



UvA-DARE (Digital Academic Repository)

Disjunctive bases: normal forms and model theory for modal logics

Enqvist, S.; Venema, Y.

DOI

[10.23638/LMCS-15\(1:30\)2019](https://doi.org/10.23638/LMCS-15(1:30)2019)

Publication date

2019

Document Version

Final published version

Published in

Logical Methods in Computer Science

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Enqvist, S., & Venema, Y. (2019). Disjunctive bases: normal forms and model theory for modal logics. *Logical Methods in Computer Science*, 15(1), [30].
[https://doi.org/10.23638/LMCS-15\(1:30\)2019](https://doi.org/10.23638/LMCS-15(1:30)2019)

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

DISJUNCTIVE BASES: NORMAL FORMS AND MODEL THEORY FOR MODAL LOGICS

SEBASTIAN ENQVIST^a AND YDE VENEMA^b

^a Stockholm University, Department of Philosophy
e-mail address: sebastian.enqvist@philosophy.su.se

^b University of Amsterdam, Institute for Logic, Language and Computation
e-mail address: y.venema@uva.nl

ABSTRACT. We present the concept of a disjunctive basis as a generic framework for normal forms in modal logic based on coalgebra. Disjunctive bases were defined in previous work on completeness for modal fixpoint logics, where they played a central role in the proof of a generic completeness theorem for coalgebraic μ -calculi. Believing the concept has a much wider significance, here we investigate it more thoroughly in its own right. We show that the presence of a disjunctive basis at the “one-step” level entails a number of good properties for a coalgebraic μ -calculus, in particular, a simulation theorem showing that every alternating automaton can be transformed into an equivalent nondeterministic one. Based on this, we prove a Lyndon theorem for the full fixpoint logic, its fixpoint-free fragment and its one-step fragment, a Uniform Interpolation result, for both the full μ -calculus and its fixpoint-free fragment, and a Janin-Walukiewicz-style characterization theorem for the μ -calculus under slightly stronger assumptions.

We also raise the questions, when a disjunctive basis exists, and how disjunctive bases are related to Moss’ coalgebraic “nabla” modalities. Nabla formulas provide disjunctive bases for many coalgebraic modal logics, but there are cases where disjunctive bases give useful normal forms even when nabla formulas fail to do so, our prime example being graded modal logic. We also show that disjunctive bases are preserved by forming sums, products and compositions of coalgebraic modal logics, providing tools for modular construction of modal logics admitting disjunctive bases. Finally, we consider the problem of giving a category-theoretic formulation of disjunctive bases, and provide a partial solution.

1. INTRODUCTION

The topic of this paper connects modal μ -calculi, coalgebra and automata. The connection between the modal μ -calculus, as introduced by Kozen [20], and automata running on infinite objects, is standard [15]. Many of the most fundamental results about the modal μ -calculus have been proved by making use of this connection, including completeness of Kozen’s axiom system [37], and model theoretic results like expressive completeness [19], uniform interpolation and a Lyndon theorem [6].

Key words and phrases: Modal logic, fixpoint logic, automata, coalgebra, graded modal logic, Lyndon theorem, uniform interpolation, expressive completeness.

The standard modal μ -calculus was generalized to a generic, coalgebraic modal μ -calculi [35], of which the modal basis was provided by Moss' original coalgebraic modality [27], now known as the *nabla* modality. From a meta-logical perspective, Moss' nabla logics and their fixpoint extensions are wonderfully well-behaved. For example, a generic completeness theorem for nabla logics by a uniform system of axioms was established [21], and this was recently extended to the fixpoint extension of the finitary Moss logic [11]. Most importantly, the automata corresponding to the fixpoint extension of Moss' finitary nabla logic always enjoy a *simulation theorem*, allowing arbitrary coalgebraic automata to be simulated by *non-deterministic* ones; this goes back to the work of Janin & Walukiewicz on μ -automata [18]. The simulation theorem provides a very strong normal form for these logics, and plays an important role in the proofs of several results for coalgebraic fixpoint logics.

The downside of this approach is that the nabla modality is rather non-standard, and understanding what concrete formulas actually say is not always easy. For this reason, another approach to coalgebraic modal logic has become popular, based on so called *predicate liftings*. This approach, going back to the work of Pattinson [31], provides a much more familiar syntax in concrete applications, but can still be elegantly formulated at the level of generality and abstraction that makes the coalgebraic approach to modal logic attractive in the first place.¹ Coalgebraic μ -calculi have also been developed as extensions of the predicate liftings based languages [3], and the resulting logics are very well behaved: for example, good complexity results were obtained in op. cit. Again, the connection between formulas and automata can be formulated in this setting [13], but a central piece is now missing: so far, no simulation theorem has been established for logics based on predicate liftings. In fact, it is not trivial even to define what a non-deterministic automaton *is* in this setting.

This problem turned up in recent work [8], by ourselves together with Seifan, where we extended our earlier completeness result for Moss-style fixpoint logics [11] to the predicate liftings setting. Our solution was to introduce the concept of a *disjunctive basis*, which formalizes in a compact way the minimal requirements that a collection of predicate liftings Λ must meet in order for the class of corresponding Λ -automata to admit a simulation theorem. Our aim in the present paper is to follow up on this conceptual contribution, which we believe is of much wider significance besides providing a tool to prove completeness results.

Exemplifying this, we shall explore some of the applications of our coalgebraic simulation theorem. Some of these transfer known results for nabla based fixpoint logics to the predicate liftings setting; for example, we show that a linear-size model property holds for our non-deterministic automata (or “disjunctive” automata as we will call them), following [35]. We also show that uniform interpolation results hold for coalgebraic fixpoint logics in the presence of a disjunctive basis, which was proved for the Moss-style languages in [26]. We prove a Lyndon theorem for coalgebraic fixpoint logics, generalizing a result for the standard modal μ -calculus proved in [6]: a formula is monotone in one of its variables if and only if it is equivalent to one in which the variable appears positively. We also prove an explicitly *one-step* version of this last result, which we believe has some practical interest for modal fixpoint logics: It is used to show that, given an expressively complete set of monotone predicate liftings, its associated μ -calculus has the same expressive power as the full μ -calculus based on the collection of all monotone predicate liftings. We also prove a Janin-Walukiewicz style characterization theorem for coalgebraic μ -calculi admitting slightly stronger form of

¹For a comparison between the two approaches, see [24].

disjunctive bases, continuing the line of work initiated in [10, 9]. Finally, we show that the sum, product and composition of two modal logics that admit a disjunctive bases also admit a disjunctive basis, thus providing a tool for modular construction of logics to which our results apply.

Next to proving these results, we compare the notion of a disjunctive basis to the nabla based approach to coalgebraic fixpoint logics. The connection will be highlighted in Section 9 where we discuss disjunctive predicate liftings via the Yoneda lemma: here the Barr lifting of the ambient functor (on which the semantics of nabla modalities are based) comes into the picture naturally. However, this is not to say that disjunctive bases are just “nablas in disguise”: it is a fundamental concept, formulated specifically to suit the approach to coalgebraic modal logics via predicate liftings, rather than nabla based languages. Furthermore, in cases where these two approaches to coalgebraic logics are not equivalent, disjunctive bases may be available even when nabla formulas fail to provide them. We shall see that there are several concrete and natural examples of this. A particularly interesting specimen is *graded modal logic*, which adds counting modalities to standard modal logic. While we will see that the standard nabla formulas of this language do not provide a disjunctive basis, nevertheless a disjunctive basis for graded modal logic does exist.

Conference version: This article is an extended version of a paper [12] presented in Ljubljana at the 2017 conference on Algebra and Coalgebra in Computer Science (CALCO 2017). Besides the material presented there, we have provided more detailed proofs, new examples and some additional results: our Janin-Walukiewicz theorem and the results on modular construction of logics admitting disjunctive bases via products, sums and compositions.

2. PRELIMINARIES

2.1. Basics of coalgebraic logic. We assume that the reader is familiar with coalgebra, coalgebraic modal logic and the basic theory of automata operating on infinite objects. The aim of this section is merely to fix some definitions and notations on notions related to coalgebraic modal logic. In an appendix to this paper we provide some basic definitions related to the theory of infinite parity games — we shall need such games for our results on coalgebraic modal *fixpoint* logics.

First of all, throughout this paper we will use the letter T to denote an arbitrary *set functor*, that is, a covariant endofunctor on the category Set having sets as objects and functions as arrows. For notational convenience we sometimes assume that T preserves inclusions; our arguments can easily be adapted to the more general case. Functors of coalgebraic interest include the identity functor Id , the *powerset functor* P , the monotone neighborhood functor M and the (finitary) bag functor B (where $\mathsf{B}S$ is the collection of *weight functions* $\sigma : S \rightarrow \omega$ with finite support). We also need the contravariant powerset functor $\check{\mathsf{P}}$.

Definition 2.1. A T -coalgebra is a pair $\mathbb{S} = (S, \sigma)$ where S is a set of objects called *states* or *points* and $\sigma : S \rightarrow \mathsf{T}S$ is the *transition* or *coalgebra map* of \mathbb{S} . A *pointed* T -coalgebra is a pair (\mathbb{S}, s) consisting of a T -coalgebra and a state $s \in S$. We call a function $f : S' \rightarrow S$ a *coalgebra homomorphism* from (S', σ') to (S, σ) if $\sigma \circ f = \mathsf{T}f \circ \sigma'$, and write $(S', \sigma') \rightrightarrows (S, \sigma)$ if there is such a coalgebra morphism mapping s' to s .

Throughout the paper we fix an (unnamed) countable supply of propositional letters (or variables), of which we often single out a finite subset X .

Definition 2.2. With X a set of proposition letters, a T -model over X is a pair (\mathbb{S}, V) consisting of a T -coalgebra $\mathbb{S} = (S, \sigma)$ and a X -valuation V on S , that is, a function $V : X \rightarrow \mathsf{PS}$. The *marking* associated with V is the transpose map $V^\flat : S \rightarrow \mathsf{PX}$ given by $V^\flat(s) := \{p \in X \mid s \in V(p)\}$. Thus the pair (\mathbb{S}, V) induces a T_X -coalgebra $(S, (V^\flat, \sigma))$, where T_X is the set functor $\mathsf{PX} \times T$.

Given T -models \mathbb{S}, \mathbb{S}' , a map $f : \mathbb{S} \rightarrow \mathbb{S}'$ is called a T -model homomorphism if it is a T_X -coalgebra homomorphism for the induced T_X -coalgebras, i.e., it is a T -coalgebra homomorphism that preserves the truth values of all propositional variables. Pointed T -models (\mathbb{S}, s) and (\mathbb{S}', s') are said to be *behaviorally equivalent*, written $(\mathbb{S}, s) \simeq (\mathbb{S}', s')$, if there is a pointed T -model (\mathbb{S}'', s'') and T -model homomorphisms $f : \mathbb{S} \rightarrow \mathbb{S}''$ and $f' : \mathbb{S}' \rightarrow \mathbb{S}''$ such that $f(s) = f'(s')$.

We will mainly follow the approach in coalgebraic modal logic where modalities are associated (or even identified) with finitary predicate liftings.

Definition 2.3. A *predicate lifting* of arity n is a natural transformation $\lambda : \check{\mathsf{P}}^n \Rightarrow \check{\mathsf{P}}T$. Such a predicate lifting is *monotone* if for every set S , the map $\lambda_S : (\mathsf{PS})^n \rightarrow \mathsf{PTS}$ preserves the subset order in each coordinate. The induced predicate lifting $\lambda^\partial : \mathsf{P}^n \Rightarrow \mathsf{PT}$, given by $\lambda^\partial_S(X_1, \dots, X_n) := \mathsf{TS} \setminus \lambda_S(S \setminus X_1, \dots, S \setminus X_n)$, is called the (*Boolean*) *dual* of λ .

Definition 2.4. A *monotone modal signature*, or briefly: *signature* for T is a set Λ of monotone predicate liftings for T , which is closed under taking boolean duals.

In this paper we will study coalgebraic modal logic with and without fixpoint operators. Given a signature Λ , the formulas of the *coalgebraic μ -calculus* $\mu\mathsf{ML}_\Lambda$ are given by the following grammar:

$$\varphi ::= p \mid \perp \mid \neg\varphi \mid \varphi_0 \vee \varphi_1 \mid \heartsuit_\lambda(\varphi_1, \dots, \varphi_n) \mid \mu q.\varphi'$$

where p and q are propositional variables, $\lambda \in \Lambda$ has arity n , and the application of the fixpoint operator μq is under the proviso that all occurrences of q in φ' are positive (i.e., under an even number of negations). We let ML_Λ denote the fixpoint-free fragment of $\mu\mathsf{ML}_\Lambda$, i.e., the “basic” coalgebraic modal logic of the signature Λ . We let $\mu\mathsf{ML}_\Lambda(X)$ denote the set of $\mu\mathsf{ML}_\Lambda$ -formulas taking free variables from X , and define the notation $\mathsf{ML}_\Lambda(X)$ similarly.

Formulas of such coalgebraic μ -calculi are interpreted in coalgebraic models, as follows. Let $\mathbb{S} = (S, \sigma, V)$ be a T -model over a set X of proposition letters. By induction on the complexity of formulas, we define a *meaning function* $\llbracket \cdot \rrbracket^\mathbb{S} : \mu\mathsf{ML}_\Lambda(X) \rightarrow \mathsf{PS}$, together with an associated *satisfaction relation* $\Vdash \subseteq S \times \mu\mathsf{ML}_\Lambda(X)$ given by $\mathbb{S}, s \Vdash \varphi$ iff $s \in \llbracket \varphi \rrbracket^\mathbb{S}$. All clauses of this definition are standard; for instance, the one for the modality \heartsuit_λ is given by

$$\mathbb{S}, s \Vdash \heartsuit_\lambda(\varphi_1, \dots, \varphi_n) \text{ if } \sigma(s) \in \lambda_S(\llbracket \varphi_1 \rrbracket^\mathbb{S}, \dots, \llbracket \varphi_n \rrbracket^\mathbb{S}). \quad (2.1)$$

For the least fixpoint operator we apply the standard description of least fixpoints of monotone maps from the Knaster-Tarski theorem and take

$$\llbracket \mu x.\varphi \rrbracket^\mathbb{S} := \bigcap \{U \in \mathsf{PS} \mid \llbracket \varphi \rrbracket^{(S, \sigma, V[x \mapsto U])} \subseteq U\},$$

where $V[x \mapsto U]$ is given by $V[x \mapsto U](x) := U$ while $V[x \mapsto U](p) := V(p)$ for $p \neq x$. A formula φ is said to be *monotone* in a variable p if, for every T -model $\mathbb{S} = (S, \sigma, V)$ and all sets $Z_1 \subseteq Z_2 \subseteq S$, we have $\llbracket \varphi \rrbracket^{(S, \sigma, V[p \mapsto Z_1])} \subseteq \llbracket \varphi \rrbracket^{(S, \sigma, V[p \mapsto Z_2])}$.

2.2. Examples. There are many well-known examples of modal logics that can be presented as coalgebraic modal logics where the modalities correspond to predicate liftings for the relevant functor. We shall not attempt to provide a complete list here, but we provide a few basic examples that will be helpful in what follows.

Next-time modality. Probably the simplest example of a non-trivial modality that can be described as a predicate lifting is the “next-time” operator of linear temporal logic. The natural way to present models of LTL coalgebraically is to take the coalgebraic type functor to be Id , the identity functor on the category of sets. Coalgebras (S, σ) for this functor just provide maps $\sigma : S \rightarrow S$, which can be thought of as providing the “next-state” function for a discrete linear flow of time.

The identity functor has a unary predicate lifting $\circ : \check{\mathbb{P}} \rightarrow \check{\mathbb{P}}$: the identity natural transformation defined by $\circ_S : Z \mapsto Z$ for $Z \subseteq S$. A little thought shows that the evaluation of formulas $\circ\varphi$ in a model turns out as expected:

$$\mathbb{S}, s \Vdash \circ\varphi \text{ iff } \mathbb{S}, \sigma(s) \Vdash \varphi.$$

The lifting \circ is monotone and dual to itself, so $\Sigma_{\text{Id}} = \{\circ\}$ is a modal signature for the identity functor. The language μML_{Λ} induced by this signature is known as the *linear-time μ -calculus*.

Basic modal logic. Kripke frames for modal logic are coalgebras for the covariant powerset functor \mathbb{P} , or rather can be represented equivalently as such: a binary relation $R \subseteq S \times S$ of a frame (S, R) can be identified with the map $R[-] : S \rightarrow \mathbb{P}S$ sending each point in S to its set of R -successors. The usual Kripkean modalities \square and \diamond come out as predicate liftings for \mathbb{P} , by setting $\square_S(Z) = \{Z' \subseteq S \mid Z' \subseteq Z\}$ and $\diamond_S(Z) = \{Z' \subseteq S \mid Z \cap Z' \neq \emptyset\}$. Unfolding the definitions we see that the naturality condition for the box modality \square says that for all $f : S \rightarrow S'$ and $Z \subseteq S$, $Z' \subseteq S'$, we have:

$$f[Z] \subseteq Z' \iff Z \subseteq f^{-1}[Z'].$$

This is the familiar adjunction between direct and inverse image. The semantics of modal formulas comes out as expected: given a Kripke model (S, R, V) represented as a \mathbb{P} -model \mathbb{S} , we have $\mathbb{S}, s \Vdash \square\varphi$ if and only if $\mathbb{S}, t \Vdash \varphi$ for every R -successor t of s . For the signature $\Sigma_{\mathbb{P}} = \{\square, \diamond\}$, the language $\mu\text{ML}_{\Sigma_{\mathbb{P}}}$ is the (standard) modal μ -calculus.

Monotone modal logic. Monotone modal logic generalizes normal modal logic by dropping the constraint that the box modality should commute with conjunctions: $\square(\varphi \wedge \psi) \iff \square\varphi \wedge \square\psi$. Semantics for monotone modal logic is given by coalgebras for the monotone neighborhood functor \mathbb{M} , which is the sub-functor of the functor $\check{\mathbb{P}} \circ \check{\mathbb{P}}$ defined by setting:

$$\mathbb{M}S = \{F \subseteq \check{\mathbb{P}}S \mid Z \in F \ \& \ Z \subseteq Z' \Rightarrow Z' \in F\}$$

The lifting $\square : \check{\mathbb{P}} \rightarrow \check{\mathbb{P}} \circ \mathbb{M}$ is defined by setting $F \in \square_S(Z)$ iff $Z \in F$. The dual lifting \diamond is defined by setting $F \in \diamond_S(Z)$ iff, for all $Z' \in F$, $Z \cap Z' \neq \emptyset$.

Monotone modal logic can be seen as a “base logic” for modal logics without distribution of the box over conjunctions, with varying interpretations, much as the modal logic \mathbf{K} can be seen as the most basic normal modal logic. Examples where such modalities appear are alternating-time temporal logic [1] and Parikh’s dynamic game logic [30]. For $\Sigma_{\mathbb{M}} = \{\square, \diamond\}$, the language $\mu\text{ML}_{\Sigma_{\mathbb{M}}}$ is known as the *monotone μ -calculus*. It stands in a similar relationship to alternating-time logic and game logic as the modal μ -calculus does to CTL and PDL.

Graded modal logic. Graded modal logic extends basic modal logic with *counting modalities* $\diamond^k\varphi$ and $\square^k\varphi$. These modalities, interpreted on Kripke models, are interpreted as “at least k successors satisfy φ ” and “there are less than k successors that do not satisfy φ ”. It is often convenient to use a slight generalization of Kripke semantics, where successors of a state are assigned “weights” from ω . Such models are based on \mathbf{B} -coalgebras, where the “bags” functor \mathbf{B} assigns to a set S the set $\mathbf{B}S$ of maps $f : S \rightarrow \omega$ such that $f(s) = 0$ for all but finitely many $s \in S$. Given a map $h : S \rightarrow S'$ and $f \in \mathbf{B}S$, the map $f' = \mathbf{B}h(f)$ is defined by:

$$f'(s') = \sum_{h(s)=s'} f(s)$$

The functor \mathbf{B} comes with an infinite supply of predicate liftings \underline{k} and \overline{k} — one pair for each $k \in \omega$ — given by:

$$\begin{aligned} \underline{k}_S &: U \mapsto \{\sigma \in \mathbf{B}S \mid \sum_{u \in U} \sigma(u) \geq k\} \\ \overline{k}_S &: U \mapsto \{\sigma \in \mathbf{B}S \mid \sum_{u \notin U} \sigma(u) < k\}. \end{aligned}$$

Over Kripke models, which can be identified with \mathbf{B} -models (S, σ, V) in which $\sigma(s)(v) \in \{0, 1\}$ for all $s, v \in S$, it is not hard to see that the formulas $\diamond^k\varphi$ and $\square^k\varphi$ get their expected meanings. For $\Sigma_{\mathbf{B}} = \{\underline{k} \mid k \in \omega\} \cup \{\overline{k} \mid k \in \omega\}$, the language $\mu\mathbf{ML}_{\Sigma_{\mathbf{B}}}$ is known as the *graded μ -calculus*.

2.3. One-step logic and one-step models. A pivotal role in our approach is filled by the *one-step versions* of coalgebraic logics. The one-step perspective on coalgebraic modal logics, developed by a number of authors over several papers [4, 31, 33, 34], has been key to proving some central results about such logics and their fixpoint extensions. In particular, it is instrumental to the theory of coalgebraic automata. Indeed our main contribution here — the concept of disjunctive bases — takes place on the one-step level. We do not assume that the reader is familiar with the framework of one-step logics, and give a self-contained introduction here.

We begin with a formal definition.

Definition 2.5. Given a signature Λ and a set A of variables, we define the set $\mathbf{Bool}(A)$ of *boolean formulas* over A and the set $\mathbf{1ML}_{\Lambda}(A)$ of *one-step Λ -formulas* over A , by the following grammars:

$$\begin{aligned} \mathbf{Bool}(A) \ni \pi &::= a \mid \perp \mid \top \mid \pi \vee \pi \mid \pi \wedge \pi \mid \neg\pi \\ \mathbf{1ML}_{\Lambda}(A) \ni \alpha &::= \heartsuit_{\lambda}\overline{\pi} \mid \perp \mid \top \mid \alpha \vee \alpha \mid \alpha \wedge \alpha \mid \neg\alpha \end{aligned}$$

where $a \in A$, $\lambda \in \Lambda$ and $\overline{\pi} = (\pi_1, \dots, \pi_n)$ is a tuple of formulas in $\mathbf{Bool}(A)$ of the same length as the arity of λ . We will denote the positive (negation-free) fragments of $\mathbf{Bool}(A)$ and $\mathbf{1ML}_{\Lambda}(A)$ as, respectively, $\mathbf{Latt}(A)$ and $\mathbf{1ML}_{\Lambda}^+(A)$:

$$\begin{aligned} \mathbf{Latt}(A) \ni \pi &::= a \mid \perp \mid \top \mid \pi \vee \pi \mid \pi \wedge \pi \\ \mathbf{1ML}_{\Lambda}^+(A) \ni \alpha &::= \heartsuit_{\lambda}\overline{\pi} \mid \perp \mid \top \mid \alpha \vee \alpha \mid \alpha \wedge \alpha \end{aligned}$$

We shall often make use of substitutions: given a finite set A , let $\vee_A : \mathbf{P}A \rightarrow \mathbf{Latt}(A)$ be the map sending B to $\bigvee B$, and let $\wedge_A : \mathbf{P}A \rightarrow \mathbf{Latt}(A)$ be the map sending B to $\bigwedge B$, and given sets A, B let $\theta_{A,B} : A \times B \rightarrow \mathbf{Bool}(A \cup B)$ be defined by mapping (a, b) to $a \wedge b$.

We need a number of properties of modal signatures, formulated in terms of the corresponding one-step logics. The first is fairly standard: expressive completeness of a modal signature Λ means that every modality that makes sense for the functor — formally, every predicate lifting — can be expressed in terms of modalities in Λ .

Definition 2.6. A monotone modal signature Λ for \mathbb{T} is *expressively complete* if, for every n -place predicate lifting λ (not necessarily in Λ) and for all variables a_1, \dots, a_n there is a formula $\alpha \in \mathbf{1ML}_\Lambda(\{a_1, \dots, a_n\})$ which is equivalent to $\heartsuit_\lambda \bar{a}$.

We will also be interested in the following variant of expressive completeness:

Definition 2.7. We say that Λ is *Lyndon complete* if, for every *monotone* n -place predicate lifting λ and variables a_1, \dots, a_n , there is a *positive* formula $\alpha \in \mathbf{1ML}_\Lambda^+(\{a_1, \dots, a_n\})$ equivalent to $\heartsuit_\lambda \bar{a}$.

We now turn to the semantics of one-step formulas, using so-called *one-step models*.

Definition 2.8. A *one-step \mathbb{T} -frame* is a pair (S, σ) with $\sigma \in \mathbb{T}S$, i.e., an object in the category $\mathcal{E}(\mathbb{T})$ of elements of \mathbb{T} . Similarly a *one-step \mathbb{T} -model* over a set A of variables is a triple (S, σ, m) such that (S, σ) is a one-step \mathbb{T} -frame and $m : S \rightarrow \mathbf{PA}$ is an A -marking on S .

A *morphism* $f : (S, \sigma) \rightarrow (S', \sigma')$ is a morphism in $\mathcal{E}(\mathbb{T})$, that is, a map from S to S' such that $\mathbb{T}f(\sigma) = \sigma'$. The map f is said to be a *morphism of one-step models* $f : (S, \sigma, m) \rightarrow (S', \sigma', m')$ if, in addition, $m = m' \circ f$.

Given a one-step model (S, σ, m) , we define the *0-step interpretation* $\llbracket \pi \rrbracket_m^0 \subseteq S$ of $\pi \in \mathbf{Bool}(A)$ by the obvious induction: $\llbracket a \rrbracket_m^0 := \{v \in S \mid a \in m(v)\}$, $\llbracket \top \rrbracket_m^0 := S$, $\llbracket \perp \rrbracket_m^0 := \emptyset$, while we use standard clauses for \wedge, \vee and \neg . Similarly, the *one-step interpretation* $\llbracket \alpha \rrbracket_m^1$ of $\alpha \in \mathbf{1ML}_\Lambda(A)$ is defined as a subset of $\mathbb{T}S$, with $\llbracket \heartsuit_\lambda(\pi_1, \dots, \pi_n) \rrbracket_m^1 := \lambda_S(\llbracket \pi_1 \rrbracket_m^0, \dots, \llbracket \pi_n \rrbracket_m^0)$, and again standard clauses apply to $\perp, \top, \wedge, \vee$ and \neg . Given a one-step model (S, σ, m) , we write $S, \sigma, m \Vdash^1 \alpha$ for $\sigma \in \llbracket \alpha \rrbracket_m^1$. Notions like one-step satisfiability, validity and equivalence are defined and denoted in the obvious way; in particular, we use \equiv^1 to denote the equivalence of one-step formulas.

For future reference we mention the following two results, the first of which states that the truth of one-step formulas is invariant under one-step morphisms.

Proposition 2.1. *Let $f : (S', \sigma', m') \rightarrow (S, \sigma, m)$ be a morphism of one-step models over A . Then for every formula $\alpha \in \mathbf{1ML}_\Lambda(A)$ we have*

$$S', \sigma', m' \Vdash^1 \alpha \text{ iff } S, \sigma, m \Vdash^1 \alpha.$$

The second proposition is a standard observation about the semantic counterpart of the syntactic notion of substitution.

Proposition 2.2. *Let (S, σ, m) be a one-step model over A , and let $\sigma : B \rightarrow \mathbf{Bool}(A)$ be a substitution. Then for every formula $\alpha \in \mathbf{1ML}_\Lambda(B)$ we have*

$$S, \sigma, m_\sigma \Vdash^1 \alpha \text{ iff } S, \sigma, m \Vdash^1 \alpha[\sigma],$$

where m_σ is the B -marking given by $m_\sigma(b) := \llbracket \sigma_b \rrbracket_m^0$.

2.4. Graph games. For readers unfamiliar with the theory of infinite games, we provide some of the basic definitions here, referring to [15] for a survey.

Definition 2.9. A *board game* is a tuple $\mathbb{G} = (G_{\exists}, G_{\forall}, E, W)$ where G_{\exists} and G_{\forall} are disjoint sets, and, with $G := G_{\exists} \cup G_{\forall}$ denoting the *board* of the game, the binary relation $E \subseteq G^2$ encodes the moves that are admissible to the respective players, and $W \subseteq G^{\omega}$ denotes the *winning condition* of the game. In a *parity game*, the winning condition is determined by a parity map $\Omega : G \rightarrow \omega$ with finite range, in the sense that the set W_{Ω} is given as the set of G -streams $\rho \in G^{\omega}$ such that the maximum value occurring infinitely often in the stream $(\Omega\rho_i)_{i \in \omega}$ is even.

Elements of G_{\exists} and G_{\forall} are called *positions* for the players \exists and \forall , respectively; given a position p for player $\Pi \in \{\exists, \forall\}$, the set $E[p]$ denotes the set of *moves* that are *legitimate* or *admissible* to Π at p . In case $E[p] = \emptyset$ we say that player Π *gets stuck* at p .

An *initialized board game* is a pair consisting of a board game \mathbb{G} and a *initial* position p , usually denoted as $\mathbb{G}@p$.

Definition 2.10. A *match* of a graph game $\mathbb{G} = (G_{\exists}, G_{\forall}, E, W)$ is nothing but a (finite or infinite) path through the graph (G, E) . Such a match ρ is called *partial* if it is finite and $E[\text{last}\rho] \neq \emptyset$, and *full* otherwise. We let PM_{Π} denote the collection of partial matches ρ ending in a position $\text{last}(\rho) \in G_{\Pi}$, and define $\text{PM}_{\Pi}@p$ as the set of partial matches in PM_{Π} starting at position p .

The *winner* of a full match ρ is determined as follows. If ρ is finite, then by definition one of the two players got stuck at the position $\text{last}(\rho)$, and so this player loses ρ , while the opponent wins. If ρ is infinite, we declare its winner to be \exists if $\rho \in W$, and \forall otherwise.

Definition 2.11. A *strategy* for a player $\Pi \in \{\exists, \forall\}$ is a map $\chi : \text{PM}_{\Pi} \rightarrow G$. A strategy is *positional* if it only depends on the last position of a partial match, i.e., if $\chi(\rho) = \chi(\rho')$ whenever $\text{last}(\rho) = \text{last}(\rho')$; such a strategy can and will be presented as a map $\chi : G_{\Pi} \rightarrow G$.

A match $\rho = (p_i)_{i < \kappa}$ is *guided* by a Π -strategy χ if $\chi(p_0 p_1 \dots p_{n-1}) = p_n$ for all $n < \kappa$ such that $p_0 \dots p_{n-1} \in \text{PM}_{\Pi}$ (that is, $p_{n-1} \in G_{\Pi}$). Given a strategy f , we say that a position p is *f-reachable* if p occurs on some f -guided partial match. A Π -strategy χ is *legitimate* in $\mathbb{G}@p$ if the moves that it prescribes to χ -guided partial matches in $\text{PM}_{\Pi}@p$ are always admissible to Π , and *winning for Π* in $\mathbb{G}@p$ if in addition all χ -guided full matches starting at p are won by Π .

A position p is a *winning position* for player $\Pi \in \{\exists, \forall\}$ if Π has a winning strategy in the game $\mathbb{G}@p$; the set of these positions is denoted as Win_{Π} . The game $\mathbb{G} = (G_{\exists}, G_{\forall}, E, W)$ is *determined* if every position is winning for either \exists or \forall .

When defining a strategy χ for one of the players in a board game, we will often confine ourselves to defining χ for partial matches that are themselves guided by χ . The following fact, independently due to Emerson & Jutla [7] and Mostowski [28], will be quite useful to us.

Fact 2.12 (Positional Determinacy). Let $\mathbb{G} = (G_{\exists}, G_{\forall}, E, W)$ be a graph game. If W is given by a parity condition, then \mathbb{G} is determined, and both players have positional winning strategies.

2.5. Automata. Given a state s in a coalgebra $\sigma : S \rightarrow \mathbb{T}S$, consider the one-step frame $(S, \sigma(s))$ — this is the local “window” into the structure of the coalgebra that is directly visible from s . The main function of one-step models for our purposes here is to provide a neat framework for *automata* running on \mathbb{T} -models. The idea is the following: at any stage in the run of an automaton \mathbb{A} on a model $\mathbb{S} = (S, \sigma, V)$, the automaton reads some point s in \mathbb{S} and takes the set $V^b(s) \subseteq \mathbf{X}$, consisting of those propositional variables that are true at s , as input. Next the automaton decides whether to continue the computation or to reject. To decide this, the automaton checks the directly visible part of the model \mathbb{S} at s , modelled as the one-step frame $\sigma(s)$, and looks for an admissible way to continue the computation one step further. The ways in which the run may continue are constrained by a one-step formula α , which depends on the current state a of the automaton and the last input $V^b(s)$ read, and is built up using states of the automaton as variables. The run continues if a marking m can be found that makes α true in the one-step model $(S, \sigma(s), m)$, and the marking m then determines which states in \mathbb{S} may be visited in the next stage of the run, and the possible next states of the automaton.

Definition 2.13. A (Λ, \mathbf{X}) -*automaton*, or more broadly, a *coalgebra automaton*, is a quadruple (A, Θ, Ω, a_I) where A is a finite set of *states*, with *initial state* $a_I \in A$, $\Theta : A \times \mathbf{P}\mathbf{X} \rightarrow \mathbf{1ML}_\Lambda^+(A)$ is the *transition map* and $\Omega : A \rightarrow \omega$ is the *priority map* of \mathbb{A} .

The semantics of such an automaton is given in terms of a two-player infinite parity game: With $\mathbb{S} = (S, \sigma, V)$ a \mathbb{T} -model over a set $\mathbf{Y} \supseteq \mathbf{X}$, the *acceptance game* $\mathcal{A}(\mathbb{A}, \mathbb{S})$ is the parity game given by the table below.

Position	Player	Admissible moves	Priority
$(a, s) \in A \times S$	\exists	$\{m : S \rightarrow \mathbf{P}A \mid (S, \sigma(s), m) \Vdash^1 \Theta(a, \mathbf{X} \cap V^b(s))\}$	$\Omega(a)$
$m : S \rightarrow \mathbf{P}A$	\forall	$\{(b, t) \mid b \in m(t)\}$	0

We say that \mathbb{A} *accepts* the pointed \mathbb{T} -model (\mathbb{S}, s) , notation: $\mathbb{S}, s \Vdash \mathbb{A}$, if (a_I, s) is a winning position for \exists in the acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{S})$.

The connection with coalgebraic modal logic is given by the following result.

Fact 2.14 [13]. There are effective constructions transforming a formula in $\mu\mathbf{ML}_\Lambda(\mathbf{X})$ into an equivalent (Λ, \mathbf{X}) -automaton, and vice versa.

3. DISJUNCTIVE FORMULAS AND DISJUNCTIVE BASES

3.1. Disjunctive formulas. In this section, we present the main conceptual contribution of the paper, and define disjunctive bases. We then consider a number of examples.

As a first step, we begin by presenting disjunctive *formulas*, originally introduced in [8], as a class of one-step formulas for a given modal signature characterized by a model-theoretic property, expressed in terms of the one-step semantics.

Definition 3.1. A one-step formula $\alpha \in \mathbf{1ML}_\Lambda^+(A)$ is called *disjunctive* if for every one-step model (S, σ, m) such that $S, \sigma, m \Vdash^1 \alpha$ there is a one-step frame (S', σ') together with a one-step frame morphism $f : (S', \sigma') \rightarrow (S, \sigma)$ and a marking $m' : S' \rightarrow \mathbf{P}A$, such that:

- (1) $S', \sigma', m' \Vdash^1 \alpha$;
- (2) $m'(s') \subseteq m(f(s'))$, for all $s' \in S'$;
- (3) $|m'(s')| \leq 1$, for all $s' \in S'$.

We sometimes refer to the one-step frame (S', σ') together with the map f as a *cover* of (S, σ) , and to the one-step model (S', σ', m') together with the map f as a *dividing cover* of (S, σ, m) for α .

The intuition behind disjunctive formulas is that, in a certain sense, they never “force” two distinct propositional variables to be true together, i.e. any one-step model for a disjunctive formula δ in $\mathbf{1ML}_\Lambda^+(A)$ can be transformed into one in which every point satisfies at most one propositional variable from A . Moreover, “transformed into” here does not just mean “replaced by”: we cannot arbitrarily change the one-step model, the output of the construction must be closely related to the one-step model that we started with.

A trivial example of a disjunctive formula is $\bigcirc a$ for $a \in A$, where we recall that \bigcirc was the next-time modality viewed as a predicate lifting for the identity functor Id . A one-step model for this functor is a triple S, s, m consisting of a set S , an element $s \in S$ and a marking $m : S \rightarrow \text{PA}$. Then $S, s, m \Vdash^{-1} \bigcirc a$ if, and only if, $a \in m(s)$. But then, no elements in S besides s are relevant to the evaluation of $\bigcirc a$, and for s we can just forget about all other variables: set $m'(s) = \{a\}$ and $m'(v) = \emptyset$ for all $s \in S \setminus \{v\}$. We have $S, s, m' \Vdash^{-1} \bigcirc a$, $m'(v) \subseteq m(v)$ and $|m'(v)| \leq 1$ for all $v \in S$.

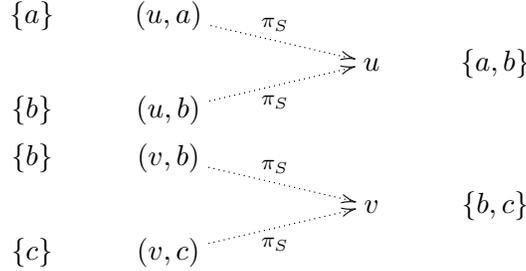
For an example of a one-step formula that is *not* disjunctive, consider $\Box a \wedge \Diamond b$ (where \Diamond and \Box are the standard modalities for Kripke structures, i.e., coalgebras for the power set functor P). Observe that a one-step model for this functor is a triple (S, σ, m) with $\sigma \subseteq S$. It should be obvious that for the formula $\Box a \wedge \Diamond b$ to hold at such a structure, σ needs to have an element s where a holds, while at the same time *every* element of σ , including s , must satisfy b . There is no escape here: we can only have $(S, \sigma, m) \Vdash^{-1} \Box a \wedge \Diamond b$ if there is an element s making *both* a and b true.

This is very different if we consider the typical disjunctive formulas for basic modal logic, which are of the form:

$$\Diamond a_1 \wedge \cdots \wedge \Diamond a_n \wedge \Box(a_1 \vee \cdots \vee a_n)$$

This is often abbreviated as $\nabla\{a_1, \dots, a_n\}$. The operator ∇ is known as the *cover modality*, or sometimes “nabla modality”. The formula ∇B says about a one-step model (S, σ, m) for P , that the following “back-and-forth” conditions hold: for all $s \in \sigma$ there is some $a \in B$ with $a \in m(s)$, and conversely, for every $a \in B$ there is some $s \in \sigma$ with $a \in m(s)$. These formulas are indeed disjunctive, but less trivially so than the next-time formulas. For example, consider the formula $\nabla\{a, b, c\}$, which is true in a one-step model (S, s, m) with $S = \{u, v\}$, $m(u) = \{a, b\}$ and $m(v) = \{b, c\}$. But there is no way to simply shrink the marking m to a dividing marking m' so that $(S, s, m) \Vdash^{-1} \nabla\{a, b, c\}$: there are too many pigeons and too few pigeon holes — so the obvious solution is to make more pigeon holes! One way to do this is to “split” the points in $\{u, v\}$ so that we make room for each variable to be witnessed at a separate point. More formally, let $S' = \{(u, a), (u, b), (v, b), (v, c)\}$ and define $m' : S' \rightarrow \{a, b, c\}$ via the projection on $\{a, b, c\}$, i.e., $m'(u, a) = \{a\}$, etc. This one-step model satisfies $\nabla\{a, b, c\}$ and it has the obvious covering map h being the projection π_S on S .

The cover modality was arguably the starting point of coalgebraic modal logic. In the seminal paper [27], Moss defined nabla modalities for all functors that preserve weak pullback squares, generalizing the cover modality. The idea is to apply the coalgebraic type functor T to sets of formulas Ψ and form new formulas from objects Γ in $\text{T}\Psi$. That is, $\nabla\Gamma$ counts as a formula if Ψ is a set of formulas and $\Gamma \in \text{T}\Psi$. The semantics is given in terms of the “Barr extension” $\bar{\text{T}}$ of the functor T , which is a relation lifting defined for $R \subseteq X \times X$

Figure 1: The cover $h : (S', m') \rightarrow (S, m)$

by setting:

$$\bar{\top}R = \{(\xi, \alpha) \in \top X \times \top A \mid \exists \rho \in \top R : \top \pi_X(\rho) = \xi \ \& \ \top \pi_A(\rho) = \alpha\}$$

We then evaluate the nabla modality by applying this relation lifting to the satisfaction relation:

$$\mathbb{S}, s \Vdash \nabla \Gamma \text{ iff } (\sigma(s), \Gamma) \in \bar{\top}(\Vdash).$$

In terms of one-step formulas, we would count $\nabla \Gamma$ as a one-step formula with variables in A , for all $\Gamma \in \top \text{Bool}(A)$. In particular, formulas of the form $\nabla \Gamma$ for $\Gamma \in \top A$ count as one-step formulas.

These formulas are in fact disjunctive:

Proposition 3.1. *Let \top be a set functor that preserves weak pullbacks. Then every formula of the form $\nabla \Gamma$, where $\Gamma \in \top A$, is disjunctive.*

To see this, suppose that a one-step model X, ξ, m satisfies the formula $\nabla \Gamma$. This means that $(\xi, \Gamma) \in \bar{\top}R$, where $R = \Vdash^0 \cap (X \times A)$ is defined by $(u, a) \in R$ iff $a \in m(u)$. By definition of the Barr extension, this means we find some $\rho \in \top R$ such that $\top \pi_A(\rho) = \Gamma$ and $\top \pi_X(\rho) = \xi$. But then we have our covering one-step model right in front of us already: the triple $(R, \rho, \eta \circ \pi_A)$ is a one-step model where $\eta : A \rightarrow \mathbb{P}A$ is the singleton map $a \mapsto \{a\}$, i.e., the unit of the powerset monad. The covering map is just the projection $\pi_X : R \rightarrow X$. It is a simple exercise to check that $(R, \rho, \eta \circ \pi_A) \Vdash^1 \nabla \Gamma$, and clearly $|\eta(\pi_A(u, a))| = 1$ and $\eta(\pi_A(u, a)) \subseteq m(u)$ for all $(u, a) \in R$. So we see that disjunctiveness of the nabla modalities is not incidental: it is hardwired into their semantics.

3.2. Disjunctive bases. We now turn to a discussion of the key notion of this paper, viz., the disjunctive bases originating with [8]. For their definition, we recall the substitutions $\wedge_A : B \mapsto \bigwedge B$ and $\theta_{A,B} : (a, b) \mapsto a \wedge b$ defined in the Preliminaries. It will also be convenient to introduce the abbreviation

$$A \boxtimes B := (A \times B) \cup A \cup B$$

for any two sets A, B .

Definition 3.2. Let \mathbb{D} be an assignment of a set of positive one-step formulas $\mathbb{D}(A) \subseteq \mathbf{1ML}_\Lambda^+(A)$ for all sets of variables A . Then \mathbb{D} is called a *disjunctive basis* for Λ if each formula in $\mathbb{D}(A)$ is disjunctive, and the following conditions hold:

- (1) $\mathsf{D}(A)$ contains \top and is closed under finite disjunctions (in particular, also $\perp = \bigvee \emptyset \in \mathsf{D}(A)$).
- (2) D is *distributive over Λ* : for every one-step formula of the form $\heartsuit_{\lambda}\bar{\pi}$ in $\mathbf{1ML}_{\Lambda}^{+}(A)$ there is a formula $\delta \in \mathsf{D}(\mathsf{P}(A))$ such that $\heartsuit_{\lambda}\bar{\pi} \equiv^1 \delta[\wedge_A]$.
- (3) D *admits a binary distributive law*: for any two formulas $\alpha \in \mathsf{D}(A)$ and $\beta \in \mathsf{D}(B)$, there is a formula $\gamma \in \mathsf{D}(A \otimes B)$ such that $\alpha \wedge \beta \equiv^1 \gamma[\theta_{A,B}]$.

As a stronger version, we say that D is a *uniform disjunctive basis* if each one-step frame (S, σ) has a single cover $f : (S', \sigma') \rightarrow (S, \sigma)$ which works for all D -formulas, in the sense that for each $\delta \in \mathsf{D}(A)$ and each marking $m : S \rightarrow \mathsf{P}A$ with $S, \sigma, m \Vdash^1 \delta$ there is a marking $m' : S' \rightarrow \mathsf{P}A$ such that (S', σ', m') together with the map f provides a dividing cover of (S, σ, m) for δ .

The key property of disjunctive bases is captured by the following normal form theorem, which is easy to derive from the definition.

Proposition 3.2. *Suppose D is a disjunctive basis for Λ . Then for every one-step formula $\alpha \in \mathbf{1ML}_{\Lambda}^{+}(A)$ there is a formula $\delta \in \mathsf{D}(\mathsf{P}(A))$ such that $\alpha \equiv^1 \delta[\wedge_A]$.*

Note that the requirement $\top \in \mathsf{D}(A)$ is needed for this proposition to hold: we have $\top \in \mathbf{1ML}_{\Lambda}^{+}(A)$, so for the proposition to be true there has to be a formula $\delta \in \mathsf{D}(\mathsf{P}(A))$ such that $\top \equiv^1 \delta[\wedge_A]$. Since $\top \in \mathsf{D}(\mathsf{P}(A))$ we can take $\delta = \top$ — without having $\top \in \mathsf{D}(\mathsf{P}(A))$ there would be no guarantee in general that an appropriate formula δ can be found.

We have seen that nabla formulas are disjunctive, and we shall soon see that they provide disjunctive bases for many modal signatures. In order for the approach of constructing disjunctive bases via nabla formulas to work however, it is necessary that the ambient functor T preserves weak pullbacks. This is one sense in which disjunctive bases generalize nabla formulas: we shall soon see an example of a functor which does not preserve weak pullbacks, but does have a signature admitting a disjunctive basis. However, we want to stress that the main motivation behind disjunctive bases is not to “get by without weak pullback preservation” — in fact, most natural examples of disjunctive bases we can think of apply to functors that do preserve weak pullbacks. Rather, the point is that the existence of a disjunctive basis is a property of *modal signatures*, not of functors. The choice of a modal signature for a functor is not arbitrary, it can substantially affect the properties of the modal languages that we get. In particular, even for functors that do preserve weak pullbacks, we can have natural choices of modal signatures for which the nabla formulas associated with the functor do not provide a disjunctive basis. Typically, this will happen when the modal signature is (a) infinite, and (b) not expressively complete. An example of this situation that will be given special attention is graded modal logic, to be presented in the next subsection. But there are much simpler cases, as we shall see.

3.3. Examples. In this subsection we present a number of examples of modal signatures admitting disjunctive bases, and one example of a modal signature that provably does not admit a disjunctive basis.

3.3.1. Disjunctive bases for weak pullback preserving functors. In the previous section we noted that disjunctive formulas generalize the Moss modalities. In many interesting cases this suffices to find a disjunctive basis.

Proposition 3.3. *Let Λ be a signature for a weak-pullback preserving functor \mathbb{T} . If Λ is Lyndon complete, then it admits a disjunctive basis.*

Proof. We shall use the (infinitary) nabla based logic for \mathbb{T} as an auxiliary language. Let $1\text{ML}_{\nabla}^{\infty}(A)$ be defined by the grammar:

$$\alpha ::= \top \mid \bigvee \Phi \mid \nabla \Gamma$$

where $\Phi \subseteq 1\text{ML}_{\nabla}^{\infty}(A)$ and $\Gamma \in \mathbb{T}A$. Semantics of formulas in $1\text{ML}_{\nabla}^{\infty}(A)$ in a one-step model is defined by the obvious recursion, where the interpretation of $\nabla \Gamma$ is as before.

We first claim that for every formula $\alpha \in 1\text{ML}_{\nabla}^{\infty}(A)$, where $A = \{a_1, \dots, a_n\}$, there is a monotone n -place predicate lifting λ for \mathbb{T} such that α is equivalent to the formula $\heartsuit_{\lambda}(a_1, \dots, a_n)$. This follows from results in [25], together with the easy observation that (monotone) predicate liftings are closed under arbitrary disjunctions. We can safely assume that $\heartsuit_{\lambda}(a_1, \dots, a_n) \in 1\text{ML}_{\Lambda}^+(A)$ since Λ was assumed to be Lyndon complete.²

With this in mind, let $\text{D}_{\nabla}(A)$ be the set of all finite disjunctions of formulas of the form \top or of the form $\heartsuit_{\lambda}(a_1, \dots, a_n)$, where λ is the n -place predicate lifting associated with some $\alpha \in 1\text{ML}_{\nabla}^{\infty}(A)$. As mentioned, all formulas of the form $\nabla \Gamma$ are disjunctive, and since disjunctivity is closed under taking arbitrary disjunctions, all formulas in $1\text{ML}_{\nabla}^{\infty}(A)$ are disjunctive too — hence all formulas in $\text{D}_{\nabla}(A)$ are disjunctive.

It remains to prove that $\text{D}_{\nabla}(A)$ is a basis for Λ , so we need to show that $\text{D}_{\nabla}(A)$ is distributive over Λ and admits a binary distributive law. For the purpose of proving distributivity over Λ it suffices to show that any formula $\beta \in 1\text{ML}_{\Lambda}^+(A)$ is equivalent to a formula $\alpha[\wedge_A]$, where $\alpha \in 1\text{ML}_{\nabla}^{\infty}(PA)$. In other words we want to prove, for an arbitrary formula $\beta \in 1\text{ML}_{\Lambda}^+(A)$:

$$\beta \equiv^1 \bigvee \{(\nabla \Gamma)[\wedge_A] \mid \Gamma \in \mathbb{T}PA \ \& \ PA, \Gamma, \text{id} \Vdash^1 \beta\}, \quad (3.1)$$

where $\text{id} : B \mapsto B$ denotes the canonical marking on the set PA . The notation here may need a bit of explanation: given $\Gamma \in \mathbb{T}PA$, we can apply the functor \mathbb{T} to the map $\wedge_A : PA \rightarrow \text{Latt}(A)$ and apply this to Γ to obtain $(\mathbb{T}\wedge_A)\Gamma \in \mathbb{T}\text{Latt}(A)$. With this established we define:

$$(\nabla \Gamma)[\wedge_A] := \nabla(\mathbb{T}\wedge_A)\Gamma.$$

For a proof of the left-to-right direction of (3.1), assume that $S, \sigma, m \Vdash^1 \beta$. From this it follows by invariance (Proposition 2.1) that $PA, (\mathbb{T}m)\sigma, \text{id} \Vdash^1 \beta$, so that $\nabla \Gamma$ where $\Gamma := (\mathbb{T}m)\sigma \in \mathbb{T}PA$ provides a candidate disjunct on the right hand side of (3.1). It remains to show that $S, \sigma, m \Vdash^1 \nabla(\mathbb{T}m)\sigma[\wedge_A]$, but this is immediate by definition of the semantics of ∇ .

For the opposite direction of (3.1), let $\Gamma \in \mathbb{T}PA$ be such that $PA, \Gamma, \text{id} \Vdash^1 \beta$. In order to show that β is a one-step semantic consequence of $\nabla \Gamma[\wedge_A]$, let (S, σ, m) be a one-step model such that $S, \sigma, m \Vdash^1 \nabla(\mathbb{T}\wedge_A)\Gamma$. Without loss of generality we may assume that $(S, \sigma, m) = (PA, \Delta, \text{id})$ for some $\Delta \in \mathbb{T}PA$.

By the semantics of ∇ it then follows from $PA, \Delta, \text{id} \Vdash^1 \nabla \Gamma[\wedge_A]$ that $(\Delta, (\mathbb{T}\wedge_A)\Gamma) \in \overline{\mathbb{T}}(\Vdash^0)$. But since $(B, \wedge_A(C)) \in \Vdash^0$ implies that $C \subseteq B$, we easily obtain that $(\Gamma, \Delta) \in \overline{\mathbb{T}}(\subseteq)$. We can now apply the following claim, the proof of which we leave as an exercise:

²Strictly speaking Lyndon completeness only guarantees there exists some formula in $1\text{ML}_{\Lambda}^+(A)$ that is equivalent to $\heartsuit_{\lambda}(a_1, \dots, a_n)$, but taking care with this distinction would only complicate notation.

CLAIM 1 . Let (S, σ, m) and (S', σ', m') be two one-step models, and let $Z \subseteq S \times S'$ be a relation such that $(\sigma, \sigma') \in \overline{T}Z$, and $m(s) \subseteq m'(s')$, for all $(s, s') \in Z$. Then for all $\alpha \in \mathbf{1ML}_\Lambda^+(A)$:

$$S, \sigma, m \Vdash^1 \alpha \text{ implies } S', \sigma', m' \Vdash^1 \alpha.$$

It is easy to see that the claim is applicable to the one-step models (PA, Γ, id) and (PA, Δ, id) , and the relation $Z = \subseteq$: we already showed that $(\Gamma, \Delta) \in \overline{T}(\subseteq)$. Furthermore, if $B \subseteq B'$ then trivially $\text{id}(B) \subseteq \text{id}(B')$. Hence it follows from $PA, \Gamma, \text{id} \Vdash^1 \beta$ that $PA, \Delta, \text{id} \Vdash^1 \beta$.

For the binary distributive law, we leave it to the reader to check that for $\Gamma \in TA$ and $\Gamma' \in TB$, the conjunction $\nabla\Gamma \wedge \nabla\Gamma'$ is equivalent to the possibly infinite disjunction of all formulas of the form $\nabla\Gamma''[\theta_{A,B}]$ such that $\Gamma'' \in T(A \times B)$, $T\pi_A(\Gamma'') = \Gamma$ and $T\pi_B(\Gamma'') = \Gamma'$. With this claim in place, the binary distributive law is easily established. \square

3.3.2. *A simple example.* The simplest non-trivial example of a disjunctive basis we can think of, that is not a special case of the nabla-based approach of the previous paragraph, is the following. Consider the functor $T = \Sigma \times \text{Id}$, where Σ is a countably infinite alphabet. For this functor, T -coalgebras are (up to unfolding) just Σ -streams, or infinite words for which the alphabet is contained in Σ . They can be viewed as triples (S, σ_1, σ_2) where $\sigma_1 : S \rightarrow \Sigma$ and $\sigma_2 : S \rightarrow S$. A simple and reasonably natural modal signature Λ for this functor is the following: we have one nullary modality $!l$ for each $l \in \Sigma$ with the interpretation: $\mathbb{S}, s \Vdash !l$ iff $\sigma_1(s) = l$, and we have a single one-place modality \circ with the interpretation $\mathbb{S}, s \Vdash \circ\varphi$ iff $\mathbb{S}, \sigma_2(s) \Vdash \varphi$. We can easily describe this language as a modal signature Λ for the functor T , and it is not hard to show that Λ has a disjunctive basis consisting of the formulas $!l$ for $l \in \Sigma$ and $\circ a$ for $a \in A$.

The one-step nabla formulas for T , which are of the form $\nabla(l, a)$, with $(S, (l', s), m) \Vdash^1 \nabla(l, a)$ iff $l' = l$ and $s \Vdash^0 a$, can easily be expressed as Λ -formulas by writing $\nabla(l, a) := !l \wedge \circ a$. However, these formulas do not form a disjunctive basis for Λ . To see why, just consider the set of variables $A = \{a\}$. The property of distributivity of disjunctive bases would then require that the formula $\circ a$ should be equivalent to a finite disjunction of formulas of the form:

$$!l \wedge \circ a$$

where $l \in \Sigma$. But the formula $\circ a$ corresponds to an infinite disjunction:

$$\bigvee \{!l \wedge \circ a \mid l \in \Sigma\}$$

It is fairly obvious that this is not equivalent to any finitary disjunction of the required shape. It is not hard to come up with other, similar examples.

3.3.3. *Graded modal logic.* A much more involved example of a modal logic that admits a disjunctive basis which cannot be reduced to nabla formulas, is graded modal logic. The bag functor (defined in Section 2.2) does preserve weak pullbacks, and so its Moss modalities are disjunctive, and so the set of all monotone liftings for \mathbf{B} does admit a disjunctive basis as an instance of Proposition 3.3. Note, however, that the latter proposition does not apply to graded modal logic, since its signature $\Sigma_{\mathbf{B}}$ is not expressively complete; this was essentially shown in [29]. It was observed in [2] that very simple formulas in the one-step language $\mathbf{1ML}_{\Sigma_{\mathbf{B}}}$ are impossible to express in the (finitary) Moss language; consequently, the Moss

modalities for the bag functor are not suitable to provide disjunctive normal forms for graded modal logic. Still, the signature $\Sigma_{\mathbf{B}}$ does have a disjunctive basis.

Definition 3.3. We say that a one-step model for the finite bag functor is *Kripkean* if all states have multiplicity 1. Note that a Kripkean one-step model (S, σ, m) can also be seen as a structure (in the sense of standard first-order model theory) for a first-order signature consisting of a monadic predicate for each $a \in A$: Simply consider the pair (S, V_m) , where $V_m : A \rightarrow \mathcal{P}S$ is the interpretation given by putting $V_m(a) := \{s \in S \mid a \in m(s)\}$. We consider special basic formulas of monadic first-order logic of the form:

$$\beta(\bar{a}, B) := \exists \bar{x}(\text{diff}(\bar{x}) \wedge \bigwedge_{i \in I} a_i(x_i) \wedge \forall y(\text{diff}(\bar{x}, y) \rightarrow \bigvee_{b \in B} b(y))),$$

where $\text{diff}(x_1, \dots, x_n)$ abbreviates the formula $\text{diff}(x_1, \dots, x_n) = \bigwedge_{1 \leq i < j \leq n} x_i \neq x_j$. To be explicit, we allow the case where \bar{a} is the empty sequence ϵ , in which case we get:

$$\beta(\epsilon, B) = \forall x \bigvee_{b \in B} b(x).$$

It is not hard to see that any Kripkean one-step \mathbf{B} -model (S, σ, m) satisfies:

$$S, \sigma, m \Vdash^1 \beta(\bar{a}, B) \text{ implies } S, \sigma, m' \Vdash^1 \beta(\bar{a}, B) \text{ for some } m' \subseteq m \text{ with } \text{Ran}(m') \subseteq \mathcal{P}_{\leq 1} A. \quad (3.2)$$

We can turn the formula $\beta(\bar{a}, B)$ into a modality $\nabla(\bar{a}; B)$ that can be interpreted in *all* one-step \mathbf{B} -models, using the observation that every one-step \mathbf{B} -frame (S, σ) (with $\sigma : S \rightarrow \omega$) has a unique Kripkean cover $(\tilde{S}, \tilde{\sigma})$ defined by putting $\tilde{S} := \bigcup \{s \times \sigma(s) \mid s \in S\}$, and $\tilde{\sigma} \in \mathcal{B}\tilde{S}$ is defined by $\tilde{\sigma}(s, i) := 1$ for all $s \in S$ and $i \in \sigma(s)$ (where we view each finite ordinal as the set of all smaller ordinals). Then we can define, for an arbitrary one-step \mathbf{B} -model (S, σ)

$$S, \sigma, m \Vdash^1 \nabla(\bar{a}; B) \text{ if } \tilde{S}, \tilde{\sigma}, m \circ \pi_S \Vdash^1 \beta(\bar{a}, B), \quad (3.3)$$

where π_S is the projection map $\pi_S : \tilde{S} \rightarrow S$. It is then an immediate consequence of (3.2) that $\nabla(\bar{a}; B)$ is a disjunctive formula.

Our main aim in this section is to show that the modalities $\nabla(\bar{a}; B)$ provide a disjunctive basis for the signature $\Sigma_{\mathbf{B}}$. As far as we know, this result is new. The hardest part in proving it is to show that the modalities $\nabla(\bar{a}; B)$ can be expressed as one-step formulas in $\mathbf{1ML}_{\Sigma_{\mathbf{B}}}^+(A)$. The reason that this is not so easy is subtle; by contrast, it is fairly straightforward to show that formulas of the form $\nabla(\bar{a}; B)$ can be expressed in $\mathbf{1ML}_{\Sigma_{\mathbf{B}}}(A)$, using Ehrenfeucht-Fraïssé games, see e.g. Fontaine & Place [14]. However, a proper disjunctive basis as we have defined it has to consist of *positive* formulas, and this will be crucial for applications to modal fixpoint logics³.

Proposition 3.4. *Every formula $\nabla(\bar{a}; B)$ is one-step equivalent to a formula in $\mathbf{1ML}_{\Sigma_{\mathbf{B}}}^+(A)$.*

Our main tool in proving this proposition will be Hall's Marriage Theorem, which can be formulated as follows. A *matching* of a bi-partite graph $\mathbb{G} = (V_1, V_2, E)$ is a subset M of E such that no two edges in M share any common vertex. M is said to *cover* V_1 if $\text{Dom}M = V_1$.

³The same subtlety appears in Janin & Lenzi [17], where the translation of the language $\mathcal{D}_{\mathbf{B}}$ into $\mathbf{1ML}_{\Sigma_{\mathbf{B}}}^+$ is required to prove that the graded μ -calculus is equivalent, over trees, to monadic second-order logic. Proposition 3.4 in fact fills a minor gap in this proof.

Fact 3.4 (Hall's Marriage Theorem). Let \mathbb{G} be a finite bi-partite graph, $\mathbb{G} = (V_1, V_2, E)$. Then \mathbb{G} has a matching that covers V_1 iff, for all $U \subseteq V_1$, $|U| \leq |E[U]|$, where $E[U]$ is the set of vertices in V_2 that are adjacent to some element of U .

Proof of Proposition 3.4. We will show this for the simple case where B is a singleton $\{b\}$. The general case is an immediate consequence of this (consider the substitution $B \mapsto \bigvee B$).

Let $\bar{a} = (a_1, \dots, a_n)$. Define $I := \{1, \dots, n\}$. For each subset $J \subseteq I$, let χ_J be the formula

$$\chi_J := \diamond^{|J|} \bigvee_{i \in J} a_i \wedge \square^{n+1-|J|} (\bigvee_{i \in J} a_i \vee b),$$

and let γ be the conjunction

$$\gamma := \bigwedge \{\chi_J \mid J \subseteq I\}.$$

What the formula χ_J says about a Kripkean (finite) one-step model is that at least $|J|$ elements satisfy the disjunction of the set $\{a_i \mid i \in J\}$, while all but at most $n - |J|$ elements satisfy the disjunction of the set $\{a_i \mid i \in J\} \cup \{b\}$. Abbreviating $\nabla(\bar{a}; b) := \nabla(\bar{a}; \{b\})$, we claim that

$$\gamma \equiv^1 \nabla(\bar{a}; b), \tag{3.4}$$

and to prove this it suffices to consider Kripkean one-step models.

It is straightforward to verify that the formula γ is a semantic one-step consequence of $\nabla(\bar{a}; b)$. For the converse, consider a Kripkean one-step model (S, σ, m) in which γ is true. Let K be an index set of size $|S| - n$, and disjoint from the index set $I = \{1, \dots, n\}$. Clearly then, $|I \cup K| = |I| + |K| = |S|$. Furthermore, let $a_k := b$, for all $k \in K$. To apply Hall's theorem, we define a bipartite graph $\mathbb{G} := (V_1, V_2, E)$ by setting $V_1 := I \cup K$, $V_2 := S$, and $E := \{(j, s) \in (I \cup K) \times S \mid a_j \in m(s)\}$. Note that this graph is finite: by definition of the bag functor only finitely many elements may have non-zero multiplicity in a one-step model, so every Kripkean one-step model has to be finite. Hence S is finite, and so both I and K are finite since $|I \cup K| = |S|$.

CLAIM 1 . The graph G has a matching that covers V_1 .

PROOF OF CLAIM We check the Hall marriage condition for an arbitrary subset $H \subseteq V_0$. In order to prove that the size of $E[H]$ is greater than that of H itself, we consider the formula $\chi_{H \cap I}$. We make a case distinction.

Case 1: $H \subseteq I$. Then $\chi_{H \cap I} = \chi_H$ implies $\diamond^{|H|} \bigvee_{i \in H} a_i$. This means that at least $|H|$ elements of S satisfy at least one variable in the set $\{a_i \mid i \in H\}$. By the definition of the graph \mathbb{G} , this is just another way of saying that $|H| \leq |E[H]|$, as required.

Case 2: $H \cap K \neq \emptyset$. Let $J := H \cap I$, then the formula $\chi_{H \cap I} = \chi_J$ implies the formula

$$\square^{n+1-|J|} (\bigvee_{j \in J} a_j \vee b).$$

Now, if $s \in S$ satisfies either b or some a_j for $j \in J$, then by the construction of \mathbb{G} we have $s \in E[H]$. We now see that $|S \setminus E[H]| \leq n - |J|$. Hence we get:

$$|E[H]| \geq |S| - (n - |J|) = |S| - n + |J|.$$

But note that $H = J \cup (H \cap K)$, so that we find

$$|H| \leq |J| + |H \cap K| \leq |J| + |K| = |J| + (|S| - n),$$

From these two inequalities it is immediate that $|H| \leq |E[H]|$, as required. \blacktriangleleft

Now consider a matching M that covers V_1 . Since the size of the set V_1 is the same as that of V_2 , any matching M of \mathbb{G} that covers V_1 is (the graph of) a bijection between these two sets. Furthermore, it easily follows that such an M restricts to a bijection between I and a subset $\{s_1, \dots, s_n\}$ of S such that $a_i \in m(s_i)$ for each $i \in I$, and that $b \in m(t)$ for each $t \notin \{u_1, \dots, u_n\}$. Hence $\nabla(\bar{a}; b)$ is true in (S, σ, m) , as required. \square

In light of this proposition, we shall continue to use the notation $\nabla(\bar{a}; B)$ for the equivalent formula in $\mathbf{1ML}_{\Sigma_{\mathbb{B}}}^+(A)$ provided in the proof.

Definition 3.5. We define $\mathbf{D}_{\mathbb{B}}(A)$ by the following grammar:

$$\delta ::= \top \mid \nabla(\bar{a}; B) \mid \delta \vee \delta$$

where \bar{a} is a tuple of elements from A and $B \subseteq A$.

Theorem 3.5. *The assignment $\mathbf{D}_{\mathbb{B}}$ provides a disjunctive basis for the signature $\Sigma_{\mathbb{B}}$.*

Proof. It remains to prove that $\mathbf{D}_{\mathbb{B}}$ is distributive over $\Sigma_{\mathbb{B}}$, and admits a binary distributive law. For the first part, consider the formula $\diamond^k \pi \in \mathbf{1ML}_{\Sigma_{\mathbb{B}}}^+(A)$. Note that \varnothing is a variable in \mathbf{PA} and that $\bigvee\{\varnothing\}[\wedge_A] = \bigwedge \varnothing = \top$. With this in mind, it is not hard to see that $\diamond^k \pi$ can be rewritten equivalently as:

$$\diamond^k \pi = \bigvee\{\nabla(B_1, \dots, B_k; \{\varnothing\}) \mid \bigwedge B_i \models \pi, \text{ for each } i\}[\wedge_A]$$

Here, \models denotes propositional consequence between formulas in $\mathbf{Bool}(A)$. Next, consider the formula $\square^k \pi \in \mathbf{1ML}_{\Sigma_{\mathbb{B}}}^+(A)$. Keeping in mind that $\bigwedge \varnothing = \top$, the reader can verify that this is equivalent to:

$$\bigvee_{m < k} \nabla(\underbrace{\varnothing, \dots, \varnothing}_{m \text{ times}}; \{B \subseteq A \mid \bigwedge B \models \pi\})[\wedge_A]$$

To establish the binary distributive law, let $\delta \in \mathbf{D}_{\mathbb{B}}(A)$ and $\delta' \in \mathbf{D}_{\mathbb{B}}(B)$. Then δ is of the form $\alpha_1 \vee \dots \vee \alpha_n$ and δ' is of the form $\beta_1 \vee \dots \vee \beta_m$, where each α_i is either equal to \top or of the form $\nabla(\bar{a}; A')$ for some \bar{a}, A' and each β_j is either equal to \top or of the form $\nabla(\bar{b}; B')$ for some \bar{b}, B' . By distributing the conjunction over disjunctions, we can rewrite the formula $\delta \wedge \delta'$ as an equivalent disjunction of formulas $\alpha_i \wedge \beta_j$ for $1 \leq i \leq n$ and $1 \leq j \leq m$. So it suffices to show that the required distributive law holds for conjunctions of this shape.

In the case where $\alpha_i = \beta_j = \top$, we have $\alpha_i \wedge \beta_j \equiv^1 \top$. But $\top \in \mathbf{D}(A \otimes B)$ and $\top = \top[\theta_{A,B}]$, so this case is done. If one and only one of α_i, β_j is equal to \top , say $\alpha_i = \nabla(\bar{a}; A')$ and $\beta_j = \top$, then:

$$\alpha_i \wedge \beta_j \equiv^1 \nabla(\bar{a}; A') \in \mathbf{D}(A) \subseteq \mathbf{D}(A \otimes B).$$

But $\nabla(\bar{a}; A')[\theta_{A,B}] = \nabla(\bar{a}; A')$, so we are done in this case as well.

Finally, the interesting case is where α_i is of the form $\nabla(a_1, \dots, a_k; A')$ and β_j is of the form $\nabla(b_1, \dots, b_l; B')$. For this case, we need some definitions: we fix the formulas $\alpha_i = \nabla(a_1, \dots, a_k; A')$ and $\beta_j = \nabla(b_1, \dots, b_l; B')$. Let an *overlap record* be a subset $O \subseteq \{1, \dots, k\} \times \{1, \dots, l\}$ such that O is the graph of a bijection from some subset of $\{1, \dots, k\}$ onto some subset of $\{1, \dots, l\}$. We denote the set of such overlap records by \mathcal{O} . Given an overlap record O , an *O-pair* is a pair (c_1, c_2) of functions $c_1 : \{1, \dots, k\} \setminus \pi_1[O] \rightarrow B'$ and $c_2 : \{1, \dots, l\} \setminus \pi_2[O] \rightarrow A'$. Finally, a *case description* is a triple (O, c_1, c_2) where O is an overlap record and (c_1, c_2) is an *O-pair*. Note that there are finitely many case descriptions.

Given a case description (O, c_1, c_2) , let $\{1, \dots, k\} \setminus \pi_1[O] = \{p_1, \dots, p_s\}$, let $\{1, \dots, l\} \setminus \pi_2[O] = \{q_1, \dots, q_t\}$, and let \vec{O} be a list of the pairs in O in some arbitrary order. Now define the formula $\chi(O, c_1, c_2)$ to be:

$$\nabla((a_{p_1}, c_1(p_1)), \dots, (a_{p_s}, c_1(p_s)), (b_{q_1}, c_2(q_1)), \dots, (b_{q_t}, c_2(q_t)), \vec{O}; A' \times B')$$

This is a formula in $\mathsf{D}_B(A \uplus B)$, in fact it is a formula in $\mathsf{D}_B(A \times B)$. It is a fairly straightforward task to verify that the formula $\alpha_i \wedge \beta_j$ is equivalent to the disjunction:

$$\bigvee \{ \chi(O, c_1, c_2) \mid (O, c_1, c_2) \text{ is a case description} \}$$

We leave this as an exercise for the reader. \square

3.3.4. An example without weak pullback preservation. Here is an example of a functor that does not preserve weak pullbacks, but still has a natural modal signature that admits a disjunctive basis. Let F be the subfunctor of P^2 given by setting $\mathsf{F}X$ to be the set of pairs $(Y, Z) \in (\mathsf{P}X)^2$ such that at least one of the sets Y, Z is finite. That is:

$$\mathsf{F}X = (\mathsf{P}_\omega X \times \mathsf{P}X) \cup (\mathsf{P}X \times \mathsf{P}_\omega X).$$

This is a well defined subfunctor of P^2 since the direct image of a finite subset $Y \subseteq X$ under any given map $f : X \rightarrow X'$ is finite.

F does not preserve weak pullbacks. To see why, consider the constant map $c : \omega \rightarrow \{0\}$, and consider the following two objects in $\mathsf{F}\omega$: $\alpha = (\{0\}, \omega)$ and $\beta = (\omega, \{0\})$. Clearly, $\mathsf{F}c(\alpha) = \mathsf{F}c(\beta) = (\{0\}, \{0\})$. But given that (P, p_1, p_2) is the pullback of the diagram $\omega \rightarrow \{0\} \leftarrow \omega$ where both maps are equal to c , if F preserves weak pullbacks then there should be some pair $\chi = (Y, Z) \in \mathsf{F}P$ such that $\mathsf{F}p_1(\chi) = \alpha$ and $\mathsf{F}p_2(\chi) = \beta$. But then $p_1[Z] = \omega$ and likewise $p_2[Y] = \omega$. So both Y and Z would have to be infinite, contradicting our definition of F .

Consider the modal signature consisting of the usual labelled diamond modalities \diamond^0 and \diamond^1 , quantifying over the left and right set in a pair $(Y, Z) \in \mathsf{F}X$ respectively, and their dual box modalities. Then we can define corresponding (finitary) nablas ∇^0 and ∇^1 in the usual way. Then, finite disjunctions of \top and conjunctions of the form $\nabla^0 A_0 \wedge \nabla^1 A_1$ for finite sets A_0, A_1 will provide a disjunctive basis for this signature. The reason that the formulas $\nabla^0 A_0 \wedge \nabla^1 A_1$ are still disjunctive is this: given a one-step model $(X, (Y, Z), m)$ satisfying the formula $\nabla^0 A_0 \wedge \nabla^1 A_1$, where $(Y, Z) \in \mathsf{F}X$ and m is a marking, we can construct an appropriate one-step frame morphism $f : (X', (Y', Z')) \rightarrow (X, (Y, Z))$ and a marking m' on X' as follows. Take:

$$\begin{aligned} X' &:= (X \times A_0) \cup (X \times A_1), \\ Y' &:= \{(u, a) \in Y \times A_0 \mid a \in m(u)\}, \\ Z' &:= \{(u, a) \in Z \times A_1 \mid a \in m(u)\}. \end{aligned}$$

The map f is of course just the projection to X , and the marking m' is defined by $m'(u, a) = \{a\}$. The crucial point is that since one of Y, Z is finite, and since A_0, A_1 are both finite sets, one of the sets Y', Z' is finite too, so $(Y', Z') \in \mathsf{F}X'$ as required. The proof that the formula $\nabla^0 A_0 \wedge \nabla^1 A_1$ also holds in the one-step model $(X', (Y', Z'), m')$ is routine.

3.3.5. *Monotone modal logic: a signature without disjunctive basis.* Finally, we provide an example of a signature that does not admit any disjunctive basis:

Proposition 3.6. *The signature $\Sigma_{\mathbf{M}}$ consisting of the box- and diamond liftings for \mathbf{M} does not have a disjunctive basis.*

Proof. Let L be the standard relation lifting for the monotone neighborhood functor. Given two one-step models X, ξ, m and X', ξ', m' over a set of variables A , we write $u \preceq u'$ if $m(u) \subseteq m'(u')$ for $u \in X$ and $u' \in X'$, and we say that X', ξ', m' *simulates* X, ξ, m if $(\xi, \xi') \in L(\preceq)$. A straightforward proof will verify the following claim.

CLAIM 1 . If X', ξ', m' simulates X, ξ, m then for every one-step formula $\alpha \in 1\text{ML}_{\Lambda}^+(A)$, $X, \xi, m \Vdash^1 \alpha$ implies $X', \xi', m' \Vdash^1 \alpha$.

Given a set A , let $\eta_A : A \rightarrow \text{PA}$ denote the map given by the unit of the powerset monad, i.e., it is the singleton map $\eta_A : a \mapsto \{a\}$. Furthermore, recall that \wedge_A is the substitution mapping $B \in \text{PA}$ to $\bigwedge B$.

CLAIM 2 . Let α be any one-step formula in $1\text{ML}_{\Lambda}(\text{PA})$ and let (X, ξ, m) be a one-step model with $m : X \rightarrow \text{PA}$. Consider the map $\eta_{\text{PA}} : \text{PA} \rightarrow \text{PPA}$, so that $\eta_{\text{PA}} \circ m$ is a marking of X with variables from PA .

- (1) If $X, \xi, \eta_{\text{PA}} \circ m \Vdash^1 \alpha$ then $X, \xi, m \Vdash^1 \alpha[\wedge_A]$.
- (2) If $X, \xi, m \Vdash^1 \alpha[\wedge_A]$ and the empty set does not appear as a variable in α , and furthermore $m(u)$ is a singleton for each $u \in X$, then $X, \xi, \eta_{\text{PA}} \circ m \Vdash^1 \alpha$.

PROOF OF CLAIM For the first part of the proposition, it suffices to note that $\llbracket B \rrbracket_{\eta_{\text{PA}} \circ m}^1 \subseteq \llbracket \bigwedge B \rrbracket_m^1$ for each $B \in \text{PA}$, and the result then follows by monotonicity of the predicate lifting corresponding to the one-step formula α .

For the second part, it suffices to note that under the additional constraint that $m(u)$ is a singleton for each $u \in X$ and the empty set does not appear as a variable in α , we have $\llbracket \bigwedge B \rrbracket_m^1 \subseteq \llbracket B \rrbracket_{\eta_{\text{PA}} \circ m}^1$ for each $B \in \text{PA}$ that appears as a variable in α . To prove this, suppose that $u \in \llbracket \bigwedge B \rrbracket_m^1$. Since B appears in α it is non-empty, and since $m(u)$ is a singleton, say $m(u) = \{b\}$, it follows that we must in fact have $B = \{b\}$. Hence:

$$B \in \{\{b\}\} = \{m(u)\} = \eta_{\text{PA}}(m(u))$$

so $u \in \llbracket B \rrbracket_{\eta_{\text{PA}} \circ m}^1$ as required. \blacktriangleleft

Now, let $A = \{a, b, c\}$ and consider the formula $\psi = \nabla\{\{a, b\}, \{c\}\}$. If 1ML_{Λ} admits a disjunctive basis, then there is a disjunctive formula δ in $1\text{ML}_{\Lambda}(\text{PA})$ such that $\psi = \delta[\wedge_A]$.

So suppose $\delta \in 1\text{ML}_{\Lambda}(\text{PA})$ is disjunctive, and suppose that $\psi = \delta[\wedge_A]$. We may in fact assume w.l.o.g. that the empty set does not appear as a variable in δ , since otherwise we just use instead the formula $\delta[\top/\emptyset]$, which is still disjunctive (this is easy to prove). We have $\delta[\top/\emptyset][\wedge_A] = \delta[\wedge_A]$ since $\wedge_A(\emptyset) = \bigwedge \emptyset = \top$.

With this in mind, consider the one-step model X, ξ, m where $X = \{x_1, x_2, x_3\}$, $\xi = \{\{x_1, x_2\}, \{x_3\}, X\}$ and $m(x_1) = \{a\}$, $m(x_2) = \{b\}$ and $m(x_3) = \{c\}$. It is easy to see that $X, \xi, m \Vdash^1 \psi$, so by assumption $X, \xi, m \Vdash^1 \delta[\wedge_A]$. But since the marking m maps every element of X to a singleton, item 2 of Claim 2 gives us that $X, \xi, \eta_{\text{PA}} \circ m \Vdash^1 \delta$.

Now, define a new one-step model X, ξ, h where as before $X = \{x_1, x_2, x_3\}$ and $\xi = \{\{x_1, x_2\}, \{x_3\}, X\}$ but where the marking $h : X \rightarrow \text{PPA}$ (with respect to variables in PA) is defined by setting $h(x_1) = \{\{a\}\}$, $h(x_2) = \{\{a\}, \{c\}\}$ and $h(x_3) = \{\{c\}\}$. It is a matter of

simple verification to check that X, ξ, h in fact simulates $X, \xi, \eta_{\mathcal{P}A} \circ m$, so by Proposition 1 we get $X, \xi, h \Vdash^1 \delta$.

Since δ is disjunctive, there should be a one-step model X', ξ', h' and a map $f : X' \rightarrow X$ such that: $X', \xi', h' \Vdash^1 \delta$, $\mathbf{M}f(\xi') = \xi$, $h'(u) \subseteq h(f(u))$ for all $u \in X'$ and $h'(u)$ is at most a singleton for each $u \in X'$. By monotonicity of δ we can in fact assume w.l.o.g. that $h'(u)$ is precisely a singleton for each $u \in X'$: if $h'(u) = \emptyset$, just pick some element e of $h(f(u))$ (since $h(v)$ is non-empty for each $v \in X$) and set $h'(u) = \{e\}$. The resulting marking still satisfies all the conditions above.

But this means that we can define a marking $n : X' \rightarrow \mathcal{P}A$ by taking each $n(u)$ for $u \in X'$ to be the unique $B \subseteq A$ such that $h'(u) = \{B\}$. Clearly, $h' = \eta_{\mathcal{P}A} \circ n$, so by the first part of Claim 2, we get $X', \xi', n \Vdash^1 \delta[\wedge_A]$, hence $X', \xi', n \Vdash^1 \psi$, i.e. $X', \xi', n \Vdash^1 \nabla\{\{a, b\}, \{c\}\}$. But from the definition of the marking h , the condition that $h'(u) \subseteq h(f(u))$ for all $u \in X'$ and from the definition of n it is clear that, for all $u \in X'$, we have $n(u) = \{a\}$ or $n(u) = \{c\}$. So to finally reach our desired contradiction, it suffices to prove the following.

CLAIM 3 . Let X, ξ, m be any one-step model such that $X, \xi, m \Vdash^1 \nabla\{\{a, b\}, \{c\}\}$. Then either there is some $u \in X$ with $\{a, c\} \subseteq m(u)$, or there is some $u \in X$ with $b \in m(u)$.

PROOF OF CLAIM Suppose there is no $u \in X$ with $b \in m(u)$. Then there is some set $Z \in \xi$ such that every $v \in Z$ satisfies a . Furthermore there must be some $B \in \xi$ such that every $l \in B$ is satisfied by some member of Z . The only choice possible for this is $\{c\}$, hence some member of Z must satisfy both a and c . \blacktriangleleft

This finishes the proof of Proposition 3.6. \square

Interestingly, we may “repair” the failure of admitting a disjunctive basis by moving to the *supported companion functor* $\underline{\mathbf{M}}$. This functor is defined as the subfunctor of $\mathbf{P} \times \mathbf{M}$, given (on objects) by $\underline{\mathbf{M}}S := \{(U, \sigma) \in \mathcal{P}S \times \mathcal{M}S \mid U \text{ supports } \sigma\}$, where we refer for the definition of the notion of support to Definition 7.1. If we add to the signature $\Sigma_{\mathbf{M}}$ of monotone modal logic a box and diamond modality accessing the support, we obtain a signature which does admit a disjunctive basis. This is connected with Proposition 3.3 and the fact that the functor $\underline{\mathbf{M}}$ preserves weak pullbacks, which was proved in [8]. The supported companion functor features prominently in our completeness proof (with Seifan) for the monotone μ -calculus [8].

4. DISJUNCTIVE AUTOMATA AND SIMULATION

We now introduce disjunctive automata, which serve as a coalgebraic generalization of non-deterministic automata for the modal μ -calculus. We refer to the sections 2.5 and 2.4 for background on, respectively, automata and their connection with μ -calculi, and the infinite games in which their semantics is formulated.

Definition 4.1. A (Λ, X) -automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$ is said to be *disjunctive* if $\Theta(c, a)$ is disjunctive, for all colors $c \in \mathcal{P}X$ and all states $a \in A$.

Given a disjunctive basis \mathbf{D} , we say that \mathbb{A} is a \mathbf{D} -automaton if $\Theta(c, a) \in \mathbf{D}(A)$, for all colors $c \in \mathcal{P}X$ and all states $a \in A$.

Definition 4.2. Let \mathbb{A} be a Λ -automaton and let (\mathbb{S}, s_I) be a pointed \mathbf{T} -model. A strategy f for \exists in $\mathcal{A}(\mathbb{A}, \mathbb{S})@(a, s)$ is *dividing* if for every t in \mathbb{S} there is at most one state b in \mathbb{A} such that the position (b, t) is *f-reachable* (i.e., occurs in some f -guided match). We say that

\mathbb{A} *strongly accepts* (\mathbb{S}, s_I) , notation: $\mathbb{S}, s_I \Vdash_s \mathbb{A}$ if \exists has a dividing winning strategy in the game $\mathcal{A}(\mathbb{A}, \mathbb{S})@(a_I, s_I)$.

Disjunctive automata are very well behaved. For instance, the following observation states a *linear-size* model property. This observation goes back to Janin & Walukiewicz [18], who proved the linear-size model property for so-called disjunctive formulas of the modal μ -calculus. Their result was generalised to the coalgebraic setting of ∇ -based automata by Kupke & Venema [22].

Theorem 4.1. *Let $\mathbb{A} = (A, \Theta, a_I, \Omega)$ be a disjunctive automaton for a set functor \mathbb{T} . If \mathbb{A} accepts some pointed \mathbb{T} -model, then it accepts one of which the carrier S satisfies $S \subseteq A$.*

The main property of disjunctive automata, which we will use throughout the remainder of this paper, is the following.

Proposition 4.2. *Let \mathbb{A} be a disjunctive Λ -automaton. Then any pointed \mathbb{T} -model which is accepted by \mathbb{A} has a pre-image model which is strongly accepted by \mathbb{A} .*

Proof. Let $\mathbb{S} = (S, \sigma, V)$ be a pointed \mathbb{T} -model, let $s_I \in S$, and let f be a winning strategy for \exists in the acceptance game $\mathcal{A} := \mathcal{A}(\mathbb{A}, \mathbb{S})@(a_I, s_I)$. By Fact 2.12 we may without loss of generality assume that f is positional. We will construct (i) a pointed \mathbb{T} -model (X, ξ, W, x_I) , (ii) a tree (X, R) which is rooted at x_I (in the sense that for every $t \in X$ there is a unique R -path from x_I to t) and supports (X, ξ) (in the sense that $\xi(x) \in \mathbb{T}R(x)$, for every $x \in X$), and (iii) a morphism $h : (X, \xi, W) \rightarrow (S, \sigma, V)$ such that $h(x_I) = s_I$. In addition (X, ξ, W, x_I) will be strongly accepted by \mathbb{A} .

More in detail, we will construct all of the above step by step, and by a simultaneous induction we will associate, with each $t \in X$ of depth k , a (partial) f -guided match Σ_t of length $2k + 1$; we will denote the final position of Σ_t as (a_t, s_t) , and will define $h(t) := s_t$.

For the base step of the construction we take some fresh object x_I , we define Σ_{x_I} to be the match consisting of the single position (a_I, s_I) , and set $h(x_I) := s_I$.

Inductively assume that we are dealing with a node $t \in X$ of depth k , and that Σ_t, a_t and s_t are as described above. Since Σ_t is an f -guided match and f is a winning strategy in \mathcal{A} , the pair (a_t, s_t) is a winning position for \exists in \mathcal{A} . In particular, the marking $m_t : S \rightarrow \text{PA}$ prescribed by f at this position satisfies

$$S, \sigma(s_t), m_t \Vdash^1 \Theta(a_t, V^b(s_t)).$$

Now by disjunctiveness of the automaton \mathbb{A} there is a set $R(t)$ (that we may take to consist of fresh objects), an object $\xi(t) \in \mathbb{T}R(t)$, an A -marking $m'_t : R(t) \rightarrow \text{PA}$ and a map $h_t : R(t) \rightarrow S$, such that⁴ $|m(u)| = 1$ and $m'_t(u) \subseteq m_t(h_t(u))$ for all $u \in R(t)$, $(\mathbb{T}h_t)\xi(t) = \sigma(s_t)$ and

$$R(t), \xi(t), m'_t \Vdash^1 \Theta(a_t, V^b(s_t)).$$

Let a_u be the unique object such that $m'_t(u) = \{a_u\}$, define $s_u := h_t(u)$, and let Σ_u be the match $\Sigma_u := \Sigma_t \cdot m_t \cdot (a_u, s_u)$.

With (X, R, x_I) the tree constructed in this way, and observing that $\xi(t) \in R(t) \subseteq X$, we let ξ be the coalgebra map on X . Taking $h : X \rightarrow S$ to be the union $\{(x_I, s_I)\} \cup \{h_t \mid t \in X\}$, we can easily verify that h is a surjective coalgebra morphism. Finally, we define the valuation $W : \mathbf{X} \rightarrow \text{PX}$ by putting $W(p) := \{x \in X \mid hx \in V(p)\}$.

⁴To simplify our construction, we strengthen clause (3) in Definition 3.1. This is not without loss of generality, but we may take care of the general case using a routine extension of the present proof.

It remains to show that \mathbb{A} strongly accepts the pointed \mathbb{T} -model (\mathbb{X}, x_I) , with $\mathbb{X} = (X, \xi, W)$; for this purpose consider the following (positional) strategy f' for \exists in $\mathcal{A}(\mathbb{A}, \mathbb{X})$. At a position $(a, t) \in A \times X$ such that $a \neq a_t$ \exists moves randomly (we may show that such a position will not occur); on the other hand, at a position of the form (a_t, t) , the move suggested by the strategy f' is the marking m'_t . Then it is obvious that f' is a dividing strategy; to see that f' is winning from starting position (a_I, x_I) , consider an infinite match Σ of $\mathcal{A}(\mathbb{A}, \mathbb{X})@_{(a_I, x_I)}$ (finite matches are left to the reader). It is not hard to see that Σ must be of the form $\Sigma = (a_0, x_0)m'_{x_0}(a_1, x_1)m'_{x_1} \cdots$, where $\Sigma^- = (a_0, h(x_0))m_{x_0}(a_1, h(x_1))m_{x_1} \cdots$ is an f -guided match of \mathcal{A} . From this observation it is immediate that Σ is won by \exists . \square

We now come to our main application of disjunctive bases, and fill in the main missing piece in the theory of coalgebraic automata based on predicate liftings: a simulation theorem. As mentioned in the introduction, Janin & Walukiewicz' simulation theorem [18] is one of the key tools in the theory of the standard modal μ -calculus, see for instance [36] for many examples. At the coalgebraic level of generality, a first simulation theorem was proved by Kupke & Venema [22] for ∇ -based automata.

Theorem 4.3 (Simulation). *Let Λ be a monotone modal signature for the set functor \mathbb{T} and assume that Λ has a disjunctive basis \mathbb{D} . Then there is an effective construction transforming an arbitrary Λ -automaton \mathbb{A} into an equivalent \mathbb{D} -automaton $\text{sim}(\mathbb{A})$.*

Proof. Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a Λ -automaton. Our definition of $\text{sim}(\mathbb{A})$ is rather standard [36], so we will confine ourselves to the definitions. The construction takes place in two steps, a ‘pre-simulation’ step that produces a disjunctive automaton $\text{pre}(\mathbb{A})$ with a non-parity acceptance condition, and a second ‘synchronization’ step that turns this nonstandard disjunctive automaton into a standard one. Both steps of the construction involve a ‘change of base’ in the sense that we obtain the transition map of the new automaton via a substitution relating its carrier to the carrier of the old automaton.

We define the pre-simulation automaton of \mathbb{A} as the structure

$$\text{pre}(\mathbb{A}) := (A^\sharp, \Theta^\sharp, NBT_{\mathbb{A}}, R_I),$$

where the carrier of the pre-simulation $\text{pre}(\mathbb{A})$ of \mathbb{A} is the collection A^\sharp of *binary* relations over A , and the initial state R_I is the singleton pair $\{(a_I, a_I)\}$. For its transition function, first define the map $\Theta^* : A \times \text{PX} \rightarrow \mathbf{1ML}_\Lambda^+(A \times A)$ by putting, for $a \in A$ and $c \in \text{PX}$:

$$\Theta^*(a, c) := \Theta(a, c)[\theta_a],$$

where $\theta_a : A \rightarrow \mathbf{Latt}(A \times A)$ is the *tagging* substitution given by $\theta_a : b \mapsto (a, b)$. Now, given a state $R \in A^\sharp$ and color $c \in \text{PX}$, take $\Theta^\sharp(R, c)$ to be an arbitrary but fixed formula in $\mathbb{D}(A^\sharp)$ such that

$$\Theta^\sharp(R, c)[\wedge_{A \times A}] \equiv \bigwedge_{a \in \text{Ran} R} \Theta^*(a, c).$$

Clearly such a formula exists by our assumption on \mathbb{D} being a disjunctive basis for Λ .

Turning to the *acceptance condition*, define a *trace* on an A^\sharp -stream $\rho = (R_n)_{0 \leq n < \omega}$ to be an A -stream $\alpha = (a_n)_{0 \leq n < \omega}$ with $R_i a_i a_{i+1}$ for all $i \leq 0$. Calling such a trace α *bad* if $\max\{\Omega(a) \mid a \text{ occurs infinitely often in } \alpha\}$ is odd, we obtain the acceptance condition of the automaton $\text{pre}(\mathbb{A})$ as the set $NBT_{\mathbb{A}} \subseteq (A^\sharp)^\omega$ of A^\sharp -streams that contain no bad trace.

Finally we produce the simulation of \mathbb{A} by forming a certain kind of product of $\text{pre}(\mathbb{A})$ with \mathbb{Z} , where $\mathbb{Z} = (Z, \delta, \Omega', z_I)$ is some deterministic parity stream automaton recognizing

the ω -regular language $NBT_{\mathbb{A}}$. More precisely, we define $\mathbf{sim}(\mathbb{A}) := (A^\sharp \times Z, \Theta'', \Omega'', (R_I, z_I))$ where:

- $\Theta''(R, z) := \Theta^\sharp(R)[(Q, \delta(R, z)/Q) \mid Q \in A^\sharp]$ and
- $\Omega''(R, z) := \Omega'(z)$.

The equivalence of \mathbb{A} and $\mathbf{sim}(\mathbb{A})$ can be proved by relatively standard means [36]. \square

5. LYNDON THEOREMS

Lyndon's classical theorem in model theory provides a syntactic characterization of a semantic property, showing that a formula is *monotone* in a predicate P if and only if it is equivalent to a formula in which P occurs only *positively*. A version of this result for the modal μ -calculus was proved by d'Agostino and Hollenberg in [6]. Here, we show that their result holds for any μ -calculus based on a signature that admits a disjunctive basis.

We first turn to the one-step version of the Lyndon Theorem, for which we need the following definition; we also recall the substitutions \wedge_A and \vee_A defined in Section 2.

Definition 5.1. A *propositional A-type* is a subset of A . For $B \subseteq A$ and $a \in A$, the formulas τ_B and τ_B^{a+} are defined by:

$$\begin{aligned} \tau_B &:= \bigwedge B \wedge \bigwedge \{\neg a \mid a \in A \setminus B\} \\ \tau_B^{a+} &:= \bigwedge B \wedge \bigwedge \{\neg b \mid b \in A \setminus (B \cup \{a\})\} \end{aligned}$$

We let τ and τ^{a+} denote the maps $B \mapsto \tau_B$ and $B \mapsto \tau_B^{a+}$, respectively.

Proposition 5.1. *Suppose Λ admits a disjunctive basis. Then for any formula α in $1\mathbf{ML}_\Lambda(A)$ there is a one-step equivalent formula of the form $\delta[\vee_{PA}][\tau]$ for some $\delta \in \mathbf{D}(\mathbf{PPA})$.*

Proof. We first check that everything is correctly typed: note that we have $\vee_{PA} : \mathbf{PPA} \rightarrow \mathbf{Bool}(PA)$ and so $\delta[\vee_{PA}] \in 1\mathbf{ML}_\Lambda(PA)$, and $\tau_{PA} : PA \rightarrow \mathbf{Bool}(A)$. So $\delta[\vee_{PA}][\tau] \in 1\mathbf{ML}_\Lambda(A)$, as required.

For the normal form proof, first note that we can use boolean duals of the modal operators to push negations down to the zero-step level. Putting the resulting formula in disjunctive normal form, we obtain a disjunction of formulas of the form $\heartsuit_{\lambda_1} \overline{\pi_1} \wedge \cdots \wedge \heartsuit_{\lambda_k} \overline{\pi_k}$, where all the π -formulas are in $\mathbf{Bool}(A)$. Repeatedly applying the distributivity of \mathbf{D} over Λ and the distributive law for \mathbf{D} , we can rewrite each such disjunct as a formula of the form $\delta[\sigma]$ where, for some set B , $\delta \in \mathbf{D}(B)$ and $\sigma : B \rightarrow \mathbf{Bool}(A)$ is some propositional substitution. Now, just apply propositional logic to rewrite each formula σ_b as a disjunction of formulas in $\tau[PA]$, and we are done. \square

Theorem 5.2 (One-step Lyndon theorem). *Let Λ be a monotone modal signature for the set functor \mathbb{T} and assume that Λ has a disjunctive basis. Any $\alpha \in 1\mathbf{ML}_\Lambda(A)$, monotone in the variable $a \in A$, is one-step equivalent to some formula in $1\mathbf{ML}_\Lambda(A)$, which is positive in a .*

Proof. By Proposition 5.1, we can assume that α is of the form $\delta[\vee_{PA}][\tau]$ for some $\delta \in \mathbf{D}(PA)$. Clearly it suffices to show that:

$$\delta[\vee_{PA}][\tau] \equiv^1 \delta[\vee_{PA}][\tau^{a+}]$$

One direction, from left to right, is easy since $\delta[\vee_{PA}]$ is a monotone formula in $1\mathbf{ML}_\Lambda(PA)$, and $[\tau_B]_m^0 \subseteq [\tau_B^{a+}]_m^0$ for each $B \subseteq A$ and each marking $m : X \rightarrow PA$.

For the converse direction, suppose $X, \xi, m \Vdash^1 \delta[\vee_{PA}][\tau^{a+}]$. We define a PA -marking $m_0 : X \rightarrow PPA$ by setting $m_0(u) := \{B \subseteq A \mid B \preceq_a m(u)\}$, where the relation \preceq_a over PA is defined by $B \preceq_a B'$ iff $B \setminus \{a\} = B' \setminus \{a\}$, and $a \notin B$ or $a \in B'$. We claim that $X, \xi, m_0 \Vdash^1 \delta[\vee_{PA}]$. Since $\delta[\vee_{PA}]$ is a monotone formula, it suffices to check that $\llbracket \tau_B^{a+} \rrbracket_m^0 \subseteq \llbracket B \rrbracket_{m_0}^0$ for each $B \subseteq A$. This follows by just unfolding definitions.

Since δ was disjunctive, so is $\delta[\vee_{PA}]$, as an easy argument will reveal. So we now find a one-step frame morphism $f : (X', \xi') \rightarrow (X, \xi)$, together with a marking $m' : X' \rightarrow PPA$ such that $|m'(u)| \leq 1$ and $m'(u) \subseteq m_0(f(u))$ for all $u \in X'$, and such that $X', \xi', m' \Vdash^1 \delta[\vee_{PA}]$. We define a new A -marking $m'' : X' \rightarrow PA$ on X' by setting $m''(u) = B$, if $m'(u) = \{B\}$, and $m''(u) = m(f(u))$ if $m'(u) = \emptyset$. Note that, for each $B \subseteq A$, we have $\llbracket B \rrbracket_{m'}^0 \subseteq \llbracket \tau_B \rrbracket_{m''}^0$, so by monotonicity of $\delta[\vee_{PA}]$ we get $X', \xi', m'' \Vdash^1 \delta[\vee_{PA}][\tau]$.

Comparing the markings m'' and $m \circ f$, we claim that $m''(u) \preceq_a m(f(u))$ for all $u \in X'$. If $m'(u) = \emptyset$, then in fact $m''(u) = m(f(u))$ by definition of m'' . If $m'(u) = \{B\}$, then $m''(u) = B \in m'(u) \subseteq m_0(f(u))$, hence $B \preceq_a m(f(u))$ by definition of m_0 . Since $\delta[\vee_{PA}][\tau]$ was monotone with respect to the variable a it follows that $X', \xi', m \circ f \Vdash^1 \delta[\vee_{PA}][\tau]$ and so $X, \xi, m \Vdash^1 \delta[\vee_{PA}][\tau]$ by naturality, thus completing the proof of the theorem. \square

A useful corollary to this theorem is that, given an expressively complete set Λ of predicate liftings for a functor \mathbb{T} , the language μML_Λ has the same expressive power as the full language $\mu ML_{\mathbb{T}}$. At first glance this proposition may seem trivial, but it is important to see that it is not: given a formula φ of $\mu ML_{\mathbb{T}}$, a naive definition of an equivalent formula in μML_Λ would be to apply expressive completeness to simply replace each subformula of the form $\heartsuit_\lambda(\psi_1, \dots, \psi_n)$ with an equivalent one-step formula α over $\{\psi_1, \dots, \psi_n\}$, using only predicate liftings in Λ . But if this subformula contains bound fixpoint variables, these must still appear positively in α in order for the translation to even produce a grammatically correct formula! We need the stronger condition of *Lyndon completeness* for Λ . Generally, we have no guarantee that expressive completeness entails Lyndon completeness, but in the presence of a disjunctive basis, we do: this is a consequence of Theorem 5.2.

Corollary 5.3. *Suppose Λ is an expressively complete set of monotone predicate liftings for \mathbb{T} . If Λ admits a disjunctive basis, then Λ is Lyndon complete and hence $\mu ML_\Lambda \equiv \mu ML_{\mathbb{T}}$.*

Proof. The simplest proof uses automata: pick a modal Λ' -automaton \mathbb{A} , where Λ' is the set of all monotone predicate liftings for \mathbb{T} , and apply expressive completeness to replace each formula α in the co-domain of the transition map Θ with an equivalent one-step formula α' using only liftings in Λ . This formula is still monotone in all the variables in A since it is equivalent to α , so we can apply the one-step Lyndon Theorem 5.2 to replace α' by an equivalent and positive one-step formula β in $1ML_\Lambda(A)$. Clearly, the resulting automaton \mathbb{A}' will be semantically equivalent to \mathbb{A} . \square

We now turn to our Lyndon Theorems for the full coalgebraic modal (fixpoint) languages. Let $(\mu ML_\Lambda)_p^M$ and $(ML_\Lambda)_p^M$ denote the fragments of respectively μML_Λ and ML_Λ , consisting of the formulas that are positive in the proposition letter p .

Theorem 5.4 (Lyndon Theorem). *There is an effective translation $(\cdot)_p^M : \mu ML_\Lambda \rightarrow (\mu ML_\Lambda)_p^M$, which restricts to a map $(\cdot)_p^M : ML_\Lambda \rightarrow (ML_\Lambda)_p^M$, and satisfies that*

$$\varphi \in \mu ML \text{ is monotone in } p \text{ iff } \varphi \equiv \varphi_p^M.$$

Proof. By the equivalence between formulas and Λ -automata and the Simulation Theorem, it suffices to prove the analogous statement for disjunctive coalgebra automata.

Given a disjunctive Λ -automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$, we define \mathbb{A}_p^M to be the automaton $(A, \Theta_p^M, \Omega, a_I)$, where

$$\Theta_p^M(a, c) := \begin{cases} \Theta(a, c) \vee \Theta(a, c \setminus \{p\}) & \text{if } p \in c \\ \Theta(a, c) & \text{if } p \notin c. \end{cases}$$

Clearly \mathbb{A}_p^M is a disjunctive automaton as well, and it is routine to show that \mathbb{A}_p^M is equivalent to a formula in μML_Λ that is positive in the variable p . The key observation is that we have the following valid equivalence of classical propositional logic, for any formulas π, Θ_1, Θ_2 :

$$(\pi \wedge \neg p \wedge \Theta_1) \vee (\pi \wedge p \wedge (\Theta_1 \vee \Theta_2)) \quad \Leftrightarrow \quad (\pi \wedge \Theta_1) \vee (\pi \wedge p \wedge \Theta_2)$$

We claim that \mathbb{A} is monotone in p iff $\mathbb{A} \equiv \mathbb{A}_p^M$. Leaving the easy direction from right to left to the reader, we prove the opposite implication. So assume that \mathbb{A} is monotone in p . Since it immediately follows from the definitions that \mathbb{A} always implies \mathbb{A}_p^M , we are left to show that \mathbb{A}_p^M implies \mathbb{A} , and since \mathbb{A}_p^M is disjunctive, by Proposition 4.2 and invariance of acceptance by coalgebra automata it suffices to prove the following:

$$\mathbb{S}, s_I \Vdash_s \mathbb{A}_p^M \text{ implies } \mathbb{S}, s_I \Vdash \mathbb{A}, \quad (5.1)$$

for an arbitrary T -model (\mathbb{S}, s_I) .

To prove (5.1), let f be a dividing winning strategy for \exists in $\mathcal{A}^M := \mathcal{A}(\mathbb{A}_p^M, \mathbb{S})@(a_I, s_I)$. Our aim is to find a subset $U \subseteq V(p)$ such that $\mathbb{S}[p \mapsto U], s_I \Vdash \mathbb{A}$; it then follows by monotonicity that $\mathbb{S}, s_I \Vdash \mathbb{A}$. Call a point $s \in S$ *f-accessible* if there is a (by assumption unique) state a_s such that the position (a_s, s) is *f-reachable* in \mathcal{A}^M . We define U as the set of *f-accessible* elements s of $V(p)$ such that:

$$S, \sigma(s), m_s \Vdash^1 \Theta(a_s, V^b(s)),$$

where m_s is the A -marking provided by f at position (a_s, s) . We let V_U abbreviate $V[p \mapsto U]$. We claim that

$$\text{if } s \text{ is } f\text{-accessible then } S, \sigma(s), m_s \Vdash^1 \Theta(a_s, V_U^b(s)). \quad (5.2)$$

To see why (5.2) holds, note that for any *f-accessible* point s , the marking m_s is a legitimate move at position (a_s, s) , since f is assumed to be winning for \exists in \mathcal{A}^M . In other words, we have $S, \sigma(s), m_s \Vdash^1 \Theta_p^M(a_s, V^b(s))$. We have to make a case distinction, between the following three cases:

Case 1: $p \notin V_U^b(s)$ and $p \notin V^b(s)$. Then $\Theta(a_s, V_U^b(s)) = \Theta_p^M(a_s, V^b(s))$, so we get $S, \sigma(s), m_s \Vdash^1 \Theta(a_s, V_U^b(s))$ as required.

Case 2: $p \notin V_U^b(s)$ but $p \in V^b(s)$. Then $\Theta(a_s, V^b(s))$ does not hold in the one-step model $S, \sigma(s), m_s$, so by definition of Θ_p^M we have $S, \sigma(s), m_s \Vdash^1 \Theta(a_s, V^b(s) \setminus \{p\})$. But $V^b(s) \setminus \{p\} = V_U^b(s)$, so again we get $S, \sigma(s), m_s \Vdash^1 \Theta(a_s, V_U^b(s))$ as required.

Case 3: $p \in V_U^b(s)$. In this case we have $S, \sigma(s), m_s \Vdash^1 \Theta(a_s, V^b(s))$ by definition of the valuation V_U . Furthermore, since $V^b(s) = V_U^b(s)$, we get $S, \sigma(s), m_s \Vdash^1 \Theta(a_s, V_U^b(s))$.

Hence, (5.2) holds in all three possible cases. Finally, it is straightforward to derive from (5.2) that f itself is a (dividing) winning strategy for \exists in the acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{S})$ initialized at (a_I, s_I) . \square

Remark 5.2. Observe that as a corollary of Theorem 5.4 and the decidability of the satisfiability problem of μML_Λ [3], it is decidable whether a given formula $\varphi \in \mu\text{ML}_\Lambda$ is monotone in p .

6. UNIFORM INTERPOLATION

Uniform interpolation is a very strong form of the interpolation theorem, first proved for the modal μ -calculus in [6]. It was later generalized to coalgebraic modal logics in [26]. However, the proof crucially relies on non-deterministic automata, and for that reason the generalization in [26] is stated for nabla-based languages. With a simulation theorem for predicate liftings based automata in place, we can prove the uniform interpolation theorem for a large class of μ -calculi based on predicate liftings.

Definition 6.1. Given a formula $\varphi \in \mu\text{ML}_\Lambda$, we let \mathbf{X}_φ denote the set of proposition letters occurring in φ . Given a set \mathbf{X} of proposition letters and a single proposition letter p , it may be convenient to denote the set $\mathbf{X} \cup \{p\}$ as $\mathbf{X}p$.

Definition 6.2. A logic \mathcal{L} with semantic consequence relation \models is said to have the property of *uniform interpolation* if, for any formula $\varphi \in \mathcal{L}$ and any set $\mathbf{X} \subseteq \mathbf{X}_\varphi$ of proposition letters, there is a formula $\varphi_{\mathbf{X}} \in \mathcal{L}(\mathbf{X})$, effectively constructible from φ , such that

$$\varphi \models \psi \text{ iff } \varphi_{\mathbf{X}} \models \psi, \quad (6.1)$$

for every formula $\psi \in \mathcal{L}$ such that $\mathbf{X}_\varphi \cap \mathbf{X}_\psi \subseteq \mathbf{X}$.

To see why this property is called *uniform interpolation*, it is not hard to prove that, if $\varphi \models \psi$, with $\mathbf{X}_\varphi \cap \mathbf{X}_\psi \subseteq \mathbf{X}$, then the formula $\varphi_{\mathbf{X}}$ is indeed an interpolant in the sense that $\varphi \models \varphi_{\mathbf{X}} \models \psi$ and $\mathbf{X}_{\varphi_{\mathbf{X}}} \subseteq \mathbf{X}_\varphi \cap \mathbf{X}_\psi$.

Theorem 6.1 (Uniform Interpolation). *Let Λ be a monotone modal signature for the set functor \mathbb{T} and assume that Λ has a disjunctive basis. Then both logics ML_Λ and μML_Λ enjoy the property of uniform interpolation.*

Following D'Agostino & Hollenberg [6], we prove Theorem 6.1 by automata-theoretic means. The key proposition in our proof is Proposition 6.2 below, which refers to the following construction on disjunctive automata.

Definition 6.3. Let \mathbf{X} be a set of proposition letters not containing the letter p . Given a disjunctive $(\Lambda, \mathbf{X}p)$ -automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$, we define the map $\Theta^{\exists p} : A \times \text{PX} \rightarrow \text{D}(A)$ by

$$\Theta^{\exists p}(a, c) := \Theta(a, c) \vee \Theta(a, c \cup \{p\}),$$

and we let $\mathbb{A}^{\exists p}$ denote the (Λ, \mathbf{X}) -automaton $(A, \Theta^{\exists p}, \Omega, a_I)$.

The following proposition shows that the operation $(\cdot)^{\exists p}$ behaves like an *existential quantifier*, but with a twist: the automaton $\mathbb{A}^{\exists p}$ accepts a pointed coalgebra model (\mathbb{S}, s_I) iff for some subset P of *some preimage model* (\mathbb{S}', s'_I) , the model $(\mathbb{S}'[p \mapsto P], s'_I)$ is accepted by \mathbb{A} .

Proposition 6.2. *Let $\mathbf{X} \subseteq \mathbf{Y}$ be sets of proposition letters, both not containing the letter p . Then for any disjunctive $(\Lambda, \mathbf{X}p)$ -automaton \mathbb{A} and any pointed \mathbb{T} -model (\mathbb{S}, s_I) over \mathbf{Y} :*

$$\mathbb{S}, s_I \Vdash \mathbb{A}^{\exists p} \text{ iff } \mathbb{S}', s'_I \Vdash_s \mathbb{A} \text{ for some } \mathbf{Y}p\text{-model } (\mathbb{S}', s'_I) \text{ such that } \mathbb{S}' \upharpoonright_{\mathbf{Y}}, s'_I \rightrightarrows \mathbb{S}, s_I. \quad (6.2)$$

Proof. We only prove the direction from left to right, leaving the other (easier) direction as an exercise to the reader. For notational convenience we assume that $\mathbf{X} = \mathbf{Y}$.

By Proposition 4.2 it suffices to assume that (\mathbb{S}, s_I) is *strongly* accepted by $\mathbb{A}^{\exists p}$ and find a subset U of S for which we can prove that $\mathbb{S}[p \mapsto U], s_I \Vdash_s \mathbb{A}$. So let f be a dividing winning strategy for \exists in $\mathcal{A}(\mathbb{A}^{\exists p}, \mathbb{S})@_{(a_I, s_I)}$ witnessing that $\mathbb{S}, s_I \Vdash_s \mathbb{A}^{\exists p}$. Call a point $s \in S$ *f-accessible* if there is a state $a \in A$ such that the position (a, s) is *f-reachable*; since this state is unique by the assumption of strong acceptance we may denote it as a_s . Clearly any position of the form (a_s, s) is winning for \exists , and hence by legitimacy of f it holds in particular that

$$S, \sigma(s), m_s \Vdash^1 \Theta^{\exists p}(a_s, V^b(s)),$$

where $m_s : S \rightarrow PA$ denotes the marking selected by f at position (a_s, s) . Recalling that $\Theta^{\exists p}(a_s, V^b(s)) = \Theta(a_s, V^b(s)) \vee \Theta(a_s, V^b(s) \cup \{p\})$, we define

$$U := \{s \in S \mid s \text{ is } f\text{-accessible and } S, \sigma(s), m_s \not\Vdash^1 \Theta(a_s, V^b(s))\}.$$

By this we ensure that, for all *f-accessible* points s :

$$s \notin U \text{ implies } S, \sigma(s), m_s \Vdash^1 \Theta(a_s, V^b(s)) \quad (6.3)$$

$$\text{while } s \in U \text{ implies } S, \sigma(s), m_s \Vdash^1 \Theta(a_s, V^b(s) \cup \{p\}) \quad (6.4)$$

Now consider the valuation $V_U := V[p \mapsto U]$, and observe that by this definition we have $V_U^b(s) = V^b(s)$ if $s \notin U$ while $V_U^b(s) = V^b(s) \cup \{p\}$ if $s \in U$. Combining this with (6.3) and (6.4) we find that

$$S, \sigma(s), m_s \Vdash^1 \Theta(a_s, V_U^b)$$

whenever s is *f-accessible*. In other words, f provides a legitimate move m_s in $\mathcal{A}(\mathbb{A}, \mathbb{S})@_{(a_s, s)}$ at any position of the form (a_s, s) . From this it is straightforward to derive that f itself is a (dividing) winning strategy for \exists in $\mathcal{A}(\mathbb{A}, \mathbb{S}[p \mapsto U])@_{(a_I, s_I)}$, and so we obtain that $\mathbb{S}[p \mapsto U], s_I \Vdash_s \mathbb{A}$ as required. \square

The remaining part of the argument follows by a fairly standard argument going back to D'Agostino & Hollenberg [6] (see also Marti et alii [26]) — a minor difference being that our quantification operation $(\cdot)^{\exists p}$ refers to pre-images rather than to bisimilar models.

Proposition 6.3. *Given any proposition letter p , there is a map $\tilde{\exists}p$ on μML_Λ , restricting to ML_Λ , such that $\mathbf{X}_{\tilde{\exists}p.\varphi} = \mathbf{X}_\varphi \setminus \{p\}$ and, for every pointed (\mathbb{S}, s_I) over a set $\mathbf{Y} \supseteq \mathbf{X}_\varphi$ with $p \notin \mathbf{Y}$:*

$$\mathbb{S}, s_I \Vdash \tilde{\exists}p.\varphi \text{ iff } \mathbb{S}', s'_I \Vdash \varphi \text{ for some } \mathbf{Y}p\text{-model } (\mathbb{S}', s'_I) \text{ such that } \mathbb{S}' \upharpoonright_{\mathbf{Y}}, s'_I \rightrightarrows \mathbb{S}, s_I. \quad (6.5)$$

Proof. Straightforward by the equivalence between formulas and Λ -automata, the Simulation Theorem, and Proposition 6.2. \square

Proof of Theorem 6.1. With p_1, \dots, p_n enumerating the proposition letters in $\mathbf{X}_\varphi \setminus \mathbf{X}$, set

$$\varphi_{\mathbf{X}} := \tilde{\exists}p_1 \tilde{\exists}p_2 \cdots \tilde{\exists}p_n.\varphi.$$

Then a relatively routine exercise shows that $\varphi \models \psi$ iff $\varphi_{\mathbf{X}} \models \psi$, for all formulas $\psi \in \mu\text{ML}_\Lambda$ such that $\mathbf{X}_\varphi \cap \mathbf{X}_\psi \subseteq \mathbf{X}$. For some detail, first assume that $\varphi \models \psi$, and take an arbitrary pointed model (\mathbb{S}_0, s_0) over some set $\mathbf{Y} \supseteq \mathbf{X}_\psi \cup \mathbf{X}$ such that $\mathbf{Y} \cap \{p_1, \dots, p_n\} = \emptyset$ and $\mathbb{S}_0, s_0 \Vdash \varphi_{\mathbf{X}}$. Then successive applications of Proposition 6.3 provide, for $i = 1, \dots, n$, pointed \mathbb{T} -models (\mathbb{S}_i, s_i) over $\mathbf{Y} \cup \{p_1, \dots, p_i\}$ such that $(\mathbb{S}_i, s_i) \Vdash \tilde{\exists}p_{i+1} \cdots \tilde{\exists}p_n.\varphi$ and $\mathbb{S}_{i+1} \upharpoonright_{(\mathbf{Y} \cup \{p_1, \dots, p_i\})}, s_{i+1} \rightrightarrows \mathbb{S}_i, s_i$, for all i . Thus in particular we have $\mathbb{S}_n, s_n \Vdash \varphi$, from which it follows by assumption that

$\mathbb{S}_n, s_n \Vdash \psi$, and since $\mathbb{S}_n \upharpoonright_Y, s_n \rightrightarrows \mathbb{S}_0, s_0$ by transitivity of \rightrightarrows , this implies that $\mathbb{S}_0, s_0 \Vdash \psi$, as required. Conversely, to show that $\varphi_X \models \psi$ implies $\varphi \models \psi$, it suffices to prove that $\varphi \models \varphi_X$; but this is almost immediate from the definitions.

Finally, it is not difficult to verify that φ_X is fixpoint-free if φ is so; that is, the uniform interpolants of a formula in ML_Λ also belong to ML_Λ . \square

7. A JANIN-WALUKIEWICZ THEOREM VIA DISJUNCTIVE BASES

Our goal in this section is to prove a coalgebraic generalization of the Janin-Walukiewicz theorem for the modal μ -calculus, using disjunctive bases. In joint work with Fatemeh Seifan [9], we have previously proved such results using automata for coalgebraic MSO and assuming some conditions on a version of the modal one-step language extended with monadic second-order quantifiers. The result that we present here uses only a condition on the modal one-step language itself, and avoids the use of automata for the MSO-language. This also means that we get a new proof of the Janin-Walukiewicz theorem for the μ -calculus, which proceeds via a strong form of bisimulation quantifiers.

In this section we will not assume that the coalgebraic type functor \mathbb{T} preserves inclusion maps, since the notion of a *support* will play an important role and so we want to be as explicit about it as possible. Instead we make the weaker assumption that \mathbb{T} preserves all monics in Set .⁵

Definition 7.1. Given $\xi \in \mathbb{T}X$, a *support* of ξ is a subset $X' \subseteq X$ such that there is some $\xi' \in \mathbb{T}X'$ with $\mathbb{T}\iota(\xi') = \xi$, where $\iota : X' \rightarrow X$ is the inclusion map. Since \mathbb{T} preserves monics, if X' is a support of $\xi \in \mathbb{T}X$ then there is a *unique* $\xi' \in \mathbb{T}X'$ with $\mathbb{T}\iota(\xi') = \xi$; we shall denote this object by $\xi|_{X'}$.

We say that \mathbb{T} *naturally admits minimal supports* if there is a natural transformation $\text{sup} : \mathbb{T} \Rightarrow \mathbb{P}$ such that $\text{sup}_X(\xi)$ is the smallest support for ξ for any $\xi \in \mathbb{T}X$.

It follows by results of Gumm [16] that a finitary functor naturally admits minimal supports iff it weakly preserves *pre-images*; thus for instance, all finitary functors preserving weak pullbacks have the property. Functors lacking this property include the monotone neighborhood functor \mathbb{M} . The reason that we require the map sup to be natural is to ensure that it gives rise to a predicate lifting for the functor \mathbb{T} . In particular, if a modal signature contains this predicate lifting then we can reason about supports within the object language. More generally:

Definition 7.2. If \mathbb{T} naturally admits minimal supports, we say that the modal signature Λ for \mathbb{T} has the *support modality* if there is a positive one-step formula $\Box a$ such that a one-step model (X, ξ, m) satisfies $\Box a$ iff $v \Vdash^0 a$ for all $v \in \text{sup}_X(\xi)$.

For example, the powerset functor has the support modality \Box witnessed by the identity natural transformation on \mathbb{P} . The multi-set functor has the support modality \Box^1 witnessed by the natural transformation mapping $\xi \mapsto \{v \in X \mid \xi(v) > 0\}$ for $\xi \in \mathbb{B}X$. More generally, it follows by the earlier mentioned results of Gumm [16] that any Lyndon complete modal signature for a finitary functor that weakly preserves pre-images has the support modality.

Coalgebraic monadic second-order logic was introduced in [9].

⁵This is a harmless assumption as we can modify \mathbb{T} into a functor \mathbb{T}' that does preserve monics, only by changing its value on the empty set. Since there is only one coalgebra on the empty set, the collection of coalgebras for \mathbb{T} and \mathbb{T}' is the same.

Definition 7.3. Given a modal signature Λ , we define the syntax of the monadic second-order logic MSO_Λ by the following grammar:

$$\varphi := \text{sr}(p) \mid p \subseteq q \mid \lambda(p, q_1, \dots, q_n) \mid \varphi \vee \psi \mid \neg\varphi \mid \exists p.\varphi$$

where λ is any n -place monotone predicate lifting in Λ and $p, q, q_1, \dots, q_n \in \text{Var}$. For the semantics, let (\mathbb{S}, s) be a pointed \mathbb{T} -model, where $\mathbb{S} = (S, \sigma, V)$. We define the satisfaction relation $\Vdash \subseteq S \times \text{MSO}_\Lambda$ as follows:

$$\begin{aligned} (\mathbb{S}, s) \Vdash \text{sr}(p) & \quad \text{iff } V(p) = \{s\} \\ (\mathbb{S}, s) \Vdash p \subseteq q & \quad \text{iff } V(p) \subseteq V(q), \\ (\mathbb{S}, s) \Vdash \lambda(p, q_1, \dots, q_n) & \quad \text{iff } \sigma(v) \in \lambda_S(V(q_1), \dots, V(q_n)) \text{ for all } v \in V(p) \\ (\mathbb{S}, s) \Vdash \exists p.\varphi & \quad \text{iff } (S, \sigma, V[p \mapsto Z], s) \Vdash \varphi, \text{ some } Z \subseteq S, \end{aligned}$$

while we use standard clauses for the Boolean connectives.

The problem we address in this section is to compare the expressive power of coalgebraic monadic second-order logic to that of the coalgebraic μ -calculus.

Definition 7.4. A formula $\varphi \in \text{MSO}_\Lambda$ is *bisimulation invariant* if for all pointed \mathbb{T} -models (\mathbb{S}, s) and (\mathbb{S}', s') it holds that $\mathbb{S}, s \Vdash \varphi$ and $(\mathbb{S}, s) \simeq (\mathbb{S}', s')$ implies $\mathbb{S}', s' \Vdash \varphi$.

Our goal is to prove the following theorem:

Theorem 7.1. *Suppose that \mathbb{T} naturally admits minimal supports, and that the modal signature Λ admits a uniform disjunctive basis, and has the support modality. Then μML_Λ is the bisimulation invariant fragment of MSO_Λ , i.e., a formula of MSO_Λ is bisimulation invariant if and only if it is equivalent to a formula in μML_Λ .*

The assumption that \mathbb{T} naturally admits minimal supports and Λ has the support modality is used to ensure that we can express hidden first-order quantifiers such as in formulas $p \subseteq q$, “up to behavioural equivalence”, in the language μML_Λ . The following observation provides the easy part of the theorem — we omit the proof, which is routine.

Proposition 7.2. *Let Λ be a set of monotone predicate liftings for the set functor \mathbb{T} . There is an inductively defined translation $(\cdot)^\diamond$ mapping any formula $\varphi \in \mu\text{ML}_\Lambda$ to a semantically equivalent formula $\varphi^\diamond \in \text{MSO}_\Lambda$.*

For the opposite direction, we need to slightly refine the construction from the proof of Proposition 4.2. Fix a uniform disjunctive basis \mathbb{D} . For every one-step frame (X, ξ) , we first make use of the assumptions that \mathbb{T} naturally admits minimal supports, and then let

$$f_\xi : (X_*, \xi_*) \rightarrow (\text{sup}_X(\xi), \xi|_{\text{sup}_X(\xi)})$$

be the one-step frame morphism witnessing that \mathbb{D} is a uniform disjunctive basis, for the case of the one-step frame $(\text{sup}_X(\xi), \xi|_{\text{sup}_X(\xi)})$. Then $\iota_{\text{sup}_X(\xi), X} \circ f_\xi : (X_*, \xi_*) \rightarrow (X, \xi)$ is also a one-step frame morphism. Note furthermore that the map $f_\xi : X_* \rightarrow \text{sup}_X(\xi)$ has to be surjective, since otherwise $f_\xi[X_*]$ is a support for $\mathbb{T}(\iota_{\text{sup}_X(\xi), X} \circ f_\xi)(\xi_*)$ in X which is smaller than $\text{sup}_X(\xi)$, which contradicts the definition.

Definition 7.5. Given a pointed \mathbb{T} -model (\mathbb{S}, s) , let the *supporting Kripke frame* be $(S, R_\mathbb{S})$ where $vR_\mathbb{S}w$ iff $w \in \text{sup}_S(\sigma(v))$.

The following proposition is similar to Proposition 4.2 in Section 4, but note that there are a few crucial differences. In particular, conditions (1) – (3) state conditions for the supporting Kripke frames involved in the construction, which only make sense when \mathbb{T}

naturally admits minimal supports. Furthermore, we are making use of the existence of a *uniform* disjunctive basis to ensure that the construction is dependent only on the input T-model (\mathbb{S}, s) , and cannot depend on an automaton \mathbb{A} as an extra parameter.

Proposition 7.3. *There is a construction that provides, for every T-model (\mathbb{S}, s) , a pre-image $f_{\mathbb{S}} : (\mathbb{S}_*, s_*) \rightarrow (\mathbb{S}, s)$ such that:*

- (1) *the structure $(\mathbb{S}_*, R_{\mathbb{S}_*})$ is a tree rooted at s_* ;*
- (2) *the map $f_{\mathbb{S}}$ is both a T-model homomorphism and a morphism of the supporting Kripke frames;*
- (3) *$f_{\mathbb{S}}$ maps S_* surjectively onto the set of points in \mathbb{S} that are reachable from s via the reflexive transitive closure of the support relation $R_{\mathbb{S}}$;*
- (4) *any disjunctive Λ -automaton \mathbb{A} accepts (\mathbb{S}, s) iff it strongly accepts (\mathbb{S}_*, s_*) .*

Proof. The construction of the pre-image $f_{\mathbb{S}} : (\mathbb{S}_*, s_*) \rightarrow (\mathbb{S}, s)$ essentially follows the construction used in the proof of Proposition 5.7 in [9], from which the first two conditions follow. Given that we have constructed the covers $\iota_{\text{sup}_X(\xi), X} \circ f_{\xi} : (X_*, \xi_*) \rightarrow (X, \xi)$ so that they factor epi-monically through the natural inclusion of supports $\iota_{\text{sup}_X(\xi), X} : \text{sup}_X(\xi) \rightarrow X$, it is easy to verify condition (3). For condition (4), the argument closely follows that in the proof of Proposition 4.2 in Section 4. We omit the details. \square

Proof of Theorem 7.1. We provide a recursive translation $t : \text{MSO}_{\Lambda}(X) \rightarrow \mu\text{ML}_{\Lambda}(X)$ such that for any formula φ and any pointed T-model (\mathbb{S}, s) , we have:

$$\mathbb{S}_*, s_* \Vdash \varphi \text{ iff } \mathbb{S}, s \Vdash t(\varphi). \quad (7.1)$$

If φ is bisimulation invariant, it then follows that $\mathbb{S}, s \Vdash \varphi$ iff $\mathbb{S}_*, s_* \Vdash \varphi$ iff $\mathbb{S}, s \Vdash t(\varphi)$, so that φ is equivalent to $t(\varphi)$ as required.

The translation is defined as follows:

φ	$t(\varphi)$
$\text{sr}(p)$	$p \wedge \nu x. \Box(x \wedge \neg p)$
$p \subseteq q$	$\nu x. (p \rightarrow q) \wedge \Box x$
$\lambda(p, q_1, \dots, q_n)$	$\nu x. (p \rightarrow \heartsuit_{\lambda}(q_1, \dots, q_n)) \wedge \Box x$
$\alpha \vee \beta$	$t(\alpha) \vee t(\beta)$
$\neg \alpha$	$\neg t(\alpha)$
$\exists p. \alpha$	$\tilde{\exists} p. t(\alpha)$

where the map $\tilde{\exists} p : \mu\text{ML}_{\Lambda}(Xp) \rightarrow \mu\text{ML}_{\Lambda}(X)$ is defined precisely as in Section 6. We can now prove the equivalence (7.1) by induction on the complexity of formulas of MSO_{Λ} . The first three atomic cases for $\text{sr}(p)$, $p \subseteq q$ and $\lambda(p, q_1, \dots, q_n)$ can be checked by routine verification, where the proofs for the first two cases uses the fact that the map $f_{\mathbb{S}}$ maps the points of \mathbb{S}_* surjectively onto the points in \mathbb{S} , the second case additionally uses the fact that \mathbb{S}_* is constructed so that its supporting Kripke frame is a tree rooted at s_* (see [9] for details), and the case for $\lambda(p, q_1, \dots, q_n)$ uses the fact that $f_{\mathbb{S}}$ is a T-model homomorphism. The induction steps for boolean connectives are trivial. Finally, the proof for the crucial case of $\exists p. \varphi$ closely follows that of Proposition 6.2, so we omit the details. \square

8. DISJUNCTIVE BASES FOR COMBINED MODAL LOGICS

An important topic in modal logic concerns methods to construct complex logics from simpler ones, in such a way that desirable metalogical properties transfer to a combined logic

from its components. For an overview of this field, see [23]. This thread has also been picked up in coalgebraic modal logic: in a particularly interesting paper [5], Cîrstea & Pattinson provide generic methods to obtain, from modal signatures for two set functors T_1 and T_2 , modal signatures for the coproduct, product, and composition of T_1 and T_2 . Since the verification that a particular modal logic has a disjunctive basis can be quite non-trivial — as witnessed here by the case of graded modal logic — it will be useful to have some methods available that guarantee the existence of a disjunctive basis being preserved by Cîrstea & Pattinson’s modular constructions of signatures. In fact, we will show that disjunctive bases can be constructed in a modular fashion as well: in each of the cases of coproduct, product and composition we will give an explicit construction of a disjunctive basis for the combined signature, based on disjunctive bases for the composing functors.

8.1. Coproduct. Let $\mathsf{T} = \mathsf{T}_1 + \mathsf{T}_2$ be the coproduct of the set functors T_1 and T_2 . For $i = 1, 2$ we will use κ_i to denote the natural transformation $\kappa_i : \mathsf{T}_i \Rightarrow \mathsf{T}$ that instantiates to the coproduct insertion map $(\kappa_i)_S : \mathsf{T}_i S \rightarrow \mathsf{T}S$, for every set S . Now suppose that we have been given signatures for T_1 and T_2 , respectively. Following Cîrstea & Pattinson [5], we define the combined signature $\Lambda_1 \oplus \Lambda_2$ for T as follows.

Definition 8.1. Where Λ_1 and Λ_2 are monotone modal signatures for T_1 and T_2 respectively, we define

$$\Lambda_1 \oplus \Lambda_2 := \{\kappa_i \circ \lambda \mid \lambda \in \Lambda_i\}.$$

In the syntax we shall write $\circ_{i,\lambda}$ rather than $\heartsuit_{\kappa_i \circ \lambda}$.

We leave it for the reader to check that $\Lambda_1 \oplus \Lambda_2$ is indeed a collection of monotone predicate liftings for T . The meaning of the $\Lambda_1 \oplus \Lambda_2$ -modalities in an arbitrary T -coalgebra \mathbb{S} is given as follows:

$$\mathbb{S}, s \Vdash \circ_{i,\lambda}(\varphi_1, \dots, \varphi_n) \text{ iff } \begin{array}{l} \sigma(s) = (\kappa_i)_S(\sigma^i) \text{ for some } \sigma^i \in \mathsf{T}_i S \\ \text{with } \sigma^i \in \lambda_S(\llbracket \varphi_1 \rrbracket^{\mathbb{S}}, \dots, \llbracket \varphi_n \rrbracket^{\mathbb{S}}). \end{array} \quad (8.1)$$

Note that the subscript i in $\circ_{i,\lambda}$ works as a *tag* indicating to which part of the coproduct $\mathsf{T}S = \mathsf{T}_1 S + \mathsf{T}_2 S$ the unfolding $\sigma(s)$ of the state s is situated.

The result on the existence of disjunctive bases that we want to prove is the following.

Theorem 8.1. *Suppose both signatures Λ_1 and Λ_2 admit a disjunctive basis. Then so does $\Lambda_1 \oplus \Lambda_2$.*

We start with giving a disjunctive basis for the combined signature.

Definition 8.2. Fix a set $A = \{a_1, \dots, a_n\}$. Given a one-step formula α in the language $\mathbf{1ML}_{\Lambda_i}(A)$, we let $\alpha^i \in \mathbf{1ML}_{\Lambda_1 \oplus \Lambda_2}(A)$ denote the result of replacing every occurrence of a modality \heartsuit_λ with $\circ_{i,\lambda}$. Define

$$(\mathsf{D}_1 \oplus \mathsf{D}_2)(A) := \{\delta_1^1 \vee \delta_2^2 \in \mathbf{1ML}_{\Lambda_1 \oplus \Lambda_2}(A) \mid \delta_i \in \mathsf{D}_i(A)\}.$$

where D_1 and D_2 are disjunctive bases for Λ_1 and Λ_2 , respectively.

It remains to show that $\mathsf{D}_1 \oplus \mathsf{D}_2$ is a disjunctive basis for the signature $\Lambda_1 \oplus \Lambda_2$.

Proof of Theorem 8.1. We first show that $\mathsf{D}_1 \oplus \mathsf{D}_2(A)$ consists of disjunctive formulas. For this purpose we fix a set A and an arbitrary formula in $\mathsf{D}_1 \oplus \mathsf{D}_2(A)$, say, δ^1 such that $\delta \in \mathsf{D}_1(A)$. Let (S, σ, m) be an arbitrary one-step T -model such that $S, \sigma, m \Vdash^{-1} \delta^1$, and make

a case distinction. If $\sigma = (\kappa_2)_S(\sigma^2)$ for some $\sigma^2 \in \mathbb{T}_2 S$ then a straightforward induction will show that $S, \sigma, m' \Vdash^1 \delta^1$, where $m'(s) := \emptyset$ for every $s \in S$.

If, on the other hand, $\sigma = (\kappa_1)_S(\sigma^1)$ for some $\sigma^1 \in \mathbb{T}_1 S$ then a routine inductive proof will reveal that the one-step \mathbb{T}_1 -model (S, σ^1, m) satisfies δ . By disjointness of \mathbb{D}_1 we then obtain a separating cover for (S, σ^1, m) , consisting of a one-step \mathbb{T}_1 -model (S', σ', m') and a map $f : S' \rightarrow S$. It is then easy to verify that the one-step \mathbb{T} -model $(S', \kappa_1(\sigma'), m')$, together with the same map f , is a separating cover for the one-step \mathbb{T} -model (S, σ, m) .

It is left to check that $\mathbb{D}_1 \oplus \mathbb{D}_2(A)$ satisfies the closure conditions of disjoint bases. We leave it to the reader to verify that (modulo equivalence) the set $\mathbb{D}_1 \oplus \mathbb{D}_2(A)$ is closed under taking disjunctions, and contains the formula \top . For condition (2), take a formula of the form $\bigcirc_{i, \lambda} \bar{\pi}$, where λ is a predicate lifting in Λ_i ; without loss of generality assume $i = 1$. Then by assumption there is a formula $\delta \in \mathbb{D}_1(\mathcal{P}A)$ such that $\heartsuit_{\lambda} \bar{\pi} \equiv^1 \delta[\wedge_A]$. It is straightforward to verify that $\alpha \equiv^1 \beta$ implies $\alpha^j \equiv^1 \beta^j$, for any pair of formulas $\alpha, \beta \in \mathbf{1ML}_{\Lambda_j}^+(A)$. But then it is immediate that $\bigcirc_{1, \lambda} \bar{\pi} \equiv^1 \delta^1[\wedge_A] \equiv^1 (\delta^1 \vee \perp)[\wedge_A]$; clearly this suffices, since $\delta^1 \vee \perp \in \mathbb{D}_1 \oplus \mathbb{D}_2(\mathcal{P}A)$.

Finally, for condition (3), consider the conjunction of two formulas in the sets $\mathbb{D}_1 \oplus \mathbb{D}_2(A)$ and $\mathbb{D}_1 \oplus \mathbb{D}_2(B)$, respectively. Using the distributive law of conjunctions over disjunctions, we may rewrite this conjunction into an equivalent disjunction of formulas of the form $\alpha^i \wedge \beta^j$, where $\alpha \in \mathbb{D}_i(A)$ and $\beta \in \mathbb{D}_j(B)$ for some $\{i, j\} \subseteq \{1, 2\}$. Clearly then it suffices to show that each conjunction of the latter form can be rewritten into the required shape. We distinguish two cases.

If $i = j$, then since \mathbb{D}_i is a disjoint basis for Λ_i , there is a formula $\gamma^i \in \mathbb{D}_i(A \uplus B)$ such that $\alpha \wedge \beta \equiv^1 \gamma^i[\theta_{A,B}]$. It is then straightforward to verify that $\alpha^i \wedge \beta^i \equiv^1 \gamma^i[\theta_{A,B}] \equiv^1 (\gamma^i \vee \perp)[\theta_{A,B}]$.

If, on the other hand, i and j are distinct, then we may without loss of generality assume that $i = 1$ and $j = 2$. We claim that in fact for *any* pair of formulas $\alpha \in \mathbf{1ML}_{\Lambda_1}^+(A)$ and $\beta \in \mathbf{1ML}_{\Lambda_2}^+(B)$ (i.e., we do not need to assume that $\alpha \in \mathbb{D}_1(A)$ or $\beta \in \mathbb{D}_2(B)$), the conjunction $\alpha^1 \wedge \beta^2$ is equivalent to a formula from the set $\{\top, \perp, \alpha^1, \beta^2\}$. The key observation in the proof of this claim is that if $\alpha = \heartsuit_{\lambda} \bar{\pi}$ and $\beta = \heartsuit_{\eta} \bar{\rho}$, then $\alpha^1 \wedge \beta^2 \equiv \perp$ — this easily follows from the observation about tagging that we made just after (8.1). Finally, note that $\{\top, \perp, \alpha^1, \beta^2\} \subseteq \mathbb{D}_1 \oplus \mathbb{D}_2(A \cup B)$. This means that every formula γ in this set belongs to $\mathbb{D}_1 \oplus \mathbb{D}_2(A \uplus B)$, and in addition satisfies that $\gamma[\theta_{A,B}] = \gamma$. But then clearly the claim suffices to find a formula $\gamma \in \mathbb{D}_1 \oplus \mathbb{D}_2(A \uplus B)$ such that $\alpha^1 \wedge \beta^2 \equiv^1 \gamma[\theta_{A,B}]$, as required. \square

8.2. Product. Given two functors $\mathbb{T}_1, \mathbb{T}_2$ and modal signatures Λ_1, Λ_2 for these functors respectively, we construct a new modal signature $\Lambda_1 \otimes \Lambda_2$ that contains the modalities of both Λ_1 and Λ_2 . We want to interpret this combined language on coalgebras that can be seen simultaneously as both \mathbb{T}_1 - and \mathbb{T}_2 -coalgebras, and the natural choice is to form the product of the two functors and consider $\mathbb{T}_1 \times \mathbb{T}_2$ -coalgebras.

Definition 8.3. Suppose Λ_1, Λ_2 are modal signatures for functors $\mathbb{T}_1, \mathbb{T}_2$ respectively. Then the modal signature $\Lambda_1 \otimes \Lambda_2$ for $\mathbb{T}_1 \times \mathbb{T}_2$ is defined by:

$$= \Lambda_1 \otimes \Lambda_2 = \{\check{\mathbf{P}}\pi_i \circ \lambda \mid \lambda \in \Lambda_i, i \in \{1, 2\}\}$$

where $\pi_1 : \mathbb{T}_1 \times \mathbb{T}_2 \rightarrow \mathbb{T}_1$ and $\pi_2 : \mathbb{T}_1 \times \mathbb{T}_2 \rightarrow \mathbb{T}_2$ are the natural projection maps.

For simplicity of notation we will allow a small bit of imprecision and just use the symbol λ to denote the predicate lifting $\check{P}\pi_i \circ \lambda$, given a lifting λ for \mathbb{T}_i . This means that we can view the languages $1\text{ML}_{\Lambda_1}(A)$ and $1\text{ML}_{\Lambda_2}(A)$ as fragments of $1\text{ML}_{\Lambda_1 \otimes \Lambda_2}(A)$. Our goal is to prove the following.

Theorem 8.2. *Suppose both signatures Λ_1 and Λ_2 admit a disjunctive basis. Then so does $\Lambda_1 \otimes \Lambda_2$.*

The disjunctive basis for the combined signature is defined as follows.

Definition 8.4. Let $\mathbb{D}_1, \mathbb{D}_2$ be disjunctive bases for Λ_1, Λ_2 respectively. Given a set of variables A we define the set of one-step formulas $(\mathbb{D}_1 \otimes \mathbb{D}_2)(A)$ to be all finite disjunctions of formulas of the form $\delta_1 \wedge \delta_2$, where $\delta_1 \in \mathbb{D}_1(A)$ and $\delta_2 \in \mathbb{D}_2(A)$.

Note that $\top \in (\mathbb{D}_1 \otimes \mathbb{D}_2)(A)$, since it is in $\mathbb{D}_1(A)$ and $\mathbb{D}_2(A)$, given that we allow a slight abuse of notation and identify the conjunction $\top \wedge \top$ with \top .

We will show that $\mathbb{D}_1 \otimes \mathbb{D}_2$ is a disjunctive basis for $\Lambda_1 \otimes \Lambda_2$ indeed. The first thing we need to check is that these formulas are indeed disjunctive.

Proposition 8.3. *Let $\delta_1 \in \mathbb{D}_1(A)$ and $\delta_2 \in \mathbb{D}_2(A)$. Then the formula $\delta_1 \wedge \delta_2$ is disjunctive.*

Proof. Let $(X, (\xi_1, \xi_2), m)$ be a one-step $\mathbb{T}_1 \times \mathbb{T}_2$ -model satisfying the formula $\delta_1 \wedge \delta_2$. Then $X, \xi_1, m \Vdash \delta_1$ and $X, \xi_2, m \Vdash \delta_2$. Since δ_1, δ_2 are disjunctive there exist two covering one-step frames $h_1 : (Y_1, \rho_1) \rightarrow (X, \xi_1)$ and $h_2 : (Y_2, \rho_2) \rightarrow (X, \xi_2)$ together with markings $m_1 : Y_1 \rightarrow \text{PX}$ and $m_2 : Y_2 \rightarrow \text{PX}$ such that for each $i \in \{1, 2\}$:

- $Y_i, \rho_i, m_i \Vdash \delta_i$,
- $m_i(u) \subseteq m(h_i(u))$ for all $u \in Y_i$,
- $|m_i(u)| \leq 1$ for all $u \in Y_i$.

We shall construct a covering one-step $\mathbb{T}_1 \times \mathbb{T}_2$ -frame based on the co-product $Y_1 + Y_2$ of the sets Y_1, Y_2 , with the covering map given by the co-tuple of the maps h_1, h_2 , as shown in Figure 2. We define the covering one-step frame to be $(Y_1 + Y_2, ((\mathbb{T}_1 i_1)\rho_1, (\mathbb{T}_2 i_2)\rho_2))$, where

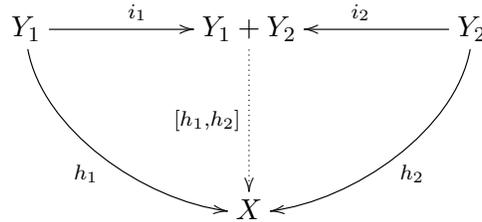


Figure 2: The covering map $[h_1, h_2]$

i_1, i_2 are the insertions of the co-product, so that $[h_1, h_2] \circ i_1 = h_1$ and $[h_1, h_2] \circ i_2 = h_2$. That the co-tuple $[h_1, h_2]$ is in fact a one-step frame morphism is shown by a straightforward calculation:

$$\begin{aligned}
 (\mathbb{T}_1 \times \mathbb{T}_2)[h_1, h_2](\mathbb{T}_1 i_1 \rho_1, \mathbb{T}_2 i_2 \rho_2) &= (\mathbb{T}_1 [h_1, h_2] \circ \mathbb{T}_1 i_1(\rho_1), \mathbb{T}_2 [h_1, h_2] \circ \mathbb{T}_2 i_2(\rho_2)) \\
 &= (\mathbb{T}_1 ([h_1, h_2] \circ i_1)(\rho_1), \mathbb{T}_2 ([h_1, h_2] \circ i_2)(\rho_2)) \\
 &= (\mathbb{T}_1 h_1(\rho_1), \mathbb{T}_2 h_2(\rho_2)) \\
 &= (\xi_1, \xi_2)
 \end{aligned}$$

The commutative diagram shown in Figure 3 may help to give an overview of the construction, in which the top middle entry shows the type of the covering one-step frame we have constructed.

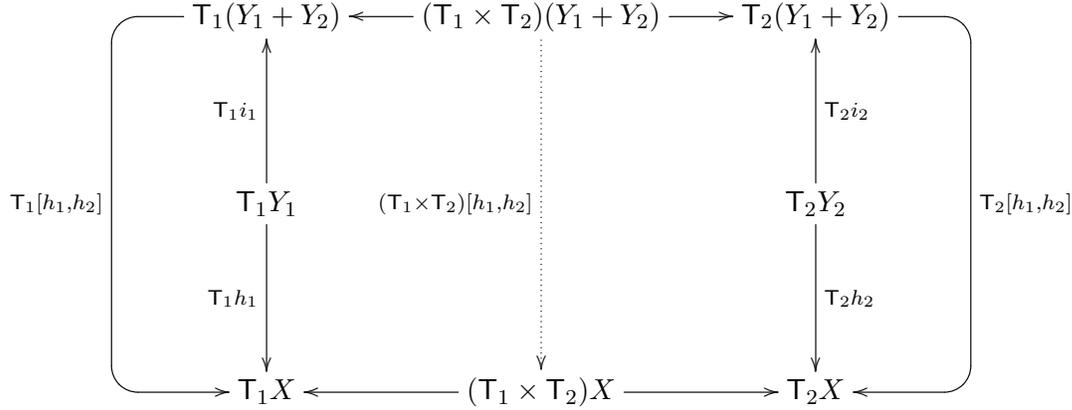


Figure 3: The map $[h_1, h_2]$ is a one-step frame morphism

The construction is completed by defining a marking on the covering one-step $\mathsf{T}_1 \times \mathsf{T}_2$ -frame by co-tupling the markings for each covering T_i -frame, to obtain the marking $[m_1, m_2]$. Since each of the markings m_1, m_2 factor through the co-tuple $[m_1, m_2]$ via the insertions, we have $[m_1, m_2](u) \subseteq m([h_1, h_2](u))$ for each $u \in Y_1 + Y_2$, and also $|[m_1, m_2](u)| \leq 1$ for each $u \in Y_1 + Y_2$.

By naturality of one-step formulas, and again since the markings m_1, m_2 factor through the co-tuple $[m_1, m_2]$ via the insertions, we have

$$Y_1 + Y_2, \mathsf{T}_1 i_1(\rho_1), [m_1, m_2] \Vdash^1 \delta_1$$

and

$$Y_1 + Y_2, \mathsf{T}_2 i_2(\rho_2), [m_1, m_2] \Vdash^1 \delta_2$$

It easily follows that:

$$Y_1 + Y_2, (\mathsf{T}_1 i_1(\rho_1), \mathsf{T}_2 i_2(\rho_2)), [m_1, m_2] \Vdash^1 \delta_1 \wedge \delta_2$$

as required. So $\delta_1 \wedge \delta_2$ is disjunctive. \square

Proof of Theorem 8.2. Since we know the formulas in $(\mathsf{D}_1 \otimes \mathsf{D}_2)(A)$ are disjunctive, we only need to check conditions (1) – (3) of Definition 3.2 one by one: condition (1) just requires the formulas in a disjunctive bases to contain \top and be closed under disjunctions, which holds by definition of $\mathsf{D}_1 \otimes \mathsf{D}_2$. For condition (2), consider a one-step formula in $\mathbf{1ML}_{\Lambda_1 \otimes \Lambda_2}(A)$ of the form $\heartsuit_{\lambda} \bar{\pi}$ where $\lambda \in \Lambda_1$ or $\lambda \in \Lambda_2$. Suppose the former is the case. Then there is a formula $\delta \in \mathsf{D}_1(\mathcal{P}A)$ such that $\heartsuit_{\lambda} \bar{\pi} \equiv^1 \delta[\wedge_A]$. But then $\delta \wedge \top \in (\mathsf{D}_1 \otimes \mathsf{D}_2)(A)$, and this formula is also one-step equivalent to $\heartsuit_{\lambda} \bar{\pi}$.

Finally, for condition (3), by rewriting positive one-step formulas into disjunctive normal form (treating modalities as atomic) we only need to consider conjunctions of the form:

$$(\delta_1 \wedge \delta_2) \wedge (\delta'_1 \wedge \delta'_2)$$

where $\delta_1 \in \mathsf{D}_1(A)$, $\delta_2 \in \mathsf{D}_2(A)$, $\delta'_1 \in \mathsf{D}_1(B)$, $\delta'_2 \in \mathsf{D}_2(B)$. Apply the distributive law of D_1 to $\delta_1 \wedge \delta'_1$ and that of D_2 to $\delta_2 \wedge \delta'_2$ to find formulas $\gamma_1 \in \mathsf{D}_1(A \times B)$ and $\gamma_2 \in \mathsf{D}_2(A \times B)$ such that

$\delta_1 \wedge \delta'_1 \equiv^1 \gamma_1[\theta_{A,B}]$ and $\delta_2 \wedge \delta'_2 \equiv^1 \gamma_2[\theta_{A,B}]$. The conjunction $\gamma_1 \wedge \gamma_2$ is in $(\mathbf{D}_1 \otimes \mathbf{D}_2)(A \times B)$ and we have:

$$(\delta_1 \wedge \delta_2) \wedge (\delta'_1 \wedge \delta'_2) \equiv^1 (\gamma_1 \wedge \gamma_2)[\theta_{A,B}]$$

as required. \square

8.3. Composition. The third and final example that we consider involves the *composition* $\mathbb{T} = \mathbb{T}_1 \circ \mathbb{T}_2$ of two set functors \mathbb{T}_1 and \mathbb{T}_2 . We first recall Cirstea & Pattinson's definition of the combined signature in the case of composition [5].

Definition 8.5. Let $\lambda \in \Lambda_1$ be an m -ary predicate lifting, and let $\alpha_1, \dots, \alpha_m$ be one-step formulas in $1\mathbf{ML}_{\Lambda_2}(A)$ for some set $A = \{a_1, \dots, a_n\}$. Then we define the n -ary predicate lifting

$$\underline{\lambda\langle\alpha_1, \dots, \alpha_m\rangle} : \check{\mathbf{P}}^n \Rightarrow \check{\mathbf{P}}\mathbb{T}$$

in the obvious way:

$$\underline{\lambda\langle\alpha_1, \dots, \alpha_m\rangle}_S : (U_1, \dots, U_n) \mapsto \lambda_{\mathbb{T}_2 S}(\widehat{\alpha_1}_S(U_1, \dots, U_n), \dots, \widehat{\alpha_m}_S(U_1, \dots, U_n)).$$

We let $\Lambda_1 \odot \Lambda_2$ denote the set of all (finitary) predicate liftings that can be obtained in this way.

In the sequel we will often abbreviate $\Lambda_1 \odot \Lambda_2$ as Λ .

Theorem 8.4. *Suppose both signatures Λ_1 and Λ_2 admit a disjunctive basis. Then so does $\Lambda_1 \odot \Lambda_2$.*

For the definition of the disjunctive basis of the combined signature, fix a set $A = \{a_1, \dots, a_n\}$ and consider the one-step language $1\mathbf{ML}_{\Lambda_1}(1\mathbf{ML}_{\Lambda_2}(A))$. Given a formula α in this language, every occurrence of an m -ary Λ_1 -modality \heartsuit_λ is of the form $\heartsuit_\lambda(\gamma_1, \dots, \gamma_m)$ with each $\gamma_i \in 1\mathbf{ML}_{\Lambda_2}(A)$. If we now replace each such subformula $\heartsuit_\lambda(\gamma_1, \dots, \gamma_m)$ with the formula $\heartsuit_{\lambda\langle\gamma_1, \dots, \gamma_m\rangle}(a_1, \dots, a_n)$, we have associated with α a unique formula $\alpha' \in 1\mathbf{ML}_\Lambda(A)$.

Definition 8.6. Define

$$(\mathbf{D}_1 \otimes \mathbf{D}_2)(A) := \{\delta' \in 1\mathbf{ML}_\Lambda(A) \mid \delta \in \mathbf{D}_1(\mathbf{D}_2(A))\}.$$

where \mathbf{D}_1 and \mathbf{D}_2 are disjunctive bases for Λ_1 and Λ_2 , respectively.

To prove that $\mathbf{D}_1 \otimes \mathbf{D}_2$ is a disjunctive basis for the signature $\Lambda_1 \odot \Lambda_2$ we first show that $\mathbf{D}_1 \otimes \mathbf{D}_2(A)$ consists of disjunctive formulas, for any set A .

Proposition 8.5. *Every formula in $\mathbf{D}_1 \otimes \mathbf{D}_2(A)$ is disjunctive.*

Proof. It suffices to show that for an arbitrary but fixed formula $\delta \in \mathbf{D}_1(\mathbf{D}_2(A))$, the formula $\delta' \in (\mathbf{D}_1 \otimes \mathbf{D}_2)(A)$ is disjunctive.

Let (S, σ, m) be an arbitrary one-step \mathbb{T} -model such that $S, \sigma, m \Vdash^1 \delta^1$, and let $\delta \in \mathbf{D}_1(B)$ and $\eta : B \rightarrow \mathbf{D}_2(A)$ be such that $\delta = \delta_1[\eta]$. Consider the one-step \mathbb{T}_1 -model $(\mathbb{T}_2 S, \sigma, m_\eta)$ where $m_\eta : \mathbb{T}_2 S \rightarrow PB$ is given by

$$m_\eta(\rho) := \{b \in B \mid (S, \rho, m) \Vdash^1 \eta(b)\}.$$

Unravelling the definitions, it is not hard to show that for any formula $\alpha_1 \in 1\mathbf{ML}_{\Lambda_1}(B)$ we have

$$(S, \sigma, m) \Vdash^1 \alpha_1[\eta] \text{ iff } (\mathbb{T}_2 S, \sigma, m_\eta) \Vdash^1 \alpha_1, \quad (8.2)$$

and as an immediate consequence of this we find that

$$(\mathbb{T}_2, \sigma, m_\eta) \Vdash^1 \delta_1.$$

Then by disjointiveness of \mathbb{D}_1 there is a one-step \mathbb{T}_1 -model (Z, ζ, n) , where $\zeta \in \mathbb{T}_1 Z$ and $n : Z \rightarrow \mathbb{P}(B)$, together with a map $f : Z \rightarrow \mathbb{T}_2 S$ such that $(\mathbb{T}_1 f)\zeta = \sigma$; $|n(z)| \leq 1$ and $n(z) \subseteq m_\eta(f(z))$, for all $z \in Z$; and $(Z, \zeta, n) \Vdash^1 \delta_1$.

Define, for $z \in Z$, $\eta_z \in \mathbb{D}(A)$ to be the formula $\eta(b)$ in case $b \in B$ is the unique element of $n(z)$; and set $\eta_z := \top$ if $n(z) = \emptyset$. Observe that for any $z \in Z$, the triple $(S, f(z), m)$ constitutes a \mathbb{T}_2 -model; we claim that

$$(S, f(z), m) \Vdash^1 \eta_z. \quad (8.3)$$

To see this, clearly we only have to pay attention to the case where $n(z) = \{b\}$ for some $b \in B$. But here it is immediate from $n(z) \subseteq m_\eta(f(z))$ that $b \in m_\eta(f(z))$, and so we obtain (8.3) by definition of m_η .

Given (8.3), we now use the disjointiveness of \mathbb{D}_2 to obtain, for each $z \in Z$, a one-step \mathbb{T}_2 -model (S_z, ρ_z, m_z) , with $\rho_z \in \mathbb{T}_2 S_z$ and $m_z : S_z \rightarrow \mathbb{P}(A)$, together with a map $g_z : S_z \rightarrow S$, such that $(\mathbb{T}_2 g_z)(\rho_z) = f(z)$; $|m_z| \leq 1$ and $m_z(t) \subseteq m(f(t))$, for all $t \in S_z$; and $(S_z, \rho_z, m_z) \Vdash^1 \eta_z$.

We are now ready to define the required separating cover for (S, σ, m) . As its domain we will take the coproduct $\coprod_{z \in Z} S_z$, and where for $z \in Z$ we let $i_z : S_z \rightarrow S'$ denote the insertion map, we define the maps $m' : S' \rightarrow \mathbb{P}A$ and $g : S' \rightarrow S$ via co-tupling; in particular, we require $m \circ i_z = m_z$ and $g \circ i_z = g_z$ for all $z \in Z$. For the definition of the distinguished element $\sigma' \in \mathbb{T}S'$, we first define the map $\rho' : Z \rightarrow \mathbb{T}_2 S'$ by setting $\rho'(z) := (\mathbb{T}_2 i_z)\rho_z$. We then put $\sigma' := (\mathbb{T}_1 \rho')\zeta$.

It is obvious that (S', σ', m) is a one-step \mathbb{T} -model; we now check that together with the map $g : S' \rightarrow S$ is indeed a separating cover for (S, σ, m) . First of all, it is obvious that $|m'(s')| \leq 1$ and $m'(z) \subseteq m(g(z))$, for all $s' \in S'$. Second, to check that

$$(\mathbb{T}g)\sigma' = \sigma \quad (8.4)$$

we first observe that $(\mathbb{T}_2 g) \circ \rho' = f$, as can easily be verified: $((\mathbb{T}_2 g) \circ \rho')(z) = (\mathbb{T}_2 g)(\rho'(z)) = (\mathbb{T}_2 g)((\mathbb{T}_2 i_z)\rho_z) = (\mathbb{T}_2 (g \circ i_z))(\rho_z) = (\mathbb{T}_2 g_z)(\rho_z) = f(z)$. We then continue with the following calculation:

$$(\mathbb{T}g)\sigma' = (\mathbb{T}_1 \mathbb{T}_2 g)((\mathbb{T}_1 \rho')\zeta) = ((\mathbb{T}_1 \mathbb{T}_2 g) \circ (\mathbb{T}_1 \rho'))(\zeta) = (\mathbb{T}_1((\mathbb{T}_2 g) \circ \rho'))(\zeta) = (\mathbb{T}_1 f)(\zeta) = \sigma.$$

Finally, we need to prove that

$$(S', \sigma', m') \Vdash^1 \delta'. \quad (8.5)$$

To see this, let $m'_\eta : \mathbb{T}_2 S' \rightarrow \mathbb{P}B$ be the marking given by

$$m'_\eta(\rho') := \{b \in B \mid (S', \rho', m') \Vdash^1 \eta(b)\},$$

and consider the marking $n' : Z \rightarrow \mathbb{P}B$ defined by $n'(z) := m'_\eta(\rho'_z)$. Our key claim is now that $n \subseteq n'$, and we prove this as follows. In case $n(z) = \emptyset$ there is nothing to prove; in

case $n(z) \neq \emptyset$, let $b \in B$ be the unique element of $n(z)$. Then we have

$$\begin{aligned}
b \in n(z) &\Rightarrow b \in m_\eta(f(z)) && \text{(assumptions on } f \text{ and } n) \\
&\Rightarrow (S, f(z), m) \Vdash^1 \eta(b) && \text{(definition } m_\eta) \\
&\Rightarrow (S_z, \rho_z, m_z) \Vdash^1 \eta(b) && ((S_z, \rho_z, m_z) \text{ is cover)} \\
&\Rightarrow (S', \rho'_z, m') \Vdash^1 \eta(b) && \text{(invariance under } i_z, \text{ Prop. 2.1)} \\
&\Rightarrow b \in m'_\eta(\rho'_z) && \text{(definition } m'_\eta) \\
&\Rightarrow b \in n'(z) && \text{(definition } n')
\end{aligned}$$

But by the monotonicity of disjunctive formulas, it follows from $n \subseteq n'$ and $(Z, \zeta, n) \Vdash^1 \delta_1$ that $(Z, \zeta, n') \Vdash^1 \delta_1$. Then by invariance (Proposition 2.1) we find that $(\top_2 S, \sigma, m'_\eta) \Vdash^1 \delta_1$, and from this we may derive (8.5), using an analogous claim to (8.2). \square

Proof of Theorem 8.4. Since we have verified the disjunctivity of all formulas in $\mathsf{D}_1 \odot \mathsf{D}_2(A)$, it remains to check that $\mathsf{D}_1 \odot \mathsf{D}_2(A)$ satisfies the closure conditions of disjunctive bases. For condition (1) this is an immediate consequence of the definitions. It is in fact not very hard to see that $\mathsf{D}_1 \odot \mathsf{D}_2$ meets the other two closure conditions as well, but full proofs are very tedious. In order to avoid convoluted syntax we confine ourselves to somewhat sketchy arguments here.

For condition (2), consider a formula of the form $\heartsuit_{\lambda \langle \alpha_1, \dots, \alpha_m \rangle} (b_1, \dots, b_n)[\pi]$, where $\lambda \in \Lambda_1$, each $\alpha_i \in \mathbf{1ML}_{\Lambda_2}(B)$ and $\pi : B \rightarrow \mathbf{Latt}(A)$. Since D_1 is a disjunctive basis for Λ_1 , we may find formulas $\delta_1 \in \mathsf{D}_1(\{1, \dots, k\})$ and $\beta_1, \dots, \beta_k \in \mathbf{1ML}_{\Lambda_2}(A)$ such that $\heartsuit_{\lambda \langle \alpha_1[\pi], \dots, \alpha_m[\pi] \rangle} \equiv^1 \delta_1(\beta_1, \dots, \beta_k)$. But, now using the fact that D_2 is a disjunctive basis for Λ_2 , we may derive from Proposition 3.2 that each formula β_i is equivalent to a formula $\gamma_i[\wedge A]$, where $\gamma_i \in \mathsf{D}_2(\mathsf{PA})$. It is then a tedious but straightforward exercise to show that $\heartsuit_{\lambda \langle \alpha_1, \dots, \alpha_m \rangle} (b_1, \dots, b_n)[\pi] \equiv^1 \delta'$, where we define $\delta := \delta_1(\gamma_1, \dots, \gamma_k)$.

Finally, for condition (3), consider two disjunctive formulas $\gamma' \in \mathsf{D}_1 \odot \mathsf{D}_2(A)$ and $\delta' \in \mathsf{D}_1 \odot \mathsf{D}_2(B)$, where $\gamma = \gamma_1[\sigma]$ and $\delta = \delta_1[\tau]$ for $\gamma_1 \in \mathsf{D}_1(A')$, $\delta_1 \in \mathsf{D}_1(B')$, and $\sigma : A' \rightarrow \mathsf{D}_2(A)$, $\tau : B' \rightarrow \mathsf{D}_2(B)$ for some sets A' and B' that without loss of generality we may take to be disjoint. D_1 being a disjunctive basis yields a formula $\beta_1 \in \mathsf{D}(A' \uplus B')$ such that $\gamma_1 \wedge \delta_1 \equiv^1 \beta_1[\theta_{A', B'}]$. By the disjointness of A' and B' we then have

$$\gamma' \wedge \delta' \equiv^1 \beta_1[\theta_{A', B'}][\sigma][\tau]. \quad (8.6)$$

Now consider an arbitrary pair $(a', b') \in A' \times B'$; since D_2 is a disjunctive basis there is a formula $\alpha_{a', b'} \in \mathsf{D}_2(A \times B)$ such that $\sigma_{a'} \wedge \tau_{b'} \equiv^1 \alpha_{a', b'}[\theta_{A, B}]$. Define the following substitution $\sigma \uplus \tau : A' \uplus B' \rightarrow \mathsf{D}_2(A \uplus B)$:

$$\sigma \uplus \tau(d) := \begin{cases} \sigma(d) & \text{if } d \in A' \\ \tau(d) & \text{if } d \in B' \\ \alpha_{a', b'} & \text{if } d = (a', b') \in A' \times B' \end{cases}$$

It then follows from the definitions that the substitutions $[\theta_{A', B'}][\sigma][\tau]$ and $[\sigma \uplus \tau][\theta_{A, B}]$ produce one-step equivalent formulas, so that combining this with (8.6) we obtain that

$$\gamma' \wedge \delta' \equiv^1 \beta_1[\sigma \uplus \tau][\theta_{A, B}].$$

This suffices, since obviously we have that $\beta_1[\sigma \uplus \tau] \in \mathsf{D}_1 \odot \mathsf{D}_2(A \uplus B)$. \square

9. YONEDA REPRESENTATION OF DISJUNCTIVE LIFTINGS

It is a well known fact in coalgebraic modal logic that predicate liftings have a neat representation via an application of the Yoneda lemma. This was explored by Schröder in [32], where it was used among other things to prove a characterization theorem for the monotone predicate liftings. Here, we apply the same idea to disjunctive liftings. We shall be working with a slightly generalized notion of predicate lifting here, taking a predicate lifting over a finite set of variables A to be a natural transformation $\lambda : \check{\mathbf{P}}^A \rightarrow \check{\mathbf{P}} \circ \mathbf{T}$. Clearly, one-step formulas in $\mathbf{1ML}_A(A)$ can then be viewed as predicate liftings over A .

Definition 9.1. Let $\lambda : \check{\mathbf{P}}^A \rightarrow \check{\mathbf{P}} \circ \mathbf{T}$ be a predicate lifting over variables $A = \{a_1, \dots, a_n\}$. The *Yoneda representation* $y(\lambda)$ of λ is the subset

$$\lambda_{\mathbf{P}A}(\text{true}_{a_1}, \dots, \text{true}_{a_n}) \in \text{PTPA}$$

where $\text{true}_{a_i} = \{B \subseteq A \mid a_i \in B\}$. We shall write simply $\lambda \subseteq \text{TPA}$ instead of $y(\lambda)$.

Definition 9.2. Given a set A , let A^\top be the set $A \cup \{\top\}$. Let $\epsilon_A \subseteq A^\top \times \mathbf{P}A$ be the relation defined by $a \epsilon_A B$ iff $a \in B$, and $\top \epsilon_A B$ for all $B \subseteq A$. Let $\eta_A : A^\top \rightarrow \mathbf{P}A$ be defined by $\eta_A(a) = \{a\}$, and $\eta_A(\top) = \emptyset$.

In the remainder of this section we assume familiarity with the Barr relation lifting $\overline{\mathbf{T}}$ associated with a functor \mathbf{T} ; see [21] for the definition and some basic properties.

Definition 9.3. A predicate lifting $\lambda \subseteq \text{TPA}$ is said to be *divisible* if, for all $\alpha \in \lambda$ there is some $\beta \in \mathbf{T}A^\top$ such that $(\beta, \alpha) \in \overline{\mathbf{T}}(\epsilon_A)$ and $\mathbf{T}\eta_A(\beta) \in \lambda$.

Proposition 9.1. *Any disjunctive lifting over A is divisible, and if \mathbf{T} preserves weak pullbacks the disjunctive liftings over A are precisely the divisible ones.*

Proof. Suppose $\lambda \subseteq \text{TPA}$ is disjunctive, and pick $\alpha \in \lambda$. Then $\mathbf{P}A, \alpha, \text{id}_{\mathbf{P}A} \Vdash^1 \lambda$, so since λ is disjunctive there are some one-step model (X, ξ, m) and map $f : X \rightarrow \mathbf{P}A$ with $m : X \rightarrow \mathbf{P}A$, $m(u) \subseteq f(u)$ for all $u \in X$, $\mathbf{T}f(\xi) = \alpha$, and $|m(u)| \leq 1$ for all $u \in X$. We define a map $g : X \rightarrow A^\top$ by setting $g : u \mapsto \top$ if $m(u) = \emptyset$, $g : u \mapsto a$ if $m(u) = \{a\}$. We tuple the maps f, g to get a map $\langle f, g \rangle : X \rightarrow A^\top \times \mathbf{P}A$. In fact, since $m(u) \subseteq f(u)$ for all $u \in X$, we have $\langle f, g \rangle : X \rightarrow \epsilon_A$. Let $\pi_1 : \epsilon_A \rightarrow A^\top$ and $\pi_2 : \epsilon_A \rightarrow \mathbf{P}A$ be the projection maps. We have the following diagram, in which the two triangles and the outer edges commute (i.e., $m = \eta_A \circ g$).

$$\begin{array}{ccc}
 & & \mathbf{P}A \\
 & \overset{m}{\curvearrowright} & \uparrow \pi_2 \\
 X & \xrightarrow{f} & \mathbf{P}A \\
 & \searrow \langle f, g \rangle & \uparrow \epsilon_A \\
 & & \epsilon_A \\
 & \xrightarrow{\langle f, g \rangle} & \downarrow \pi_1 \\
 & & A^\top \\
 & \searrow g & \curvearrowright \eta_A
 \end{array}$$

Now apply \mathbf{T} to this diagram and define $\beta \in \mathbf{T}A^\top$ to be $\mathbf{T}(\pi_1 \circ \langle f, g \rangle)(\xi) = \mathbf{T}g(\xi)$. First, we have $(\beta, \alpha) \in \overline{\mathbf{T}}(\epsilon_A)$, witnessed by $\mathbf{T}(\langle f, g \rangle)(\xi) \in \mathbf{T}\epsilon_A$. We claim that $\mathbf{T}\eta_A(\beta) \in \lambda$. But since $X, \xi, m \Vdash^1 \lambda$ and $m = \eta_A \circ g$, naturality of λ applied to the map $g : X \rightarrow A^\top$, gives $A^\top, \beta, \eta_A \Vdash^1 \lambda$. Another naturality argument, applied to $\eta_A : (A^\top, \beta, \eta_A) \rightarrow (\mathbf{P}A, \mathbf{T}\eta_A(\beta), \text{id}_{\mathbf{P}A})$ gives $\mathbf{P}A, \mathbf{T}\eta_A(\beta), \text{id}_{\mathbf{P}A} \Vdash^1 \lambda$, i.e., $\mathbf{T}\eta_A(\beta) \in \lambda$.

For the converse direction, under the assumption that \mathbf{T} preserves weak pullbacks, suppose that λ is divisible, and suppose $X, \xi, m \Vdash^1 \lambda$. We get $\mathbf{T}m(\xi) \in \lambda$ and so we find some

$\beta \in \mathsf{TA}^\top$ with $\beta(\overline{\mathsf{T}\epsilon_A})\mathsf{T}m(\xi)$ and $\mathsf{T}\eta_A(\beta) \in \lambda$. Pick some $\beta' \in \mathsf{T}\epsilon_A$ with $\mathsf{T}\pi_2(\beta') = \mathsf{T}m(\xi)$ and $\mathsf{T}\pi_1(\beta') = \beta$. Let R, g_1, g_2 be the pullback of the diagram $X \rightarrow \mathsf{PA} \leftarrow \epsilon_A$, shown in the diagram:

$$\begin{array}{ccc}
 X & \xrightarrow{m} & \mathsf{PA} \\
 g_1 \uparrow & & \uparrow \pi_2 \\
 R & \xrightarrow{g_2} & \epsilon_A \\
 & & \downarrow \pi_1 \\
 & & A^\top
 \end{array}
 \quad \eta_A$$

By weak pullback preservation there is $\rho \in \mathsf{TR}$ with $\mathsf{T}g_1(\rho) = \xi$ and $\mathsf{T}g_2(\rho) = \beta'$. The map $g_1 : (R, \rho) \rightarrow (X, \xi)$ is thus a cover, and we have a marking m' on R defined by $\eta_A \circ \pi_1 \circ g_2$ (follow the bottom-right path in the previous diagram). It is now routine to check that $R, \rho, m' \Vdash^1 \lambda$, and $|m'(u)| \leq 1$ and $m'(u) \subseteq m(g_1(u))$ for all $u \in R$, so we are done. \square

For the moment, we leave the question open, whether a similar characterization of disjunctive predicate liftings can be proved without weak pullback preservation. We also leave it as an open problem to characterize the functors that admit a disjunctive basis.

Acknowledgement. We thank the two anonymous referees for their detailed and helpful comments.

REFERENCES

- [1] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM (JACM)*, 49(5):672–713, 2002.
- [2] J. Bergfeld. Moss’s coalgebraic logic: Examples and completeness results. Master’s thesis, Institute for Logic, Language and Computation, University of Amsterdam, 2009.
- [3] C. Cirstea, C. Kupke, and D. Pattinson. EXPTIME tableaux for the coalgebraic μ -calculus. In E. Grädel and R. Kahle, editors, *Computer Science Logic (CSL 2009)*, volume 5771 of *Lecture Notes in Computer Science*, pages 179–193. Springer, 2009.
- [4] C. Cirstea and D. Pattinson. Modular construction of modal logics. In Ph. Gardner and N. Yoshida, editors, *Proceedings of the 15th International Conference on Concurrency Theory (CONCUR 2004)*, volume 3170 of *LNCS*, pages 258–275. Springer, 2004.
- [5] C. Cirstea and D. Pattinson. Modular construction of complete coalgebraic logics. *Theoretical Computer Science*, 388(1):83 – 108, 2007.
- [6] G. D’Agostino and M. Hollenberg. Logical questions concerning the μ -calculus. *Journal of Symbolic Logic*, 65:310–332, 2000.
- [7] E.A. Emerson and C.S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *Proceedings of the 32nd Symposium on the Foundations of Computer Science*, pages 368–377. IEEE Computer Society Press, 1991.
- [8] S. Enqvist, F. Seifan, and Y. Venema. Completeness for μ -calculi: a coalgebraic approach. Technical Report PP-2017-04, Institute for Logic, Language and Computation, Universiteit van Amsterdam. *Annals of Pure and Applied Logic*, to appear.
- [9] S. Enqvist, F. Seifan, and Y. Venema. An expressive completeness theorem for coalgebraic modal mu-calculi. *Logical Methods in Computer Science*, Volume 13, Issue 2. An earlier version of this paper appeared in LICS (2015).

- [10] S. Enqvist, F. Seifan, and Y. Venema. Monadic second-order logic and bisimulation invariance for coalgebras. In *Proceedings of the 30th Annual ACM/IEEE Symposium on Logic in Computer Science, (LICS 2015)*, pages 353–365, 2015.
- [11] S. Enqvist, F. Seifan, and Y. Venema. Completeness for coalgebraic fixpoint logic. In *Proceedings of the 25th EACSL Annual Conference on Computer Science Logic (CSL 2016)*, volume 62 of *LIPICs*, pages 7:1–7:19, 2016.
- [12] S. Enqvist and Y. Venema. Disjunctive bases: Normal forms for modal logics. In F. Bonchi and B. König, editors, *Proceedings of the 7th Conference on Algebra and Coalgebra in Computer Science (CALCO 2017)*, volume 72 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 11:1–11:16, 2017.
- [13] G. Fontaine, R. Leal, and Y. Venema. Automata for coalgebras: An approach using predicate liftings. In *Automata, Languages and Programming: 37th International Colloquium ICALP'10*, volume 6199 of *LNCS*, pages 381–392. Springer, 2010.
- [14] G. Fontaine and T. Place. Frame definability for classes of trees in the μ -calculus. In *Proceedings of the 35th International Symposium on Mathematical Foundations of Computer Science (MFCS 2010)*, pages 381–392. Springer, 2010.
- [15] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logic, and Infinite Games*, volume 2500 of *LNCS*. Springer, 2002.
- [16] H.P. Gumm. From T-coalgebras to filter structures and transition systems. In J. L. Fiadeiro, editor, *Algebra and Coalgebra in Computer Science (CALCO 2005)*, volume 3629 of *LNCS*, pages 194–212, 2005.
- [17] D. Janin and G. Lenzi. Relating levels of the μ -calculus hierarchy and levels of the monadic hierarchy. In *Proceedings of the 16th Annual IEEE Symposium on Logic in Computer Science (LICS 2001)*, pages 347–356, 2001.
- [18] D. Janin and I. Walukiewicz. Automata for the modal μ -calculus and related results. In J. Wiedermann and P. Hájek, editors, *Mathematical Foundations of Computer Science 1995, 20th International Symposium (MFCS'95)*, volume 969 of *LNCS*, pages 552–562. Springer, 1995.
- [19] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ -calculus w.r.t. monadic second-order logic. In *Proceedings CONCUR '96*, 1996.
- [20] D. Kozen. Results on the propositional μ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
- [21] C. Kupke, A. Kurz, and Y. Venema. Completeness for the coalgebraic cover modality. *Logical Methods in Computer Science*, 8(3), 2010.
- [22] C. Kupke and Y. Venema. Coalgebraic automata theory: basic results. *Logical Methods in Computer Science*, 4, 2008. (An earlier version of this paper appeared in LICS 2005.).
- [23] A. Kurucz. Combining modal logics. In J. van Benthem, P. Blackburn, and F. Wolter, editors, *Handbook of Modal Logic*, pages 721–756. Elsevier, 2006.
- [24] A. Kurz and R. Leal. Modalities in the stone age: a comparison of coalgebraic logics. *Theoretical Computer Science*, 430:88–116, 2012.
- [25] R.A. Leal. Predicate liftings versus nabla modalities. *Electronic Notes in Theoretical Computer Science*, 203:195–220, 06 2008.
- [26] J. Marti, F. Seifan, and Y. Venema. Uniform interpolation for coalgebraic fixpoint logic. In L.S. Moss and P. Sobocinski, editors, *Proceedings of the Sixth Conference on Algebra and Coalgebra in Computer Science (CALCO 2015)*, volume 35 of *LIPICs*, pages 238–252. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
- [27] L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999. (Erratum published *APAL* 99:241–259, 1999).
- [28] A.M. Mostowski. Games with forbidden positions. Technical Report 78, Instytut Matematyki, Uniwersytet Gdański, Poland, 1991.
- [29] E. Pacuit and S. Salame. Majority logic. In *Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning (KR2004)*, pages 598–605, 2004.
- [30] R. Parikh. The logic of games and its applications. *Annals of Discrete Mathematics*, 24:111–139, 1985.
- [31] D. Pattinson. Coalgebraic modal logic: soundness, completeness and decidability of local consequence. *Theoretical Computer Science*, 309(1–3):177–193, 2003.
- [32] L. Schröder. Expressivity of coalgebraic modal logic: the limits and beyond. *Theoretical Computer Science*, pages 230–247, 2008.
- [33] L. Schröder and D. Pattinson. Pspace bounds for rank-1 modal logics. *ACM Transactions on Computational Logic (TOCL)*, 10(2):13, 2009.

- [34] L. Schröder and D. Pattinson. Rank-1 modal logics are coalgebraic. *Journal of Logic and Computation*, 20(5):1113–1147, 2010.
- [35] Y. Venema. Automata and fixed point logic: a coalgebraic perspective. *Information and Computation*, 204:637–678, 2006.
- [36] Y. Venema. Lectures on the modal μ -calculus. Lecture Notes, ILLC, University of Amsterdam, 2012.
- [37] I. Walukiewicz. Completeness of Kozen’s axiomatisation of the propositional μ -calculus. *Information and Computation*, 157:142–182, 2000.