal logic, to have uniform interpolation: In the general case, we have established a reduction to the one-step logic; and in the case where the modalities are separating, we have given a simple semantic criterion, namely preservation of (finite) surjective weak pullbacks, which in the monotone case is in fact also necessary for one-step interpolation. We have thus reproved uniform interpolation for the relational modal logics K and KD and for monotone (neighbourhood) modal logic, and newly established uniform interpolation for coalition logic, neighbourhood logic (i.e. classical modal logic), and various logics of finite-monoid-weighted transition systems. All proofs are entirely semantic; we leave a proof-theoretic treatment, in generalization of tentative results based on cut-free sequent systems [28], for future work. In particular, such a treatment will hopefully lead to practically feasible algorithms for the computation of interpolants. Another open question is how our results relate to definability of bisimulation quantifiers.

Acknowledgments The authors wish to thank Tadeusz Litak and Sebastian Enqvist for helpful discussions.

References

- 1 J. Adámek, H. Herrlich, and G. E. Strecker. *Abstract and Concrete Categories*. Wiley Interscience, 1990. Republished in *Reprints in Theor. Appl. Cat.* 17 (2006).
- 2 M. Bílková. Uniform interpolation and propositional quantifiers in modal logics. *Stud. Log.*, 85:1–31, 2007.
- 3 B. Chellas. *Modal Logic*. Cambridge University Press, 1980.
- 4 D. Coumans and B. Jacobs. Scalars, monads, and categories. arXiv 1003.0585, 2010.

- 5 W. Craig. Linear reasoning. A new form of the Herbrand-Gentzen theorem. *J. Symb. Log.*, 22:250–268, 1957.
- 6 G. D’Agostino and M. Hollenberg. Logical questions concerning the μ -calculus: Interpolation, Lyndon and Łoś-Tarski. *J. Symb. Log.*, 65:310–332, 2000.
- 7 S. Demri and D. Lugiez. Complexity of modal logics with Presburger constraints. *J. Appl. Log.*, 8:233–252, 2010.
- 8 D. Gabbay. Craig’s interpolation theorem for modal logics. In *Conference in Mathematical Logic—London’70*, vol. 255 of *LNM*, pp. 111–127. Springer, 1972.
- 9 D. Gabbay. Craig interpolation theorem for intuitionistic logic and extensions Part III. *J. Symb. Log.*, 42:269–271, 1977.
- 10 S. Ghilardi. An algebraic theory of normal forms. *Ann. Pure Appl. Log.*, 71:189–245, 1995.
- 11 S. Ghilardi and M. Zawadowski. Undefinability of propositional quantifiers in the modal system S4. *Stud. Log.*, 55:259–271, 1995.
- 12 H. P. Gumm and T. Schröder. Coalgebraic structure from weak limit preserving functors. In *Coalgebraic Methods in Computer Science, CMCS 2000*, vol. 33 of *ENTCS*, pp. 111–131. Elsevier, 2000.
- 13 H. P. Gumm and T. Schröder. Monoid-labeled transition systems. In *Coalgebraic Methods in Computer Science, CMCS 2001*, vol. 44 of *ENTCS*, pp. 185–204. Elsevier, 2001.
- 14 H. H. Hansen and C. Kupke. A coalgebraic perspective on monotone modal logic. In *Coalgebraic Methods in Computer Science, CMCS 2004*, vol. 106 of *ENTCS*, pp. 121–143. Elsevier, 2004.
- 15 H. H. Hansen, C. Kupke, and E. Pacuit. Neighbourhood structures: Bisimilarity and basic model theory. *Log. Meth. Comput. Sci.*, 5, 2009.
- 16 L. Henkin. An extension of the Craig-Lyndon interpolation theorem. *J. Symb. Log.*, 28:201–216, 1963.
- 17 D. Kozen. Results on the propositional μ -calculus. *Theor. Comput. Sci.*, 27:333–354, 1983.
- 18 A. Kurz and R. Leal. Modalities in the Stone age: A comparison of coalgebraic logics. *Theor. Comput. Sci.*, 430:88–116, 2012.
- 19 C. Lutz and F. Wolter. Foundations for uniform interpolation and forgetting in expressive description logics. In *International Joint Conference on Artificial Intelligence, IJCAI 2011*, pp. 989–995. IJCAI/AAAI, 2011.
- 20 J. Marti. Relation liftings in coalgebraic modal logic. Master’s thesis, Universiteit van Amsterdam, 2011.
- 21 J. Marti, F. Seifan, and Y. Venema. Uniform interpolation for coalgebraic fixpoint logic. In *Algebra and Coalgebra in Computer Science, CALCO 2015*, vol. 35 of *LIPICs*, pp. 238–252. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
- 22 J. Marti and Y. Venema. Lax extensions of coalgebra functors and their logic. *J. Comput. Syst. Sci.*, 81:880–900, 2015.
- 23 L. Moss. Coalgebraic logic. *Ann. Pure Appl. Logic*, 96:277–317, 1999.
- 24 R. Myers, D. Pattinson, and L. Schröder. Coalgebraic hybrid logic. In L. de Alfaro, ed., *Foundations of Software Science and Computation Structures, FoSSaCS 2009*, vol. 5504 of *LNCS*, pp. 137–151. Springer, 2009.
- 25 D. Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoret. Comput. Sci.*, 309:177–193, 2003.
- 26 D. Pattinson. Expressive logics for coalgebras via terminal sequence induction. *Notre Dame J. Formal Log.*, 45:19–33, 2004.
- 27 D. Pattinson. The logic of exact covers: Completeness and uniform interpolation. In *Logic in Computer Science, LICS 2013*, pp. 418–427. IEEE, 2013.
- 28 D. Pattinson and L. Schröder. Cut elimination in coalgebraic logics. *Inf. Comput.*, 208:1447–1468, 2010.

- 29 M. Pauly. A modal logic for coalitional power in games. *J. Log. Comput.*, 12:149–166, 2002.
- 30 A. Pitts. On an interpretation of second order quantification in first order intuitionistic propositional logic. *J. Symb. Log.*, 57:33–52, 1992.
- 31 J. Rutten. Universal coalgebra: A theory of systems. *Theor. Comput. Sci.*, 249:3–80, 2000.
- 32 L. Santocanale and Y. Venema. Uniform interpolation for monotone modal logic. In *Advances in Modal Logic, AiML 2010*, pp. 350–370. College Publications, 2010.
- 33 L. Schröder. A finite model construction for coalgebraic modal logic. *J. Log. Algebr. Prog.*, 73:97–110, 2007.
- 34 L. Schröder. Expressivity of coalgebraic modal logic: The limits and beyond. *Theor. Comput. Sci.*, 390:230–247, 2008.
- 35 L. Schröder and D. Pattinson. PSPACE bounds for rank-1 modal logics. *ACM Trans. Comput. Log.*, 10(2:13):1–33, 2009.
- 36 L. Schröder and D. Pattinson. Strong completeness of coalgebraic modal logics. In *Symposium on Theoretical Aspects of Computer Science, STACS 2009*, vol. 3 of *LIPIcs*, pp. 673–684. Schloss Dagstuhl – Leibniz-Center of Informatics, 2009.
- 37 L. Schröder and D. Pattinson. Rank-1 modal logics are coalgebraic. *J. Log. Comput.*, 20(5):1113–1147, 2010.
- 38 D. Turi. *Functorial operational semantics and its denotational dual*. PhD thesis, Vrije Universiteit, Amsterdam, 1996.
- 39 A. Visser. Uniform interpolation and layered bisimulation. In *Gödel '96 (Brno, 1996)*, vol. 6 of *Lect. Notes Log.*, pp. 139–164. Springer, 1996.
- 40 K. Wang, Z. Wang, R. Topor, J. Pan, and G. Antoniou. Concept and role forgetting in *ALC*-ontologies. In *The Semantic Web, ISWC 2009*, vol. 5823 of *LNCS*. Springer, 2009.

A Omitted Proofs

Proof of Lemma 2.9

We note that for $A \in \mathfrak{A}$, $\text{can}_{\mathfrak{A}}(A) = \{U \in \mathcal{S}(\mathfrak{A}) \mid U \subseteq A\}$.

We have a map $p : X \rightarrow \mathcal{S}(\mathfrak{A})$ that maps $x \in X$ to the unique atom $U \in \mathcal{S}(\mathfrak{A})$ such that $x \in U$. We show more generally that for the interpretations of ϕ and $\phi \text{can}_{\mathfrak{A}}$, $\phi = (Fp)^{-1} \phi \text{can}_{\mathfrak{A}}$. We use induction over ϕ , with trivial Boolean steps. For the modal case, we note that for $A \in \mathfrak{A}$, $A = p^{-1}[\text{can}_{\mathfrak{A}}(A)]$; the claim then follows by naturality of predicate liftings. \blacktriangleleft

Proof of Lemma 2.10

We note that the canonical projection $f : \mathcal{S}(\mathfrak{A}_1) \rightarrow \mathcal{S}(\mathfrak{A}_0)$ maps $U \in \mathcal{S}(\mathfrak{A}_1)$ to the unique \mathfrak{A}_0 -atom V such that $U \subseteq V$.

We prove the claim by induction over ϕ , with trivial Boolean steps. For the modal case, we note that for $A \in \mathfrak{A}_0$, $\text{can}_{\mathfrak{A}_1}(A) = f^{-1}[\text{can}_{\mathfrak{A}_0}(A)]$. The claim follows by naturality of predicate liftings. \blacktriangleleft

Proof of Lemma 2.13

Define the natural transformation $\eta : \mathbf{F} \rightarrow M_{\mathbf{F}}^{\Lambda}$ by $\eta_X(t) = \{\phi \in \text{Prop}(\Lambda(\mathcal{P}(X))) \mid t \models \phi\}$, i.e., $t \in \mathbf{F}X$ is mapped to its *theory* in $M_{\mathbf{F}}^{\Lambda}(X)$. Given a set X , η_X is surjective by satisfiability of sets in $M_{\mathbf{F}}^{\Lambda}(X)$, and injective since Λ is separating and each $t \in \mathbf{F}X$ is uniquely determined by the set $\{\phi \in \Lambda(\mathcal{P}(X)) \mid t \models \phi\}$. Hence \mathbf{F} and $M_{\mathbf{F}}^{\Lambda}$ are naturally isomorphic. \blacktriangleleft

Proof of Lemma 3.2

We prove only the claim for surjective pullbacks. So let

$$\begin{array}{ccc} P & \xrightarrow{\pi_1} & X \\ \pi_2 \downarrow & & \downarrow f \\ Y & \xrightarrow{g} & Z \end{array}$$

be a surjective pullback, and let $s \in \mathbf{F}X$, $t \in \mathbf{F}Y$ such that $\mathbf{F}f(s) = \mathbf{F}g(t)$. We have to show that there exists $u \in \mathbf{F}Z$ such that $\mathbf{F}\pi_1(u) = s$ and $\mathbf{F}\pi_2(u) = t$. Since \mathbf{F} is finitary (and we assume that \mathbf{F} preserves inclusions), there exist finite subsets $X_0 \subseteq X$, $Y_0 \subseteq Y$ such that $s \in \mathbf{F}X_0 \subseteq \mathbf{F}X$, $t \in \mathbf{F}Y_0 \subseteq \mathbf{F}Y$. Since f and g are surjective, we can moreover ensure, by suitably enlarging X_0 and Y_0 , that $f[X_0] = g[Y_0] =: Z_0$. Take $P_0 = P \cap \pi_1^{-1}[X_0] \cap \pi_2^{-1}[Y_0]$. Then

$$\begin{array}{ccc} P_0 & \xrightarrow{\pi_1} & X_0 \\ \pi_2 \downarrow & & \downarrow f \\ Y_0 & \xrightarrow{g} & Z_0 \end{array}$$

(where we abuse f, g, π_1, π_2 to denote their restrictions) is a finite surjective pullback, hence weakly preserved by \mathbf{F} ; it follows that $u \in \mathbf{F}P_0 \subseteq \mathbf{F}P$ exists as required. \blacktriangleleft

Proof of Lemma 3.4

Let

$$\begin{array}{ccc}
P & \xrightarrow{\pi_1} & X \\
\pi_2 \downarrow & & \downarrow f \\
Y & \xrightarrow{g} & Z
\end{array}$$

be a surjective pullback, and let $\alpha_1 \in \mathcal{M}X$, $\alpha_2 \in \mathcal{M}Y$ such that $\mathcal{M}f(\alpha_1) = \mathcal{M}g(\alpha_2)$. We put

$$\beta = \text{Up}(\{\pi_1^{-1}[U] \mid U \in \alpha_1\} \cup \{\pi_2^{-1}[V] \mid V \in \alpha_2\}) \in \mathcal{M}P,$$

and claim that $\mathcal{M}\pi_1(\beta) = \alpha_1$ and $\mathcal{M}\pi_2(\beta) = \alpha_2$. By symmetry, it suffices to prove $\mathcal{M}\pi_1(\beta) = \alpha_1$. Here, ‘ \supseteq ’ is clear by construction of β ; we prove ‘ \subseteq ’. We distinguish the following cases for $U \in \mathcal{M}\pi_1(\beta)$:

If $\pi_1^{-1}[U'] \subseteq \pi_1^{-1}[U]$ for some $U' \in \alpha_1$, then $U' = \pi_1\pi_1^{-1}[U'] \subseteq \pi_1\pi_1^{-1}[U] = U$ by surjectivity of π_1 , and hence $U \in \alpha_1$.

Otherwise, $\pi_2^{-1}[V] \subseteq \pi_1^{-1}[U]$ for some $V \in \alpha_2$. Then $\pi_1[\pi_2^{-1}[V]] \in \alpha_1$ by the compatibility lemma, and $\pi_1[\pi_2^{-1}[V]] \subseteq \pi_1[\pi_1^{-1}[U]] = U$ by surjectivity of π_1 , so $U \in \alpha_1$. ◀

Proof of Lemma 3.5

‘*Only if*’: For $U \in \alpha_1$, we have $f^{-1}[f[U]] \supseteq U$ and therefore $f[U] \in \mathcal{M}f(\alpha_1) = \mathcal{M}g(\alpha_2)$, so $g^{-1}[f[U]] \in \alpha_2$; but $g^{-1}[f[U]] = \pi_2[\pi_1^{-1}[U]]$ by the pullback property. The other direction is symmetric.

‘*If*’: For the inclusion $\mathcal{M}f(\alpha_1) \subseteq \mathcal{M}g(\alpha_2)$, let $U \in \mathcal{M}f(\alpha_1)$. Then $f[f^{-1}[U]] = U$ by surjectivity, and $f^{-1}[U] \in \alpha_1$, hence $g^{-1}[U] = g^{-1}[f[f^{-1}[U]]] = \pi_2[\pi_1^{-1}[f^{-1}[U]]] \in \alpha_2$ by compatibility, so $U \in \mathcal{M}g(\alpha_2)$. The reverse inclusion is shown symmetrically. ◀

Proof of Lemma 3.7

It only remains to show ‘if’. It is easy to see that refinability is exactly the condition that ensures that \mathbf{S}_M maps non-empty binary products to weak pullbacks. Now \mathbf{S}_M is generally *additive* [4], i.e. maps finite coproducts to products. Since every finite set is a coproduct of singletons, every finite surjective pullback is a sum of non-empty finite products. Thus, \mathbf{S}_M maps finite surjective pullbacks to products of weak pullbacks, which are again weak pullbacks. ◀

Proof of Lemma 3.9

‘*If*’: If M is refinable and positive, then \mathbf{S}_M preserves all weak pullbacks [13] and hence has a separating set of monotone predicate liftings [18].

‘*Only if*’: Let $a + b = 0$, $a \neq 0 \neq b$ in M . We can understand the elements of $\mathbf{S}_M(X)$ as formal linear combinations over X with coefficients from M . Recall that a monotone n -ary predicate lifting λ for \mathbf{S}_M is equivalently determined by a subset of $\mathbf{S}_M(2^n)$ that is closed under the relation \lesssim given by taking $t \lesssim s$ if t, s can be written as linear combinations $s = \sum_{i=1}^n a_i x_i$, $t = \sum_{i=1}^n a_i y_i$ such that $x_i \leq y_i$ for $i = 1, \dots, n$ in the pointwise ordering on 2^n [34]. Now let X be a set with at least two distinct elements $x \neq y$. We show that $ax, ay \in \mathbf{S}_M X$ are indistinguishable under n -ary monotone predicate liftings by showing that for any $f : X \rightarrow 2^n$, we have $Tf(ax) \lesssim Tf(ay) \lesssim Tf(ax)$. Namely: Writing $\perp^n \in 2^n$ for the n -tuple (\perp, \dots, \perp) , we have

$$Tf(ax) = af(x) = af(x) + a\perp^n + b\perp^n \lesssim af(x) + af(y) + bf(x) = af(y) = Tf(ay),$$

and analogously $Tf(ay) \lesssim Tf(ax)$. ◀

Proof of Lemma 4.2

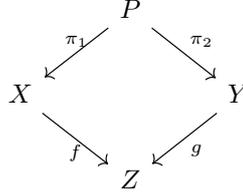
‘Only if’ is trivial; we prove ‘if’. So let $\phi \in \text{Prop}(\Lambda(\text{Prop}(V_1)))$, and let $V_0 \subseteq V_1$. We put

$$i(\phi) = \bigwedge \{ \rho \in \text{Prop}(\Lambda(\text{Prop}(V_0))) \mid \models \phi \rightarrow \rho \},$$

effectively a finite formula because Λ is finite and Prop preserves finite sets up to logical equivalence, and claim that $i(\phi)$ is a uniform V_0 -interpolant of ϕ . We clearly have $\models \phi \rightarrow i(\phi)$. Now let $\psi \in \text{Prop}(\Lambda(\text{Prop}(V_2)))$ where $V_1 \cap V_2 \subseteq V_0$, and let $\models \phi \rightarrow \psi$. By one-step interpolation, we have $\rho \in \text{Prop}(V_0)$ such that $\models \phi \rightarrow \rho$ and $\models \rho \rightarrow \psi$. But then ρ is a conjunct of $i(\phi)$, so $\models i(\phi) \rightarrow \psi$, as required. \blacktriangleleft

Proof of Lemma 4.6

Let



be a finite surjective pullback. As indicated in Section 2, we can identify F with its MSS functor, i.e. we assume that FX consists of maximally satisfiable subsets $\Phi \subseteq \text{Prop}(\Lambda(\mathcal{P}(X)))$.

In this reading, we are given $\Phi_1 \in FX$ and $\Phi_2 \in FY$ such that $Ff(\Phi_1) = Fg(\Phi_2)$, which by a straightforward generalization of Lemma 3.5 means that Φ_1 and Φ_2 are *compatible*, i.e.

$$\phi \in \Phi_2 \text{ implies } \phi \sigma_{\pi_2} \sigma_{\pi_1} \in \Phi_1$$

and symmetrically. We have to show that there exists $\Phi \in FP$ such that $F\pi_1(\Phi) = \Phi_1$ and $F\pi_2(\Phi) = \Phi_2$, i.e.

$$\phi \in \Phi_1 \iff \phi \sigma_{\pi_1} \in \Phi$$

and correspondingly for Φ_2 . Observe that in the above condition, ‘ \implies ’ is sufficient, because the logic has negation.

That is, we have to show that the set

$$\{ \phi \sigma_{\pi_1} \mid \phi \in \Phi_1 \} \cup \{ \phi \sigma_{\pi_2} \mid \phi \in \Phi_2 \}$$

is one-step satisfiable. Since Φ_1 and Φ_2 are effectively finite and closed under conjunctions, it thus suffices to show that whenever $\phi_1 \in \Phi_1$ and $\phi_2 \in \Phi_2$, then

$$\phi = \phi_1 \sigma_{\pi_1} \wedge \phi_2 \sigma_{\pi_2}$$

is one-step satisfiable. Assume the contrary; then $\phi_1 \sigma_{\pi_1} \rightarrow \neg \phi_2 \sigma_{\pi_2}$ is one-step valid. Now let \mathfrak{A}_1 denote the Boolean subalgebra of $\mathcal{P}(P)$ consisting of the π_1 -invariant sets, correspondingly \mathfrak{A}_2 for the π_2 -invariant sets. Since $\mathcal{L}(\Lambda)$ has one-step interpolation, we therefore find $\rho \in \text{Prop}(\Lambda(\mathfrak{A}_1 \cap \mathfrak{A}_2))$ such that $P \models \phi_1 \sigma_{\pi_1} \rightarrow \rho$ and $P \models \rho \rightarrow \neg \phi_2 \sigma_{\pi_2}$. Now by Lemma 4.7,

$$X \models \phi_1 \sigma_{\pi_1} \sigma_{\pi_1} \rightarrow \rho \sigma_{\pi_1}.$$

Since π_1 is surjective, $\phi_1 \sigma_{\pi_1^{-1}} \sigma_{\pi_1} = \phi_1$, and since $\phi_1 \in \Phi_1$, it follows by maximal satisfiability that $\rho \sigma_{\pi_1} \in \Phi_1$, so by compatibility,

$$\rho \sigma_{\pi_1} \sigma_{\pi_1^{-1}} \sigma_{\pi_2} \in \Phi_2. \quad (3)$$

We claim that for $A \in \mathfrak{A}_1$,

$$(\sigma_{\pi_1} \sigma_{\pi_1^{-1}} \sigma_{\pi_2})(A) = \sigma_{\pi_2}(A).$$

Indeed, keeping in mind that composition of substitutions is diagrammatic, we have

$$\pi_2[\pi_1^{-1}[\pi_1[A]]] = \pi_2[A]$$

by π_1 -invariance of A . Thus, from (3), we obtain $\rho \sigma_{\pi_2} \in \Phi_2$. Since $\rho \rightarrow \neg \phi_2 \sigma_{\pi_2^{-1}} \in \text{Prop}(\Lambda(\mathfrak{A}_2))$, we have by Lemma 4.7

$$X \models \rho \sigma_{\pi_2} \rightarrow \neg \phi_2 \sigma_{\pi_2^{-1}} \sigma_{\pi_2},$$

and since, again, $\phi_2 \sigma_{\pi_2^{-1}} \sigma_{\pi_2} = \phi_2$ because π_2 is surjective, it follows that $\neg \phi_2 \in \Phi_2$, in contradiction to satisfiability of Φ_2 . \blacktriangleleft

Proof of Lemma 4.7

Induction over ϕ , with trivial Boolean steps. For the modal step, i.e. by the above for atoms of the form $\heartsuit f^{-1}[B]$, note that Ff is also surjective, and for $t \in FX$,

$$s \models \heartsuit f^{-1}[B] \iff Ff(s) \models \heartsuit B = \heartsuit f[f^{-1}[B]],$$

where the last equality is by surjectivity. \blacktriangleleft

Proof of Lemma 5.3

Let $\phi \in \text{Prop}(\Lambda(\text{Prop}(V_1)))$ and $\psi \in \text{Prop}(\Lambda(\text{Prop}(V_2)))$ such that $\models \phi \rightarrow \psi$. Since $\mathcal{L}(\Lambda)$ has propositional variables, we can see ϕ and ψ as formulas in $\mathcal{F}(\Lambda)$, so ϕ and ψ have an interpolant $\theta \in \mathcal{F}(\Lambda, V_1 \cap V_2)$

We proceed to show that we can modify θ to obtain an interpolant in $\text{Prop}(\Lambda(\text{Prop}(V_1 \cap V_2)))$. We begin by eliminating occurrences of propositional variables at the top level of θ , i.e. outside the scope of modal operators. Let θ' arise from θ by replacing all such occurrences with \top . Then, of course, $\theta' \in \mathcal{F}(\Lambda, V_1 \cap V_2)$. Moreover, we still have $\models \phi \rightarrow \theta'$: Let x be a state in a coalgebraic model C such that $C, x \models \phi$; we have to show $C, x \models \theta'$. Since coalgebraic modal logic has the tree model property [35], we can assume that C is a tree. Under this assumption, we can change the valuation of all propositional variables at x to \top without affecting satisfaction of ϕ , obtaining C' such that $C', x \models \phi$. Then $C', x \models \theta$, which by the choice of valuation at x implies $C', x \models \theta'$. Since θ' does not contain propositional variables at the top level and C is tree-shaped, it follows that $C, x \models \theta'$. The proof that $\models \theta' \rightarrow \psi$ is entirely analogous.

We can thus assume that $\theta \in \text{Prop}(\Lambda(\mathcal{F}(\Lambda)))$. That is, θ has the form

$$\theta = \theta_0 \sigma$$

where

$$\theta_0 \in \text{Prop}(\Lambda(\text{Prop}(V_0 \cup (V_1 \cap V_2))))$$

where V_0 is disjoint from V_1 and V_2 , and

$$\sigma : V_0 \rightarrow \Lambda(\mathcal{F}(\Lambda)).$$

Informally speaking, V_0 and σ represent the parts of θ that stick out of rank 1. Now let y be a state in some coalgebraic model D , and define $\sigma' : V_0 \rightarrow 2$ by

$$\sigma'(v) = \top \quad \text{iff} \quad D, x \models \sigma(v).$$

Then $\theta' := \theta_0 \sigma'$ is rank 1 and mentions only variables in $V_1 \cap V_2$; we claim that it interpolates between ϕ and ψ .

Again, we prove only that $\models \phi \rightarrow \theta'$. So let x be a state in a coalgebraic model C such that $C, x \models \phi$; we have to show $C, x \models \theta'$. As above, we can assume that C is tree-shaped. We now construct a coalgebraic model C' by replacing every successor state of x and the associated subtree (according to the supporting Kripke frame of C) with a copy of y and D , keeping however the original valuation for the propositional variables. Then

$$C \text{ and } C' \text{ satisfy the same rank-1 formulas,} \tag{4}$$

in particular $C', x \models \phi$ and therefore $C', x \models \theta$. By the definition of σ' it follows that $C', x \models \theta'$, and since θ' is rank 1, this implies $C, x \models \theta'$, again by (4). \blacktriangleleft

Details for Example 5.4

We use liberally that the rule $a \vee b / \Box a \vee \Box b$ is *one-step cutfree complete* for the logic of \mathcal{N}_\vee [36], i.e. every valid clause over $\{\Box\}(\text{Prop}(V))$ must contain $\Box\phi$ and $\Box\psi$ such that $\phi \vee \psi$ is valid.

Assume that an interpolant in $\text{Prop}(\{\Box\}(\text{Prop}(\{p\})))$ exists. Then by the same argument as in Lemma 4.2, the formula χ obtained as the essentially finite conjunction of all clauses ρ over $\{\Box\}(\text{Prop}(\{p\}))$ such that $\Box(p \vee q) \vee \rho$ is (one-step) valid is also an interpolant. In the definition of χ , we can clearly omit clauses ρ that are themselves valid. Then the only way that $\Box(p \vee q) \vee \rho$ is valid is that ρ contains (and then w.l.o.g. equals) a literal $\Box(\theta)$ with $\theta \in \text{Prop}(\{p\})$ such that $(p \vee q) \vee \theta$ is valid. Up to logical equivalence, θ is one of $\top, \perp, p, \neg p$; validity of $(p \vee q) \vee \theta$ leaves only $\theta = \neg p$. But $\Box(\neg p) \rightarrow \Box(\neg p \vee r)$ is not valid. \blacktriangleleft

Proof of Theorem 5.5

Let $\phi \in \mathcal{L}_n(\Lambda, V_1)$, and let $V_0 \subseteq V_1$. We claim that

$$i(\phi) = \bigwedge \{ \phi' \in \mathcal{F}_n(\Lambda, V_0) \mid \models \phi \rightarrow \phi' \}$$

(by Lemma 5.1, effectively a finite formula) is a uniform V_0 -interpolant for ϕ . Clearly, $\models \phi \rightarrow i(\phi)$, so it remains to show that $\models i(\phi) \rightarrow \psi$ given $\psi \in \mathcal{F}(\Lambda, V_2)$ such that $\models \phi \rightarrow \psi$ and $V_1 \cap V_2 \subseteq V_0$. We can clearly assume that $i(\phi)$ is satisfiable; since $\perp \in \mathcal{F}_n(\Lambda, V_0)$, this implies that \perp is not a conjunct of $i(\phi)$, i.e. ϕ is satisfiable. We can then dualize the goal and prove that whenever $i(\phi) \wedge \psi$ is satisfiable for some $\psi \in \mathcal{F}(\Lambda, V_2)$ where $V_1 \cap V_2 \subseteq V_0$, then $\phi \wedge \psi$ is satisfiable. So we are given two models, $D = (Y, \zeta, \theta_2)$ such that $y_0 \models_D i(\phi) \wedge \psi$ for some $y_0 \in Y$, and $C = (X, \xi, \theta_1)$ such that $x_0 \models_C \phi$ for some $x_0 \in X$. Let $\rho = \text{Th}_{V_0}(y_0)$. Then $\phi \wedge \rho$ is satisfiable (otherwise, $\models \phi \rightarrow \neg\rho$, so $i(\phi) \rightarrow \neg\rho$ and hence $y_0 \models \neg\rho$, contradiction), so we assume that x_0 in fact satisfies $\phi \wedge \rho$.

Using a minor variation of standard model constructions in coalgebraic modal logic [33, 36, 24], we can assume that both these models are finite dags in which all states have a

well-defined *height* (distance from any initial state in a smallest supporting Kripke frame, which exists because the models are finite), with x_0 and y_0 being initial states whose depth (length of the longest path starting at x_0 and y_0 , respectively) equals the rank of the relevant formulas, and in which every state x of height $n - k$ in C is uniquely determined (among the states of height $n - k$) by $\text{Th}_{V_1}^k(x)$, correspondingly for $y \in D$ and $\text{Th}_{V_2}^k(x)$. Moreover, we can assume that the models are *canonical*, i.e. every maximally satisfiable subset of $\mathcal{F}_k(\Lambda, V_1)$ is indeed satisfied at a unique state of height $n - k$ of C , as by Lemma 5.1, there are only finitely many such sets; correspondingly for D and $\mathcal{F}_k(\Lambda, V_2)$.

We now construct a model $E = (Z, \gamma, \theta)$ of $\phi \wedge \psi$ as follows. We put

$$Z = \{(x, y) \in X \times Y \mid n \geq \text{ht}(x) = \text{ht}(y) =: k, \text{Th}_{V_0}^{n-k}(x) = \text{Th}_{V_0}^{n-k}(y)\} \\ \cup \{y \in Y \mid \text{ht}(y) > n\},$$

denoting the first part by Z_0 and the second by Z_1 , and their height- k levels by Z_0^k, Z_1^k, Z^k , respectively. Note that $(x_0, y_0) \in Z_0$.

By construction, states $z \in Z$ have a well-defined height $\text{ht}(z)$. We define a $V_1 \cup V_2$ -valuation θ on Z by

$$Z_0 \ni (x, y) \in \theta(p) \iff \begin{cases} x \in \theta_1(p) & (p \in V_1) \\ y \in \theta_2(p) & (p \in V_2) \end{cases} \\ Z_1 \ni y \in \theta(p) \iff y \in \theta_2(p),$$

which is well-defined by construction of Z_0 . Moreover, we define a coalgebra structure $\gamma : Z \rightarrow \mathbf{FZ}$ such that $\gamma(z) \in \mathbf{FZ}^{k+1}$ for $z \in Z^k$. We put

$$\gamma(y) = \zeta(y) \in \mathbf{FZ}_1 \subseteq \mathbf{FZ}$$

for $y \in Z_1$ (using preservation of inclusions by \mathbf{F}), and on states $(x, y) \in Z_0$ of maximal height n by

$$\gamma(x, y) = \zeta(y).$$

On the rest of Z_0 we define γ by a *coherence* requirement: we define a *pseudo-satisfaction* relation \vDash^0 between

- states in Z of height k and $\mathcal{F}(\Lambda, V_1)$ -formulas of rank at most $n - k$, as well as
- states in Z of height k and $\mathcal{F}(\Lambda, V_2)$ -formulas of rank at most $\text{rk}(\psi) - k$.

by

$$(x, y) \vDash^0 \rho \iff \begin{cases} x \vDash_C \rho & (\rho \in \mathcal{F}(\Lambda, V_1)) \\ y \vDash_D \rho & (\rho \in \mathcal{F}(\Lambda, V_2)) \end{cases} \\ y \vDash^0 \rho \iff y \vDash_D \rho.$$

This is well-defined by construction of Z_0 . For a formula $\rho \in \mathcal{F}_n(\Lambda, V_1) \cup \mathcal{F}_{\text{rk}(\psi)}(\Lambda, V_2)$, we then have the *pseudo-extension* $\hat{\rho} \subseteq Z$ defined by

$$\hat{\rho} = \{z \in Z \mid \text{rk}(\rho) \leq n - \text{ht}(z), z \vDash^0 \rho\}.$$

Then we say that γ is *coherent* if for $\heartsuit \rho \in \mathcal{F}_{n-k}(\Lambda, V_1) \cup \mathcal{F}_{n-k}(\Lambda, V_2)$, $k \leq n$, and $(x, y) \in Z_0^k$,

$$\gamma(x, y) \vDash \heartsuit(\hat{\rho} \cap Z_0^{k+1}) \iff (x, y) \vDash^0 \heartsuit \rho.$$

21:22 Uniform Interpolation in Monotone Coalgebraic Modal Logic

For $i = 0, 1, 2$, let \mathfrak{A}_i be the Boolean subalgebra of $\mathcal{P}Z_0^{k+1}$ of sets of the form $\hat{\rho} \cap Z_0^{k+1}$ for $\rho \in \mathcal{F}_{n-k-1}(\Lambda, V_i)$. Then, of course, $\mathfrak{A}_1 \cap \mathfrak{A}_2 \subseteq \mathfrak{A}_0$, and by induction, $\mathfrak{A}_1, \mathfrak{A}_2$ are interpolable.

We take W to be the set of variables a_ρ where ρ ranges over $\mathcal{F}_{n-k-1}(\Lambda, V_1 \cup V_2)$, and $W_i = \{a_\rho \mid \rho \in \mathcal{F}_{n-k-1}(\Lambda, V_i)\}$ for $i = 0, 1, 2$. We define $\phi_0 \in \mathbf{Prop}(\Lambda(W_1))$ and $\psi_0 \in \mathbf{Prop}(\Lambda(W_2))$ by

$$\begin{aligned}\phi_0 &= \bigwedge \{ \epsilon \heartsuit a_\rho \mid (x, y) \models^0 \epsilon \heartsuit \rho, \rho \in \mathcal{F}_{n-k-1}(\Lambda, V_1), \epsilon \in \{ \cdot, \neg \} \} \\ \psi_0 &= \bigwedge \{ \epsilon \heartsuit a_\rho \mid (x, y) \models^0 \epsilon \heartsuit \rho, \rho \in \mathcal{F}_{n-k-1}(\Lambda, V_2), \epsilon \in \{ \cdot, \neg \} \}.\end{aligned}$$

We let τ denote the $\mathcal{P}(Z_0^{k+1})$ -valuation given by

$$\tau(a_\rho) = \hat{\rho} \cap Z_0^{k+1}.$$

It suffices to show that $(\phi_0 \wedge \psi_0)\tau$ is τ -satisfiable. Now $\rho \in \mathbf{Prop}(\Lambda(\mathfrak{A}_0))$ be the uniform one-step \mathfrak{A}_0 -interpolant of $\phi_0\tau \in \mathbf{Prop}(\Lambda(\mathfrak{A}_1))$, and let $i(\phi) \in \mathbf{Prop}(\Lambda(W_0))$ such that $i(\phi)\tau = \rho$. Then it suffices to show that $(i(\phi_0) \wedge \psi_0)\tau$ is satisfiable. Now let $Y^k \subseteq Y$ be the set of states of height k in D , and let τ_1 denote the $\mathcal{P}(Y^{k+1})$ -valuation given by

$$\tau_1(a_\rho) = \llbracket \rho \rrbracket_D \cap Y^{k+1}.$$

Then

$$\tau(a_\rho) = \pi_2^{-1}[\tau_1(a_\rho)] \quad \text{for } a_\rho \in W_2$$

by the definition of $\hat{\rho}$, where $\pi_2 : Z_0 \rightarrow Y$ is the second projection, so since $i(\phi_0) \wedge \psi_0 \in \mathbf{Prop}(\Lambda(\mathbf{Prop}(W_2)))$ it suffices to show that $(i(\phi_0) \wedge \psi_0)\tau_1$ is satisfiable. In fact we claim that $\zeta(y) \models (i(\phi_0) \wedge \psi_0)\tau_1$. Since for $\rho \in \mathcal{F}_{n-k-1}(\Lambda, V_2)$, $(x, y) \models^0 \epsilon \heartsuit \rho$ iff $y \models \epsilon \heartsuit \rho$ iff $\zeta(y) \models \epsilon \heartsuit a_\rho \tau_1$, we have $\zeta(y) \models \psi_0 \tau_1$ by definition of ψ_0 . It remains to show that $\zeta(y) \models i(\phi_0)\tau_1$. Let σ be the substitution defined by $\sigma(a_\rho) = \rho$. Then we have

$$\begin{aligned}\zeta(y) \models i(\phi_0)\tau_1 & \\ \iff y \models i(\phi_0)\sigma & \quad (\text{semantics}) \\ \iff x \models i(\phi_0)\sigma & \quad (i(\phi_0)\sigma \in \mathcal{F}_{n-(k-1)}(\Lambda, V_0)) \\ \iff x \models \phi_0\sigma, & \end{aligned}$$

and the last condition holds by construction of ϕ_0 because for $\rho \in \mathcal{F}_{n-k}(\Lambda, V_0)$, $(x, y) \models \epsilon \heartsuit \rho$ iff $x \models \epsilon \heartsuit \rho$. This shows that a coherent γ exists. Then, we have by induction on $\rho \in \mathcal{F}_n(\Lambda, V_1) \cup \mathcal{F}_{\text{rk}(\psi)}(\Lambda, V_2)$ that $z \models_E \rho$ iff $z \models^0 \rho$ for $\text{ht}(z) = k$ and $\text{rk}(\rho) \leq n - k$; so in particular $z_0 = (x_0, y_0) \models \phi$ and $z_0 \models \psi$, as required. \blacktriangleleft