Disjunctive Bases: Normal Forms for Modal Logics

Sebastian Enqvist
sebastian.enqvist@fil.lu.se

Yde Venema
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 mu-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 mu-calculus, in particular, a simulation 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, and a Uniform Interpolation result, for both the full mu-calculus and its fixpoint-free fragment.

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.

Finally, we consider the problem of giving a category-theoretic formulation of disjunctive bases, and provide a partial solution.

Keywords Modal logic, fixpoint logic, automata, coalgebra, graded modal logic, Lyndon theorem, uniform interpolation.

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 [12], and automata running on infinite objects, is standard [8]. 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 [22], and model theoretic results like expressive completeness [11], uniform interpolation and a Lyndon theorem [3].

The standard modal μ-calculus was generalized to a generic, coalgebraic modal μ-calculi [20], of which the modal basis was provided by Moss' original coalgebraic modality [16], 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 [13], and this was recently extended to the fixpoint extension of the finitary Moss logic [4]. 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 [10]. 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 [18], 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. (For a comparison between the two approaches, see [14].) Coalgebraic μ-calculi have also been developed as extensions of the predicate liftings based languages [2], 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 [6], 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 [5], by ourselves together with Seifan, where we extended our earlier completeness result for Moss-style fixpoint logics [4] 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 [20]. 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 [15]. Finally, we prove a Lyndon theorem for coalgebraic fixpoint logics, generalizing a result for the standard modal μ-calculus proved in [3]: 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.

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 7 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. This is not to say that disjunctive bases are just "nablas in disguise": it is a fundamental concept, and in some cases it is the *right* concept as *opposed* to nabla formulas. As a clear example of this, we consider *graded modal logic*, which adds counting modalities to modal logic. While we will see that this language has a disjunctive basis, at the same time we will prove that no such basis can be based on the nabla modalities.

2 Preliminaries

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 to fix some definitions and notations.

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* $\sigma : S \to \omega$ with finite support). We also need the contravariant powerset functor $\breve{\mathsf{P}}$.

A T-*coalgebra* is a pair $\mathbb{S} = (S, \sigma)$ where S is a set of objects called *states* or *points* and $\sigma : S \to \mathsf{T}S$ 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 $f : S' \to S$ a *coalgebra homomorphism* from (S', σ') to (S, σ) if $\sigma \circ f = \mathsf{T}f \circ \sigma'$, and write $(\mathbb{S}', s') \rightrightarrows (\mathbb{S}, s)$ if there is such a coalgebra morphism mapping s' to s.

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

We will mainly follow the approach in coalgebraic modal logic where modalities are associated (or even identified) with finitary predicate liftings. A *predicate lifting* of arity n is a natural transformation $\lambda : \breve{\mathsf{P}}^n \Rightarrow \breve{\mathsf{P}}\mathsf{T}$. Such a predicate lifting is *monotone* if for every set S, the map $\lambda_S : (\mathsf{P}S)^n \to \mathsf{PT}S$ preserves the subset order in each coordinate. The induced predicate lifting $\lambda^\partial : \mathsf{P}^n \Rightarrow \mathsf{PT}$, given by $\lambda^\partial_S(X_1, \ldots, X_n) := \mathsf{T}S \setminus \lambda_S(S \setminus X_1, \ldots, S \setminus X_1)$, is called the *(Boolean) dual* of λ. 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.

Given a signature Λ, the formulas of the *coalgebraic μ-calculus* $\mu\mathtt{ML}_\Lambda$ are given by the following grammar:

$$\varphi ::= p \mid \bot \mid \neg\varphi \mid \varphi_0 \vee \varphi_1 \mid \heartsuit_\lambda(\varphi_1, \ldots, \varphi_n) \mid \mu x.\varphi'$$

where p and x are propositional variables, $\lambda \in \Lambda$ has arity n, and the application of the fixpoint operator μx is under the proviso that all occurrences of x in φ' are positive (i.e., under an even number of negations). We let $\mathtt{ML}_\Lambda$ and $\mu\mathtt{ML}_\Lambda(\mathtt{X})$ denote, respectively, the fixpoint-free fragment of $\mu\mathtt{ML}_\Lambda$ and the set of $\mu\mathtt{ML}_\Lambda$-formulas taking free variables from $\mathtt{X}$.

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 $\mathtt{X}$ of proposition letters. By induction on the complexity of formulas, we define a *meaning function* $[\![\cdot]\!]^\mathbb{S} : \mu\mathtt{ML}_\Lambda(\mathtt{X}) \to \mathsf{P}S$, together with an associated *satisfaction relation* $\Vdash \subseteq S \times \mu\mathtt{ML}_\Lambda(\mathtt{X})$ given by $\mathbb{S}, s \Vdash \varphi$ iff $s \in [\![\varphi]\!]^\mathbb{S}$. All clauses of this definition are standard; for instance, the one for the modality $\heartsuit_\lambda$ is given by

$$\mathbb{S}, s \Vdash \heartsuit_\lambda(\varphi_1, \ldots, \varphi_n) \text{ if } \sigma(s) \in \lambda_S([\![\varphi_1]\!]^\mathbb{S}, \ldots, [\![\varphi_n]\!]^\mathbb{S}). \tag{1}$$

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

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

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 formulas 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 $[\![\varphi]\!]^{(S,\sigma,V[p\mapsto Z_1])} \subseteq [\![\varphi]\!]^{(S,\sigma,V[p\mapsto Z_2])}$.

Well-known examples of coalgebraic modalities include the next-time operator $\bigcirc$ of linear time temporal logic, the standard Kripkean modalities $\Box$ and $\Diamond$, the more general modalities of monotone modal logic, and the *counting modalities* $\Diamond^k$ and $\Box^k$ of graded modal logic, which can be interpreted over B-coalgebras using the predicate liftings $\underline{k}$ and $\overline{k}$ given by

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

A pivotal role in our approach is filled by the one-step versions of coalgebraic logics. Given a signature Λ and a set A of variables, we define the set $\mathtt{Bool}(A)$ of *boolean formulas* over A and the set

$\mathtt{1ML}_\Lambda(A)$ of *one-step* Λ*-formulas* over A, by the following grammars:

$$\begin{array}{rcl} \mathtt{Bool}(A) \ni \pi & ::= & a \mid \bot \mid \top \mid \pi \vee \pi \mid \pi \wedge \pi \mid \neg\pi \\ \mathtt{1ML}_\Lambda(\mathtt{X}, A) \ni \alpha & ::= & \heartsuit_\lambda \overline{\pi} \mid \bot \mid \top \mid \alpha \vee \alpha \mid \alpha \wedge \alpha \mid \neg\alpha \end{array}$$

where $a \in A$ and $\lambda \in \Lambda$. We will denote the positive (negation-free) fragments of $\mathtt{Bool}(A)$ and $\mathtt{1ML}_\Lambda(A)$ as, respectively, $\mathtt{Latt}(A)$ and $\mathtt{1ML}^+_\Lambda(A)$.

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

A monotone modal signature Λ for T is *expressively complete* if, for every n-place predicate lifting λ and variables $a_1, \ldots, a_n$ there is a formula $\alpha \in \mathtt{1ML}_\Lambda(\{a_1, \ldots, a_n\})$ which is equivalent to $\heartsuit_\lambda \overline{a}$. We will also be interested in the following strengthening of expressive completeness: we say that Λ is *Lyndon complete* if, for every *monotone* n-place predicate lifting λ and variables $a_1, \ldots, a_n$, there is a *positive* formula $\alpha \in \mathtt{1ML}^+_\Lambda(\{a_1, \ldots, a_n\})$ equivalent to $\heartsuit_\lambda \overline{a}$.

One-step formulas are naturally interpreted in the following structures. A *one-step* T*-frame* is a pair (S, σ) with $\sigma \in \mathsf{T}S$, i.e., an object in the category $\mathcal{E}(\mathsf{T})$ of elements of T. Similarly a *one-step* T*-model* over a set A of variables is a triple (S, σ, m) such that (S, σ) is a one-step T-frame and $m : S \to \mathsf{P}A$ is an A-marking on S. *Morphism* of one-step frames and of one-step models are defined in the obvious way.

Given a one-step model (S, σ, m), we define the 0*-step interpretation* $[\![\pi]\!]^0_m \subseteq S$ of $\pi \in \mathtt{Bool}(A)$ by the obvious induction: $[\![a]\!]^0_m := \{v \in S \mid a \in m(v)\}$, $[\![\top]\!]^0_m := S$, $[\![\bot]\!]^0_m := \varnothing$, and the standard clauses for $\wedge, \vee$ and $\neg$. Similarly, the *one-step interpretation* $[\![\alpha]\!]^1_m$ of $\alpha \in \mathtt{1ML}_\Lambda(A)$ is defined as a subset of $\mathsf{T}S$, with $[\![\heartsuit_\lambda(\pi_1, \ldots, \pi_n)]\!]^1_m := \lambda_S([\![\pi_1]\!]^0_m, \ldots, [\![\pi_n]\!]^0_m)$, and standard clauses for $\bot, \top, \wedge, \vee$ and $\neg$. Given a one-step modal (S, σ, m), we write $S, \sigma, m \Vdash^1 \alpha$ for $\sigma \in [\![\alpha]\!]^1_m$. Notions like one-step satisfiability, validity and equivalence are defined in the obvious way.

A $(\Lambda, \mathtt{X})$*-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* $a_I \in A$, $\Theta : A \times \mathsf{P}\mathtt{X} \to \mathtt{1ML}^+_\Lambda(\mathtt{X}, A)$ is the *transition map* and $\Omega : A \to \omega$ is the *priority map* of $\mathbb{A}$. Its semantics is given in terms of a two-player infinite parity game: With $\mathbb{S} = (S, \sigma, V)$ a T-model over a set $\mathtt{Y} \supseteq \mathtt{X}$, the *acceptance game* $\mathcal{A}(\mathbb{A}, \mathbb{S})$ is the parity game given by the table below.

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

We say that $\mathbb{A}$ *accepts* the pointed 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 $\mathcal{A}(\mathbb{A}, \mathbb{S})$.

Fact 2.1 *There are effective constructions transforming a formula in* $\mu\mathtt{ML}_\Lambda(\mathtt{X})$ *into an equivalent* $(\Lambda, \mathtt{X})$*-automaton, and vice versa.*

3 Disjunctive formulas and disjunctive bases

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

Definition 3.1 A one-step formula $\alpha \in \mathtt{1ML}^+_\Lambda(\mathtt{X}, A)$ 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 morphism $f : (S', \sigma') \to (S, \sigma)$ and a marking $m' : S' \to \mathsf{P}A$ 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 $s' \in S'$. ◁

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

(1) $\mathtt{D}(A)$ is closed under finite disjunctions (in particular, it contains $\top = \bigvee \varnothing$).

(2) $\mathtt{D}$ is *distributive over* Λ: for every one-step formula of the form $\heartsuit_\lambda \overline{\pi}$ there is a formula $\delta \in \mathtt{D}(\mathsf{P}(A))$ such that $\heartsuit_\lambda \overline{\pi} \equiv^1 \delta[\wedge_A]$. (3) $\mathtt{D}$ *admits a binary distributive law*: for any two formulas $\alpha \in \mathtt{D}(A)$ and $\beta \in \mathtt{D}(B)$, there is a formula $\gamma \in \mathtt{D}(A \times B)$ such that $\alpha \wedge \beta \equiv^1 \gamma[\theta_{A,B}]$. ◁

Disjunctive bases for weak pullback preserving functors It is not hard to prove that disjunctive formulas generalize the Moss modalities, which are tightly connected to weak pullback preservation of the coalgebraic type functor. (Due to space limitations we refer to [13] for the details on the syntax and semantics of 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 T. If Λ is Lyndon complete, then it admits a disjunctive basis.*

Proof. Let $\mathtt{D}_\nabla(A)$ be the set of all (finite and infinite) disjunctions of formulas of the form $\nabla\beta$, with $\beta \in \mathsf{T}A$. Such disjunctions can be regarded as n-ary predicate liftings, where $|A| = n$, so we can apply expressive completeness and treat them as one-step formulas in $\mathtt{1ML}^+_\Lambda(A)$. As mentioned, it is easy to verify that all formulas of the form $\nabla\beta$ are disjunctive, and since disjunctivity is closed under taking disjunctions, all formulas in $\mathtt{D}_\nabla(A)$ are disjunctive. It remains to show that $\mathtt{D}_\nabla(A)$ is a basis for Λ.

It remains to prove that any formula $\alpha \in \mathtt{1ML}^+_\Lambda(A)$ is equivalent to a (possibly infinite) disjunction of formulas of the form $\nabla\Gamma[\chi_A]$, with $\Gamma \in \mathsf{TP}A$. Note that any such formula can be written as $\nabla\Gamma[\chi_A] = \nabla(\mathsf{T}\chi_A)\Gamma$ (where we remind the reader that the substitution $\chi_A : \mathsf{P} \to \mathtt{Latt}(A)$ is the function mapping a set $B \subseteq A$ to its conjunction $\bigwedge B$). This means that it suffices to prove, for an arbitrary formula $\alpha \in \mathtt{1ML}^+_\Lambda(A)$:

$$\alpha \equiv^1 \bigvee \{\nabla(\mathsf{T}\chi_A)\Gamma \mid \mathsf{P}A, \Gamma, \mathsf{id} \Vdash^1 \alpha\}, \tag{2}$$

where $(\mathsf{P}A, \Gamma, \mathsf{id})$ denotes the canonical one-step A-model on the set $\mathsf{P}A$.

For a proof of the left-to-right direction of (2), assume that $S, \sigma, m \Vdash^1 \alpha$. It is easy to derive from this that $\mathsf{P}A, (\mathsf{T}m)\sigma, \mathsf{id} \Vdash^1 \alpha$, so that $\Gamma := (\mathsf{T}m)\sigma \in \mathsf{TP}A$ provides a candidate disjunct on the right hand side of (2). It remains to show that $S, \sigma, m \Vdash^1 \nabla(\mathsf{T}\chi_A)(\mathsf{T}m)\sigma$, but this is immediate by definition of the semantics of ∇.

For the opposite direction of (2), let $\Gamma \in \mathsf{TP}A$ be such that $\mathsf{P}A, \Gamma, \mathsf{id} \Vdash^1 \alpha$. In order to show that $\nabla(\mathsf{T}\chi_A)\Gamma \vDash^1 \alpha$, let (S, σ, m) be a one-step model such that $S, \sigma, m \Vdash^1 \nabla(\mathsf{T}\chi_A)\Gamma$. Without loss of generality we may assume that $(S, \sigma, m) = (\mathsf{P}A, \Delta, \mathsf{id})$ for some $\Delta \in \mathsf{TP}A$.

By the semantics of ∇ it then follows from $\mathsf{P}A, \Delta, \mathsf{id} \Vdash^1 \nabla(\mathsf{T}\chi_A)\Gamma$ that $(\Delta, (\mathsf{T}\chi_A)\Gamma) \in \overline{\mathsf{T}}(\Vdash^0)$. But since $(B, \chi_A(C)) \in \Vdash^0$ implies that $C \subseteq B$, we easily obtain that $(\Gamma, \Delta) \in \overline{\mathsf{T}}(\subseteq)$.

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 $(\sigma, \sigma') \in \overline{\mathsf{T}}Z$, and $m(s) \subseteq m'(s')$, for all $(s, s') \in Z$. Then for all $\alpha \in \mathtt{1ML}^+_\Lambda(A)$:

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

Finally, it is easy to see that the claim is applicable to the one-step models $(\mathsf{P}A, \Gamma, \mathsf{id})$ and $(\mathsf{P}A, \Delta, \mathsf{id})$, and the relation $\subseteq$. Hence it follows from $\mathsf{P}A, \Gamma, \mathsf{id} \Vdash^1 \alpha$ that $\mathsf{P}A, \Delta, \mathsf{id} \Vdash^1 \alpha$. QED

Graded modal logic Our main motivating example to introduce disjunctive bases is graded modal logic. The bag functor does preserve weak pullbacks, and so its Moss modalities are disjunctive, and the set of all monotone liftings for B does admit a disjunctive basis as an instance of Proposition 3.3. Note, however, that this proposition does not apply to graded modal logic, since the signature Σ_{B} is not expressively complete; this was essentially shown in [17]. It was observed already in [1] that very simple formulas in the one-step language $\mathtt{1ML}_{\Sigma_{\mathsf{B}}}$ 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 Σ_{B} does have a disjunctive basis.

Definition 3.4 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 $V_m : A \to \mathsf{P}S$ 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:

$$\gamma(\overline{a}, B) := \exists \overline{x}(\mathsf{diff}(\overline{x}) \wedge \bigwedge_{i \in I} a_i(x_i) \wedge \forall y(\mathsf{diff}(\overline{x}, y) \to \bigvee_{b \in B} b(y)))$$

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

$$S, \sigma, m \Vdash^1 \gamma(\overline{a}, B) \text{ implies } S, \sigma, m' \Vdash^1 \gamma(\overline{a}, B) \text{ for some } m' \subseteq m \text{ with } \mathsf{Ran}(m') \subseteq \mathsf{P}_{\leq 1}A. \qquad (3)$$

We can turn the formula $\gamma(\overline{a}, B)$ into a modality $\nabla(\overline{a}; B)$ that can be interpreted in *all* one-step B-models, using the observation that every one-step B-frame (S, σ) has a unique Kripkean cover $(\widetilde{S}, \widetilde{\sigma})$ defined by putting $\widetilde{S} := \bigcup\{s \times \sigma(s) \mid s \in S\}$, and $\widetilde{\sigma}(s, i) := 1$ 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 B-model (S, σ)

$$S, \sigma, m \Vdash^1 \nabla(\overline{a}; B) \text{ if } \widetilde{S}, \widetilde{\sigma}, m \circ \pi_S \Vdash^1 \gamma(\overline{a}, B), \qquad (4)$$

where π_S is the projection map $\pi_S : \widetilde{S} \to S$. It is then an immediate consequence of (3) that $\nabla(\overline{a}; B)$ is a disjunctive formula. $\lhd$

Theorem 3.5 *The collection* D_{B} *provides a disjunctive basis for the signature* Σ_{B}.

As far as we know, this result is new. The hardest part in proving it is actually not to show that the language D_{B} is distributive over Σ_{B} or that it admits a distributive law (these are easy exercises that we leave to the reader), but to show that formulas in $\mathsf{D}_{\mathsf{B}}(A)$ can be expressed as one-step formulas in $\mathtt{1ML}^+_{\Sigma_{\mathsf{B}}}(A)$. The reason that this is not so easy is subtle; by contrast, it is fairly straightforward to show that formulas in $\mathsf{D}_{\mathsf{B}}(A)$ can be expressed in $\mathtt{1ML}_{\Sigma_{\mathsf{B}}}(A)$, using Ehrenfeucht-Fraïssé games, see e.g. Fontaine & Place [7]. 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 logics[1].

Proposition 3.6 *Every formula* $\nabla(\overline{a}; B) \in \mathsf{D}_{\mathsf{B}}$ *is one-step equivalent to a formula in* $\mathtt{1ML}_{\Sigma_{\mathsf{B}}}(A)$.

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 $\mathsf{Dom}M = V_1$.

[1]The same subtlety appears in Janin & Lenzi [9], where the translation of the language D_{B} into $\mathtt{1ML}^+_{\Sigma_{\mathsf{B}}}$ is required to prove that the graded μ-calculus is equivalent, over trees, to monadic second-order logic. Proposition 3.6 in fact fills a minor gap in this proof.

Fact 3.7 (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 3.6 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 $B \mapsto \bigvee B$).

Where $\overline{a} = (a_1, \ldots, a_n)$, define $I := \{1, \ldots, n\}$. For each subset $J \subseteq I$, let χ_J be the formula

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

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 $\nabla(\overline{a};b) := \nabla(\overline{a};\{b\})$, we claim that

$$\gamma \equiv^1 \nabla(\overline{a};b), \tag{5}$$

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 $\nabla(\overline{a};b)$. 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 I. 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 $E := \{(j, s) \in (I \cup K) \times S \mid a_j \in m(s)\}$.

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

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

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. ◀

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 $\{s_1, ..., s_n\}$ of S such that $a_i \in m(s_i)$ for each $i \in I$, and that $b \in m(t)$ for each $t \notin \{u_1, ..., u_n\}$. Hence $\nabla(\overline{a};b)$ is true in (S, σ, m), as required. ◀

This concludes the proof of Theorem 3.5.

An example without weak pullback preservation There are also functors that do not preserve weak pullbacks, but do have a disjunctive basis. As an example of this, consider the subfunctor $\mathsf{P}^{2/3}$ of P^3 given by:

$$\mathsf{P}_{2/3}S = \{(Z_0, Z_1, Z_2) \mid Z_0 \cap Z_1 \neq \varnothing \text{ or } Z_1 \cap Z_2 \neq \varnothing\}.$$

While it is easy to show that this functor does not preserve weak pullbacks, The signature Σ_{P^3} (regarded as a set of liftings for $\mathsf{P}_{2/3}$ rather than P^3) still admits a disjunctive basis.

A non-example Finally, we provide an example of a signature that does not admit any disjunctive basis:

Proposition 3.8 *The signature Σ consisting of the box- and diamond liftings for 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 $\alpha \in \mathtt{1ML}^+_\Lambda(A)$, $X, \xi, m \Vdash^1 \alpha$ implies $X', \xi', m' \Vdash^1 \alpha$.

Given a set A, let $\eta_A : A \to \mathsf{P}A$ 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 subsitution mapping $B \in \mathsf{P}A$ to $\bigwedge B$.

Claim 2 Let α be any one-step formula in $\mathtt{1ML}_\Lambda(\mathsf{P}A)$ and let (X, ξ, m) be a one-step model with $m : X \to \mathsf{P}A$. Consider the map $\eta_{\mathsf{P}A} : \mathsf{P}A \to \mathsf{P}\mathsf{P}A$, so that $\eta_{\mathsf{P}A} \circ m$ is a marking of X with variables from $\mathsf{P}A$.

1. If $X, \xi, \eta_{\mathsf{P}A} \circ m \Vdash^1 \alpha$ 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 $X, \xi, \eta_{\mathsf{P}A} \circ m \Vdash^1 \alpha$.

Proof of Claim For the first part of the proposition, it suffices to note that $[\![B]\!]^1_{\eta_{\mathsf{P}A} \circ m} \subseteq [\![\bigwedge B]\!]^1_m$ for each $B \in \mathsf{P}A$, 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 $[\![\bigwedge B]\!]^1_m \subseteq [\![B]\!]^1_{\eta_{\mathsf{P}A} \circ m}$ for each $B \in \mathsf{P}A$ that appears as a variable in α. To prove this, suppose that $u \in [\![\bigwedge B]\!]^1_m$. 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:

$$B \in \{\{b\}\} = \{m(u)\} = \eta_{\mathsf{P}A}(m(u))$$

so $u \in [\![B]\!]^1_{\eta_{\mathsf{P}A} \circ m}$ as required. ◀

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

So suppose $\delta \in \mathtt{1ML}_\Lambda(\mathsf{P}A)$ 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 $\delta[\top/\varnothing]$, which is still disjunctive (this is easy to prove). We have $\delta[\top/\varnothing][\wedge_A] = \delta[\wedge_A]$ since $\wedge_A(\varnothing) = \bigwedge \varnothing = \top$.

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 $X, \xi, \eta_{\mathsf{P}A} \circ m \Vdash^1 \delta$.

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 $h : X \to \mathsf{PP}A$ (with respect to variables in $\mathsf{P}A$) 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 $X, \xi, \eta_{\mathsf{P}A} \circ m$, 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 $f : X' \to X$ such that: $X', \xi', h' \Vdash^1 \delta$, $\mathsf{M}f(\xi') = \xi$, $h'(u) \subseteq h(f(u))$ 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 $h'(u) = \varnothing$, 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 $n : X' \to \mathsf{P}A$ by taking each $n(u)$ for $u \in X'$ to be the unique $B \subseteq A$ such that $h'(u) = \{B\}$. Clearly, $h' = \eta_{\mathsf{P}A} \circ n$, 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. ◀

This finishes the proof of Proposition 3.8. QED

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.

Definition 4.1 A $(\Lambda, \mathtt{X})$-automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$ is said to be *disjunctive* (relative to a disjunctive basis $\mathtt{D}$) if $\Theta(c, a) \in \mathtt{D}(A)$, for all colors $c \in \mathsf{P}\mathtt{X}$ 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 T-model. A strategy f for $\exists$ in $\mathcal{A}(\mathbb{A}, \mathbb{S})@(a, s)$ is *separating* if for every s in $\mathbb{S}$ there is at most one state a in $\mathbb{A}$ such that the position (a, s) 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 separating winning strategy in the game $\mathcal{A}(\mathbb{A}, \mathbb{S})@(a, s)$. ◁

Disjunctive automata are very well behaved. For instance, the following result, which can be proved using essentially the same argument as in [20], states a *linear-size* model property.

Theorem 4.3 *Let $\mathbb{A} = (A, \Theta, a_I, \Omega)$ be a disjunctive automaton for a set functor T. If $\mathbb{A}$ accepts some pointed 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 4.4 *Let $\mathbb{A}$ be a disjunctive Λ-automaton. Then any pointed 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 T-model, let $s_I \in S$, and let f be a winning strategy for $\exists$ in the acceptance game $\mathcal{A} := \mathcal{A}(\mathbb{A}, \mathbb{S})@(a_I, s_I)$; without loss of generality we may assume that f is positional. We will construct (i) a pointed 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 x) and supports (X, ξ) (in the sense that $\xi(x) \in \mathsf{T}R(x)$, for every $x \in X$), (iii) a morphism $h : (X, \xi, W) \to (S, \sigma, V)$ 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 $m_t : S \to \mathsf{P}A$ prescribed by f at this position satisfies

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

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 $\xi(t) \in \mathsf{T}R(t)$, an A-marking $m'_t : R(t) \to \mathsf{P}A$ and a map $h_t : R(t) \to S$, such that[2] $|m(u)| = 1$ and $m'_t(u) \subseteq m_t(h_t(u))$ for all $u \in R(t)$, $(\mathsf{T}h_t)\xi(t) = \sigma(s_t)$ and

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

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

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 $h : X \to S$ 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 $W : \mathtt{X} \to \mathsf{P}X$ by putting $W(p) := \{x \in X \mid hx \in V(p)\}$.

It remains to show that $\mathbb{A}$ strongly accepts the pointed 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 separating 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 $\Sigma^- = (a_0, h(s_0))m_{x_0}(a_1, h(s_1))m_{x_1} \cdots$ is an f-guided match of $\mathcal{A}$. From this observation it is immediate that Σ is won by $\exists$. QED

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.

Theorem 4.5 (Simulation) *Let Λ be a monotone modal signature for the set functor T and assume that Λ has a disjunctive basis. Then there is an effective construction transforming an arbitrary Λ-automaton $\mathbb{A}$ into an equivalent disjunctive Λ-automaton $\mathsf{sim}(\mathbb{A})$.*

[2] To 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.

Proof. Assume that $\mathtt{D}$ is a disjunctive basis for Λ, and let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a Λ-automaton. Our definition of $\mathsf{sim}(\mathbb{A})$ is rather standard [21], so we will confine ourselves to the definitions. The construction takes place in two steps, a 'pre-simulation' step that produces a disjunctive automaton $\mathsf{pre}(\mathbb{A})$ with a non-parity acceptance condition, and a second 'synchronization' step that turns this nonstandard disjunctive automaton into a standard one.

We define the pre-simulation automaton of $\mathbb{A}$ as the structure $\mathsf{pre}(\mathbb{A}) := (A^\sharp, \Theta^\sharp, NBT_{\mathbb{A}}, R_I)$, where the carrier of the pre-simulation $\mathsf{pre}(\mathbb{A})$ 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 $\Theta^\star : A \times \mathtt{P}\mathtt{X} \to \mathtt{1ML}^+_\Lambda(A \times A)$ by putting, for $a \in A$ and $c \in \mathtt{P}\mathtt{X}$:

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

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

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

Clearly such a formula exists by our assumption on $\mathtt{D}$ 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 $\mathsf{pre}(\mathbb{A})$ 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 $\mathsf{pre}(\mathbb{A})$ 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 $\mathsf{sim}(\mathbb{A}) := (A^\sharp \times Z, \Theta'', \Omega'', (R_I, z_I))$ 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 $\mathsf{sim}(\mathbb{A})$ can be proved by relatively standard means [21]. QED

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 [3]. 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^{a+}_B are defined by:

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

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

Proposition 5.2 *Suppose Λ admits a disjunctive basis. Then for any formula α in $\mathtt{1ML}_\Lambda(A)$ there is a one-step equivalent formula of the form $\delta[\vee_{\mathsf{P}A}][\tau]$ for some $\delta \in \mathtt{D}(\mathsf{P}A)$.*

Proof. Let's first check that everything is correctly typed: note that we have $\vee_{\mathsf{P}A} : \mathsf{P}A \to \mathtt{Bool}(\mathsf{P}A)$ and so $\delta[\vee_{\mathsf{P}A}] \in \mathtt{1ML}_\Lambda(\mathsf{P}A)$, and $\tau_{\mathsf{P}A} : \mathsf{P}A \to \mathtt{Bool}(A)$. So $\delta[\vee_{\mathsf{P}A}][\tau] \in \mathtt{1ML}_\Lambda(A)$, 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 $\heartsuit_{\lambda_1}\pi_1 \wedge ... \wedge \heartsuit_{\lambda_k}\pi_k$, where $\pi_1, ..., \pi_k \in \mathtt{Bool}(A)$. Repeatedly applying the distributivity of $\mathtt{D}$ over Λ and the distributive law for $\mathtt{D}$, we can rewrite each such disjunct as a formula of the form $\delta[\sigma]$ where $\delta \in \mathtt{D}(\{1, ..., k\})$ and $\sigma : \{1, ..., k\} \to \mathtt{Bool}(A)$ is defined by setting $i \mapsto \pi_i$. Now, just apply propositional logic to rewrite each formula π_i as a disjunction of formulas in $\tau[\mathsf{P}A]$, and we are done. QED

Theorem 5.3 (One-step Lyndon theorem) *Let Λ be a monotone modal signature for the set functor T and assume that Λ has a disjunctive basis. Any $\alpha \in \mathtt{1ML}_\Lambda(A)$, monotone in the variable $a \in A$, is one-step equivalent to some formula in $\mathtt{1ML}_\Lambda(A)$, which is positive in a.*

Proof. By Proposition 5.2, we can assume that α is of the form $\delta[\vee_{\mathsf{P}A}][\tau]$ for some $\delta \in \mathtt{D}(\mathsf{P}A)$. Clearly it suffices to show that :

$$\delta[\vee_{\mathsf{P}A}][\tau] \equiv^1 \delta[\vee_{\mathsf{P}A}][\tau^{a+}]$$

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

For the converse direction, suppose $X, \xi, m \Vdash^1 \delta[\vee_{\mathsf{P}A}][\tau^{a+}]$. We define a $\mathsf{P}A$-marking $m_0 : X \to \mathsf{PP}A$ by setting $m_0(u) := \{B \subseteq A \mid B \preceq_a m(u)\}$, where the relation $\preceq_a$ over $\mathsf{P}A$ 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 $X, \xi, m_0 \Vdash^1 \delta[\vee_{\mathsf{P}A}]$. Since $\delta[\vee_{\mathsf{P}A}]$ is a monotone formula, it suffices to check that $[\![\tau_B^{a+}]\!]^0_m \subseteq [\![B]\!]^0_{m_0}$ for each $B \subseteq A$. This follows by just unfolding definitions.

Since δ was disjunctive, so is $\delta[\vee_{\mathsf{P}A}]$, as an easy argument will reveal. So we now find a one-step frame morphism $f : (X', \xi') \to (X, \xi)$, together with a marking $m' : X' \to \mathsf{PP}A$ such that $|m'(u)| \leq 1$ and $m'(u) \subseteq m_0(f(u))$ for all $u \in X'$, and such that $X', \xi', m' \Vdash^1 \delta[\vee_{\mathsf{P}A}]$. We define a new A-marking $m'' : X' \to \mathsf{P}A$ 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 $[\![B]\!]^0_{m'} \subseteq [\![\tau_B]\!]^0_{m''}$, so by monotonicity of $\delta[\vee_{\mathsf{P}A}]$ we get $X', \xi', m'' \Vdash^1 \delta[\vee_{\mathsf{P}A}][\tau]$.

If we compare the markings m'' and $m \circ f$, we see 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 $\delta[\vee_{\mathsf{P}A}][\tau]$ was monotone with respect to the variable a it follows that $X', \xi', m \circ f \Vdash^1 \delta[\vee_{\mathsf{P}A}][\tau]$ and so $X, \xi, m \Vdash^1 \delta[\vee_{\mathsf{P}A}][\tau]$ by naturality, thus completing the proof of the theorem. QED

A useful corollary to this theorem is that, given an expressively complete set Λ of predicate liftings for a functor T, the language $\mu\mathtt{ML}_\Lambda$ has the same expressive power as the full language $\mu\mathtt{ML}_\mathsf{T}$. At first glance this proposition may seem trivial, but it is important to see that it is not: given a formula φ of $\mu\mathtt{ML}_\mathsf{T}$, a naive definition of an equivalent formula in $\mu\mathtt{ML}_\Lambda$ would be to apply expressive completeness to simply replace each subformula of the form $\heartsuit_\lambda(\psi_1, ..., \psi_n)$ with an equivalent one-step formula α over $\{\psi_1, ..., \psi_n\}$, 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 5.3.

Corollary 5.4 *Suppose Λ is an expressively complete set of monotone predicate liftings for T. If Λ admits a disjunctive basis, then Λ is Lyndon complete and hence $\mu\mathtt{ML}_\Lambda \equiv \mu\mathtt{ML}_\mathsf{T}$.*

Proof. The simplest proof uses automata: pick a modal Λ'-automaton $\mathbb{A}$, where Λ' is the set of all monotone predicate liftings for 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 is formula still monotone in all the variables in A since it is equivalent to α, so we can apply the one-step Lyndon Theorem 5.3 to replace α' by an equivalent and positive one-step formula β in $\mathtt{1ML}_\Lambda(A)$. Clearly, the resulting automaton $\mathbb{A}'$ will be semantically equivalent to $\mathbb{A}$. QED

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

Theorem 5.5 (Lyndon Theorem) *There is an effective translation $(\cdot)_p^M : \mu\mathtt{ML}_\Lambda \to (\mu\mathtt{ML}_\Lambda)_p^M$, which restricts to a map $(\cdot)_p^M : \mathtt{ML}_\Lambda \to (\mathtt{ML}_\Lambda)_p^M$, and satisfies that*

$$\varphi \in \mu\mathtt{ML} \textit{ is monotone in } p \textit{ iff } \xi \equiv \xi_p^M.$$

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

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

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 $\mu\mathtt{ML}_\Lambda$ that is positive in the variable p.

We claim that $\mathbb{A}$ is monotone in p iff $\mathbb{A} \equiv \mathbb{A}_p^M$. Leaving the direction from right to left to the reader, we prove the opposite implication. So assume that $\mathbb{A}$ is monotone in p. Since it is easy to see that $\mathbb{A}_p^M$ always implies $\mathbb{A}$, we are left to show that $\mathbb{A}_p^M$ implies $\mathbb{A}$, and since $\mathbb{A}_p^M$ is disjunctive, by Proposition 4.4 and invariance of acceptance by coalgebra automata it suffices to prove the following:

$$\mathbb{S}, s_I \Vdash_s \mathbb{A}_p^M \text{ implies } \mathbb{S}, s_I \Vdash \mathbb{A}, \tag{6}$$

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

To prove (6), let f be a separating 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 mononotonicity 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 accessible elements of $V(p)$, and let V_U abbreviate $V[p \mapsto U]$. We claim that

$$\text{if } s \text{ is } f\text{-accessbible then } S, \sigma(s), m_s \Vdash^1 \Theta(V_U^\flat(s), a_s), \tag{7}$$

where m_s is the A-marking provided by f at position (a_s, s). To see why (7) holds, note that for any f-accessbible 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 $S, \sigma(s), m_s \Vdash^1 \Theta_p^M(V^\flat(s), a_s)$. But then (7) is immediate by the definitions of Θ_p^M and U.

Finally, it is straightforward to derive from (7) that f itself is a (separating) winning strategy for $\exists$ in the acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{S})$ initialized at (a_I, s_I). QED

Remark 5.6 Observe that as a corollary of Theorem 5.5 and the decidability of the satisfiability problem of $\mu\mathtt{ML}_\Lambda$ [2], it is decidable whether a given formula $\varphi \in \mu\mathtt{ML}$ is monotone in p. $\lhd$

6 Uniform Interpolation

Uniform interpolation is a very strong form of the interpolation theorem, first proved for the modal μ-calculus in [3]. It was later generalized to coalgebraic modal logics in [15]. However, the proof crucially relies on non-deterministic automata, and for that reason the generalization in [15] 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 $\varphi \in \mu\mathtt{ML}_\Lambda$, we let $\mathtt{X}_\varphi$ denote the set of proposition letters occurring in φ. Given a set $\mathtt{X}$ of proposition letters and a single proposition letter p, it may be convenient to denote the set $\mathtt{X} \cup \{p\}$ as $\mathtt{X}p$. ◁

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 $\mathtt{X} \subseteq \mathtt{X}_\varphi$ of proposition letters, there is a formula $\varphi_\mathtt{X} \in \mathcal{L}(\mathtt{X})$, effectively constructible from φ, such that

$$\varphi \models \psi \text{ iff } \varphi_\mathtt{X} \models \psi, \tag{8}$$

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

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

Theorem 6.3 (Uniform Interpolation) *Let Λ be a monotone modal signature for the set functor T and assume that Λ has a disjunctive basis. Then both logics $\mathtt{ML}_\Lambda$ and $\mu\mathtt{ML}_\Lambda$ enjoy the property of uniform interpolation.*

Following D'Agostino & Hollenberg [3], we prove Theorem 6.3 by automata-theoretic means. The key proposition in our proof is Proposition 6.5 below, which refers to the following construction on disjunctive automata.

Definition 6.4 Let $\mathtt{X}$ be a set of proposition letters not containing the letter p. Given a disjunctive $(\Lambda, \mathtt{X}p)$-automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$, we define the map $\Theta^{\exists p} : A \times \mathsf{P}\mathtt{X} \to \mathtt{D}(A)$ by

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

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

Proposition 6.5 *Let $\mathtt{X} \subseteq \mathtt{Y}$ be sets of proposition letters, both not containing the letter p. Then for any disjunctive $(\Lambda, \mathtt{X}p)$-automaton $\mathbb{A}$ and any pointed T-model $(\mathbb{S}, s_I)$ over $\mathtt{Y}$:*

$$\mathbb{S}, s_I \Vdash \mathbb{A}^{\exists p} \text{ iff } \mathbb{S}', s'_I \Vdash_s \mathbb{A} \text{ for some } \mathtt{Y}p\text{-model } (S', s'_I) \text{ such that } \mathbb{S}'{\upharpoonright}_\mathtt{Y}, s'_I \rightrightarrows \mathbb{S}, s_I. \tag{9}$$

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 $\mathtt{X} = \mathtt{Y}$.

By Proposition 4.4 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 separating 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

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

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

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

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

$$s \notin U \text{ implies } S, \sigma(s), m_s \Vdash^1 \Theta(V^\flat(s), a_s) \tag{10}$$

$$\text{while } s \in U \text{ implies } S, \sigma(s), m_s \Vdash^1 \Theta(V^\flat(s) \cup \{p\}, a_s) \tag{11}$$

Now consider the valuation $V_U := V[p \mapsto U]$, and observe that by this definition we have $V_U^\flat(s) = V^\flat(s)$ if $s \notin U$ while $V_U^\flat(s) = V^\flat(s) \cup \{p\}$ if $s \in U$. Combining this with (10) and (11) we find that

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

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 (separating) 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. QED

The remaining part of the argument follows by a fairly standard argument going back to D'Agostino & Hollenberg [3] (see also Marti et alii [15]), with a twist provided by the fact that the 'bisimulation quantifier' here refers to pre-images rather than to bisimilar models.

Proposition 6.6 *Given any proposition letter p, there is a map $\exists p$ on $\mu\mathtt{ML}_\Lambda$, restricting to $\mathtt{ML}_\Lambda$, such that $\mathtt{X}_{\exists p.\varphi} = \mathtt{X}_\varphi \setminus \{p\}$ and, for every pointed $(\mathbb{S}, s_I)$ over a set $\mathtt{Y} \supseteq \mathtt{X}_\varphi$ with $p \notin \mathtt{Y}$:*

$$\mathbb{S}, s_I \Vdash \exists p.\varphi \textit{ iff } \mathbb{S}', s'_I \Vdash \varphi \textit{ for some } \mathtt{Y}p\textit{-model } (S', s'_I) \textit{ such that } \mathbb{S}'{\restriction_{\mathtt{Y}}}, s'_I \rightrightarrows \mathbb{S}, s_I. \tag{12}$$

Proof. Straightforward by the equivalence between formulas and Λ-automata, the Simulation Theorem, and Proposition 6.5. QED

Proof of Theorem 6.3 With $p_1, \ldots, p_n$ enumerating the proposition letters in $\mathtt{X}_\varphi \setminus \mathtt{X}$, set

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

Then a relatively routine exercise shows that $\varphi \models \psi$ iff $\varphi_{\mathtt{Y}} \models \psi$, for all formulas $\psi \in \mu\mathtt{ML}_\Lambda$ such that $\mathtt{X}_\varphi \cap \mathtt{X}_\psi \subseteq \mathtt{X}$. Finally, it is not difficult to verify that $\varphi_{\mathtt{Y}}$ is fixpoint-free if φ is so; that is, the uniform interpolants of a formula in $\mathtt{ML}_\Lambda$ also belong to $\mathtt{ML}_\Lambda$. ◀

7 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 [19], 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 $\lambda : \breve{\mathsf{P}}^A \to \breve{\mathsf{P}} \circ \mathsf{T}$. Clearly, one-step formulas in $1\mathtt{ML}_\Lambda(A)$ can then be viewed as predicate liftings over A.

Definition 7.1 Let $\lambda : \breve{\mathsf{P}}^A \to \breve{\mathsf{P}} \circ \mathsf{T}$ be a predicate lifting over variables $A = \{a_1, ..., a_n\}$. The *Yoneda representation* $y(\lambda)$ of λ is the subset

$$\lambda_{\mathsf{P}A}(\mathsf{true}_{a_1}, ..., \mathsf{true}_{a_n}) \in \mathsf{PTP}A$$

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

Definition 7.2 Given a set A, let $A^\top$ be the set $A \cup \{\top\}$. Let $\epsilon_A \subseteq A^\top \times \mathsf{P}A$ be the relation defined by $a \epsilon_A B$ iff $a \in B$, and $\top \epsilon_A B$ for all $B \subseteq A$. Let $\eta_A : A^\top \to \mathsf{P}A$ be defined by $\eta_A(a) = \{a\}$, and $\eta_A(\top) = \varnothing$. $\lhd$

In the remainder of this section we assume familiarity with the Barr relation lifting $\overline{\mathsf{T}}$ associated with a functor T; see [13] for the definition and some basic properties.

Definition 7.3 A predicate lifting $\lambda \subseteq \mathsf{TP}A$ is said to be *divisible* if, for all $\alpha \in \lambda$ there is some $\beta \in \mathsf{T}A^\top$ such that $(\beta, \alpha) \in \overline{\mathsf{T}}(\epsilon_A)$ and $\mathsf{T}\eta_A(\beta) \in \lambda$. $\lhd$

Proposition 7.4 *Any disjunctive lifting over A is divisible, and if T preserves weak pullbacks the disjunctive liftings over A are precisely the divisible ones.*

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

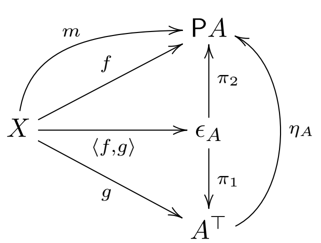

Now apply T to this diagram and define $\beta \in \mathsf{T}A^\top$ to be $\mathsf{T}(\pi_1 \circ \langle f, g \rangle)(\xi) = \mathsf{T}g(\xi)$. First, we have $(\beta, \alpha) \in \overline{\mathsf{T}}(\epsilon_A)$, witnessed by $\mathsf{T}(\langle f, g \rangle)(\xi) \in \mathsf{T}\epsilon_A$. We claim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

For the converse direction, under the assumption that T preserves weak pullbacks, suppose that λ is divisible, and suppose $X, \xi, m \Vdash^1 \lambda$. We get $\mathsf{T}m(\xi) \in \lambda$ and so we find some $\beta \in \mathsf{T}A^\top$ 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 $X \to \mathsf{P}A \leftarrow \epsilon_A$, shown in the diagram.

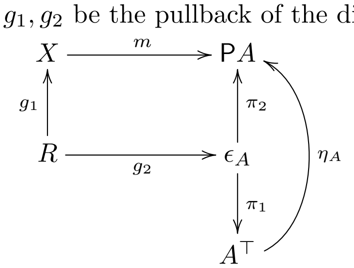

By weak pullback preservation there is $\rho \in \mathsf{T}R$ with $\mathsf{T}g_1(\rho) = \xi$ and $\mathsf{T}g_2(\rho) = \beta'$. The map $g_1 : (R, \rho) \to (X, \xi)$ is thus a cover, and we have a marking m' on R defined by $\eta_A \circ \pi_1 \circ g_2$ (follow the bottom-right path in the previous 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. QED

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

References

[1] J. Bergfeld. Moss's coalgebraic logic: Examples and completeness results. Master's thesis, Institute for Logic, Language and Computation, University of Amsterdam, 2009.

[2] C. Cîrstea, C. Kupke, and D. Pattinson. EXPTIME tableaux for the coalgebraic μ-calculus. In *Computer Science Logic (CSL 2009)*, volume 5771 of *Lecture Notes in Computer Science*, pages 179–193. Springer, 2009.

[3] G. D'Agostino and M. Hollenberg. Logical questions concerning the μ-calculus. *Journal of Symbolic Logic*, 65:310–332, 2000.

[4] 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.

[5] 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, 2017.

[6] 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.

[7] G. Fontaine and T. Place. Frame definability for classes of trees in the mu-calculus. In *Proceedings of the 35th International Symposium on Mathematical Foundations of Computer Science (MFCS 2010)*, pages 381–392. Springer, 2010.

[8] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logic, and Infinite Games*, volume 2500 of *LNCS*. Springer, 2002.

[9] D. Janin and G. Lenzi. Relating levels of the mu-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.

[10] D. Janin and I. Walukiewicz. Automata for the modal μ-calculus and related results. In *Mathematical Foundations of Computer Science 1995, 20th International Symposium (MFCS'95)*, volume 969 of *LNCS*, pages 552–562. Springer, 1995.

[11] 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.

[12] D. Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27:333–354, 1983.

[13] C. Kupke, A. Kurz, and Y. Venema. Completeness for the coalgebraic cover modality. *Logical Methods in Computer Science*, 8(3), 2010.

[14] A. Kurz and R. Leal. Modalities in the stone age: a comparison of coalgebraic logics. *Theoretical Computer Science*, 430:88–116, 2012.

[15] J. Marti, F. Seifan, and Y. Venema. Uniform interpolation for coalgebraic fixpoint logic. In *Proceedings of the Sixth Conference on Algebra and Coalgebra in Computer Science (CALCO 2015)*, volume 35 of *LIPIcs*, pages 238–252, 2015.

[16] L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999. (Erratum published *APAL* 99:241–259, 1999).

[17] E. Pacuit and S. Salame. Majority logic. In *Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning (KR 2004)*, pages 598–605, 2004.

[18] D. Pattinson. Coalgebraic modal logic: soundness, completeness and decidability of local consequence. *Theoretical Computer Science*, 309(1–3):177–193, 2003.

[19] L. Schröder. Expressivity of coalgebraic modal logic: the limits and beyond. *Theoretical Computer Science*, pages 230–247, 2008.

[20] Y. Venema. Automata and fixed point logic: a coalgebraic perspective. *Information and Computation*, 204:637–678, 2006.

[21] Y. Venema. Lectures on the modal μ-calculus. Lecture Notes, ILLC, University of Amsterdam, 2012.

[22] I. Walukiewicz. Completeness of Kozen's axiomatisation of the propositional μ-calculus. *Information and Computation*, 157:142–182, 2000.

A Graph games

For reader unfamiliar with the theory of infinite games, we provide some of basic definitions here, referring to [8] for a survey.

Definition A.1 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 $\Omega : G \to \omega$ 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 $E[p] = \varnothing$ 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 $\mathbb{G}@p$. $\lhd$

Definition A.2 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 $E[\mathsf{last}\rho] \neq \varnothing$, and *full* otherwise. We let PM_Π denote the collection of partial matches ρ ending in a position $\mathsf{last}(\rho) \in G_\Pi$, and define $\mathrm{PM}_\Pi@p$ 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 $\mathsf{last}(\rho)$, and so this player looses ρ, while the opponent wins. If ρ is infinite, we declare its winner to be $\exists$ if $\rho \in W$, and $\forall$ otherwise. $\lhd$

Definition A.3 A *strategy* for a player $\Pi \in \{\exists, \forall\}$ is a map $\chi : \mathrm{PM}_\Pi \to G$. A strategy is *positional* if it only depends on the last position of a partial match, i.e., if $\chi(\rho) = \chi(\rho')$ whenever $\mathsf{last}(\rho) = \mathsf{last}(\rho')$; such a strategy can and will be presented as a map $\chi : G_\Pi \to G$.

A match $\rho = (p_i)_{i<\kappa}$ is *guided* by a Π-strategy χ if $\chi(p_0p_1 \dots p_{n-1}) = p_n$ for all $n < \kappa$ such that $p_0 \dots p_{n-1} \in \mathrm{PM}_\Pi$ (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 $\mathrm{PM}_\Pi@p$ 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 $\forall$. $\lhd$

When defining a strategy χ for one of the players in a board game, we can and in practice will confine ourselves to defining χ for partial matches that are themselves guided by χ.

Fact A.4 (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.*