Annals of Pure and Applied Logic 170 (2019) 578–641

Contents lists available at ScienceDirect

Annals of Pure and Applied Logic

www.elsevier.com/locate/apal

Completeness for μ-calculi: A coalgebraic approach

Sebastian Enqvist [a,*], Fatemeh Seifan [b,**,1], Yde Venema [c,**]

[a] Department of Philosophy, Stockholm University, Universitetsvägen 10D, Frescati, Stockholm, Sweden
[b] Department of Computer Science, University of Erlangen–Nürnberg, Martensstraat 3, 91058 Erlangen, Germany
[c] Institute for Logic, Language and Computation, University of Amsterdam, Science Park 105, 1098 XG Amsterdam, Netherlands

A R T I C L E I N F O

Article history:
Received 8 May 2017
Received in revised form 23 August 2018
Accepted 2 December 2018
Available online 28 December 2018

MSC:
03B45
03B70
68Q60
91A43

Keywords:
Modal fixpoint logic
Completeness
Coalgebra
Automata
Graded modal mu-calculus
Monotone modal mu-calculus

A B S T R A C T

We set up a generic framework for proving completeness results for variants of the modal mu-calculus, using tools from coalgebraic modal logic. We illustrate the method by proving two new completeness results: for the graded mu-calculus (which is equivalent to monadic second-order logic on the class of unranked tree models), and for the monotone modal mu-calculus.

Besides these main applications, our result covers the Kozen–Walukiewicz completeness theorem for the standard modal mu-calculus, as well as the linear-time mu-calculus and modal fixpoint logics on ranked trees. Completeness of the linear-time mu-calculus is known, but the proof we obtain here is different and places the result under a common roof with Walukiewicz' result.

Our approach combines insights from the theory of automata operating on potentially infinite objects, with methods from the categorical framework of coalgebra as a general theory of state-based evolving systems. At the interface of these theories lies the notion of a coalgebraic modal one-step language. One of our main contributions here is the introduction of the novel concept of a disjunctive basis for a modal one-step language. Generalizing earlier work, our main general result states that in case a coalgebraic modal logic admits such a disjunctive basis, then soundness and completeness at the one-step level transfer to the level of the full coalgebraic modal mu-calculus.

* Principal corresponding author.
** Corresponding authors.
E-mail addresses: thesebastianenqvist@gmail.com (S. Enqvist), fateme.sayfan@gmail.com (F. Seifan), y.venema@uva.nl (Y. Venema).

[1] The research of this author has been made possible by *Vrije Competitie* grant 612.001.115 of the Netherlands Organisation for Scientific Research (NWO).

https://doi.org/10.1016/j.apal.2018.12.004
0168-0072/

1. Introduction

1.1. Modal μ-calculi

Over the past fifty years, the formalism of modal logic has developed into what is undoubtedly the most widely applied branch of logic. Phenomena from a wide spectrum of application areas, ranging from metaphysics in philosophy to game theory in economics, and from arithmetic in mathematics to the semantics of natural language in linguistics, have been modeled in some version or variant of modal logic. This success is largely due to the fine balance that modal formalisms strike between expressiveness and computational feasibility, but also to the well-behaved (and well-understood) unifying meta-logical theory of modal logic [5].

Many applications of modal logic require that the basic modal language is extended to express some kind of recursion. This can be taken care of in the form of fixpoint connectives (such as the common knowledge operator in epistemic logic, or the until operator in linear time temporal logic), or via explicit (least- and greatest) fixpoint operators. In the latter case we will speak of a *μ-calculus* extending the more basic modal logic, the prime example being the 'standard' modal μ-calculus as introduced by Kozen [22]. Other examples include the linear time μ-calculus [4], the μ-calculus on ranked trees, the graded modal μ-calculus [23], and the monotone μ-calculus [10].

Our earlier remark on the balance between expressiveness and computational feasibility still applies to such propositional modal μ-calculi. In the setting of specification and verification of various kinds of processes, adding this powerful yet tractable form of recursion to the language of basic modal logic enables us to express and reason about the behavior of state-based evolving systems in a manner that goes far beyond the more local properties that can be expressed in the basic formalism, while on the other hand, this additional expressive behavior comes at a very low computational cost: the EXPTIME complexity of the satisfiability problem for the full modal μ-calculus [9] is no worse than that of virtually any extension of basic modal logic.

Given this importance of modal fixpoint logics (which include logics like LTL of CTL that are obtained from more basic modal logics by adding fixpoint connectives rather than explicit fixpoint operators), there is a clear need to study and further develop their general theory.

1.2. Completeness

The question that we address here concerns the axiomatization problem for modal μ-calculi. That is, our goal is to find, for each member of the above-mentioned 'family of μ-calculi', a (finite) set of axioms and derivation rules that generate the class of *valid* formulas in the associated class of models. For the time being we take this 'associated class' to consist of *all* models for the language, that is, we do not impose additional conditions on the models such as, in the case of standard Kripke models, reflexivity or transitivity. Even without such additional constraints the axiomatization problem for modal fixpoint logics is notoriously difficult,[2] and there seems to be very little in the way of general results (as an exception we mention results on so-called flat fixpoint logics [33,37]). In fact, while many results are known about axiomatizations for concrete logics based on fixpoint connectives, until recently, only two completeness results for μ-calculi were known: the Kozen–Walukiewicz completeness theorem for the standard μ-calculus [39], and Kaivola's completeness result for the linear time μ-calculus [21].

Note that in both cases, the axiomatization is as simple and natural as the μ-calculus itself: add, to a sound and complete axiomatization of the basic (i.e., fixpoint-free) language, a single axiom schema and a rule schema. Together these capture the least fixpoint operator in the sense that the *pre-fixpoint axiom schema*

[2] We refer to the introduction of our earlier work [12] for a more detailed analysis.

$$\varphi[\mu p.\varphi/p] \to \mu p.\varphi \tag{1}$$

simply states that $\mu x.\varphi$ is a pre-fixpoint of φ, while the *Kozen–Park induction rule*:

$$\frac{\varphi[\psi/p] \to \psi}{\mu p.\varphi \to \psi} \tag{2}$$

expresses that $\mu x.\varphi$ is indeed its *least* pre-fixpoint.

This naturally raises the question whether other μ-calculi can be axiomatized in an equally simple way, and our paper will provide a positive answer. Our goal, in fact, is to set up a *general* framework for proving completeness for variants of the standard modal μ-calculus. This framework will be founded on two pillars, viz., the theories of *coalgebra* and *automata* operating on (possibly infinite) objects, respectively.

1.3. Coalgebras, modal logic & automata

A suitable abstraction level for studying various μ-calculi in a unified framework is provided by the theory of (universal) *coalgebra* [31,18], which has found a place in theoretical computer science as a natural mathematical environment for modeling various sorts of state-based evolving systems, such as, indeed, streams, labeled transition systems, Markov chains, etc. The attraction of the coalgebraic approach lies in its combination of mathematical simplicity with wide applicability: many features of (computational) processes, such as nondeterminism, input/output or probability, can be elegantly and naturally encoded in the coalgebraic type T (which formally is an endofunctor on some suitable category). This makes the theory of universal coalgebra well-equipped for a uniform study of various notions that are salient in the study of (possibly infinite) behavior, such as invariance or behavioral equivalence.

Almost since its emergence in logic and computer science, coalgebra has been firmly linked to modal logic: Aczel [1] already noted that Kripke models are natural examples of co-algebras, and Moss [26] initiated the application of modal-type languages for reasoning about coalgebras of arbitrary type. The idea is that the role of equations in algebra is played by modal formulas in coalgebra — and in case infinite behavior is to be specified, modal *fixpoint* formulas are called for. Note that the link works in both directions: the theory of coalgebraic modal logic can be applied to design suitable modal languages for the specification and verification of coalgebraic behavior, but it can also be instrumental in the study of modal logic, by providing modal formalisms with a coalgebraic semantics. Currently, the most common approach to coalgebraic modal logic, going back to the work of Pattinson [30] and others, is based on a categorical analysis of the semantics of modalities in terms of so-called *predicate liftings* for the type functor T (see section 2.2 for the details). That is, in line with the uniform and parametric approach of universal coalgebra, a generic coalgebraic modal logic may be given as a pair consisting of a functor T (providing the semantics of T-coalgebras), together with a set Λ of predicate liftings for T (providing the modalities and their interpretation). As we will see in section 2, all of the mentioned μ-calculi are instances of the *coalgebraic μ-calculus* introduced by Cîrstea, Kupke & Pattinson [6], that is, they are extensions of such coalgebraic modal logics with explicit fixpoint operators.

All our proofs involve *automata* in an essential way. This should not come as a surprise, as the use of automata (more specifically: finite state devices operating on potentially infinite objects such as infinite words, trees, and Kripke models) is well established in the study of fixpoint logics [16]. Pertinent to our work here is the realization that much of the theory of modal (fixpoint) logic and automata is essentially coalgebraic in nature. The coalgebra automata that we will employ here were developed by Fontaine, Leal & the third author [14] as the automata-theoretic counterpart of the coalgebraic μ-calculi that we just discussed.

The key observation underlying the links between coalgebra, modal logic and automata is that many of the properties of modal fixpoint logic are already manifest at the *one-step level*, that is, at the level of formulas of modal depth one and one-step unfoldings of coalgebra states. For instance, this observation was the guiding principle in the authors' work on Janin–Walukiewicz style expressive completeness results for coalgebraic μ-calculi [10]. Here our approach will follow the same track: a pivotal role in our proofs will be played by the notion of a *one-step logic* associated with a pair (T, Λ), stemming from the work on coalgebraic logic by Cîrstea, Pattinson, Schröder and others [7,30,35,36]. Generalizing earlier results on specific coalgebraic fixpoint logics (viz., the ones based on Moss' coalgebraic modality [11]), our main aim will be to show that, under some conditions, the completeness of a coalgebraic μ-calculus is already determined by the completeness of the associated one-step logic.

1.4. Contribution

The contribution of this paper is threefold. First of all, our coalgebraic analysis of one-step logic for nondeterministic automata has led us to isolate the concept of a *disjunctive one-step formula* (Definition 3.15), and the related notion of a *disjunctive basis* for a set of modalities (Definition 3.20). Disjunctivity is the property of one-step formulas that ensures *nondeterministic* behavior of the corresponding automata; essentially, a one-step formula is disjunctive if it only admits special, 'harmless' conjunctions. A set of modalities (predicate liftings) admits a disjunctive basis if there are sufficiently many disjunctive formulas; intuitively, what this achieves is that we may eliminate conjunctions, by proving a simulation theorem stating that every alternating Λ-automaton can be transformed into an equivalent nondeterministic one. Our approach here can be seen as a continuation and generalization of work by Muller & Schupp [28], Janin & Walukiewicz [20], and, in particular, Arnold & Niwiński [3]. In the main result of this paper we will see an important application of disjunctivity, but we believe there to be many more (a first study by the first and the third author can be found in [13]).

Our second contribution comprises a general completeness theorem for modal μ-calculi. Formulated in coalgebraic terminology, it states that, in case a coalgebraic modal logic admits a disjunctive basis, then soundness and completeness at the one-step level transfers to the level of the full coalgebraic modal mu-calculus. Note that we may speak of such a transfer since every one-step axiomatization $\mathbf{H}$ naturally induces an axiom system $\mu\mathbf{H}$ for the corresponding μ-calculus (Definition 4.5).

Theorem 1.1. *Let* T *be a set functor, let* Λ *be a monotone modal signature for* T*, and let* $\mathbf{H}$ *be a one-step axiomatization for* Λ *and* T*. If* $\mathbf{H}$ *is one-step sound and complete and* Λ *admits a disjunctive basis, then* $\mu\mathbf{H}$ *is a sound and complete axiom system for the* $\mu\mathtt{ML}_\Lambda$*-formulas that are valid in the class of all* T*-coalgebras.*

For a *proof* of this theorem: much of the technical ground-work was carried out in [12], where we provided a fully automata-theoretic proof of the Kozen–Walukiewicz completeness theorem for the standard modal μ-calculus, and in [11], where the authors extended this approach to coalgebraic μ-calculi based on Moss-style modalities. While Theorem 1.1 significantly generalizes the latter result, its proof is fairly similar to that of the earlier results. Because of this, and for reasons of space limitations, we confine ourselves to a high-level proof sketch in section 9.

As a direct corollary to Theorem 1.1, we obtain the following completeness result that directly transfers soundness and completeness from a coalgebraic modal logic to its fixpoint extension.

Corollary 1.2. *Let* T *be a set functor, let* Λ *be a monotone modal signature for* T *which admits a disjunctive basis. If* $\mathbf{L}$ *is a sound and complete axiomatization for the (fixpoint-free)* $\mathtt{ML}_\Lambda$*-formulas that are valid in the class of all* T*-coalgebras, then so is* $\mu\mathbf{L}$ *for the set of* $\mu\mathtt{ML}_\Lambda$*-validities.*

Third, as corollaries of Theorem 1.1 we obtain concrete completeness results, for various modal μ-calculi. Some of these are well known, such as the Kozen–Walukiewicz result for the standard modal μ-calculus, or

Kaivola's completeness theorem for the linear-time μ-calculus. Others are, as far as we are aware, new; as explicit examples we mention our results on *graded* and *monotone* modal logic.

In the case of graded modal logic, our completeness result is a fairly direct consequence of the general theorem, since graded modal logic corresponds to a coalgebraic modal logic for the *bag* functor B (Example 2.6(d)), and we will show that this similarity type admits a disjunctive basis.

Theorem 1.3. *Let* $\mathbf{B}$ *be the axioms for graded modal logic given in Definition 4.3. Then the induced axiomatization* $\mu\mathbf{B}$ *is sound and complete for the valid formulas of the graded modal* μ*-calculus.*

Axiomatizing the validities of the monotone modal μ-calculus is more challenging, since the monotone neighborhood functor M interpreting this system (cf. Example 2.6(c)) does not admit a disjunctive basis itself. Fortunately, we may take a detour via its so-called *supported companion* $\underline{\mathsf{M}}$, which *does* allow a disjunctive basis. Analyzing the relation between the two functors and their associated μ-calculi, in the final section we will prove the following completeness result. Following our definitions, $\mu\mathbf{M}$ is the axiomatization for monotone modal logic given by the monotonicity and duality axioms for $\Diamond$ and $\Box$ (cf. Definition 4.1).

Theorem 1.4. *The axiomatization* $\mu\mathbf{M}$ *is sound and complete for the valid formulas of the monotone modal* μ*-calculus.*

2. A coalgebraic approach to μ-calculi

In this section we introduce a coalgebraic framework for modal μ-calculi. The presentation here can be seen as a summary of previous work done in coalgebraic fixpoint logic and automata theory (see [6,14] and references therein).

We assume familiarity with basic notions from category theory, not going beyond categories, functors, natural transformations, and simple operations on these. (Some more information is provided in the appendix.) We let Set denote the category with sets as objects and functions as arrows. An endofunctor on Set will simply be called a *set functor*.[3] Three functors that feature prominently in this paper are the identity functor Id, and the *co-* and *contravariant power set functor*, P and $\breve{\mathsf{P}}$, respectively. Both act on objects by mapping a set S to its power set $\mathsf{P}S = \breve{\mathsf{P}}S$; a function $f : S' \to S$ is mapped by P to the direct image function $\mathsf{P}f : \mathsf{P}S' \to \mathsf{P}S$ given by $(\mathsf{P}f)X' := \{fs' \in S \mid s' \in X'\}$, and by $\breve{\mathsf{P}}$ to the inverse image function $\breve{\mathsf{P}}f : \mathsf{P}S \to \mathsf{P}S'$ given by $(\breve{\mathsf{P}}f)X := \{s' \in S' \mid fs' \in X\}$.

2.1. Coalgebra

In the introduction we described coalgebra as a mathematical framework for modeling various kinds of state-based evolving systems. Formally, this is captured by letting the *transition type* of such a system be determined by an *endofunctor* on some suitable category. For our purposes, we can restrict attention to the category Set.

Definition 2.1. Let $\mathsf{T} : \mathsf{Set} \to \mathsf{Set}$ be a set functor. A T-*coalgebra*, or *coalgebra of type* T, 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 $\mathbb{S}$. We will call $\sigma(s)$ the *(one-step) unfolding* of the state s. A *pointed* T-coalgebra is a pair $(\mathbb{S}, s)$ consisting of a T-coalgebra and a state $s \in S$.

[3] Without loss of generality and for technical convenience, we will assume in this paper that every set functor preserves inclusions, see [14].

We call a function $f : S' \to S$ a *coalgebra homomorphism*

$$\begin{array}{ccc} S' & \xrightarrow{f} & S \\ {\scriptstyle \sigma'}\downarrow & & \downarrow{\scriptstyle \sigma} \\ \mathsf{T}S' & \xrightarrow{\mathsf{T}f} & \mathsf{T}S \end{array}$$

from (S', σ') to (S, σ) if the above diagram commutes. ◁

Many mathematical structures featuring in computer science and in modal logic can be naturally presented as coalgebras. The following list is by no means exhaustive.

Example 2.2. Throughout this example we let $\mathtt{X}$ denote a fixed set of proposition letters.

(a) Streams (infinite words) over an alphabet or color set C are coalgebras for the functor $\mathsf{Id}_C := \mathsf{Id} \times C$, which maps a set S to the product $S \times C$. A stream $(a_n)_{n\in\omega}$ can then be modeled as the coalgebra (ω, σ) where σ maps a state $n \in \omega$ to the pair consisting of its successor $succ(n)$ and its color a_n. As a special case, a (natural-numbers based) linear time model over a set $\mathtt{X}$ of proposition letters can be identified with a $\mathsf{P}\mathtt{X}$-stream, and hence, with a coalgebra for the functor $\mathsf{Id}_{\mathsf{P}\mathtt{X}}$.

(b) Kripke frames are coalgebras for the power set functor P. That is, a Kripke frame (S, R) can be represented as the P-coalgebra (S, σ_R), where $\sigma_R : s \mapsto R[s]$ maps a state s to its successor set. It is not hard to verify that the notion of a bounded morphism between two Kripke frames coincides with that of a coalgebra morphism for P-coalgebras.

(c) With L denoting a set of atomic actions, we may see a transition system $(S, (R_\ell)_{\ell\in L})$, where each atomic action ℓ is interpreted as a binary relation $R_\ell \subseteq S \times S$, as a coalgebra for the functor P^L.

(d) For $k \in \omega$ with $k > 1$, the *k-ary tree* is the structure $(k^*, (succ_i)_{i<k}$, where k^* is the set of all finite sequences of natural numbers smaller than k, and $succ_i$ is the i-th successor function mapping a sequence $s \in k^*$ to the sequence $s \cdot i$. We may present this structure as a coalgebra for the functor Id^k.

(e) Define the *neighborhood* functor $\mathsf{N} : \mathsf{Set} \to \mathsf{Set}$ as the composition of the contravariant power set with itself, $\mathsf{N} := \breve{\mathsf{P}} \circ \breve{\mathsf{P}}$. Coalgebras for this functor correspond to the so-called neighborhood frames in modal logic, but they do not play an important role here.

However, restrictions of this functor also yield various interesting classes of structures. In particular, we will consider the *monotone neighborhood functor* M given by $\mathsf{M}S := \{\mathcal{U} \in \mathsf{N}S \mid \mathcal{U}$ is upward closed with respect to $\subseteq$ $\}$ and $\mathsf{M}f := \mathsf{N}f$. M-coalgebras are well known in modal logic as monotone neighborhood frames.

(f) Of significant interest here is the finitary *multiset* of *bag* functor B. This functor takes an object S to the collection $\mathsf{B}S$ of *weight functions* $\sigma : S \to \omega$ with finite support (that is, for which the set $\{s \in S \mid \sigma(s) > 0\}$ is finite). Its action on arrows is as follows: given a map $f : S \to S'$ and a weight function $\sigma \in \mathsf{B}S$, we define the weight function $(\mathsf{B}f)\sigma : S' \to \omega$ by setting $((\mathsf{B}f)\sigma)(s') := \sum\{\sigma(s) \mid f(s) = s'\}$.

Coalgebras for this functor are *weighted* transition systems, where each transition from one state to another carries a weight given by a natural number. Note that a finitely branching Kripke frame (S, R) can be seen as a B-coalgebra (S, ρ_R), if we define, for any state s, a weight function $\rho_R(s)$ on S given by $\rho_R(s)(t) = 1$ if Rst and $\rho_R(s)(t) = 0$ otherwise. ◁

We can generalize the distinction, of Kripke *models* as opposed to Kripke *frames*, to coalgebras of arbitrary type.

Definition 2.3. Let T be a set functor and let $\mathtt{X}$ be a set of proposition letters. We define the set functor $\mathsf{T}_{\mathtt{X}} := \mathsf{P}\mathtt{X} \times \mathsf{T}$. 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

$$V^\flat(s) := \{p \in \mathtt{X} \mid s \in V(p)\}.$$

Hence the pair $(\mathbb{S}, V)$ induces a $\mathsf{T}_\mathtt{X}$-coalgebra $(S, (V^\flat, \sigma))$. ◁

2.2. Modalities as predicate liftings

The most common approach to coalgebraic modal logic these days proceeds from a formal analysis of what a "modality" is, in a very generally setting. The idea is to view a modal operator as a proposition (dependent on a number of variables), about a single unfolding step of a state in a coalgebra.

Example 2.4. Using the notation of Example 2.2, we may formulate the semantics of the standard modal operators $\Diamond$ and $\Box$ in a Kripke model $\mathbb{S} = (S, R, V)$ as follows:

$$\begin{array}{lll} \mathbb{S}, s \Vdash \Diamond\varphi & \text{iff} & \sigma_R(s) \cap [\![\varphi]\!]^\mathbb{S} \neq \varnothing \\ \mathbb{S}, s \Vdash \Box\varphi & \text{iff} & \sigma_R(s) \subseteq [\![\varphi]\!]^\mathbb{S}, \end{array}$$

where $[\![\varphi]\!]^\mathbb{S} := \{s \in S \mid \mathbb{S}, s \Vdash \varphi\}$. Thus the coalgebraic perspective on standard modal logic is that the modalities $\Diamond$ and $\Box$ express statements about the unfolding $\sigma_R(s)$ of s. We can make this more explicit by defining the following maps $\lambda^\Diamond, \lambda^\Box : \mathsf{P}S \to \mathsf{P}\mathsf{P}S$:

$$\begin{array}{ll} \lambda^\Diamond : & U \mapsto \{T \in \mathsf{P}S \mid T \cap U \neq \varnothing\} \\ \lambda^\Box : & U \mapsto \{T \in \mathsf{P}S \mid T \subseteq U\}. \end{array}$$

Now we may formulate the semantics of $\Diamond$ via the map $\lambda^\Diamond$:

$$\mathbb{S}, s \Vdash \Diamond\varphi \text{ iff } \sigma_R(s) \in \lambda^\Diamond([\![\varphi]\!]^\mathbb{S}), \tag{3}$$

and similarly for $\Box$ and $\lambda^\Box$. ◁

Generalizing this to coalgebras of arbitrary type, the idea underlying coalgebraic modal logic is that (the semantics of) modalities are given by so-called predicate liftings.

Definition 2.5. Given a set functor T and $n \in \omega$, an *n-place predicate lifting* λ for T is an assignment[4] of a map

$$\lambda_S : (\mathsf{P}S)^n \to \mathsf{P}\mathsf{T}S,$$

to each set S, subject to the constraint that for any map $f : S' \to S$ and any n-tuple $\overline{Z} = (Z_1, ..., Z_n) \in (\mathsf{P}S)^n$ we have, for all $\sigma \in \mathsf{T}S$:

$$\sigma \in \lambda_{S'}(f^{-1}[\overline{Z}]) \text{ iff } \mathsf{T}f(\sigma) \in \lambda_S(\overline{Z}) \tag{4}$$

where $f^{-1}[\overline{Z}]$ abbreviates $(f^{-1}[Z_1], \ldots, f^{-1}[Z_n])$. ◁

[4] In categorical terms, an n-ary predicate lifting is simply a natural transformation $\lambda : \check{\mathsf{P}}^n \Rightarrow \check{\mathsf{P}}\mathsf{T}$, see Remark 2.7.

To obtain a suitable modal language for describing coalgebraic behavior, with each predicate lifting λ we associate a modality $\heartsuit_\lambda$ with the same arity as λ. The semantics of $\heartsuit_\lambda$ in a T-model $\mathbb{S} = (S, \sigma, V)$ is given by the following generalization of (3):

$$\mathbb{S}, s \Vdash \heartsuit_\lambda(\overline{\varphi}) \text{ if } \sigma(s) \in \lambda_S([\![\varphi_1]\!]^{\mathbb{S}}, \dots, [\![\varphi_n]\!]^{\mathbb{S}}). \tag{5}$$

The reason to impose condition (4) on predicate liftings is to ensure that, generalizing bisimulation invariance of modal logic, every modality $\heartsuit_\lambda$ will be invariant under coalgebra morphisms.

Example 2.6. Besides the standard diamond and box operators of Kripke models, the operators of many well-known variants of modal logic are in fact instances of modalities that are induced by predicate liftings.

(a) The next-time operator $\bigcirc$ of linear temporal logic can be obtained as the modality associated with the *identity map*, seen as a unary predicate lifting $\lambda^{\bigcirc} : U \mapsto U$ for the identity functor Id.

(b) Let $\bigcirc_i$ be the modality that, interpreted over tree models of branching degree k, has the following meaning: $\mathbb{S}, s \Vdash \bigcirc_i \varphi$ iff $\mathbb{S}, succ_i(s) \Vdash \varphi$. This modality is induced by the unary predicate lifting $\lambda^{\bigcirc_i}_S : \mathsf{P}S \to \mathsf{P}(S^k)$ given by

$$\lambda^{\bigcirc_i} : U \mapsto \{(s_0, \dots, s_{k-1}) \in S^k \mid s_i \in U\}.$$

(c) With respect to the monotone neighborhood functor M, we define two unary predicate liftings, ϵ and ϵ^∂:

$$\begin{aligned} \epsilon_S : \quad & U \mapsto \{\alpha \in \mathsf{M}S \mid U \in \alpha\} \\ {\epsilon^\partial}_S : \quad & U \mapsto \{\alpha \in \mathsf{M}S \mid S \setminus U \notin \alpha\}. \end{aligned}$$

It is now easy to verify that the induced operators $\heartsuit_\epsilon$ and $\heartsuit_{\epsilon^\partial}$ coincide with the standard monotone modalities $\Box$ and $\Diamond$:

$$\begin{aligned} \mathbb{S}, s \Vdash \Box\varphi \quad & \text{iff} \quad U \subseteq [\![\varphi]\!]^{\mathbb{S}}, \text{ for some } U \in \sigma(s) \\ \mathbb{S}, s \Vdash \Diamond\varphi \quad & \text{iff} \quad U \cap [\![\varphi]\!]^{\mathbb{S}} \neq \varnothing, \text{ for all } U \in \sigma(s). \end{aligned}$$

(d) Finally, we consider the bag functor B. Given a natural number k, we define the predicate liftings $\underline{k}$ and $\overline{k}$ by putting

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

Interpreted over standard Kripke models (seen as B-coalgebras as specified in Example 2.2(f)), the modalities associated with these liftings are the *counting modalities* of graded modal logic:

$$\begin{aligned} \mathbb{S}, s \Vdash \heartsuit_{\underline{k}}\varphi \quad & \text{iff} \quad s \text{ has } \geq k \text{ successors } t \text{ with } \mathbb{S}, t \Vdash \varphi \\ \mathbb{S}, s \Vdash \heartsuit_{\overline{k}}\varphi \quad & \text{iff} \quad s \text{ has } < k \text{ successors } t \text{ with } \mathbb{S}, t \not\Vdash \varphi. \end{aligned}$$

In the sequel we use the standard notation for these modalities, i.e., $\Diamond^k$ and $\Box^k$ for $\heartsuit_{\underline{k}}$ and $\heartsuit_{\overline{k}}$, respectively. ◁

Remark 2.7. In categorical terms, an n-ary predicate lifting is a natural transformation $\lambda : \breve{\mathsf{P}}^n \Rightarrow \breve{\mathsf{P}}\mathsf{T}$: (4) simply means that the following diagram commutes:

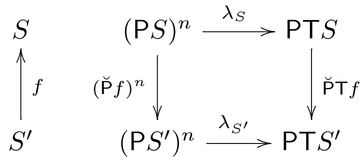

for every function $f : S' \to S$. ◁

2.3. Moss' modalities

As mentioned in the introduction, an important role in this paper is played by so-called *disjunctive* formulas, and a key example of such formulas is provided by the so-called *cover modality* from standard modal logic. It is a slightly non-standard connective that takes a finite set of formulas as its argument.

Definition 2.8. Given a finite set Φ, we let $\nabla\Phi$ abbreviate the formula

$$\nabla\Phi := \bigwedge \Diamond\Phi \wedge \Box \bigvee \Phi,$$

where $\Diamond\Phi$ denotes the set $\{\Diamond\varphi \mid \varphi \in \Phi\}$. ◁

As a primitive operator, this modality was independently introduced by Janin & Walukiewicz [20] in automata theory (with a different notation), and by Moss [26] in coalgebraic logic, where in fact it provided the starting point of the use of modal logic for coalgebras. Here we provide the basic syntactic and semantic definitions for these generalized, coalgebraic modalities; for a more detailed discussion we refer to Kupke, Kurz & Venema [24]. The key concept needed to work with the ∇ modalities is that of a *relation lifting*. (For notation related to binary relations we refer to the appendix.)

Definition 2.9. Let T be a set functor. Given a binary relation R between two sets X_1 and X_2, we define the T-*lifting* of R as the relation $\overline{\mathsf{T}}R \subseteq \mathsf{T}X_1 \times \mathsf{T}X_2$ given as:

$$\overline{\mathsf{T}}R := \{((\mathsf{T}\pi_1^R)\rho, (\mathsf{T}\pi_2^R)\rho) \mid \rho \in \mathsf{T}R\}.$$

Here $\pi_i : R \to S_i$ for $i = 1, 2$ are the projection functions. ◁

Fact 2.10. *The relation lifting $\overline{\mathsf{T}}$ associated with a set functor T has the following properties:*

(1) $\overline{\mathsf{T}}$ extends T: $\overline{\mathsf{T}}f = \mathsf{T}f$ for all functions $f : X_1 \to X_2$;
(2) $\overline{\mathsf{T}}$ preserves the diagonal: $\overline{\mathsf{T}}\mathsf{Id}_X = \mathsf{Id}_{\mathsf{T}X}$ for any set X;
(3) $\overline{\mathsf{T}}$ is monotone: $R \subseteq Q$ implies $\overline{\mathsf{T}}R \subseteq \overline{\mathsf{T}}Q$ for all relations $R, Q \subseteq X_1 \times X_2$;
(4) $\overline{\mathsf{T}}$ commutes with taking converse: $\overline{\mathsf{T}}R^\circ = (\overline{\mathsf{T}}R)^\circ$ for all relations $R \subseteq X_1 \times X_2$;
(5) $\overline{\mathsf{T}}$ distributes over relation composition: $\overline{\mathsf{T}}(R \,;\, Q) = \overline{\mathsf{T}}R \,;\, \overline{\mathsf{T}}Q$, for all relations $R \subseteq X_1 \times X_2$ and $Q \subseteq X_2 \times X_3$, provided the functor T preserves weak pullbacks.

We can now introduce the coalgebraic cover modality ∇_{T}, for an arbitrary set functor T.

Definition 2.11. Let T be some set functor. For any finite set $\mathcal{L}_0$ of formulas, and any element $\Gamma \in \mathsf{T}\mathcal{L}_0$, we let $\nabla_{\mathsf{T}}\Gamma$ denote a new formula.

For the semantics of this formula in a T-model $\mathbb{S} = (S, \sigma, V)$, we define

$$\mathbb{S}, s \Vdash \nabla_{\mathsf{T}}\Gamma \text{ iff } (\sigma(s), \Gamma) \in \overline{\mathsf{T}}(\Vdash),$$

where we inductively assume that the satisfaction relation $\Vdash \subseteq S \times \mathcal{L}_0$ has been defined. ◁

2.4. Coalgebraic μ-calculi

Associating a modality $\heartsuit_\lambda$ with each predicate lifting λ, we obtain a modal language $\mathtt{ML}_\Lambda$ geared towards T-coalgebras, for any set Λ of predicate liftings for T. In fact, the relation between predicate liftings and modalities is so tight that in parlance we will often be sloppy and blur the distinctions between the two notions. Here we are interested in *coalgebraic μ-calculi*, that is, extensions of such coalgebraic modal logics with fixpoint operators.

Definition 2.12. Given a set Λ of predicate liftings, the formulas of the modal fixpoint language $\mu\mathtt{ML}_\Lambda$ are given by the following grammar:

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

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 will employ various syntactic notions such as *subformulas*, *free* and *bound* variables, *substitutions* etc. All of these admit standard definitions and notations, and due to space limitations we refrain from giving details.

Definition 2.13. Given a set Λ of predicate liftings and a set $\mathtt{X}$ of proposition letters, we let $\mu\mathtt{ML}_\Lambda(\mathtt{X})$ denote the set of $\mu\mathtt{ML}_\Lambda$-formulas φ of which all free variables belong to $\mathtt{X}$. ◁

Turning to the semantics of these languages, in order to guarantee well-definedness we restrict attention to predicate liftings that are monotone.

Definition 2.14. A predicate lifting $\lambda : \mathsf{P}^n \Rightarrow \mathsf{PT}$ is *monotone* if for every set S, the map $\lambda_S : (\mathsf{P}S)^n \to \mathsf{PT}S$ is order-preserving in each coordinate (with respect to the subset order). The induced predicate lifting $\lambda^\partial : \mathsf{P}^n \Rightarrow \mathsf{PT}$, given by

$$\lambda^\partial_S(X_1, \dots, X_n) := \mathsf{T}S \ \setminus \ \lambda_S(S \setminus X_1, \dots, S \setminus X_1),$$

is called the *(Boolean) dual* of λ. ◁

All predicate liftings discussed in Example 2.6 are monotone, and come in dual pairs (note that the operators $\bigcirc$ and $\bigcirc_i$ are self-dual).

Definition 2.15. A *monotone modal signature*, or briefly: a *signature* for a set functor T is a set Λ of monotone predicate liftings for T, that is closed under taking boolean duals. In this setting we refer to the triple $(\mathsf{T}, \Lambda, \mu\mathtt{ML}_\Lambda)$ as the *(coalgebraic) μ-calculus* associated with Λ and T. ◁

We will often work with fixpoint formulas in negation normal form.

Definition 2.16. Let Λ be a monotone modal signature for a set functor T. A $\mu\mathtt{ML}_\Lambda$-formula is in *negation normal form* if it can be generated by the following grammar:

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

where p and x are propositional variables, $\lambda \in \Lambda$ has arity n, and the application of the fixpoint operators μx and νx is under the proviso that all occurrences of x in φ' are positive (i.e., not in the scope of a negation). ◁

Formulas of such coalgebraic μ-calculi are interpreted in coalgebraic models, as follows.

Definition 2.17. 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 $s \in [\![\varphi]\!]^{\mathbb{S}}$. Most clauses of this definition are standard; the one for the modality $\heartsuit_\lambda$ is given by (5). For the least fixpoint operator we apply the standard description of least fixpoints of monotone maps from the Knaster–Tarski theorem and take

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

where the valuation $V[x \mapsto U]$ is given by $V[x \mapsto U](x) = U$ while $V[x \mapsto U](p) = V(p)$ for $p \neq x$. ◁

Example 2.18. The table below shows how the standard modal μ-calculus and some of its variants can be presented in this format. We also take this table as canonically defining a fixed signature Σ_T for the functors listed.

T	Σ_T-modalities	Name
Id	$\{\bigcirc\}$	linear time μ-calculus
Id^k	$\{\bigcirc_i \mid 0 \leq i < k\}$	tree μ-calculus
P	$\{\Diamond, \Box\}$	standard (mono-)modal μ-calculus
P^L	$\{\Diamond_\ell, \Box_\ell \mid \ell \in L\}$	standard (poly-)modal μ-calculus
B	$\{\Diamond^k, \Box^k \mid k \in \omega\}$	graded μ-calculus
M	$\{\Diamond, \Box\}$	monotone (mono-)modal μ-calculus
M^L	$\{\Diamond_\ell, \Box_\ell \mid \ell \in L\}$	monotone (poly-)modal μ-calculus

In the case of the graded μ-calculus, it is not hard to prove that a $\mu\mathtt{ML}_{\Sigma_\mathsf{B}}$-formula (where Σ_B is the signature of the counting modalities) is satisfiable in a Kripke model iff it is satisfiable in a finitely branching Kripke model iff it is satisfiable in a B-coalgebra model. This justifies us referring to the coalgebraic μ-calculus for B and Σ_B as the graded modal μ-calculus. ◁

Remark 2.19. Some of these μ-calculi have a very tight connection with monadic second-order logic on trees: the μ-calculi based on modalities $\bigcirc_i$, $0 \leq i < k$ are expressively equivalent to monadic second-order logic on ranked trees of branching degree k [9]. The graded mu-calculus is expressively equivalent to monadic second-order logic on unranked trees with arbitrary branching, see [19]. ◁

3. One-step logic

3.1. One-step syntax and semantics

As mentioned in the introduction, a pivotal role in our approach is filled by the so-called one-step versions of our coalgebraic logics.

Definition 3.1. Given a set of predicate liftings Λ, and two disjoint sets $A, \mathtt{X}$ of variables, we define the set $\mathtt{Bool}(A)$ of *boolean formulas* over A and the set $1\mathtt{ML}_\Lambda(\mathtt{X}, A)$ of *one-step Λ-formulas* over A and parameters $\mathtt{X}$, by the following grammars:

$$\mathtt{Bool}(A) \ni \pi ::= a \mid \bot \mid \top \mid \pi \vee \pi \mid \pi \wedge \pi \mid \neg\pi$$

$$1\mathtt{ML}_\Lambda(\mathtt{X}, A) \ni \alpha ::= p \mid \bot \mid \top \mid \heartsuit_\lambda\overline{\pi} \mid \alpha \vee \alpha \mid \alpha \wedge \alpha \mid \neg\alpha$$

where $a \in A$, $p \in \mathtt{X}$ and $\lambda \in \Lambda$. The *A-positive fragment* of $\mathtt{1ML}_\Lambda(\mathtt{X}, A)$, denoted $\mathtt{1ML}^+_\Lambda(\mathtt{X}, A)$, consists of those formulas in $\mathtt{1ML}_\Lambda(\mathtt{X}, A)$ in which no $a \in A$ appears in the scope of a negation. We will denote the negation-free fragment of $\mathtt{Bool}(A)$ as $\mathtt{Latt}(A)$ and refer to its elements as *lattice formulas* over A.

In case $\mathtt{X} = \varnothing$ we will write $\mathtt{1ML}_\Lambda(A)$ and $\mathtt{1ML}^+_\Lambda(A)$ rather than $\mathtt{1ML}_\Lambda(\varnothing, A)$ and $\mathtt{1ML}^+_\Lambda(\varnothing, A)$, respectively. ◁

A significant part of our work revolves around connections between one-step languages that are based on distinct (but related) sets of variables. Most of these connections are given by substitutions.

Definition 3.2. Given two sets A and B of variables, any substitution $\rho : A \to \mathtt{Bool}(B)$ naturally induces a translation $[\rho]$ mapping $\mathtt{1ML}_\Lambda(\mathtt{X}, A)$-formulas to $\mathtt{1ML}_\Lambda(\mathtt{X}, B)$-formulas. For this translation we shall use postfix notation, $\alpha[\rho] \in \mathtt{1ML}_\Lambda(B)$ denoting the result of applying the substitution $\rho : A \to \mathtt{Bool}(B)$ to the formula $\alpha \in \mathtt{1ML}_\Lambda(A)$. In case $\rho : A \to \mathtt{Latt}(B)$ maps variables to lattice formulas, we can and will assume that $\alpha[\rho] \in \mathtt{1ML}^+_\Lambda(\mathtt{X}, B)$ whenever $\alpha \in \mathtt{1ML}^+_\Lambda(\mathtt{X}, A)$.
We fix notation for the following concrete substitutions:

- $\chi_A : \mathsf{P}A \to \mathtt{Latt}(A)$ will denote the map $B \mapsto \bigwedge B$;
- $\theta_{A,B} : A \times B \to \mathtt{Latt}(A \cup B)$ will denote the map $(a, b) \mapsto a \wedge b$;
- given $a \in A$, $\tau_a : A \to A \times A \subseteq \mathtt{Latt}(A \times A)$ is the *tagging* substitution given by $b \mapsto (a, b)$. ◁

One-step formulas are naturally interpreted in *one-step models*, which consist of a one-step frame together with a marking.

Definition 3.3. A *one-step* $\mathsf{T}_\mathtt{X}$*-frame* is a pair (S, σ) with $\sigma \in \mathsf{T}_\mathtt{X}S$. A *one-step* $\mathsf{T}_\mathtt{X}$*-model* over a set A of variables is a triple (S, σ, m) such that (S, σ) is a one-step $\mathsf{T}_\mathtt{X}$-frame and $m : S \to \mathsf{P}A$ is an A-marking on S. ◁

Definition 3.4. Given a marking $m : S \to \mathsf{P}A$, 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 1*-step interpretation* $[\![\alpha]\!]^1_m$ of $\alpha \in \mathtt{1ML}_\Lambda(\mathtt{X}, A)$ is defined as a subset of $\mathsf{T}_\mathtt{X}S$, with $[\![p]\!]^1_m := \{(\mathtt{Y}, \tau) \mid p \in \mathtt{Y}\}$,

$$[\![\heartsuit_\lambda(\pi_1, \dots, \pi_n)]\!]^1_m := \{(\mathtt{Y}, \tau) \mid \tau \in \lambda_S([\![\pi_1]\!]^0_m, \dots, [\![\pi_n]\!]^0_m)\},$$

and standard clauses for $\bot, \top, \wedge, \vee$ and $\neg$. Given a one-step model (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.

Definition 3.5. Let α and α' be one-step formulas. The formula α is *one-step satisfiable* if there is a one-step model (S, σ, m) such that $S, \sigma, m \Vdash^1 \alpha$, and *one-step valid* if $S, \sigma, m \Vdash^1 \alpha$ for all one-step models (S, σ, m). We say that α' is a *one-step consequence* of α (written $\alpha \vDash^1 \alpha'$) if $S, \sigma, m \Vdash^1 \alpha$ implies $S, \sigma, m \Vdash^1 \alpha'$, for all one-step models (S, σ, m), and that α and α' are *one-step equivalent*, notation: $\alpha \equiv^1 \alpha'$, if $\alpha \vDash^1 \alpha'$ and $\alpha' \vDash^1 \alpha$. ◁

The framework of one-step logic facilitates a concise definition of the following notion.

Definition 3.6. A monotone modal signature Λ for T is *expressively complete* if, for every monotone n-place predicate lifting $\lambda \notin \Lambda$ and variables $a_1, \dots, a_n$ there is a formula $\alpha \in \mathtt{1ML}^+_\Lambda(\{a_1, \dots, a_n\})$ which is equivalent to $\heartsuit_\lambda \overline{a}$. ◁

Note that expressive completeness requires every monotone predicate lifting to be definable by a *positive* one-step formula. In [13] we used the term *Lyndon completeness* for this concept, to distinguish it from the weaker notion that we called expressive completeness there, which does not require defining formulas to be positive. Here, we will not be concerned with this subtle distinction and the technical problems that are related to it, so we will stick to the term expressive completeness throughout the paper.

We also need morphisms between one-step frames and models.

Definition 3.7. A *one-step frame morphism* between two one-step frames (S', σ') and (S, σ) is a map $f : S' \to S$ such that $(\mathsf{T}_{\mathtt{X}} f)\sigma' = \sigma$. In case such a map satisfies $m' = m \circ f$,

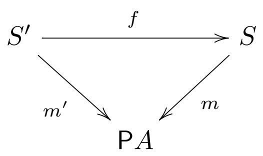

for some markings m and m' on S and S', respectively, we say that f is a *one-step model morphism* from (S', σ', m') to (S, σ, m). ◁

The following proposition, stating that the truth of one-step formulas is invariant under one-step morphisms, is fundamental. We will occasionally refer to this proposition as *naturality*, since this invariance essentially boils down to the naturality of the predicate liftings in Λ.

Proposition 3.8. *Let $f : (S', \sigma', m') \to (S, \sigma, m)$ be a morphism of one-step models over A. Then for every formula $\alpha \in \mathtt{1ML}_\Lambda(A)$ we have*

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

Formulating it differently, for any one-step frame (S', σ'), any marking $m : S \to \mathsf{P}A$, and any map $f : S' \to S$, we have

$$S', \sigma', m \circ f \Vdash^1 \alpha \text{ iff } S, (\mathsf{T}_{\mathtt{X}} f)\sigma', m \Vdash^1 \alpha.$$

As a specific instance of this invariance result we obtain the following corollary which we mention explicitly for future reference.

Corollary 3.9. *Let (S, σ, m) be a one-step A-model, and let $T \subseteq S$ be a subset of S such that $\sigma \in \mathsf{T}_{\mathtt{X}} T$. Then for every formula $\alpha \in \mathtt{1ML}_\Lambda(A)$ we have*

$$S, \sigma, m \Vdash^1 \alpha \text{ iff } T, \sigma, m{\upharpoonright}_T \Vdash^1 \alpha.$$

Proof. Immediate from Proposition 3.8 by the observation that the inclusion map $\iota : T \hookrightarrow S$ is a one-step model morphism. □

The following proposition states that the meaning of a one-step formula only depends on the variables occurring in it.

Proposition 3.10. *Let (S, σ, m) be a one-step model over A, and let $\alpha \in \mathtt{1ML}_\Lambda(A)$ be a one-step formula which belongs to the set $\mathtt{1ML}_\Lambda(B)$, for some subset $B \subseteq A$. Then we have*

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

where m^B is the B-marking given by $m^B(s) := m(s) \cap B$.

For positive one-step formulas we have the following monotonicity property.

Proposition 3.11. *Let (S, σ) be a one-step frame, and let $m, m' : S \to \mathsf{P}A$ be two markings such that $m(s) \subseteq m'(s)$, for all $s \in S$. Then we have*

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

for any formula $\alpha \in \mathtt{1ML}^+_\Lambda(A)$.

Finally, we need a coalgebraic notion of bisimulation which is inspired by an idea of Gorín & Schröder [15].

Definition 3.12. Let (S, σ) and (S', σ') be one-step frames, and Λ a signature. If $Z \subseteq S \times S'$ satisfies:

- for all $U_1, \ldots, U_n \subseteq S$ and $\lambda \in \Lambda$: $\sigma \in \lambda_S(U_1, \ldots, U_n)$ implies $\sigma' \in \lambda_{S'}(Z[U_1], \ldots, Z[U_n])$,
- for all $U'_1, \ldots, U'_n \subseteq S'$ and $\lambda \in \Lambda$: $\sigma' \in \lambda_{S'}(U'_1, \ldots, U'_n)$ implies $\sigma \in \lambda_S(Z^{-1}[U'_1], \ldots, Z^{-1}[U'_n])$,

then we call Z a *one-step Λ-bisimulation* between these one-step frames, denoted as $Z : (S, \sigma) \leftrightarrows^1_\Lambda (S', \sigma')$. In case $\mathsf{Dom}Z = S$ and $\mathsf{Ran}Z = S'$, we call Z *full*, and write $Z : (S, \sigma) \leftrightarrows^1_{\Lambda,f} (S', \sigma')$. ◁

Proposition 3.13. *Let $\leftrightarrows^1_{\Lambda,*}$ denote either $\leftrightarrows^1_\Lambda$ or $\leftrightarrows^1_{\Lambda,f}$, and let (S, σ), (S', σ') and (S'', σ'') be one-step frames. Then*

(1) $\mathsf{Id}_S : (S, \sigma) \leftrightarrows^1_{\Lambda,*} (S, \sigma)$;

(2) if $Z : (S, \sigma) \leftrightarrows^1_{\Lambda,*} (S', \sigma')$ then $Z^\circ : (S', \sigma') \leftrightarrows^1_{\Lambda,*} (S, \sigma)$;

(3) if $Y : (S, \sigma) \leftrightarrows^1_{\Lambda,*} (S', \sigma')$ and $Z : (S', \sigma') \leftrightarrows^1_{\Lambda,*} (S'', \sigma'')$ then $Y ; Z : (S, \sigma) \leftrightarrows^1_{\Lambda,*} (S'', \sigma'')$;

(4) If $f : (S, \sigma) \to (S', \sigma')$ is a one-step frame morphism, then $f : (S, \sigma) \leftrightarrows^1_\Lambda (S', \sigma')$, with $f : (S, \sigma) \leftrightarrows^1_{\Lambda,f} (S', \sigma')$ holding iff f is surjective.

The following observation, generalizing the Propositions 3.8 and 3.11, states that at the level of models, the truth of positive one-step formulas is transferred under one-step bisimulations, provided these interact properly with the markings.

Proposition 3.14. *Let $Z : (S, \sigma) \leftrightarrows^1_\Lambda (S', \sigma')$ and let m and m' be A-markings on S and S' such that*

$$m(s) \subseteq m'(s'),$$

whenever $(s, s') \in Z$. Then for all $\alpha \in \mathtt{1ML}^+_\Lambda(\mathtt{X}, A)$:

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

3.2. Disjunctive formulas

Definition 3.15. 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 $s' \in S'$. ◁

Intuitively, these conditions express that if a disjunctive formula is satisfiable, then it is satisfiable in a closely linked model where no point satisfies more than one variable in A simultaneously (and hence, no proper conjunction over A). Note that the map f mentioned in the above definition is not necessarily a one-step model morphism, since in clause (2) of the definition we do not require equality, and because of clause (3), the inclusion in clause (2) will generally be strict.

Remark 3.16. (1) Using Propositions 3.13 and 3.14, it is not difficult to show that for every one-step frame morphism $f : (S', \sigma') \to (S, \sigma)$ such that $m'(s') \subseteq m(fs')$ for all $s' \in S'$, then $S', \sigma', m' \Vdash^1 \alpha$ implies $S, \sigma, m \Vdash^1 \alpha$, for all one-step formulas $\alpha \in \mathtt{1ML}^+_\Lambda(\mathtt{X}, A)$.

From this it follows that we could have defined disjunctivity of a formula α equivalently by requiring, for an arbitrary one-step model (S, σ, m), that $S, \sigma, m \Vdash^1 \alpha$ *if and only if* there is a one-step frame morphism f satisfying the conditions of Definition 3.15.

(2) Consider two formulas $\alpha \in \mathtt{1ML}^+_\Lambda(A)$ and $\pi \in \mathtt{Bool}(X)$. Provided π is consistent, it is easy to see that α is disjunctive iff $\pi \wedge \alpha$ is so. ◁

Example 3.17. (a) The formula $\bigcirc a$ of linear time logic is easily seen to be disjunctive, as are the tree formulas $\bigcirc_i a$.

(b) The canonical example of a disjunctive formula is given by the *cover modality* ∇ of standard modal logic:

$$\nabla\{a_1, \dots, a_n\} \equiv \Diamond a_1 \wedge \dots \Diamond a_n \wedge \Box(a_1 \vee \dots \vee a_n).$$

(c) The above two examples can be generalized to arbitrary functors that preserve weak pullbacks. In fact, one may show that Moss' modality ∇_T (cf. section 2.3) provides disjunctive formulas, for *every* weak-pullback preserving functor T. To see this, suppose that $S, \sigma, m \Vdash^1 \nabla_\mathsf{T}\gamma$, for some $\gamma \in \mathsf{T}A$. Define $S' := S \times A$, let $f : S \times A \to S$ be the left projection map, and let $m : S' \to \mathsf{P}A$ be given by

$$m'(s, a) := \begin{cases} \{a\} & \text{if } a \in m(s) \\ \varnothing & \text{otherwise.} \end{cases}$$

Let Z denote the relation $Z := \{(s, a) \in S \times A \mid a \in m(s)\}$, and similarly define $Z' := \{(s', a) \in S' \times A \mid a \in m'(s')\}$. It is easy to see that $Z = f^\circ \,;\, Z'$, and so by Fact 2.10(5) we find $\overline{\mathsf{T}}Z = (\mathsf{T}f)^\circ \,;\, \overline{\mathsf{T}}Z'$ (here we use the fact that T preserves weak pullbacks). But from $S, \sigma, m \Vdash^1 \nabla_\mathsf{T}\gamma$ it follows that $(\sigma, \gamma) \in \overline{\mathsf{T}}Z = (\mathsf{T}f)^\circ \,;\, \overline{\mathsf{T}}Z'$, and so there must be an object $\sigma' \in \mathsf{T}S'$ such that $\sigma = (\mathsf{T}f)\sigma'$ and $(\sigma', \gamma) \in \overline{\mathsf{T}}Z'$, which means that $S', \sigma', m' \Vdash^1 \nabla_\mathsf{T}\gamma$, as required. Finally, it is obvious from its definition that m' satisfies the conditions (2) and (3) of Definition 3.15.

(d) An interesting example is provided by the bag functor. We say that a one-step model for the finite multi-set functor is *Kripkean* if all states have multiplicity 1 or 0. 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 $(\mathsf{Base}(\sigma), V_m)$, where $\mathsf{Base}(\sigma) := \{s \in S \mid \sigma(s) > 0\}$ and $V_m : A \to \mathsf{P}(\mathsf{Base}(\sigma))$ is the interpretation given by putting $V_m(a) := \{s \in \mathsf{Base}(\sigma) \mid a \in m(s)\}$. We consider special basic formulas of monadic first-order logic of the form:

$$\gamma(\overline{a}, B) := \exists \overline{x}(\mathsf{diff}(\overline{x}) \wedge \bigwedge_{1 \leq i \leq n} a_i(x_i) \wedge \forall y(\mathsf{diff}(\overline{x}, y) \to \bigvee_{b \in B} b(y)))$$

where $\overline{a} = (a_1, \dots, a_n)$ and $\overline{x} = (x_1, \dots, x_n)$ are n-tuples of, respectively, elements of A and individual variables, and $\mathsf{diff}(\overline{x})$ abbreviates the formula $\bigwedge_{1 \leq i \neq j \leq n} \neg x_i = x_j$.

It is not hard to see that if the formula $\gamma(\overline{a}, B)$ holds in a Kripkean one-step B-model (S, σ, m), then it will continue to hold if we shrink m to a marking $m' \subseteq m$ such that $|m(a)| \leq 1$, for all $a \in A$:

$$S, \sigma, m \Vdash^1 \gamma(\overline{a}, B) \text{ implies } S, \sigma, m' \Vdash^1 \gamma(\overline{a}, B) \text{ for some } m' \subseteq m \text{ with } \mathsf{Ran}(m') \subseteq \mathsf{P}_{\leq 1} A. \tag{6}$$

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 $i \in \sigma(s)$ (here, we have viewed each finite ordinal in the standard manner as the set of all the smaller ordinals, so in particular 0 is defined to be the empty set). Then we can define, for an arbitrary one-step B-model (S, σ):

$$S, \sigma, m \Vdash^1 \nabla(\overline{a};B) \text{ if } \widetilde{S}, \widetilde{\sigma}, m \circ \pi_S \Vdash^1 \gamma(\overline{a}, B), \tag{7}$$

where π_S is the projection map $\pi_S : \widetilde{S} \to S$. It is then an immediate consequence of (6) that $\nabla(\overline{a};B)$ is a disjunctive formula. ◁

The next two, rather technical results, will be needed further on, when we work with games associated with coalgebra automata.

Proposition 3.18. *Let $\alpha \in \mathtt{1ML}^+_\Lambda(A)$ be disjunctive, let (S, σ, m) be a one-step model over A such that $S, \sigma, m \Vdash^1 \alpha$, and let $T \subseteq S$ be such that $\sigma \in \mathsf{T}_\mathtt{X} T$. Then there is a frame homomorphism $f : (S', \sigma') \to (S, \sigma)$ and some marking m' satisfying, next to the clauses (1)–(3) in Definition 3.15, the condition that $Ran(f) = T$.*

Proof. Let α, (S, σ, m) and T be as in the formulation of the proposition. Since $\sigma \in \mathsf{T}_\mathtt{X} T$, the inclusion map $\iota : T \hookrightarrow S$ is a one-step model morphism:

$$\iota : (T, \sigma, m{\restriction_T}) \to (S, \sigma, m).$$

Then by naturality it follows that $T, \sigma, m{\restriction_T} \Vdash^1 \alpha$, so by disjunctivity of α we obtain a one-step model (S', σ', m') and a one-step frame morphism $g : (S', \sigma') \to (T, \sigma)$ satisfying the clauses (1)–(3) in Definition 3.15. It is then easy to verify that the map $f := \iota \circ g$ is a frame homomorphism $f : (S', \sigma') \to (S, \sigma)$ that meets the requirements (1)–(3) of Definition 3.15, and satisfies $\mathsf{Ran}(f) \subseteq T$.

In case the inclusion $\mathsf{Ran}(f) \subseteq T$ is proper, we extend S' to a set $S'' := S' \uplus (T \setminus \mathsf{Ran} F)$ by adding dummy elements to S', we define an A-marking m'' on S'' by putting $m''(u) := m'(u)$ if $u \in S''$ and $m''(u) := \varnothing$ otherwise, and we define a map $f' : S'' \to S$ by putting $f'(u) := f(u)$ if $u \in S'$, and $f'(u) := u$ otherwise. It is then a routine exercise to check that the one-step model (S'', σ', m''), together with the map f', satisfies all the mentioned requirements. □

Proposition 3.19. *Let $\alpha \in \mathtt{1ML}^+_\Lambda(\mathcal{G})$ be disjunctive, where $\mathcal{G} \subseteq \mathsf{P}A$ is a collection of subsets of A. Then for every one-step model (S, σ, m) over A such that $S, \sigma, m \Vdash^1 \alpha[\chi_A]$ there is a frame homomorphism $f : (S', \sigma') \to (S, \sigma)$ and an A-marking $m' : S' \to \mathsf{P}A$ such that:*

(1) $S', \sigma', m' \Vdash^1 \alpha[\chi_A]$;

(2) $m'(s') \subseteq m(f(s'))$*, for all* $s' \in S'$*;*
(3) $m'(s') \in \mathcal{G}$*, for all* $s' \in S'$*.*

In case $T \subseteq S$ *is such that* $\sigma \in \mathsf{T}_{\mathtt{X}} T$*, we may additionally assume that*
(4) $\mathsf{Ran} f = T$*.*

Proof. Fix $\alpha_0 \in \mathtt{1ML}^+_\Lambda(\mathcal{G})$, and assume that $S, \sigma, m \Vdash^1 \alpha_0[\chi_A]$ for some one-step A-model (S, σ, m).

Our first step is to turn (S, σ, m) into a $\mathcal{G}$-model by defining the $\mathcal{G}$-marking $m_{\mathcal{G}}$ by

$$m_{\mathcal{G}}(s) := \{B \in \mathcal{G} \mid B \subseteq m(s)\}.$$

Claim 1. *For all* $\alpha \in \mathtt{1ML}^+_\Lambda(\mathcal{G})$ *we have*

$$S, \sigma, m \Vdash^1 \alpha[\chi] \text{ iff } S, \sigma, m_{\mathcal{G}} \Vdash^1 \alpha. \tag{8}$$

Proof of Claim. First we prove by induction on the complexity of formulas that

$$[\![\pi[\chi]]\!]^0_m = [\![\pi]\!]^0_{m_{\mathcal{G}}} \tag{9}$$

for all $\pi \in \mathtt{Latt}(\mathcal{G})$. For the base case of (9) we take an arbitrary $\pi = B \in \mathcal{G}$, and we reason as follows. Unraveling the definitions on the left hand side of (9) we find that

$$[\![B[\chi]]\!]^0_m = [\![\bigwedge B]\!]^0_m = \bigcap_{b \in B} [\![b]\!]^0_m = \{s \in S \mid b \in m(s) \text{ for all } b \in B\} = \{s \in S \mid B \subseteq m(s)\}.$$

For the right hand side we find

$$[\![B]\!]^0_{m_{\mathcal{G}}} = m_{\mathcal{G}}(B) = \{s \in S \mid B \subseteq m(s)\},$$

and so (9) is immediate. The inductive steps are trivial and left for the reader.

The claim itself is also proved by a straightforward formula induction. The base case of this induction, where α is a formula of the form $\heartsuit_\lambda \pi$, is proved as follows:

$$\begin{array}{llll} S, \sigma, m \Vdash^1 \heartsuit_\lambda \pi & \text{iff} & \sigma \in \lambda\big([\![\pi[\chi]]\!]^0_m\big) & (\text{definition } \Vdash^1) \\ & \text{iff} & \sigma \in \lambda\big([\![\pi]\!]^0_{m_{\mathcal{G}}}\big) & (9) \\ & \text{iff} & S, \sigma, m_{\mathcal{G}} \Vdash^1 \heartsuit_\lambda \pi & (\text{definition } \Vdash^1) \end{array}$$

We omit the routine induction steps of the proof (8). ◀

From our assumption that $S, \sigma, m \Vdash^1 \alpha_0[\chi_A]$ it follows directly by Claim 1 that

$$S, \sigma, m_{\mathcal{G}} \Vdash^1 \alpha_0. \tag{10}$$

By the disjunctivity of α_0 we then obtain a cover $f : (S', \sigma') \to (S, \sigma)$ and a $\mathcal{G}$-marking $m'_{\mathcal{G}}$ such that $(S', \sigma', m'_{\mathcal{G}}) \Vdash^1 \alpha_0$ and, for all $s' \in S'$, $m'_{\mathcal{G}}(s') \subseteq m_{\mathcal{G}} \circ f(s')$ and $|m'_{\mathcal{G}}(s')| \leq 1$. Furthermore, observe that in case $T \subseteq S$ is such that $\sigma \in \mathsf{T}_{\mathtt{X}} T$, by Proposition 3.18 we may take f to be such that $\mathsf{Ran}(f) = T$, taking care of clause (4) in the proposition.

Now define an A-marking m' on S' by putting

$$m'(s') := \begin{cases} B & \text{if } m'_{\mathcal{G}}(s') = \{B\} \\ \varnothing & \text{if } m'_{\mathcal{G}}(s') = \varnothing. \end{cases}$$

Claim 2. *For all* $\alpha \in \mathtt{1ML}^+_\Lambda(\mathcal{G})$ *we have*

$$S', \sigma', m'_{\mathcal{G}} \Vdash^1 \alpha \text{ only if } S', \sigma', m' \Vdash^1 \alpha[\chi]. \tag{11}$$

Proof of Claim. As in the previous claim, we first look at zero-step formulas. By induction on the complexity of formulas we will prove that

$$[\![\pi]\!]^0_{m'_{\mathcal{G}}} \subseteq [\![\pi[\chi]]\!]^0_{m'} \tag{12}$$

for all $\pi \in \mathtt{Latt}(\mathcal{G})$. For the base case of (12) we calculate, for an arbitrary $\pi = B \in \mathcal{G}$:

$$\begin{array}{rcll} [\![B]\!]^0_{m'_{\mathcal{G}}} & = & \{s' \in S' \mid B \in m'_{\mathcal{G}}(s')\} & (\text{definition } [\![\cdot]\!]^0) \\ & = & \{s' \in S' \mid \{B\} = m'_{\mathcal{G}}(s')\} & (|m'_{\mathcal{G}}(s')| \leq 1) \\ & \subseteq & \{s' \in S' \mid B = m'(s')\} & (\text{definition } m') \\ & \subseteq & \{s' \in S' \mid B \subseteq m'(s')\} & (\text{obvious}) \\ & = & [\![B[\chi]]\!]^0_{m'} & (\text{as in proof Claim 1}) \end{array}$$

This proves the base case of (12). As usual, we omit the trivial inductive steps.

Turning to the claim itself, we observe that in the base case of the inductive proof, where α is a formula of the form $\heartsuit_\lambda \pi$, we may reason as follows:

$$\begin{array}{llll} S', \sigma', m'_{\mathcal{G}} \Vdash^1 \heartsuit_\lambda \pi & \text{iff} & \sigma' \in \lambda\big([\![\pi[\chi]]\!]^0_{m'_{\mathcal{G}}}\big) & (\text{definition } \Vdash^1) \\ & \text{only if} & \sigma' \in \lambda\big([\![\pi]\!]^0_{m'}\big) & ((12), \text{ monotonicity of } \lambda) \\ & \text{iff} & S', \sigma', m' \Vdash^1 \heartsuit_\lambda \pi & (\text{definition } \Vdash^1) \end{array}$$

Since the inductive steps of the proof are routine, this establishes the Claim. ◀

As an immediate consequence of Claim 2 and (10) we obtain that $S', \sigma', m' \Vdash^1 \alpha_0[\chi]$, which establishes the first part of Proposition 3.19. The second part follows by the definitions of the respective markings $m_{\mathcal{G}}, m'_{\mathcal{G}}$ and m': let $B := m'(s')$, then $m'_{\mathcal{G}}(s') = \{B\}$, so $B \in m_{\mathcal{G}}(fs')$ which then implies that $B \subseteq m(s)$. The third and last part of the proposition is immediate by the definition of m' and the fact that $m'_{\mathcal{G}}$ is a $\mathcal{G}$-marking. □

3.3. Disjunctive bases

A key concept in our approach is that of a *disjunctive basis*, which we define now. It is very close in spirit to Arnold & Niwiński's notion of *distributivity* [3, Chapter 9]. For the definition, it will be convenient to introduce the abbreviation:

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

for any two sets A, B.

Definition 3.20. 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) $\mathtt{D}(A)$ contains $\top$ and is closed under finite disjunctions (in particular, also $\bot = \bigvee \varnothing \in \mathtt{D}(A)$).

(2) $\mathtt{D}$ is *distributive over* Λ: for every one-step formula of the form $\heartsuit_\lambda \overline{\pi}$ in $\mathtt{1ML}^+_\Lambda(A)$ there is a formula $\delta \in \mathtt{D}(\mathsf{P}(A))$ such that $\heartsuit_\lambda \overline{\pi} \equiv^1 \delta[\chi_A]$.

(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 $\gamma \in \mathtt{D}(A \uplus B)$ such that $\alpha \wedge \beta \equiv^1 \gamma[\theta_{A,B}]$. ◁

Intuitively, what a disjunctive basis achieves is to allow us to eliminate conjunctions in a certain sense.

Proposition 3.21. *Let* $\mathtt{D}$ *be an assignment of a set of disjunctive one-step formulas* $\mathtt{D}(A) \subseteq \mathtt{1ML}^+_\Lambda(A)$ *for all finite sets* A*, satisfying clauses (1) and (3) from Definition 3.20. Then* $\mathtt{D}$ *is a disjunctive basis for* Λ *iff for any formula* $\alpha \in \mathtt{1ML}^+_\Lambda(A)$*, there is a formula* $\delta \in \mathtt{D}(\mathsf{P}A)$ *such that* $\alpha \equiv^1 \delta[\chi_A]$.

In passing we note the following consequence of the binary distributive law. We use the term "distributive law" informally here, leaving the question aside whether it allows a category-theoretic formulation.

Proposition 3.22. *Any binary distributive law* δ *for* $\mathtt{D}$ *induces a distributive law* $\widehat{\delta} : \mathsf{P}_\omega\mathtt{D} \to \mathtt{D}\mathsf{P}_\omega$*, in the sense that*

$$\bigwedge \Delta \equiv_1 \widehat{\delta}_A(\Delta)[\chi_A]$$

for any finite set Δ *of formulas in* $\mathtt{D}(A)$*.*

There is a wealth of functors that admit a disjunctive basis; several of these are presented and examined in detail in the related paper [13] by the first and third authors, where the relationship between disjunctive formulas and nabla formulas is also clarified. Here, we only present a brief summary of the results in [13].

Disjunctive bases via cover modalities. Let Λ be an expressively complete signature for a weak-pullback preserving functor T. Then Λ admits a disjunctive basis. This is due to the fact that, when Λ is expressively complete, the one-step language can express predicate liftings that correspond to arbitrary disjunctions of Moss-style ∇-formulas. A more precise explanation can be found in [13].

It follows that the signatures we have associated in Example 2.18 with the identity functor Id, the tree functor Id^k, and the functors P and P^L, all admit disjunctive bases. Furthermore, whenever the functor T preserves weak pullbacks and restricts to finite sets, (the finitary version of) Moss' language for T is expressively complete. This means that our main result in [11] fits into the present framework as a special case.

An example without expressive completeness. Consider the functor $\mathsf{T} = L \times \mathsf{Id}$, where L is a countably infinite set of labels. For this functor, T-coalgebras are (up to unfolding) just L-streams, or infinite words for which the alphabet is contained in L. They can be viewed as triples (S, σ_1, σ_2) where $\sigma_1 : S \to \Sigma$ and $\sigma_2 : S \to S$. A 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 $\bigcirc$ with the interpretation $\mathbb{S}, s \Vdash \bigcirc\varphi$ iff $\mathbb{S}, \sigma_2(s) \Vdash \varphi$. The signature Λ is not expressively complete, but still has a disjunctive basis.

An example without weak pullback preservation. We now give 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, and F does not preserve weak pullbacks. Consider the modal signature consisting of the usual labeled 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. This signature has a disjunctive basis, in fact the usual cover modalities still give a disjunctive basis even though the functor does not preserve weak pullbacks.

Graded modal logic. Finally, we consider the example of graded modal logic. This case provides another example of a signature that is not expressively complete, yet has a disjunctive basis. To start with, it is not hard to see that the signature Σ_{B} of the counting modalities for the bag functor B (which does preserve weak pullbacks) is not expressively complete. For a simple example showing this, just consider the (monotone) predicate lifting maj given by:

$$\mathsf{maj}_X(Z) = \{\xi \in \mathsf{B}X \mid \sum_{v \in Z} \xi(v) \geq \sum_{v \in X \setminus Z} \xi(v)\}.$$

It was shown by Pacuit & Salame [29] that the corresponding formula $\heartsuit_{\mathsf{maj}}\varphi$ (which in a finitely branching Kripke model states that at least half the successors satisfy φ) cannot be expressed in the language of graded modal logic. Nevertheless, Σ_{B} does admit a disjunctive basis [13]. The proof is not trivial and we believe the result has independent interest, so we list it as a theorem:

Theorem 3.23. *The signature Σ_{B} has a disjunctive basis.*

4. Derivation systems

In this section we introduce our one-step derivation systems, and we discuss their relation with the derivation systems for coalgebraic μ-calculi. The idea of one-step logics and one-step completeness, however, has been studied extensively in the literature on coalgebraic modal logic by various authors, including Cîrstea, Pattinson, and Schröder, see [30,34,35,6] for some selected references.

4.1. One-step soundness and completeness

In this subsection we will see that there is really *logic* to be done at the level of one-step formulas. Recall that in Definition 3.5 we introduced some standard *semantic* notions pertaining to one-step formulas. With these in place, we now consider *derivation systems* for one-step logics.

Definition 4.1. Given a signature Λ for T, a *one-step axiomatization* $\mathbf{H}$ is just a set of formulas $\mathbf{H} \subseteq \mathtt{1ML}_\Lambda(\mathtt{Var})$, where $\mathtt{Var}$ is a fixed countable set of propositional variables.

The *one-step derivation system* $\mathbf{H}^1$ associated with $\mathbf{H}$ consists of the following axioms and rules.

- ($\mathbf{H}$) All formulas in $\mathbf{H}$ are axioms of $\mathbf{H}^1$.
- (MP) From $\alpha \to \beta$ and α, derive β, where $\alpha, \beta \in \mathtt{1ML}_\Lambda(\mathtt{Var})$.
- (CT) All substitution instances $\alpha \in \mathtt{1ML}_\Lambda(\mathtt{Var})$ of propositional tautologies are axioms.
- (Cg) For all $\overline{\pi}, \overline{\rho} \in \mathtt{Bool}(\mathtt{Var})$, if each $\pi_i \leftrightarrow \rho_i$ is a substitution instance of a propositional tautology then $\heartsuit_\lambda \overline{\pi} \leftrightarrow \heartsuit_\lambda \overline{\rho}$ is an axiom.
- (US) Given any substitution $\tau : \mathtt{Var} \to \mathtt{Bool}(\mathtt{Var})$ and $\alpha \in \mathtt{1ML}_\Lambda(\mathtt{Var})$, derive $\alpha[\tau]$ from α.
- (Du) The formula $\heartsuit_{\lambda^\partial}(a_0, \ldots, a_{n-1}) \leftrightarrow \neg\heartsuit_\lambda(\neg a_0, \ldots, \neg a_{n-1})$ is an axiom, for all $\lambda \in \Lambda$ and $\overline{a} \in \mathtt{Var}$.
- (Mon) For all $\lambda \in \Lambda$ and $\overline{a}, \overline{b} \in \mathtt{Var}$, the formula $\heartsuit_\lambda(a_0, \ldots, a_{n-1}) \to \heartsuit_\lambda(a_0 \vee b_0, \ldots, a_{n-1} \vee b_{n-1})$ is an axiom.

Table 1
Axioms for various μ-calculi.

H	T	Axioms
$\mathbf{I}$	Id	a. $\neg\bigcirc a \leftrightarrow \bigcirc\neg a$ b. $\bigcirc\top$
$\mathbf{I}^k$	Id^k	a. $\neg\bigcirc_i a \leftrightarrow \bigcirc_i\neg a$ b. $\bigcirc_i\top$
$\mathbf{K}$	P	a. $\Box(a \wedge b) \leftrightarrow (\Box a \wedge \Box b)$ b. $\Box\top$
$\mathbf{K}^L$	P^L	a. $[l](a \wedge b) \leftrightarrow ([l]a \wedge [l]b)$ b. $[l]\top$
$\mathbf{B}$	B	a. $\Diamond^{n+1}a \to \Diamond^n a$ b. $\Box^1(a \to b) \to (\Diamond^n a \to \Diamond^n b)$ c. $\Diamond^0!(a \wedge b) \wedge \Diamond^{k_1}!a \wedge \Diamond^{k_2}!b \to \Diamond^{k_1+k_2}!(a \vee b)$ d. $\Box^1\top$

We write $\vdash^1_{\mathbf{H}} \alpha$ and say that α is *one-step* **H**-*derivable* if α is provable in the Hilbert-style system consisting of the axioms and rules of $\mathbf{H}^1$. We write $\alpha \vdash^1_{\mathbf{H}} \beta$ for $\vdash^1_{\mathbf{H}} \alpha \to \beta$. We also write $\alpha \equiv^1_{\mathbf{H}} \beta$ for $\alpha \vdash^1_{\mathbf{H}} \beta$ and $\beta \vdash^1_{\mathbf{H}} \alpha$. ◁

We now introduce one of the central ingredients of our framework:

Definition 4.2. A one-step axiomatization **H** is said to be *one-step sound* if $\vDash^1 \alpha$ whenever $\vdash^1_{\mathbf{H}} \alpha$, for $\alpha \in \mathtt{1ML}_\Lambda(A)$. The system **H** is said to be *one-step complete* if $\vdash^1_{\mathbf{H}} \alpha$ whenever $\vDash^1 \alpha$, for $\alpha \in \mathtt{1ML}_\Lambda(A)$. ◁

Definition 4.3. Table 1 presents one-step axiomatizations for a number of coalgebraic signatures, associated with the functors in the table as presented in Example 2.18.

Here, $\Diamond^{k+1}!\pi$ abbreviates $\Diamond^k\pi \wedge \neg\Diamond^{k+1}\pi$, and $\Diamond^0!\pi$ abbreviates $\neg\Diamond^0\pi$. ◁

Proposition 4.4. *All of the axiomatizations given in Definition 4.3 are one-step sound and complete.*

With one exception, we omit the proof of one-step completeness for these systems; the proofs for **I** and $\mathbf{I}^k$ are very easy, and the other cases are more or less just re-stating results from [35].

Proof. We focus on the most difficult case, the system **B** for graded modal logic.

First, given a subset B of some fixed finite set A, we define the *full type of* B to be the propositional formula

$$\tau_B := \bigwedge_{a \in B} a \wedge \bigwedge_{a \in A \setminus B} \neg a,$$

and we define a *simple conjunction* over A to be a formula of the shape:

$$(*) \quad \Diamond^{k_1}\tau_1 \wedge \ldots \wedge \Diamond^{k_n}\tau_n \wedge \Diamond^{k'_1}!\tau'_1 \wedge \ldots \wedge \Diamond^{k'_m}!\tau'_m$$

where each τ_i and each τ'_i is a full type.

Claim 1. *Any consistent simple conjunction is one-step satisfiable.*

Proof of Claim. Given a consistent simple conjunction γ of the shape $(*)$, it follows by definition of the operator $\Diamond^k!$ and the axiom $\mathbf{B}(a)$ that $\tau'_i \neq \tau'_j$ whenever $k'_i \neq k'_j$, for $1 \leq i \leq j \leq m$, and that $\tau_i = \tau'_j$ implies $k_i \leq k'_j$, for $1 \leq i \leq n$ and $1 \leq j \leq m$.

We now consider the one-step A-model $(\mathsf{P}A, \Gamma_\gamma, \mathsf{id})$ on the power set $\mathsf{P}A$ of A, where the marking is the canonical marking given by the identity map on $\mathsf{P}A$, and $\Gamma_\gamma \in \mathsf{BP}A$ is the weight function given by $\Gamma_\gamma(B) = k'_j$ if τ_B is the full type τ'_j for $1 \leq j \leq n$; otherwise set $\Gamma_\gamma(\tau)$ to be the largest k such that $\Diamond^k \tau_B$ is a conjunct of γ (where we may think of $\Diamond^0 \varphi = \top$ as a conjunct of every formula γ of shape (*)). It is then straightforward to check that the conjunction $(*)$ is true in $(\mathsf{P}A, \Gamma_\gamma, \mathsf{id})$, as required. ◀

Claim 2. *Every one-step formula in* $\mathtt{1ML}_{\Sigma_\mathsf{B}}(A)$ *is provably equivalent in* **B** *to a disjunction of simple conjunctions.*

Proof of Claim. First, a simple disjunctive normal form argument, together with the observation that every formula in $\mathtt{Bool}(A)$ is equivalent to a disjunction of full types and applying the axioms $\mathbf{B}(b\&d)$, we can write any one-step formula as a disjunction of conjunctions of the shape:

$$\Diamond^{k_1} \bigvee \Phi_1 \wedge ... \wedge \Diamond^{k_n} \bigvee \Phi_n \wedge \Box^{k'_1} \pi_1 \wedge ... \wedge \Box^{k'_m} \pi_m,$$

where each Φ_i is a set of full types. It now suffices to show that each conjunct $\Diamond^{k_i} \bigvee \Phi_i$ and each $\Box^{k'_j} \bigvee \Phi'_j$ can be replaced by equivalent disjunctions of the right shape, and then distribute conjunctions over disjunctions to put the formula back in disjunctive normal form. This is proved using the following two claims:

(I) Let π_1, π_2 be mutually inconsistent formulas in $\mathtt{Bool}(A)$. Then the formula $\Diamond^k(\pi_1 \vee \pi_2)$ is provably equivalent to the disjunction:

$$\bigvee \{ \Diamond^{k_1} \pi_1 \wedge \Diamond^{k_2} \pi_2 \mid k_1 + k_2 = k \}$$

(II) Let π be any formula in $\mathtt{Bool}(A)$. Then $\Box^k \pi$ is provably equivalent to the disjunction of all formulas of the form:

$$\Diamond^{k_1} ! \tau_1 \wedge ... \wedge \Diamond^{k_n} ! \tau_n$$

such that $k_1 + ... + k_n < k$ and $\{\tau_1, ..., \tau_n\}$ is the set of all full types that are inconsistent with π.

In each of the proofs of these two claims the central role is played by the axiom $\mathbf{B}(c)$. We omit the details. ◀

Finally, the completeness result directly follows from these two claims. □

4.2. Linked derivation systems

With a one-step axiomatization **H** we may not only associate a *one-step derivation system* $\mathbf{H}^1$, **H** also induces an axiom system for the *μ-calculus* based on the signature of **H**.

Definition 4.5. Let **H** be any one-step axiomatization. We define the Hilbert system $\mu\mathbf{H}$ as follows: as axioms we take all axioms in **H**, the axioms (Du) and (Mon), all substitution instances of propositional tautologies, and the *pre-fixpoint schema* (1) given in the introduction. As rules, we take modus ponens, the uniform substitution rule (derive $\varphi[\tau]$ from φ, where $\tau : \mathtt{Var} \to \mu\mathtt{ML}_\Lambda$), the *congruence rule*:

$$\frac{\varphi \leftrightarrow \psi}{\heartsuit_\lambda \varphi \leftrightarrow \heartsuit_\lambda \psi}$$

and, finally, the *Kozen–Park induction rule* (2) discussed in the introduction.

We write $\vdash_\mathbf{H} \varphi$ to say that φ is provable in the system $\mu\mathbf{H}$, $\varphi \vdash_\mathbf{H} \psi$ for $\vdash_\mathbf{H} \varphi \to \psi$ and $\varphi \equiv_\mathbf{H} \psi$ for $\vdash_\mathbf{H} \varphi \leftrightarrow \psi$. ◁

The following proposition will provide a crucial link between the associated derivation systems at the one-step level and at the μ-calculus level, in our completeness proof.

Proposition 4.6 *(Consistency reduction). Suppose that* $\mathtt{D}$ *is a disjunctive basis for* Λ*. Furthermore, suppose* $\mathbf{H}$ *is a one-step sound and complete axiomatization, and let* $\sigma : A \to \mu\mathtt{ML}_\Lambda$ *be a map assigning some formula in* $\mu\mathtt{ML}_\Lambda$ *to every variable in* A*. If* α *is a formula in* $\mathtt{1ML}^+_\Lambda(A)$ *such that* $\nvdash_{\mathbf{H}} \neg\alpha[\sigma]$*, then there exists a one-step model* $X, \xi, m \Vdash^1 \alpha$ *(where* $\xi \in \mathsf{T}_{\mathtt{X}} X$*) such that for each* $u \in X$*, we have* $\nvdash_{\mathbf{H}} \neg \bigwedge \sigma[m(u)]$*.*

Proof. To keep notation simple we take all predicate liftings to be unary. Using expressive completeness of the disjunctive fragment $\mathtt{D}(A)$ and applying distributivity for $\mathtt{D}$ as supplied by Proposition 3.21, we can rewrite the formula α as a disjunction ξ of formulas of the form $\delta[\chi_A]$ for $\delta \in \mathtt{D}(\mathsf{P}A)$.

Pick a disjunct $\delta[\chi_A]$ of ξ such that $\delta[\chi_A][\sigma]$ is consistent in $\mu\mathbf{H}$, which must exist since otherwise the whole disjunction $\xi[\sigma]$ is inconsistent and hence $\alpha[\sigma]$ is inconsistent contrary to assumption. It can be checked that:

$$\delta[\chi_A][\sigma] \equiv_{\mathbf{H}} \delta[\tau][\chi_A][\sigma] \tag{13}$$

where the map $\tau : \mathsf{P}A \to \mu\mathtt{ML}_\Lambda$ is defined by:

$$\tau(B) = \begin{cases} B & \text{if } \bigwedge \sigma[B] \text{ is } \mu\mathbf{H}\text{-consistent} \\ \bot & \text{otherwise.} \end{cases}$$

To see this, we first prove by induction on the complexity of a lattice formula π over $\mathsf{P}A$ that:

$$\pi[\chi_A][\sigma] \equiv_{\mathbf{H}} \pi[\tau][\chi_A][\sigma] \tag{14}$$

Using this we can prove by induction on one-step formulas α over A that:

$$\alpha[\chi_A][\sigma] \equiv_{\mathbf{H}} \alpha[\tau][\chi_A][\sigma]$$

We only consider the case where α is of the form $\heartsuit_\lambda \pi$, and since π is a lattice formula over $\mathsf{P}A$ we can reason as follows:

$$\begin{aligned} (\heartsuit_\lambda \pi)[\chi_A][\sigma] &= \heartsuit_\lambda(\pi[\chi_A][\sigma]) \\ &\equiv_{\mathbf{H}} \heartsuit_\lambda(\pi[\tau][\chi_A][\sigma]) \\ &= (\heartsuit_\lambda \pi)[\tau][\chi_A][\sigma]. \end{aligned}$$

For the second step here we have used the congruence rule. This finishes the proof of (14). We now see that $\delta[\tau][\chi_A][\sigma]$ is consistent in $\mu\mathbf{H}$ (since $\delta[\chi_A][\sigma]$ was consistent), and it follows immediately that $\nvdash^1_{\mathbf{H}} \neg\delta[\tau][\chi_A]$ by contraposition.

Using the substitution property for $\mathbf{H}$ (contrapositively) we find that $\nvdash^1_{\mathbf{H}} \neg\delta[\tau]$. From one-step completeness we get $\nVdash^1 \neg\delta[\tau]$, so we find a set X and a marking $m : X \to \mathsf{PP}A$ such that $[\![\delta]\!]^1_m \neq \varnothing$. Hence we find $\xi \in \mathsf{T}_{\mathtt{X}} X$ such that $X, \xi, m \Vdash^1 \delta[\tau]$.

We now change the marking m to a new marking n as follows: for $u \in X$ we set

$$n(u) := \{B \subseteq A \mid B \in m(u) \; \& \; \tau(B) \neq \bot\}.$$

Then for each $B \subseteq A$ we clearly have $[\![\tau(B)]\!]^0_n = [\![B]\!]^0_m$, and we get for all positive one-step formulas β over $\mathsf{P}A$ that:

$$X, \xi, m \Vdash^1 \beta[\tau] \text{ iff } X, \xi, n \Vdash^1 \beta.$$

Hence, in particular, we get:

$$X, \xi, n \Vdash^1 \delta.$$

By disjunctivity of δ we can now pick a cover $f : (X', \xi') \to (X, \xi)$ and a marking $n' : X' \to \mathsf{PP}A$ with $n'(v) \subseteq n(f(v))$ for each $v \in X'$, where each $n'(v)$ for $v \in X'$ is either empty or a singleton, and such that $X', \xi', n' \Vdash^1 \delta$. Define a new marking $n^\dagger : X' \to \mathsf{P}A$ by setting:

$$n^\dagger(u) := \begin{cases} B & \text{if } n'(u) = \{B\} \\ \varnothing & \text{if } n'(u) = \varnothing \end{cases}$$

Then one can check that for each $B \subseteq A$ we have:

$$[\![B]\!]^0_{n'} \subseteq [\![\chi_A(B)]\!]^0_{n^\dagger}$$

So by a monotonicity argument we get for all formulas $\beta \in \mathtt{1ML}^+_\Lambda(\mathsf{P}A, \mathtt{X})$ that $X', \xi', n' \Vdash^1 \beta$ implies $X', \xi', n^\dagger \Vdash^1 \beta[\chi_A]$. In particular, we get $X', \xi', n^\dagger \Vdash^1 \delta[\chi_A]$. It follows that $X', \xi', n^\dagger \Vdash^1 \xi$, hence $X', \xi', n^\dagger \Vdash^1 \alpha$, and it can be checked that $\bigwedge \sigma[n^\dagger(u)]$ is consistent for each $u \in X$. $\square$

5. Coalgebra automata

5.1. Λ-automata

As mentioned in the introduction, our approach is essentially automata-theoretic in nature. In this section we introduce the specific kind of *coalgebra automata* that we will use in this paper — these originate with Fontaine, Leal & Venema [14].

Throughout this section we fix a set $\mathtt{X}$ of proposition letters.

Definition 5.1. A $\mathtt{X}$-*automaton structure for* Λ, or briefly, a Λ-*automaton structure*, is a triple (A, Θ, Ω) where A is a finite set of *states*, $\Omega : A \to \omega$ is the *priority map* of the automaton, while the *transition map*

$$\Theta : A \to \mathtt{1ML}^+_\Lambda(\mathtt{X}, A)$$

maps states to one-step formulas. We turn such a structure into a *modal* $\mathtt{X}$-*automaton for* Λ, or briefly, a Λ-*automaton* by expanding the structure with a starting state $a_I \in A$. In case we discuss automata for an arbitrary or unknown signature Λ, we will use the term *coalgebra automata* rather than Λ-automata.

The *underlying structure* of an automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$ is the triple (A, Θ, Ω). With $b \in A$, let $\mathbb{A}\langle b\rangle$ denote the variant of $\mathbb{A}$ that takes b as its starting state, i.e., $\mathbb{A}\langle b\rangle = (A, \Theta, \Omega, b)$. $\lhd$

The semantics of coalgebra automata is given in terms of a two-player infinite parity game [16].

Definition 5.2. Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a Λ-automaton, and let $\mathbb{S} = (S, \sigma, V)$ be a T-model, both over the set $\mathtt{X}$ of proposition letters. The *acceptance game* $\mathcal{A}(\mathbb{A}, \mathbb{S})$ for $\mathbb{A}$ with respect to $\mathbb{S}$ is defined as in the following table:

Position	Player	Admissible moves
(a, s)	$\exists$	$\{m : S \to \mathsf{P}A \mid (S, \sigma(s), m) \Vdash^1 \Theta(a)\}$
m	$\forall$	$\{(b, t) \mid b \in m(t)\}$

The winning conditions are as usual for parity games. That is, the loser of a finite play is the player who got stuck. An infinite play $(a_1, s_1)m_1(a_2, s_2)m_2(a_3, s_3)m_3 \dots$ induces a stream $a_1a_2a_3 \dots$ over the alphabet A, and we declare the winner of this play to be $\exists$ if the highest priority that appears infinitely often in the stream $\Omega(a_1)\Omega(a_2)\Omega(a_3)\dots$ is even, and $\forall$ is the winner otherwise.

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 $\mathcal{A}(\mathbb{A}, \mathbb{S})$. The *language* $L(\mathbb{A})$ recognized by $\mathbb{A}$ is the class of pointed T-models accepted by $\mathbb{A}$. $\lhd$

To gain some intuitions, note that the acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{S})$ moves in *rounds* from one basic position of the form (a, s) to another. Each round starts with $\exists$ picking an A-marking m on S that turns the one-step unfolding of s into a one-step model $(S, \sigma(s), m)$ that is supposed to satisfy the one-step formula $\Theta(a)$. Looking at this marking m as a binary relation of *witnesses*, $\forall$ then finishes by picking a new basic position from this set.

Definition 5.3. Let $\mathbb{A}$ and $\mathbb{A}'$ be two modal automata. We say that $\mathbb{A}$ *(semantically) implies* $\mathbb{A}'$, notation: $\mathbb{A} \models \mathbb{A}'$, if $L(\mathbb{A}) \subseteq L(\mathbb{A}')$, and that $\mathbb{A}$ and $\mathbb{A}'$ are *equivalent*, notation: $\mathbb{A} \equiv \mathbb{A}'$, if they recognize the same language, i.e., if $L(\mathbb{A}) = L(\mathbb{A}')$. The two automata are *one-step equivalent*, notation: $\mathbb{A} \equiv^1 \mathbb{A}'$, if $A = A'$, $\Omega = \Omega'$, $a_I = a'_I$, and $\Theta(a) \equiv^1 \Theta(a)$ for all $a \in A$. A Λ-automaton $\mathbb{A}$ is *equivalent* to a formula $\varphi \in \mu\mathtt{ML}_\Lambda$ if any pointed T-model $(\mathbb{S}, s)$ is accepted by $\mathbb{A}$ iff $\mathbb{S}, s \Vdash \varphi$. $\lhd$

It is obvious that one-step equivalence implies equivalence.

In the remainder of this subsection we introduce various concepts and notations pertaining to Λ-automata and automaton structures.

Definition 5.4. The *(directed) graph* of an automaton structure $\mathbb{A} = (A, \Theta, \Omega)$ is the pair $(G, \rightsquigarrow_\mathbb{A})$, where $a \rightsquigarrow_\mathbb{A} b$ if b occurs in the formula $\Theta(a)$, and we let $\rhd_\mathbb{A}$ denote the transitive closure of $\rightsquigarrow_\mathbb{A}$. If $a \rhd_\mathbb{A} b$ we say that b is *active* in a. We write $a \bowtie_\mathbb{A} b$ if $a \lhd_\mathbb{A} b$ and $b \lhd_\mathbb{A} a$.

A *cluster* of $\mathbb{A}$ is a cell of the equivalence relation generated by $\bowtie_\mathbb{A}$ (i.e., the smallest equivalence relation on A containing $\bowtie_\mathbb{A}$). A cluster C is *degenerate* if it is of the form $C = \{a\}$ with $a \not\bowtie_\mathbb{A} a$; by extension we will also call the state a *degenerate*. The unique cluster to which a state $a \in A$ belongs is denoted as C_a.

The concepts introduced here in fact pertain to pairs of the form (A, Θ), where $\Theta : A \to \mathtt{1ML}^+_\Lambda(\mathtt{X}, A)$; for such structures we will use the same terminology, writing $\lhd_\Theta, \rhd_\Theta$ rather than $\lhd_\mathbb{A}, \rhd_\mathbb{A}$, etc. $\lhd$

Definition 5.5. Fix a Λ-automaton structure $\mathbb{A} = (A, \Theta, \Omega)$. The *size* $|\mathbb{A}|$ of $\mathbb{A}$ is defined as the cardinality of its carrier A.

We write $a \sqsubset_\mathbb{A} b$ if $\Omega(a) < \Omega(b)$, and $a \sqsubseteq_\mathbb{A} b$ if $\Omega(a) \leq \Omega(b)$. When clear from context we sometimes write $\sqsubset$ and $\sqsubseteq$ instead, dropping the explicit reference to $\mathbb{A}$.

Given a state a of $\mathbb{A}$, we write $\eta_a = \mu$ if $\Omega(a)$ is odd, and $\eta_a = \nu$ if $\Omega(a)$ is even, and we call a an η_a-state. The sets of μ- and ν-states are denoted with A^μ and A^ν, respectively.

We say that $\mathbb{A}$ is *positive* in a proposition letter $p \in \mathtt{X}$ if each occurrence of p in each formula $\Theta(a)$ is positive, that is, not in the scope of a negation.

A state $a \in A$ is called a *true* state of $\mathbb{A}$ if $\Theta(a) = \top$. $\lhd$

5.2. From formulas to automata

Generalizing the automata-theoretic perspective on the modal μ-calculus as in [40], Λ-automata are the counterpart of the coalgebraic μ-calculus associated with Λ, in the sense that there are effective constructions transforming $\mu\mathtt{ML}_\Lambda$-formulas into equivalent Λ-automata, and vice versa [14]. In this section and the next, we have a closer look at these transformations. For some more detail and motivation of these definitions we refer the reader to [12].

First we consider some operations on automata that correspond to the connectives of our language. For the definition of the complementation operation on automata, we need the following auxiliary definition.

Definition 5.6. The *(boolean) dual* α^∂ of a one-step formula $\alpha \in \mathtt{1ML}^+_\Lambda(\mathtt{X}, A)$ is the formula we obtain from α by simultaneously replacing all occurrences of $p \in \mathtt{X}$ with $\neg p$, $\wedge$ with $\vee$, $\heartsuit_\lambda$ with $\heartsuit_{\lambda^\partial}$, and vice versa. ◁

Definition 5.7. Let $\mathbb{A} = (A, \Theta_A, \Omega_A, a_I)$ and $\mathbb{B} = (B, \Theta_B, \Omega_B, b_I)$ be two Λ-automata over $\mathtt{X}$.

(1) With $\odot \in \{\wedge, \vee\}$, we let $\mathbb{A} \odot \mathbb{B}$ denote the automaton $(C, \Theta_C, \Omega_C, i_{\mathbb{A}\odot\mathbb{B}})$, where $i_{\mathbb{A}\odot\mathbb{B}}$ is some arbitrarily chosen object, $C := A \uplus B \uplus \{i_{\mathbb{A}\odot\mathbb{B}}\}$, Θ_C and Ω_C agree with, respectively, Θ_A and Ω_A on A and with, respectively Θ_B and Ω_B on B, whereas for the initial state $i_{\mathbb{A}\odot\mathbb{B}}$ we define

$$\begin{array}{lll} \Theta_C(i_{\mathbb{A}\odot\mathbb{B}}) & := & \Theta_A(a_I) \odot \Theta_B(b_I) \\ \Omega_C(i_{\mathbb{A}\odot\mathbb{B}}) & := & k+1, \end{array}$$

where k is the maximum priority of $\mathbb{A}$, $\mathbb{B}$.

(2) We let $\neg\mathbb{A}$ denote the automaton $(A, \Theta^\partial_A, \Omega_{\neg\mathbb{A}}, a_I)$, where Θ^∂_A maps a state a to the boolean dual of $\Theta(a)$ (see Definition 5.6), and $\Omega_{\neg\mathbb{A}}$ is given by

$$\Omega_{\neg\mathbb{A}}(a) := 1 + \Omega_A(a).$$

(3) For $\lambda \in \Lambda$ (assumed to be unary, for simplicity) we define $\heartsuit_\lambda \mathbb{A} = (C, \Theta_C, \Omega_C, i_C)$ as the automaton given by $C := A \uplus \{i_C\}$, Θ_C and Ω_C agree with, respectively, Θ_A and Ω_A on A, whereas for the initial state i_C we define

$$\begin{array}{lll} \Theta_C(i_C) & := & \heartsuit_\lambda a_I \\ \Omega_C(i_C) & := & k+1, \end{array}$$

where k is the maximum priority of $\mathbb{A}$. We leave it to the reader to carry out the straightforward generalization of this construction to arbitrary, n-ary predicate liftings. ◁

Next we define a substitution operation on automata.

Definition 5.8. Let $\mathbb{A} = (A, \Theta_A, \Omega_A, a_I)$ and $\mathbb{B} = (B, \Theta_B, \Omega_B, b_I)$ be two Λ-automata over the sets $\mathtt{X} \uplus \{p\}$ and $\mathtt{X}$, respectively, and assume that $\mathbb{A}$ is positive in p. We define $\mathbb{A}[\mathbb{B}/p] = (C, \Theta_C, \Omega_C, i_C)$ as the Λ-automaton over $\mathtt{X}$ defined by $C := A \uplus B$, whereas Θ_C is given by

$$\Theta_C(c) := \begin{cases} \Theta_A(c)[\Theta_B(b_I)/p] & \text{if } c \in A \\ \Theta_B(b) & \text{if } c \in B. \end{cases}$$

Finally, we set $\Omega_C(b) := \Omega_B(b)$ for $b \in B$ and $\Omega(a) := n + \Omega(a)$ for $a \in A$, where n is the least even number greater than any priority in $\mathbb{B}$. ◁

In order to define least and greatest fixpoint operators on automata we need the following proposition, where we recall that $\equiv^1_{\mathbf{H}}$ denotes the relation of one-step *provable* equivalence with respect to the ambient one-step derivation system $\mathbf{H}$, cf. Definition 4.1.

Proposition 5.9. *For every Λ-automaton $\mathbb{A}$ positive in x, and any state $a \in A$, there are formulas θ^a_0 and θ^a_1 in which x does not appear, such that*

$$\Theta(a) \equiv^1_{\mathbf{H}} (x \wedge \theta^a_0) \vee \theta^a_1.$$

Table 2
The automata $\mathbb{A}^x$, $\mu x.\mathbb{A}$ and $\nu x.\mathbb{A}$.

Automaton	$\Theta(a_i)$	$\Theta(\underline{x})$	$\Omega(a_i)$	$\Omega(\underline{x})$	i
$\mathbb{A}^x$	$\theta_i^a[\kappa]$	x	$\Omega_A(a)$	0	$\underline{x}$
$\mu x.\mathbb{A}$	$\theta_i^a[\kappa]$	$\theta_1^{a_I}[\kappa]$	$\Omega_A(a)$	$m+1$	$\underline{x}$
$\nu x.\mathbb{A}$	$\theta_i^a[\kappa]$	$\theta_0^{a_I}[\kappa] \vee \theta_1^{a_I}[\kappa]$	$\Omega_A(a)$	$m+2$	$\underline{x}$

Definition 5.10. Let $\mathbb{A} = (A, \Theta_A, \Omega_A, a_I)$ be a Λ-automaton over the set $\mathtt{X} \uplus \{x\}$, and assume that $\mathbb{A}$ is positive in x. By Proposition 5.9 for each $a \in A$ we may fix formulas $\theta_0^a, \theta_1^a \in \mathtt{1ML}_\Lambda^+(\mathtt{X}, A)$ such that $\Theta(a) \equiv_{\mathbf{H}}^1 (x \wedge \theta_0^a) \vee \theta_1^a$. We now define automata $\mathbb{A}^x$, $\mu x.\mathbb{A}$ and $\nu x.\mathbb{A}$. All three structures are based on the same carrier, viz., the set $(A \times \{0,1\}) \uplus \{\underline{x}\}$; we will denote states of the form (a,i) as a_i, if no confusion is likely. Of all these three automata, we specify their transition map Θ, priority map Ω and initial state i in Table 2. In this table, κ denotes the substitution

$$\kappa : a \mapsto (\underline{x} \wedge a_0) \vee a_1,$$

while m is the smallest even number that is greater than the maximum priority of $\mathbb{A}$. $\lhd$

Definition 5.11. By induction on the complexity of a modal μ-formula $\varphi \in \mu\mathtt{ML}_\Lambda$ we define a Λ-automaton $\mathbb{A}_\varphi$.

First of all, we need to consider atomic formulas: given any propositional variable p, we take some arbitrary object a distinct from p to be the one and only state of $\mathbb{A}_p$, and define $\Theta_p(a) = p$, and $\Omega_p(a) = 0$.

With this in place, we can complete the translation as follows:

$$\begin{array}{lcl} \mathbb{A}_{\neg\varphi} & := & \neg\mathbb{A}_\varphi \\ \mathbb{A}_{\varphi\vee\psi} & := & \mathbb{A}_\varphi \vee \mathbb{A}_\psi \\ \mathbb{A}_{\heartsuit_\lambda\varphi} & := & \heartsuit_\lambda\mathbb{A}_\varphi \\ \mathbb{A}_{\mu x.\varphi} & := & \mu x.\mathbb{A}_\varphi, \end{array}$$

i.e., by applying the operations we have defined above to handle the various connectives of the coalgebraic μ-calculus. $\lhd$

5.3. From automata to formulas

In the opposite direction we will need an actual map transforming an initialized modal automaton into an equivalent μ-calculus formula. For our definition of such a map, which is a variation of the one found in [16], we need some preparations. For a proper inductive formulation of this definition it is convenient to extend the class of automata, allowing states of the automaton to appear in the scope of a modality in a one-step formula.

Definition 5.12. A *generalized automaton structure* over $\mathtt{X}$ is a triple $\mathbb{A} = (A, \Theta, \Omega)$ such that A is a finite set of states, $\Omega : A \to \omega$ is a priority map, and $\Theta : A \to \mathtt{1ML}_\Lambda^+(\mathtt{X}, A \cup \mathtt{X})$ maps states of A to *generalized one-step formulas.* $\lhd$

Whenever possible, we will apply concepts that have been defined for automata structures to these generalized structures without explicit notification. For the operational semantics of generalized modal automata we may extend the notion of a one-step model in the obvious way. Readers who are interested in the details may consult [12].

Definition 5.13. A (generalized) automaton structure $\mathbb{A} = (A, \Theta, \Omega)$ is called *linear* if the relation $\sqsubset_\mathbb{A}$ is a linear order (i.e., the priority map Ω is injective), and satisfies $\Omega(a) > \Omega(b)$ in case b is active in a but not vice versa. A *linearization* of $\mathbb{A}$ is a linear automaton $\mathbb{A}' = (A, \Theta, \Omega')$ such that (1) for all $a \in A$, $\Omega'(a)$ has the same parity as $\Omega(a)$, and (2) for all $a, b \in A$ that belong to the same cluster we have $\Omega'(a) < \Omega'(b)$ iff $\Omega(a) < \Omega(b)$. ◁

Our focus on linear automaton structures is justified by Proposition 5.14; for the definitions of the satisfiability and consequence games involved in this definition, see Section 6. It is easy to prove this Proposition, based on the observation that, given a linearization $\mathbb{A}' = (A, \Theta, \Omega')$ of an automaton $\mathbb{A} = (A, \Theta, \Omega)$, any stream $\alpha \in A^\omega$ is winning for $\exists/\forall$ relative to Ω iff it is winning for the same player relative to Ω'.

Proposition 5.14. *Every automaton structure $\mathbb{A}$ has a linearization $\mathbb{A}^l$ such that, for all $a \in A$,*

(1) $\mathbb{A}\langle a\rangle \vDash_{\mathsf{G}} \mathbb{A}^l\langle a\rangle$ and $\mathbb{A}^l\langle a\rangle \vDash_{\mathsf{G}} \mathbb{A}\langle a\rangle$;

(2) each player $\Pi \in \{\exists, \forall\}$ has a winning strategy in $\mathcal{S}(\mathbb{A}\langle a\rangle)$ (resp. $\mathcal{S}_{thin}(\mathbb{A}\langle a\rangle)$) iff she/he has a winning strategy in $\mathcal{S}(\mathbb{A}^l\langle a\rangle)$ (resp. $\mathcal{S}_{thin}(\mathbb{A}^l\langle a\rangle)$).

Definition 5.15. We introduce a map

$$\mathtt{tr}_{\mathbb{A}} : A \to \mu\mathtt{ML}_\Lambda(\mathtt{X})$$

for any linear generalized $\mathtt{X}$-automaton structure $\mathbb{A} = (A, \Theta, \Omega)$. These maps are defined by induction on the size of $\mathbb{A}$.

In case $|\mathbb{A}| = 1$, we set

$$\mathtt{tr}_{\mathbb{A}}(a) := \eta_a a.\Theta(a),$$

where a is the unique state of $\mathbb{A}$.

In case $|\mathbb{A}| > 1$, by linearity there is a unique state m reaching the maximal priority of $\mathbb{A}$, that is, with $\Omega(m) = \max(\mathsf{Ran}(\Omega))$. Let $\mathbb{A}^- = (A^-, \Theta^-, \Omega^-)$ be the $\mathtt{X} \cup \{m\}$-automaton structure given by $A^- := A \setminus \{m\}$, while Θ^- and Ω^- are defined as the restrictions of, respectively, Θ and Ω to A^-. Since $|\mathbb{A}^-| < |\mathbb{A}|$, inductively[5] we may assume a map $\mathtt{tr}_{\mathbb{A}^-} : A \to \mu\mathtt{ML}_\Lambda(X \cup \{m\})$.

Now we first define

$$\mathtt{tr}_{\mathbb{A}}(m) := \eta_m m.\Theta(m)[\mathtt{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-],$$

and then set

$$\mathtt{tr}_{\mathbb{A}}(a) := \mathtt{tr}_{\mathbb{A}^-}(a)[\mathtt{tr}_{\mathbb{A}}(m)/m]$$

for the states $a \neq m$. ◁

We now turn to the translation map for arbitrary automaton structures. We already saw that every automaton structure has at least one linearization. Furthermore, by the following result the translation maps of different linearizations of the same structure are provably equivalent.

---

[5] Observe that since m is a proposition letter and not a variable in $\mathbb{A}^-$, the latter structure need not be a Λ-automaton, even if $\mathbb{A}$ is. It is for this reason that we introduced the notion of a *generalized* Λ-automaton.

Proposition 5.16. *Let $\mathbb{A}' = (A, \Theta, \Omega')$ and $\mathbb{A}'' = (A, \Theta, \Omega'')$ be two linearizations of the automaton structure $\mathbb{A} = (A, \Theta, \Omega)$. Then*

$$\mathtt{tr}_{\mathbb{A}'}(a) \equiv_{\mathbf{H}} \mathtt{tr}_{\mathbb{A}''}(a)$$

for all $a \in A$.

Proposition 5.16 ensures that modulo provable equivalence the following definition of $\mathtt{tr}(\mathbb{A})$ for an arbitrary automaton $\mathbb{A}$ does not depend on the particular choice of a linearization for the underlying automaton structure of $\mathbb{A}$.

Definition 5.17. With each automaton structure $\mathbb{A} = (A, \Theta, \Omega)$ we associate an arbitrary but fixed linearization $\mathbb{A}^l$ of $\mathbb{A}$ (with the understanding that $\mathbb{A}^l = \mathbb{A}$ in case $\mathbb{A}$ itself is linear). We then define $\mathtt{tr}_{\mathbb{A}} := \mathtt{tr}_{\mathbb{A}^l}$. Finally, given an arbitrary Λ-automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$, we let

$$\mathtt{tr}(\mathbb{A}) := \mathtt{tr}_{\mathbb{A}^l}(a_I)$$

define the translation of the automaton $\mathbb{A}$ itself. ◁

Proposition 5.18. *The following claims hold, for all Λ-automata $\mathbb{A}, \mathbb{B}$:*
(1) $\mathtt{tr}(\mathbb{A} \odot \mathbb{B}) \equiv_{\mathbf{H}} \mathtt{tr}(\mathbb{A}) \odot \mathtt{tr}(\mathbb{B})$ *for* $\odot \in \{\wedge, \vee\}$*;*
(2) $\mathtt{tr}(\neg\mathbb{A}) \equiv_{\mathbf{H}} \neg\mathtt{tr}(\mathbb{A})$*;*
(3) $\mathtt{tr}(\heartsuit_\lambda \mathbb{A}) \equiv_{\mathbf{H}} \heartsuit_\lambda \mathtt{tr}(\mathbb{A})$ *for all* $\lambda \in \Lambda$*;*
(4) if $\mathbb{A}$ is positive in p then $\mathtt{tr}(\eta p.\mathbb{A}) \equiv_{\mathbf{H}} \eta p.\mathtt{tr}(\mathbb{A})$ *for* $\eta \in \{\mu, \nu\}$*;*
(5) if $\mathbb{A}$ is positive in p then $\mathtt{tr}(\mathbb{A}[\mathbb{B}/p]) \equiv_{\mathbf{H}} \mathtt{tr}(\mathbb{A})[\mathtt{tr}(\mathbb{B})/p]$*;*
(6) $\mathtt{tr}(\mu x.\mathbb{A}) \equiv_{\mathbf{H}} \mu x.\mathtt{tr}(\mathbb{A}^x)$*.*

The following theorem establishes the central property of the translations from formulas to automata and back that links these constructions with the proof theory of coalgebraic μ-calculi in the appropriate way:

Theorem 5.19. *For every formula φ, we have $\varphi \equiv_{\mathbf{H}} \mathtt{tr}(\mathbb{A}_\varphi)$.*

Remark 5.20. This theorem allows us pass freely between $\mu\mathtt{ML}_\Lambda$-formulas and Λ-automata without losing information about consistency or provability, and to apply proof-theoretic concepts to automata. For example we say that the automaton $\mathbb{A}$ is *consistent* if $\mu\mathbf{H} \nvdash \neg\mathtt{tr}(\mathbb{A})$, we may write $\mathbb{A} \vdash \mathbb{B}$ to abbreviate $\mu\mathbf{H} \vdash \mathtt{tr}(\mathbb{A}) \to \mathtt{tr}(\mathbb{A})$ etc. ◁

6. Games for coalgebra automata

Our completeness result is based on a number of automata-theoretic concepts, specifically, two *games* played with automata that we call the *satisfiability* game and the *consequence game*. The satisfiability game related to an automaton $\mathbb{A}$ is played between players $\exists$ ("Eloise") and $\forall$ ("Abélard"), and the aim of Eloise is to construct a model accepted by $\mathbb{A}$ step by step. The consequence game related to two automata, $\mathbb{A}$ and $\mathbb{B}$, is also played between two players, now prosaically called 'player I' and 'player II'; here the aim of the second player is to systematically show that the first automaton implies the second one, in some strong, structural sense. Both games proceed in rounds, moving from one basic position to another, and these moves all involve one-step models over the collection $A^\sharp$ of binary relations over the carrier set of the automaton $\mathbb{A}$ (and of the collection $B^\sharp$ of binary relations over the carrier set of the second automaton, in case of the consequence game). Furthermore, for both kinds of games, infinite plays naturally induce streams of binary relations, and the winning conditions of both games are expressed in terms of the collection of *traces* through such streams. For a more detailed introduction of these games, in the setting of the standard modal μ-calculus, we refer to Enqvist, Seifan & Venema [12].

6.1. Traces and canonical one-step models

We first introduce some terminology and notation for the auxiliary notions of traces and canonical one-step models.

Definition 6.1. Fix a set A. We let $A^\sharp$ denote the set of binary relations over A, that is, $A^\sharp := \mathsf{P}(A \times A)$.

Given a finite word $\Sigma = R_1 R_2 R_3 \dots R_k$ over the set $A^\sharp$, a *trace* through Σ is a finite A-word $\alpha = a_0 a_1 a_2 \dots a_k$ such that $a_i R_{i+1} a_{i+1}$ for all $i < k$. A *trace* through an $A^\sharp$-stream $\Sigma = R_1 R_2 R_3 \dots$ is an A-stream $\alpha = a_0 a_1 a_2 \dots$, such that $a_i R_{i+1} a_{i+1}$ for all $i < \omega$). In both cases we denote the set of traces through Σ as Tr_Σ.

Given a stream $\Sigma = R_1 R_2 R_3 \dots$ over $A^\sharp$ we denote by $\Sigma|_k$ the word $R_1 \dots R_k$, and for a trace $\tau = a_0 a_1 a_2 \dots$ on Σ we denote by $\tau|_k$ the restricted trace $a_0 \dots a_k$ on $\Sigma|_k$. We use similar notation for restrictions of finite words over $A^\sharp$. ◁

Definition 6.2. Fix a finite set A and a priority map $\Omega : A \to \omega$. We let NBT_Ω denote the set of $A^\sharp$-streams that contain no bad trace, that is, no trace $\tau = a_0 a_1 \dots$ such that $\max(\Omega[\mathit{Inf}(\tau)])$, the highest priority occurring infinitely often on τ, is odd. In case Ω is the priority map of a coalgebra automaton $\mathbb{A}$, we will usually write $NBT_\mathbb{A}$ instead of NBT_Ω. ◁

It is easy to define a stream automaton that recognizes exactly the *complement* of the set $NBT_\mathbb{A}$, for any parity automaton $\mathbb{A}$; by closure under complementation it then follows that $NBT_\mathbb{A}$ itself is also an ω-regular language; details can be found in [38].

Proposition 6.3. *Given a finite set A and a priority map $\Omega : A \to \omega$, there is a parity stream automaton recognizing the set NBT_Ω, seen as a stream language over $A^\sharp$.*

Now we consider the one-step models based on the set $A^\sharp$ of binary relations over A.

Definition 6.4. Given a set A and a state $a \in A$, the *natural* or *canonical a-marking* on the set $A^\sharp$ is defined as the map $n_a^A : A^\sharp \to \mathsf{P}A$ given by

$$n_a^A : R \mapsto R[a].$$

In case A is known from context, we will usually write n_a rather than n_a^A, and define, for a one-step formula $\alpha \in \mathtt{1ML}^+_\Lambda(\mathtt{X}, A)$, $[\![\alpha]\!]^1_a := \{\Gamma \in \mathsf{T}_\mathtt{X} A^\sharp \mid A^\sharp, \Gamma, n_a \Vdash^1 \alpha\}$. ◁

Remark 6.5. The notation $[\![\alpha]\!]^1_a$ may seem to be somewhat ambiguous, since it does not refer to the ambient variable set A. However, by Proposition 3.10 and Corollary 3.9 it follows that, for any pair of sets A, B such that $\alpha \in \mathtt{1ML}^+_\Lambda(\mathtt{X}, A) \cap \mathtt{1ML}^+_\Lambda(\mathtt{X}, B)$ we have

$$\{\Gamma \in \mathsf{T}_\mathtt{X} A^\sharp \mid A^\sharp, \Gamma, n_a^A \Vdash^1 \alpha\} = \{\Gamma \in \mathsf{T}_\mathtt{X} B^\sharp \mid B^\sharp, \Gamma, n_a^B \Vdash^1 \alpha\}.$$

As another instance of Corollary 3.9, for any subset $\mathcal{R} \subseteq A^\sharp$ and for any object $\Gamma \in \mathsf{T}_\mathtt{X}\mathcal{R}$ we have

$$A^\sharp, \Gamma, n_a \Vdash^1 \alpha \text{ iff } \mathcal{R}, \Gamma, n_a{\upharpoonright}_\mathcal{R} \Vdash^1 \alpha,$$

where $n_a{\upharpoonright}_{\mathcal{R}}$ is the natural a-marking on $A^\sharp$, restricted to $\mathcal{R}$. If no confusion is likely, we will often denote the marking $n_a{\upharpoonright}_{\mathcal{R}}$ simply by n_a. ◁

Remark 6.6. We may think of any object $\Gamma \in \mathsf{T}_{\mathtt{X}}A^\sharp$ as a *family* $\{(A^\sharp, \Gamma, n_a) \mid a \in A\}$ of one-step models *on the same one-step frame* $(A^\sharp, \Gamma)$. It may occasionally be useful, however, to consider this 'family of one-step models' as one single model. To do so, we involve, for each $a \in A$, the substitution $\tau_a : A \to A \times A$ that *tags* each variable $b \in A$ with its 'origin' a, that is, $\tau_a : b \mapsto (a, b)$. One may verify, on the basis of a straightforward formula induction, that

$$A^\sharp, \Gamma, n_a \Vdash^1 \alpha \text{ iff } A^\sharp, \Gamma, \mathsf{id}_{A^\sharp} \Vdash^1 \alpha[\tau_a]$$

for each one-step formula $\alpha \in \mathtt{1ML}^+_\Lambda(\mathtt{X}, A)$. In particular, it follows that

$$\Gamma \in \bigcap_{a \in B} [\![\Theta(a)]\!]^1_a \text{ iff } A^\sharp, \Gamma, \mathsf{id}_{A^\sharp} \Vdash^1 \bigwedge_{a \in B} \alpha[\tau_a],$$

for any family $\{\Theta(a) \mid a \in B\}$ of formulas. ◁

The following rather technical lemma will be needed to ensure that we can make simplifying assumptions on the strategies that players use in the games that we are about to introduce. Recall that $\leftrightarroweq^1_{\Lambda,f}$ is the one-step bisimulation relation introduced in Definition 3.12, and that $\lhd_\Theta$ is the relation of one state being active in another relative to Θ, as introduced in Definition 5.4.

Proposition 6.7. *Let $\Theta : A \to \mathtt{1ML}^+_\Lambda(\mathtt{X}, A)$ be some map, and fix some $R \in A^\sharp$ and some $\mathcal{Q} \subseteq A^\sharp$, $\Gamma \in \mathsf{T}_{\mathtt{X}}\mathcal{Q}$ such that*

$$\Gamma \in \bigcap_{a \in \mathsf{Ran}R} [\![\Theta(a)]\!]^1_a.$$

(1) There are $\mathcal{Q}' \subseteq A^\sharp$ and $\Gamma' \in \mathsf{T}_{\mathtt{X}}\mathcal{Q}'$ such that $\Gamma' \in \bigcap_{a \in \mathsf{Ran}R} [\![\Theta(a)]\!]^1_a$, $\subseteq : (\mathcal{Q}', \Gamma') \leftrightarroweq^1_{\Lambda,f} (\mathcal{Q}, \Gamma)$, and for each $Q \in \mathcal{Q}'$: $\mathsf{Dom}Q \subseteq \mathsf{Ran}R$.

(2) There are $\mathcal{Q}' \subseteq A^\sharp$ and $\Gamma' \in \mathsf{T}_{\mathtt{X}}\mathcal{Q}'$ such that $\Gamma' \in \bigcap_{a \in \mathsf{Ran}R} [\![\Theta(a)]\!]^1_a$, $\subseteq : (\mathcal{Q}', \Gamma') \leftrightarroweq^1_{\Lambda,f} (\mathcal{Q}, \Gamma)$, and for each $Q \in \mathcal{Q}'$: $b \lhd_\Theta a$ whenever $(a, b) \in Q$.

(3) Let there be, for some subset $B \subseteq A$, a collection $\{\mathcal{G}_b \subseteq \mathsf{P}A \mid b \in B\}$ such that for every $C \in \mathsf{P}A$ there is a $C' \in \mathcal{G}_b$ such that $C' \subseteq C$. Furthermore, assume that, for each $b \in B$:

$$\Theta(b) \in \{\alpha[\chi] \mid \alpha \in \mathtt{D}(\mathcal{G}_b)\}.$$

Then there are $\mathcal{Q}' \subseteq A^\sharp$ and $\Gamma' \in \mathsf{T}_{\mathtt{X}}\mathcal{Q}'$ such that $\Gamma' \in \bigcap_{a \in \mathsf{Ran}R} [\![\Theta(a)]\!]^1_a$, $\subseteq : (\mathcal{Q}', \Gamma') \leftrightarroweq^1_{\Lambda,f} (\mathcal{Q}, \Gamma)$, and for each $Q \in \mathcal{Q}'$: $Q[b] \in \mathcal{G}_b$, for all $b \in B$.

Proof. For part (1), consider the map $F : A^\sharp \to A^\sharp$ given by

$$F(Q) := Q \cap (\mathsf{Ran}R \times A).$$

We leave it for the reader to verify that the pair $(\mathsf{Ran}(F), (\mathsf{T}_{\mathtt{X}}F)\Gamma)$ meets the requirements. Part (2) is proved similarly, using the map $Q \mapsto Q \cap \rhd$.

For part (3), we will prove the statement for the special case where B is a singleton $B = \{b\}$, while we show that $\mathcal{Q}'$ additionally satisfies

$$\{Q[a] \mid Q \in \mathcal{Q}'\} \subseteq \{Q[a] \mid Q \in \mathcal{Q}\} \tag{15}$$

for all $a \neq b$. The general case can then be obtained from the special one by a straightforward iteration, taking care of B's elements one by one. The role of (15) is to ensure that new iterations do not spoil the progress booked in earlier rounds.

So let $b \in A$ be such that $\Theta(b)$ is of the form $\alpha_b[\chi]$ for some $\alpha_b \in \mathtt{D}(\mathcal{G}_b)$. By assumption on Γ and Corollary 3.9 we have $\mathcal{Q}, \Gamma, n_b \Vdash^1 \alpha_b[\chi]$. Applying Proposition 3.19 we obtain a one-step model (S, σ, m) and a map $F : S \to \mathcal{Q}$ such that $(\mathsf{T_X}F)\sigma = \Gamma$, $\mathsf{Ran}(F) = \mathcal{Q}$, $S, \sigma, m \Vdash^1 \alpha_b[\chi]$ and, for all $s \in S$ we have $m(s) \subseteq n_b(F_s) = F_s[b]$ and $m(s) \in \mathcal{G}_b$.

Now define the map $G : S \to A^\sharp$ by setting

$$G_s[a] := \begin{cases} m(s) & \text{if } a = b \\ F_s[a] & \text{if } a \neq b. \end{cases}$$

We claim that the object $\Gamma' \in \mathsf{T_X}A^\sharp$, given as

$$\Gamma' := (\mathsf{T_X}G)\sigma,$$

together with the set $\mathcal{Q}' := \mathsf{Ran}G$, has all the desired properties.

To start with, it is easy to see that $F : (S, \sigma) \to (\mathcal{Q}, \Gamma)$ and $G : (S, \sigma) \to (\mathcal{Q}', \Gamma')$ are surjective one-step homomorphisms, so that it follows from Proposition 3.13 and the fact that $G(s) \subseteq F(s)$ for all $s \in S$ that $\subseteq : (\mathcal{Q}', \Gamma') \leftrightarrows^1_{\Lambda,f} (\mathcal{Q}, \Gamma)$.

Our next step is to prove that $\Gamma' \in \bigcap_{a \in \mathsf{Ran}R} [\![\Theta(a)]\!]^1_a$, or equivalently, that

$$A^\sharp, \Gamma', n_a \Vdash^1 \Theta(a) \text{ for all } a \in \mathsf{Ran}R. \tag{16}$$

To see this, make a case distinction. If $a = b$, it follows from the definitions that $(n_b \circ G)(s) = n_b(G_s) = G_s[b] = m(s)$, so that G is a one-step model homomorphism

$$G : (S, \sigma, m) \to (A^\sharp, \Gamma', n_b).$$

From this (16) is immediate by $S, \sigma, m \Vdash^1 \alpha_b[\chi]$.

In case $a \neq b$ we have to do a bit more work. Define the A-marking $m_a : S \to \mathsf{P}A$ by putting $m_a(s) := F_s[a]$. It is easy to check that this turns F into a one-step model homomorphism

$$F : (S, \sigma, m_a) \to (A^\sharp, \Gamma, n_a)$$

and G into a one-step model homomorphism

$$G : (S, \sigma, m_a) \to (A^\sharp, \Gamma', n_a).$$

But then by naturality we immediately obtain that

$$\begin{aligned} A^\sharp, \Gamma, n_a \Vdash^1 \alpha \quad & \text{iff} \quad S, \sigma, m_a \Vdash^1 \alpha \\ & \text{iff} \quad A^\sharp, \Gamma', n_a \Vdash^1 \alpha \end{aligned}$$

for all one-step formulas $\alpha \in \mathtt{1ML}^+_\Lambda(A)$, so in particular for $\alpha = \Theta(a)$. Thus (16) follows by the assumption that $A^\sharp, \Gamma, n_a \Vdash^1 \Theta(a)$.

Table 3
Admissible moves in the satisfiability game $\mathcal{S}(\mathbb{A})$.

Position	Player	Admissible moves
$R \in A^\sharp$	$\exists$	$\{(\mathcal{R},\Gamma) \in \mathsf{P}A^\sharp \times \mathsf{T}_{\mathsf{X}}A^\sharp \mid \Gamma \in \mathsf{T}_{\mathsf{X}}\mathcal{R} \cap \bigcap_{a\in\mathsf{Ran}R}[\![\Theta(a)]\!]^1_a\}$
$(\mathcal{R},\Gamma)$	$\forall$	$\{R \in A^\sharp \mid R \subseteq R'$ for some $R' \in \mathcal{R}\}$

Having established (16) we continue with proving that

$$Q'[b] \in \mathcal{G}_b \tag{17}$$

for each $Q' \in \mathcal{Q}'$. This is in fact easy, since each such Q' is by definition of the form G_s, for some $s \in S$. Hence $Q'[b] = m(s) \in \mathcal{G}_b$ by the assumptions on the one-step model (S, σ, m).

This leaves (15) to take care of. Let $a \in A$ be distinct from b, and take an arbitrary $Q' \in \mathcal{Q}'$, say, $Q' = G_s$ for $s \in S$. Then by definition of $G : S \to A^\sharp$ we have $G_s[a] = F_s[a]$, and since $\mathsf{Ran}(F) \subseteq \mathcal{Q}$ we are done. □

6.2. The satisfiability game

We now turn to the definition and technical details of the satisfiability game $\mathcal{S}(\mathbb{A})$ associated with an automaton $\mathbb{A}$. As we shall see, it is rather similar to a *tableau* in that the aim of $\exists$ is to step-by-step construct a model that is accepted by $\mathbb{A}$.

Definition 6.8. Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a modal automaton. Then the *satisfiability game* $\mathcal{S}(\mathbb{A})$ is the graph game of which the moves are given by Table 3. Positions of the form $R \in A^\sharp$ are called *basic*.

The winner of an infinite play of the satisfiability game is given by the induced stream $\Sigma = R_0R_1 \ldots \in (A^\sharp)^\omega$ of basic positions. This winner is $\exists$ if Σ belongs to the set NBT_Ω, that is, if Σ contains no bad traces, and it is $\forall$ otherwise. A winning strategy of $\forall$ in $\mathcal{S}(\mathbb{A})$ may be called a *refutation* of $\mathbb{A}$. ◁

The satisfiability game is sound and complete in the following sense. For a proof of this Proposition we refer to [14].

Proposition 6.9 *(Adequacy). Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a modal automaton $\mathbb{A}$. Then $\exists$ has a winning strategy in $\mathcal{S}(\mathbb{A})$ iff the language recognized by $\mathbb{A}$ is non-empty.*

The purpose of the following Proposition is to justify some simplifying assumptions on the strategies employed by $\exists$ in the satisfiability game. Here $\overline{\mathsf{P}}$ denotes the relation lifting associated to P as in Definition A.3.

Proposition 6.10. *Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a Λ-automaton, and let $\mathcal{N} \subseteq A^\sharp$ be a set of relations. Assume that for every basic position $R \in A^\sharp$ of the satisfiability game, and every legitimate move $(\mathcal{R}, \Gamma)$ of $\exists$ there is a legitimate move $(\mathcal{R}', \Gamma')$ such that $\mathcal{R}' \subseteq \mathcal{N}$ and $\mathcal{R}' \mathrel{\overline{\mathsf{P}}{\subseteq}} \mathcal{R}$. Then for any winning position in $\mathcal{S}(\mathbb{A})$ $\exists$ has a winning strategy that restricts her moves to pairs $(\mathcal{R}, \Gamma)$ with $\mathcal{R} \subseteq \mathcal{N}$.*

Proof. Assume that $\exists$ has a winning strategy f in the game $\mathcal{S}(\mathbb{A})$ initialized at position R_0. We need to provide her with a winning $\mathcal{N}$-strategy, that is, a strategy $\overline{f}$ that always selects moves $(\mathcal{R}, \Gamma)$ with $\mathcal{R} \subseteq \mathcal{N}$.

We will define this strategy $\overline{f}$ by induction on the length of partial $\mathcal{S}(\mathbb{A})$-plays. Simultaneously, for any such play

$$\Sigma = R_0(\mathcal{R}_0, \Gamma_0)R_1(\mathcal{R}_1, \Gamma_1) \ldots R_k$$

which is $\overline{f}$-guided, we will define a parallel play

Table 4
Admissible moves in the alternative satisfiability game.

Position	Player	Admissible moves
$R \in A^\sharp$	$\exists$	$\bigcap_{a \in \mathsf{Ran} R} [\![\Theta(a)]\!]^1_a$
$\Gamma \in \mathsf{T_X} A^\sharp$	$\forall$	$\mathsf{Base}(\Gamma)$

$$\Sigma^* = R_0(\mathcal{R}^*_0, \Gamma^*_0)R_1(\mathcal{R}^*_1, \Gamma^*_1)\dots R_k$$

which is guided by $\exists$'s winning strategy f. If we can maintain such a shadow play infinitely long, it is routine to prove that $\overline{f}$ is winning for $\exists$.

For the case where $k = 0$ there is nothing to prove, so assume inductively that there are partial plays Σ and Σ^* as above. Observe that since the last positions of Σ and Σ^* are identical, the set of $\exists$'s legitimate moves in Σ and Σ^* are the same. Let $(\mathcal{R}, \Gamma)$ be the move prescribed by $\exists$'s winning strategy f in the partial play Σ^*, then by assumption there is a legitimate move $(\mathcal{R}', \Gamma')$ such that $\mathcal{R}' \subseteq \mathcal{N}$ and $\mathcal{R}' \overline{\mathsf{P}}{\subseteq} \mathcal{R}$. Then we let

$$\overline{f}(\Sigma) := (\mathcal{R}', \Gamma')$$

be $\exists$'s move in Σ. This defines the strategy $\overline{f}$.

To finish the inductive step, consider an arbitrary continuation of the play $\Sigma \cdot (\mathcal{R}', \Gamma')$, say, where $\forall$ plays some relation Q. By definition, Q is a subset of some $Q' \in \mathcal{R}'$, while by $\mathcal{R}' \overline{\mathsf{P}}{\subseteq} \mathcal{R}$ we may find some $Q'' \in \mathcal{R}$ such that $Q' \subseteq Q''$. But then it follows from $Q \subseteq Q''$ that Q is also a legitimate move for $\forall$ in $\Sigma^* \cdot (\mathcal{R}, \Gamma)$. In other words, the two $k+1$-length plays $\Sigma \cdot (\mathcal{R}', \Gamma') \cdot Q$ and $\Sigma^* \cdot (\mathcal{R}, \Gamma) \cdot Q$ satisfy the required conditions. □

Remark 6.11. As a consequence of Proposition 6.10, we can always make some minimality assumptions on $\exists$'s strategy in the satisfiability game. In particular, suppose that $\exists$, at some position $R \in A^\sharp$ in a play of $\mathcal{S}(\mathbb{A})$, picks a move $(\mathcal{R}, \Gamma) \in \mathsf{P}A^\sharp \times \mathsf{T_X}A^\sharp$. Then by the Propositions 6.7 and 6.10 we can assume without loss of generality that, for all $Q \in \mathcal{R}$:

(1) $\mathsf{Dom}(Q) \subseteq \mathsf{Ran}(R)$;
(2) b occurs in $\Theta(a)$, for all $(a, b) \in Q$;
(3) $|Q[a]| \leq 1$, whenever $\Theta(a)$ is a disjunctive formula. ◁

Remark 6.12. We remark in passing that the moves made by $\exists$ can always be assumed without loss of generality to be of the form $(\mathcal{R}, \Gamma)$ where $\mathcal{R}$ is the unique *smallest* subset of $A^\sharp$ with $\Gamma \in \mathsf{T_X}\mathcal{R}$; this set is called the *base* of Γ, and denoted as $\mathsf{Base}(\Gamma)$ (cf. Definition A.11). (That this set exists follows from standard results in coalgebra, together with the assumption that T (and hence $\mathsf{T_X}$) preserves inclusion maps.) Based on this, an alternative but equivalent formulation of the satisfiability game (more compatible with the versions used in [11,12]) is given in Table 4. ◁

6.3. Consequence game

The *consequence game* $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ is played between two players I (female) and II (male). To provide some intuitions, the aim of the second player is to provide "step-by-step" a construction that systematically turns any winning strategy for $\exists$ in $\mathcal{S}(\mathbb{A})$ into a winning strategy in $\mathcal{S}(\mathbb{A}')$. A strategy for player II thus provides a tight structural connection between the two automata. More in detail, the basic positions of the game are pairs $(R, R') \in A^\sharp \times A'^\sharp$, and at such a position Player I picks an admissible move $(\mathcal{R}, \Gamma)$ for $\exists$ in $\mathcal{S}(\mathbb{A})$ at the position R. After this Player II must respond with an admissible move $(\mathcal{R}', \Gamma')$ for $\exists$ in $\mathcal{S}(\mathbb{A}')$ at the position R', but also, crucially, with a full one-step bisimulation $Z \subseteq \mathcal{R} \times \mathcal{R}'$ linking Γ and Γ'. This round of the play finishes with player I picking an element (Q, Q') of Z as the next basic position.

Table 5
Admissible moves in the consequence game $\mathcal{C}(\mathbb{A}, \mathbb{A}')$.

Position	P	Moves
(R, R')	I	$\{((\mathcal{R}, \Gamma), R') \mid \Gamma \in \mathsf{T}_{\mathsf{X}}\mathcal{R} \cap \bigcap_{a \in \mathsf{Ran} R} [\![\Theta(a)]\!]^1_a\}$
$((\mathcal{R}, \Gamma), R')$	II	$\{((\mathcal{R}, \Gamma), (\mathcal{R}', \Gamma')) \mid \Gamma' \in \mathsf{T}_{\mathsf{X}}\mathcal{R}' \cap \bigcap_{b \in \mathsf{Ran} R'} [\![\Theta(b)]\!]^1_b\}$
$((\mathcal{R}, \Gamma), (\mathcal{R}', \Gamma'))$	II	$\{Z \mid Z : (\mathcal{R}, \Gamma) \leftrightarrows^1_{\Lambda, f} (\mathcal{R}', \Gamma')\}$
$Z \subseteq A^\sharp \times A'^\sharp$	I	Z

We can now provide the formal definition of the consequence game $\mathcal{C}(\mathbb{A}, \mathbb{A}')$. Note that this version of the game differs from the one given in [11,12] in that it explicitly refers to support sets for the objects Γ and Γ', and in that it allows player II to come up with a binary relation rather than with a partial function.

Definition 6.13. Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ and $\mathbb{A}' = (A', \Theta', \Omega', a'_I)$ be Λ-automata. The rules of the *consequence game* $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ are given by Table 5. Positions of the form $(R, R') \in A^\sharp \times A'^\sharp$ are called *basic*. For the *winning conditions* of this game, consider an infinite play Σ of $\mathcal{C}(\mathbb{A}, \mathbb{A}')$, and let

$$(R_0, R'_0)(R_1, R'_1)(R_2, R'_2) \dots$$

be the induced stream of basic positions in Σ. Then player I is the winner of Σ if $R_0 R_1 \dots \in NBT_\Omega$ but $R'_0 R'_1 \dots \notin NBT_{\Omega'}$; that is, if there is a bad trace on the $\mathbb{A}'$-side but not on the $\mathbb{A}$-side.

If the position $\big(\{(a_I, a_I)\}, \{(a'_I, a'_I)\}\big)$ is a winning position for player II in $\mathcal{C}(\mathbb{A}, \mathbb{A}')$, we say that $\mathbb{A}'$ is a *game consequence of* $\mathbb{A}$, notation: $\mathbb{A} \models_{\mathsf{G}} \mathbb{A}'$. ◁

We have the following soundness result for this game.

Proposition 6.14. *For any two modal automata $\mathbb{A}$ and $\mathbb{A}'$ it holds that*

$$\mathbb{A} \models_{\mathsf{G}} \mathbb{A}' \textit{ implies } \mathbb{A} \models \mathbb{A}'. \tag{18}$$

For future reference we give the following proposition, stating that the consequence relation $\models_{\mathsf{G}}$ is reflexive and transitive.

Proposition 6.15. *Let $\mathbb{A}$, $\mathbb{A}'$ and $\mathbb{A}''$ be modal automata.*

(1) $\mathbb{A} \models_{\mathsf{G}} \mathbb{A}$;
(2) if $\mathbb{A} \models_{\mathsf{G}} \mathbb{A}'$ *and* $\mathbb{A}' \models_{\mathsf{G}} \mathbb{A}''$ *then* $\mathbb{A} \models_{\mathsf{G}} \mathbb{A}''$.

Proof. Clearly, the proof of the first item is trivial. Concerning the transitivity of $\models_{\mathsf{G}}$, it is a routine exercise to verify that player II can compose any two winning strategies in the games $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ and $\mathcal{C}(\mathbb{A}', \mathbb{A}'')$, respectively, to obtain a winning strategy in the game $\mathcal{C}(\mathbb{A}, \mathbb{A}'')$. □

Remark 6.16. Note that by Proposition 3.13 we always have $F : (\mathcal{R}, \Gamma) \leftrightarrows^1_{\Lambda, f} (F[\mathcal{R}], \mathsf{T}_{\mathsf{X}} F(\Gamma))$, for any map F having $\mathcal{R}$ as its domain. A strategy for player II in the consequence game $\mathcal{C}(\mathbb{A}, \mathbb{B})$ is said to be *functional* if his response to any play ending in a position $((\mathcal{R}, \Gamma), R')$ is of the form $(F[\mathcal{R}], \mathsf{T}_{\mathsf{X}} F(\Gamma))$ followed by (the graph of) F for some map $F : \mathcal{R} \to B^\sharp$. ◁

Similar to the satisfiability game, we will often want to make certain assumptions on the strategy of player I in the consequence game. These assumptions will be justified by the following analog of Proposition 6.10.

Proposition 6.17. *Let $\mathbb{A} = (A, \Theta_A, \Omega_A, a_I)$ and $\mathbb{B} = (B, \Theta_B, \Omega_B, b_I)$ be Λ-automata, and let $\mathcal{N} \subseteq A^\sharp$ be a set of relations. Assume that for every basic position $(Q, R) \in A^\sharp \times B^\sharp$ of the consequence game, and every* *legitimate move* $(\mathcal{Q}', \Gamma')$ *of player I, she has a legitimate move* $(\mathcal{Q}, \Gamma)$ *such that* $\mathcal{Q} \subseteq \mathcal{N}$ *and* $\subseteq : (\mathcal{Q}, \Gamma) \leftrightarrows^1_{\Lambda,f} (\mathcal{Q}', \Gamma')$.

Then for any winning position in $\mathcal{C}(\mathbb{A}, \mathbb{B})$*, player I has a winning strategy that restricts her moves to pairs* $(\mathcal{Q}, \Gamma)$ *with* $\mathcal{Q} \subseteq \mathcal{N}$.

Proof. We write $Q_0 := \{(a_I, a_I)\}$, $R_0 := \{(b_I, b_I)\}$, and abbreviate $\mathcal{C} := \mathcal{C}(\mathbb{A}, \mathbb{B})@(Q_0, R_0)$ (check Definition A.13 for notation). Let f be a winning strategy for player I in $\mathcal{C}$. In the same game we will provide I with a winning strategy $\overline{f}$, that restricts her moves to pairs $(\mathcal{Q}, \Gamma)$ with $\mathcal{Q} \subseteq \mathcal{N}$. This strategy $\overline{f}$ will be defined by induction on the length of a partial $\overline{f}$-guided play, while by a simultaneous induction we will (†) associate with each $\overline{f}$-guided play $\Sigma = (Q_n, R_n)_{n \leq k}$ an f-guided shadow play $\Sigma' = (Q'_n, R_n)_{n \leq k}$ such that $Q_n \subseteq Q'_n$ for all $n \leq k$.

Clearly this holds at the start of every $\mathcal{C}$-play if we take $Q'_0 := Q_0$. For the inductive step of the definition, fix a partial $\overline{f}$-guided play $\Sigma = (Q_n, R_n)_{n \leq k}$, and let $\Sigma' = (Q'_n, R_n)_{n \leq k}$ be the inductively given shadow play. In order to provide player I with a move in Σ, first consider the move $(\mathcal{Q}', \Gamma') \in \mathsf{P}A^\sharp \times \mathsf{T}_\mathsf{X} A^\sharp$ provided by f in the shadow play Σ'. By assumption there is a legitimate move $(\mathcal{Q}, \Gamma)$ at position (Q'_k, R_k) such that $\mathcal{Q} \subseteq \mathcal{N}$ and $\subseteq : (\mathcal{Q}, \Gamma) \leftrightarrows^1_{\Lambda,f} (\mathcal{Q}', \Gamma')$. Since $Q_k \subseteq Q'_k$ (and hence, $\mathsf{Ran} Q_k \subseteq \mathsf{Ran} Q'_k$), it is easy to see that this move $(\mathcal{Q}, \Gamma)$ is also legitimate at the last position (Q_k, R_k) of Σ. Hence we may take this pair $(\mathcal{Q}, \Gamma)$ to be the move suggested by the strategy $\overline{f}$.

Continuing the inductive definition, suppose that player II's answers to I's move $(\mathcal{Q}, \Gamma)$ are, successively, $(\mathcal{R}, \Delta) \in \mathsf{P}B^\sharp \times \mathsf{T}_\mathsf{X} B^\sharp$ and $\mathcal{Z} \subseteq A^\sharp \times B^\sharp$. Now consider the relation $\mathcal{Z}' \subseteq A^\sharp \times B^\sharp$ defined by $\mathcal{Z}' := \supseteq ; \mathcal{Z}$. We claim that

$$(\mathcal{R}, \Delta) \text{ and } \mathcal{Z}' \text{ are legitimate moves for II at position } ((\mathcal{Q}', \Gamma'), R) \tag{19}$$

and

$$\text{for all } (Q', R) \in \mathcal{Z}' \text{ there is a } (Q, R) \in \mathcal{Z} \text{ such that } Q \subseteq Q'. \tag{20}$$

For a proof of (19), observe that the legitimacy of $(\mathcal{R}, \Delta)$ is obvious. For the legitimacy of $\mathcal{Z}'$ we have to prove that $\mathcal{Z}' : (\mathcal{Q}', \Gamma') \leftrightarrows^1_{\Lambda,f} (\mathcal{R}, \Delta)$; but by Proposition 3.13 this follows from $\supseteq : (\mathcal{Q}', \Gamma') \leftrightarrows^1_{\Lambda,f} (\mathcal{Q}, \Gamma)$ and $\mathcal{Z} : (\mathcal{Q}, \Gamma) \leftrightarrows^1_{\Lambda,f} (\mathcal{R}, \Delta)$. The claim (20) is immediate from the definitions.

Based on the statements (19) and (20), we can finish our inductive definition: Suppose that in the play $\Sigma' \cdot ((\mathcal{Q}', \Gamma'), R_k) \cdot ((\mathcal{Q}', \Gamma'), (\mathcal{R}, \Delta)) \cdot \mathcal{Z}'$, player I's winning strategy f tells her to pick a pair $(Q', R) \in \mathcal{Z}$; then in the play $\Sigma \cdot ((\mathcal{Q}, \Gamma), R_k) \cdot ((\mathcal{Q}, \Gamma), (\mathcal{R}, \Delta)) \cdot \mathcal{Z}$ we let the strategy $\overline{f}$ pick a pair $(Q, R) \in \mathcal{Z}$ as given by (20). Clearly this is a legitimate move for player I. Finally, where $\Sigma \cdot (Q, R)$ is the continuation of Σ in terms of basic positions, the associated continuation of the shadow play is $\Sigma' \cdot (Q', R)$, and so it is obvious that player I has been able to maintain the constraint (†).

It should be clear that the thus defined strategy $\overline{f}$ always picks legitimate moves of the right type. It remains to check that it is a winning strategy in $\mathcal{C}$.

It is straightforward to verify that player I will never get stuck in an $\overline{f}$-guided play, so we confine our attention to infinite plays. Let $\Sigma = (Q_n, R_n)_{n<\omega}$ be an infinite $\overline{f}$-guided play, then clearly there is an infinite f-guided shadow play $\Sigma' = (Q'_n, R_n)_{n<\omega}$ such that $Q_n \subseteq Q'_n$ for all $n < \omega$. By our assumption that f is a winning strategy in $\mathcal{C}$, the play Σ' is a win for player I. That is, all traces through $(Q'_n)_{n<\omega}$ are good, while there is a bad trace through $(R_n)_{n<\omega}$. Obviously then, all traces through $(Q_n)_{n<\omega}$ are good, and so the existence of a bad trace through $(R_n)_{n<\omega}$ means that Σ as well is a win for player I. □

7. Taming traces

7.1. Disjunctive and semi-disjunctive automata

We now introduce two classes of special Λ-automata, for which the satisfiability game simplifies significantly. Essentially, these automata are designed to prevent that the number of traces that we need to consider in a play of the satisfiability game multiplies uncontrollably, making the combinatorics involved in keeping track of these traces un-manageable.

Definition 7.1. A Λ-automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$ with free variables $\mathtt{X}$ is said to be *disjunctive* (relative to a disjunctive basis $\mathtt{D}$) if for all $a \in A$, the formula $\Theta(a) \in \mathtt{1ML}_\Lambda(\mathtt{X}, A)$ is of the form $\pi \wedge \delta$ with $\pi \in \mathtt{Bool}(\mathtt{X})$ and $\delta \in \mathtt{D}(A)$. ◁

Note that if $\mathbb{A} = (A, \Theta, \Omega, a_I)$ is disjunctive, every one-step formula $\Theta(a)$ is disjunctive indeed (cf. Remark 3.16).

The weaker concept of a *semi-disjunctive automaton*, which is similar to Walukiewicz' *weakly aconjunctive formulas*, is more subtle. They are designed to control the branching of traces in the satisfiability game, within each given cluster of the automaton.

Definition 7.2. Given an automaton $\mathbb{A}$ and a state $a \in A$, a subset $B \subseteq A$ is called *a-safe* if, for all $b \neq b'$ in B, at least one of b, b' either belongs to a different cluster than a, or has an even priority, which is higher than all odd properties that are reached in the cluster of a. We let $\mathsf{Sf}_a \subseteq \mathsf{P}(A)$ denote the set of a-safe subsets of A.

The automaton $\mathbb{A}$ is said to be *semi-disjunctive* if, for all $a \in A$, $\Theta(a)$ is of the form $\pi \wedge \delta[\chi_A]$ with $\pi \in \mathtt{Bool}(\mathtt{X})$ and $\delta \in \mathtt{D}(\mathsf{Sf}_a)$. ◁

Semi-disjunctive automata are tightly related to what we call the *thin* satisfiability game, in which the moves of ∀ are restricted in order to control the branching of traces. As we showed in [12] (cf. Proposition 6.6), in a play of the thin satisfiability game essentially there will be at most finitely many bad traces.

Definition 7.3. Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a Λ-automaton. Given a state $a \in A$, we call a relation $R \in A^\sharp$ *thin* with respect to $\mathbb{A}$ and a, or *a-thin (with respect to* $\mathbb{A}$*)*, if:

(1) for all $b \in A$ with aRb, we have $b \lhd a$;
(2) $R[a] \subseteq A$ is C_a-safe.

Given a subset $B \subseteq A$, we call R *B-thin* if it is b-thin for all $b \in B$. We denote the collection of A-thin relations in $A^\sharp$ by ${A^\sharp}_{thin}$. ◁

Definition 7.4. The *thin satisfiability game* $\mathcal{S}_{thin}(\mathbb{A})$ is defined just as $\mathcal{S}(\mathbb{A})$, except that admissible moves R of ∀ are subject to the additional *thinness* constraint: $R \in {A^\sharp}_{thin}$. ◁

Proposition 7.5. *Let* $\mathbb{A}$ *be semi-disjunctive. Then for each player* $\Pi \in \{\exists, \forall\}$*, a position is winning for* Π *in* $\mathcal{S}(\mathbb{A})$ *iff it is winning for* Π *in* $\mathcal{S}_{thin}(\mathbb{A})$.

We note the following closure properties for disjunctive and semi-disjunctive automata. Here we say that an automaton is (semi-)disjunctive *modulo provable equivalence* if it is provably equivalent to a (semi-)disjunctive automaton. We omit the proof of this Proposition, since it is completely similar to that of Proposition 6.15 in [12].

Proposition 7.6. *Let* $\mathbb{A}$ *and* $\mathbb{B}$ *be two* Λ*-automata, where the modal signature* Λ *admits a disjunctive basis.*

(1) If $\mathbb{A}$ is disjunctive, then it is also semi-disjunctive.
(2) If $\mathbb{A}$ and $\mathbb{B}$ are disjunctive, then so is $\mathbb{A} \vee \mathbb{B}$.
(3) If $\mathbb{A}$ and $\mathbb{B}$ are semi-disjunctive, then so is $\mathbb{A} \vee \mathbb{B}$.
(4) If $\mathbb{A}$ and $\mathbb{B}$ are semi-disjunctive, then so is $\mathbb{A} \wedge \mathbb{B}$, modulo provable equivalence.
(5) If $\mathbb{A}$ and $\mathbb{B}$ are semi-disjunctive, then so is $\mathbb{A}[\mathbb{B}/x]$, modulo provable equivalence.
(6) If $\mathbb{A}$ is disjunctive and positive in x, then $\mathbb{A}^x$ and $\nu x.\mathbb{A}$ are semi-disjunctive, modulo provable equivalence.

7.2. A key lemma

The following lemma is central in that it links together our two main automata-theoretic tools, the satisfiability game and the consequence game:

Proposition 7.7. *Let $\mathbb{A}$ and $\mathbb{D}$ be respectively a semi-disjunctive and an arbitrary Λ-automaton, and assume that $\mathbb{A} \vDash_{\mathsf{G}} \mathbb{D}$. Then the automaton $\mathbb{A} \wedge \neg\mathbb{D}$ has a thin refutation.*

Before we prove this proposition, we formulate an auxiliary lemma. Recall that the transition map of the automaton $\neg\mathbb{D}$ is defined by taking boolean duals of the formulas assigned by the transition map of $\mathbb{D}$, and the priority map is defined by simply raising all priorities by 1. We shall need the following fact on boolean duals, which is a straightforward consequence of the definitions.

Proposition 7.8. *Let (S, σ) be a one-step $\mathsf{T}_{\mathtt{X}}$-frame, let α be a one-step formula in $\mathtt{1ML}^{+}_{\Lambda}(\mathtt{X}, A)$ and let $m, m' : S \to \mathsf{P}(A)$ be two markings such that $S, \sigma, m \Vdash^1 \alpha$ and $S, \sigma, m' \Vdash^1 \alpha^{\partial}$. Then for some $a \in A$ and some $s \in S$ we have $a \in m(s) \cap m'(s)$.*

Proof of Proposition 7.7. To fix notation, let $\mathbb{A} = (A, \Theta_{\mathbb{A}}, \Omega_{\mathbb{A}}, a_I)$, $\mathbb{D} = (D, \Theta_{\mathbb{D}}, \Omega_{\mathbb{D}}, d_I)$ and let $\mathbb{B}$ denote the automaton $\mathbb{A} \wedge \neg\mathbb{D}$. We write $\mathbb{B} = (B, \Theta_{\mathbb{B}}, \Omega_{\mathbb{B}}, b_I)$ and recall that $B = A \uplus D \uplus \{b_I\}$.

Assume that player II has a winning strategy χ in the consequence game $\mathcal{C}(\mathbb{A}, \mathbb{D})$ starting at position $\big(\{(a_I, a_I)\}, \{(d_I, d_I)\}\big)$. Our aim is to provide a thin refutation for the automaton $\mathbb{B}$, that is, a winning strategy for player $\forall$ in the thin satisfiability game for the automaton $\mathbb{A} \wedge \neg\mathbb{D}$. It will be convenient to make some simplifying assumptions on $\exists$'s strategy in this game. The proof of this claim follows from Proposition 6.7 and Proposition 6.10.

Claim 1. *Without loss of generality we may assume that in any play of $\mathcal{S}_{thin}(\mathbb{A} \wedge \neg\mathbb{D})$, $\exists$ only picks moves $(\mathcal{Q}, \Gamma)$ such that each $R \in \mathcal{Q}$ is A-thin and, after two rounds of the play, satisfies $R = \mathsf{Res}_A R \cup \mathsf{Res}_D R$.*

We will now define a strategy σ for $\forall$ in $\mathcal{S}(\mathbb{B})$, inductively making sure that the following two conditions are maintained, for any σ-guided partial play $\Sigma = R_0 \dots R_n$:

(†) R_n is thin, and for $n \geq 1$ satisfies $|\mathsf{Ran}(R_n) \cap D| = 1$;
(‡) There is a χ-guided shadow $\mathcal{C}(\mathbb{A}, \mathbb{D})$-play of the form $(S_0, S'_0)(S_1, S'_1)\dots(S_n, S'_n)$, where
  (a) $S_0 = \{(a_I, a_I)\}$ and $S'_0 = \{(d_I, d_I)\}$;
  (b) $S_1 = \{(a_I, a) \in A \times A \mid (b_I, a) \in R_1\}$ and $\{(d_I, d) \in D \times D \mid (b_I, d) \in R_1\} \subseteq S'_1$;
  (c) for each $i > 1$ we have $R_i \cap (A \times A) = S_i$ and $R_i \cap (D \times D)$ is a singleton $\{(d, d')\}$ with $d \in \mathsf{Ran}(R_{i-1}) \cap D$ and $(d, d') \in S'_i$.

For $n = 0$ by definition we have $R_0 = \{b_I, b_I\}$, $S_0 = \{(a_I, a_I)\}$ and $S'_0 = \{(d_I, d_I)\}$, so that the conditions (†) and (‡) hold. We leave it for the reader to verify that the case where $n = 1$ can be seen as a notational variant of the general case, and focus on showing how $\forall$ can extend the play $R_0 \dots R_n$ to $R_0 \dots R_n R_{n+1}$ and maintain the above two conditions in the case that $n > 1$.

Suppose that the inductive hypothesis has been maintained for the partial play Σ consisting of the positions $R_0 R_1 \dots R_n$ where $n > 1$, with shadow play $(S_0, S'_0)(S_1, S'_1) \dots (S_n, S'_n)$, and let $(\mathcal{Q}, \Gamma) \in \mathsf{P}B^\sharp \times \mathsf{T_x} B^\sharp$ be an arbitrary move chosen by $\exists$ at Σ. By Claim 1 we may assume that each member of $\mathcal{Q}$ is thin relative to $\mathbb{A}$. By legitimacy of $(\mathcal{Q}, \Gamma)$ as a move for $\exists$ we have

$$B^\sharp, \Gamma, n_b^B \Vdash^1 \Theta_B(b) \text{ for all } b \in \mathsf{Ran} R_n, \tag{21}$$

where we recall that $n_b^B : B^\sharp \to \mathsf{P}B$ denotes the natural B-marking on $B^\sharp$, given by $n_b^B : R \mapsto R[b]$. Then by Corollary 3.9 and Proposition 3.10 we obtain that

$$\mathcal{Q}, \Gamma, n_d^D \Vdash^1 \Theta_D(d)^\partial, \tag{22}$$

where d is the unique element of $\mathsf{Ran}(R_n) \cap D$, and

$$B^\sharp, \Gamma, n_a^A \Vdash^1 \Theta_A(a) \text{ for all } a \in A \cap \mathsf{Ran} R_n. \tag{23}$$

Recall that $\mathsf{Res}_A : B^\sharp \to A^\sharp$ is the map sending a relation R to its restriction $R \cap (A \times A)$. By Proposition 3.8 we may infer from (23) that

$$A^\sharp, (\mathsf{T_x}\mathsf{Res}_A)\Gamma, n_a^A \Vdash^1 \Theta_A(a), \text{ for all } a \in A \cap \mathsf{Ran} R_n, \tag{24}$$

while it follows from Proposition A.10 that $(\mathsf{T_x}\mathsf{Res}_A)\Gamma \in \mathsf{T_x}(\mathsf{Res}_A[\mathcal{Q}])$. But then by Proposition 3.10 we have that

$$\mathsf{Res}_A[\mathcal{Q}], (\mathsf{T_x}\mathsf{Res}_A)\Gamma, n_a^A \Vdash^1 \Theta_A(a), \text{ for all } a \in A \cap \mathsf{Ran} R_n, \tag{25}$$

and hence the pair $((\mathsf{T_x}\mathsf{Res}_A)\Gamma, \mathsf{Res}_A[\mathcal{Q}])$ is admissible as a move for player I in the consequence game at position (S_n, S'_n). Thus Player II's winning strategy χ in $\mathcal{C}(\mathbb{A}, \mathbb{D})$ provides a pair $(\mathcal{Q}', \Gamma') \in \mathsf{P}D^\sharp \times \mathsf{T_x} D^\sharp$ such that $\Gamma' \in \mathsf{T_x}\mathcal{Q}'$ and

$$\mathcal{Q}', \Gamma', n_d^D \Vdash^1 \Theta_D(d), \tag{26}$$

followed by a relation $\mathcal{Z} \subseteq \mathsf{Res}_A[\mathcal{Q}] \times \mathcal{Q}'$ such that $\mathsf{Res}_A[\mathcal{Q}], \Gamma \leftrightarroweq^1_{\Lambda,f} \mathcal{Q}', \Gamma'$.

We shall prove the following claim:

Claim 2. *There are $S \in \mathcal{Q}$, $S' \in \mathcal{Q}'$, and $c \in D$ with $(\mathsf{Res}_A S, S') \in \mathcal{Z}$ and $(d, c) \in S' \cap \mathsf{Res}_D S$.*

Proof of Claim. It follows from Proposition 3.13 that the composition $\mathcal{Z}'$ of (the graph of) the map Res_A and $\mathcal{Z}$ is a full one-step Λ-bisimulation $\mathcal{Z}' : \mathcal{Q}, \Gamma \leftrightarroweq^1_{\Lambda,f} \mathcal{Q}', \Gamma'$. Hence, if we define a marking $m : \mathcal{Q} \to \mathsf{P}(D)$ by setting

$$m(S) := \bigcup \{ S'[d] \mid (\mathsf{Res}_A S, S') \in \mathcal{Z} \},$$

then we may apply Proposition 3.14 to (26) and obtain

$$\mathcal{Q}, \Gamma, m \Vdash^1 \Theta_D(d). \tag{27}$$

But then by Proposition 7.8, it follows from (22) and (27) that there is some $c \in D$ and some $S \in \mathcal{Q}$ such that $c \in n_d^D(S) \cap m(S)$. Unraveling the definitions of n_d^D and m we find that, respectively, $(d, c) \in \mathsf{Res}_D S$ and $(d, c) \in S'$ for some S' with $(Res_A S, S') \in \mathcal{Z}$, as required. ◀

With this claim in place, we define the next move for ∀ prescribed by the strategy σ to be the relation $R_{n+1} := \mathsf{Res}_A S \cup \{(d, c)\}$, where $S \in \mathcal{Q}$ and $c \in D$ are as described in the claim, so that $(d, c) \in F(\mathsf{Res}_A S) \cap \mathsf{Res}_D S$. Note that this is a legitimate move in response to $(\mathcal{Q}, \Gamma)$ since $R_{n+1} \subseteq S \in \mathcal{Q}$. The shadow play is then extended by the pair $(S_{n+1}, S'_{n+1}) := (\mathsf{Res}_A S, F(\mathsf{Res}_A S))$ so that condition (‡c) of the induction hypothesis holds as an immediate consequence of the claim. For condition (†), it is obvious that $|\mathsf{Ran}(R_{n+1}) \cap D| = 1$; thinness of the relation R_{n+1} follows from the assumption that $S \in \mathcal{Q}$ was thin relative to $\mathbb{A}$.

To show that the thus defined strategy σ is winning for ∀, first observe that he never gets stuck, so that we may focus on infinite plays. It suffices to prove that every infinite σ-guided play contains a bad trace, so consider an arbitrary such play $\Sigma = (R_i)_{i \geq 0}$.

Clearly we may assume that all initial parts of Σ, corresponding to the partial plays $(R_i)_{0 \leq i \leq n}$, satisfy the conditions (†) and (‡). From this it follows that Σ itself has an infinite χ-guided shadow play $(S_i, S'_i)_{i \geq 0}$ satisfying the condition (‡a–c). In addition, it follows from (†) that Σ will contain a *unique* trace in D, which by (‡) will also be a trace on the right side of the shadow play in the consequence game. That is, the play $R_0 R_1 R_n \ldots$ contains a unique trace of the form $b_I d_1 d_2 d_3 \ldots$ with each d_i in D, and this is a trace through the stream $S'_0 S'_1 S'_2 \ldots$ as well. If this trace is bad, then we are done. If not, then given the priorities assigned to states in $\neg\mathbb{D}$ it must be bad as a trace in $\mathbb{D}$ since parities are swapped in $\neg\mathbb{D}$. Hence there must be a bad trace $b_I a_1 a_2 a_3 \ldots$ on the left side $S_0 S_1 S_2 \ldots$ of the shadow play in the consequence game, since this shadow play was guided by the winning strategy χ of Player II. But then this trace $b_I a_1 a_2 a_3 \ldots$ is also a bad trace in the play $R_0 R_1 R_2 \ldots$ of the satisfiability game. Summarizing, we see that either the unique trace through D in Σ is bad or there is some bad trace through A in Σ. In either case, Σ is a loss for ∃ as required. □

8. A strong simulation theorem for Λ-automata

The goal of this section is to prove a strengthened simulation theorem for coalgebra automata: we will provide a construction $\mathsf{sim}(\cdot)$ transforming an arbitrary Λ-automaton $\mathbb{A}$ into a disjunctive automaton $\mathsf{sim}(\mathbb{A})$ that is not only semantically equivalent to $\mathbb{A}$, but in fact game-equivalent to $\mathbb{A}$ in the strong sense as stated in Theorem 8.2 below. The definition of $\mathsf{sim}(\mathbb{A})$ essentially uses the disjunctive basis of the signature.

The construction of $\mathsf{sim}(\mathbb{A})$ takes place in two steps, a ‘pre-simulation’ step that produces a disjunctive automaton $\mathbb{A}^\sharp$ with a non-standard acceptance condition, and a second ‘synchronization’ step that turns this acceptance condition into a parity condition. 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.

The construction of the pre-simulation of an automaton $\mathbb{A}$ is very closely related to the satisfiability game for Λ-automata; in particular, states of the pre-simulation of $\mathbb{A}$ are the same as the basic positions of $\mathcal{S}(\mathbb{A})$, namely binary relations in $A^\sharp$, and the initial state R_I is $\{(a_I, a_I)\}$. For the definition of the transition map $\Theta^\sharp$ of the pre-simulation automaton, we remind the reader of Remark 6.6, where we showed how to think of the admissibility criterion of ∃'s moves in the satisfiability game in terms of the satisfaction of a single formula:

$$\Gamma \in \bigcap_{a \in \mathsf{Ran} R} [\![\Theta(a)]\!]^1_a \text{ iff } A^\sharp, \Gamma, \mathsf{id}_{A^\sharp} \Vdash^1 \bigwedge_{a \in \mathsf{Ran} R} \alpha[\tau_a].$$

The acceptance condition $NBT_{\mathbb{A}} \subseteq (A^\sharp)^\omega$ consists of the streams over $A^\sharp$ that do not contain any bad traces. Finally, the simulation $\mathsf{sim}(\mathbb{A})$ is produced by forming a certain kind of product of the pre-simulation of $\mathbb{A}$

with a deterministic stream automaton that recognizes the stream language $NBT_{\mathbb{A}}$; we refer to [12] for the details.

Definition 8.1. Assume that $\mathtt{D}$ is expressively complete for Λ, and let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a Λ-automaton. Define $\Theta^\star : A \to \mathtt{1ML}^+_\Lambda(\mathtt{X}, A \times A)$ by putting, for each $a \in A$,

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

where $\tau_a : A \to \mathtt{Latt}(A \times A)$ is the *tagging* substitution given in Remark 6.6 by

$$\tau_a : b \mapsto (a, b).$$

Now consider a binary relation $R \in A^\sharp$; as an easy consequence of Proposition 3.21 we may pick formulas $\pi_R \in \mathtt{Bool}(\mathtt{X})$ and $\delta_R \in \mathtt{D}(\mathsf{P}(A \times A)) = \mathtt{D}(A^\sharp)$ such that

$$\pi_R \wedge \delta_R[\chi_{A\times A}] \equiv_1 \bigwedge_{a \in \mathsf{Ran} R} \Theta^\star(a).$$

Then, using these formulas for the definition of the following map $\Theta^\sharp : A^\sharp \to \mathtt{D}(\mathtt{X}, A^\sharp)$:

$$\Theta^\sharp(R) := \pi_R \wedge \delta_R,$$

we obtain the *pre-simulation* of $\mathbb{A}$ as the automaton $\mathsf{pre}(\mathbb{A}) = (A^\sharp, \Theta^\sharp, NBT_{\mathbb{A}}, R_I)$, where $R_I := \{(a_I, a_I)\}$.

Since the acceptance condition $NBT_{\mathbb{A}}$ is an ω-regular language with alphabet $A^\sharp$ as we noted in Section 6, we may pick some deterministic parity automaton $\mathbb{Z} = (Z, \delta, \Omega', z_I)$ that recognizes $NBT_{\mathbb{A}}$. Finally we define $\mathsf{sim}(\mathbb{A})$ to be the structure $(D, \Theta'', \Omega'', d_I)$ where:

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

We also define a “forgetful” map $G_{\mathbb{A}} : D \to A^\sharp$ by mapping (R, z) to R. $\lhd$

Theorem 8.2. *The map* $\mathsf{sim}(\cdot)$ *assigns to each modal automaton* $\mathbb{A}$ *a disjunctive modal automaton* $\mathsf{sim}(\mathbb{A})$ *such that*

(1) $\mathbb{A} \models_{\mathsf{G}} \mathsf{sim}(\mathbb{A})$ *and* $\mathsf{sim}(\mathbb{A}) \models_{\mathsf{G}} \mathbb{A}$*;*

(2) $\mathbb{B}[\mathsf{sim}(\mathbb{A})/p] \models_{\mathsf{G}} \mathbb{B}[\mathbb{A}/p]$*, for any modal* $\mathtt{X}$*-automaton* $\mathbb{B}$ *which is positive in* $p \in \mathtt{X}$*.*

Proof. To show that $\mathbb{A} \models_{\mathsf{G}} \mathsf{sim}(\mathbb{A})$ is easy: fix the stream automaton $\mathbb{Z}$ that recognizes $NBT_{\mathbb{A}}$. Then every finite word $R_0 \ldots R_k$ over $A^\sharp$ determines an associated state of $\mathbb{Z}$ by simply running $\mathbb{Z}$ on the word $R_0 \ldots R_k$; so for R_0 the associated state is z_I, for $R_0 R_1$ the associated state is $\zeta(R_0, z_I)$ etc. Since every k-length partial play Σ of the consequence game $\mathcal{C}(\mathbb{A}, \mathsf{sim}(\mathbb{A}))$ determines a word $R_0 \ldots R_k$ over $A^\sharp$ in the obvious way, we can associate a state z_Σ of $\mathbb{Z}$ with each such partial play. If Player I continues the play Σ consisting of basic positions $(R_0, R'_0) \ldots (R_k, R'_k)$ by choosing the move $(\mathcal{R}, \Gamma) \in \mathsf{P}A^\sharp \times \mathsf{T}_{\mathtt{X}} A^\sharp$, then we let Player II respond with the function $F : \mathcal{R} \to (A^\sharp \times Z)^\sharp$ that is defined by mapping $R \in \mathcal{R}$ to the singleton $\{((R_k, z_\Sigma), (R, \zeta(R_k, z_\Sigma))\}$. It can be checked that this defines a functional winning strategy for Player II, and we leave the details to the reader.

The direction $\mathsf{sim}(\mathbb{A}) \models_{\mathsf{G}} \mathbb{A}$ of clause (1), which can be seen as a simple special case of clause (2), will follow from the Propositions 8.5 and 8.6, as will clause (2) itself. $\square$

The difficult part of Theorem 8.2 is to prove clause (2), and this will be the focus of the rest of this section. It will be convenient to state more abstractly what the crucial properties are of the automaton $\mathsf{sim}(\mathbb{A})$ that we have associated with an arbitrary automaton $\mathbb{A}$. First we need an auxiliary definition, for which we recall the notion of a *true state* from Definition 5.5.

Definition 8.3. Given a disjunctive automaton $\mathbb{D} = (D, \Theta, \Omega, d_I)$, and a fixed true state $d_\top$ of $\mathbb{D}$, we let

$$\mathsf{sing}_\top(d) := \begin{cases} \varnothing & \text{if } d = d_\top \\ \{d\} & \text{if } d \neq d_\top \end{cases}$$

define the D-marking $\mathsf{sing}_\top : D \to \mathsf{P}D$. ◁

Definition 8.4. Let $\mathbb{A} = (A, \Theta_A, \Omega_A, a_I)$ and $\mathbb{D} = (D, \Theta_D, \Omega_D, d_I)$ be an arbitrary and a disjunctive Λ-automaton, respectively. We say that $\mathbb{D}$ is a *disjunctive companion* of $\mathbb{A}$ if $\mathbb{D}$ has a true state $d_\top$, and there is a map $G : D \to A^\sharp$ satisfying the following conditions:

(DC1) $G(d_I) = \{(a_I, a_I)\}$ and $G(d_\top) = \varnothing$.

(DC2) Let $\delta \in \mathsf{T_X} D$ be such that $D, \delta, \mathsf{sing}_\top \Vdash^1 \Theta_{\mathbb{D}}(d)$. Then $(\mathsf{T_X} G)\delta \in \bigcap_{a \in \mathsf{Ran}(Gd)} [\![\Theta_A(a)]\!]^1_a$.

(DC3) If $G(d_i)_{i\in\omega} \in (A^\sharp)^\omega$ contains a bad $\mathbb{A}$-trace, then $(d_i)_{i\in\omega}$ is itself a bad $\mathbb{D}$-trace. ◁

Proposition 8.5. *The simulation map* $\mathsf{sim}(\cdot)$ *assigns a disjunctive companion to any modal automaton.*

Proof. It is fairly straightforward to check that the projection map $G_{\mathbb{A}} : D \to A^\sharp$ specified in Definition 8.1, which simply forgets the states of the stream automaton used in the product construction, has all the properties required to witness that $\mathsf{sim}(\mathbb{A})$ is a disjunctive companion of $\mathbb{A}$. □

Proposition 8.6. *Let $\mathbb{A}$ and $\mathbb{B}$ be arbitrary modal automata, let $\mathbb{D}$ be a disjunctive companion of $\mathbb{A}$, and assume that $\mathbb{B}$ is positive in p. Then*

$$\mathbb{B}[\mathbb{D}/p] \vDash_{\mathsf{G}} \mathbb{B}[\mathbb{A}/p].$$

Before we set out to prove this proposition we prove an auxiliary result, which we need to make a simplifying assumption on the moves player I makes in the consequence game associated with $\mathbb{B}[\mathbb{D}/p]$ and $\mathbb{B}[\mathbb{A}/p]$.

Proposition 8.7. *Let Θ_{BD} denote the transition map of the automaton $\mathbb{B}[\mathbb{D}/p]$, where $\mathbb{B}$ is an arbitrary Λ-automaton (positive in p), and $\mathbb{D}$ is a disjunctive Λ-automaton. Fix some $R \in A^\sharp$, $\mathcal{Q} \subseteq (B \cup D)^\sharp$, some $C \subseteq \mathsf{Ran}R$ and $\Gamma \in \mathsf{T_X}\mathcal{Q}$ such that*

$$\Gamma \in \bigcap_{a \in \mathsf{Ran}R} [\![\Theta(a)]\!]^1_a.$$

Then there are $\mathcal{Q}' \subseteq A^\sharp$ and $\Gamma' \in \mathsf{T_X}\mathcal{Q}'$ such that $\Gamma' \in \bigcap_{a\in\mathsf{Ran}R}[\![\Theta(a)]\!]^1_a$, $\subseteq : (\mathcal{Q}', \Gamma') \leftrightarrows^1_{\Lambda,f} (\mathcal{Q}, \Gamma)$, and for each $Q \in \mathcal{Q}'$ and $c \in C$, we have $|Q[c] \cap D| \leq 1$.

Proof. As in the proof of Proposition 6.7(3), we will prove the statement for the special case where C is a singleton, $C = \{c\}$, while we show that $\mathcal{Q}'$ additionally satisfies

$$\{Q[a] \mid Q \in \mathcal{Q}'\} \subseteq \{Q[a] \mid Q \in \mathcal{Q}\} \tag{28}$$

for all $a \neq c$. The general case can then be obtained from the special one by a straightforward iteration, taking care of C's elements one by one. The role of (28) is to ensure that new iterations do not spoil the progress booked in earlier rounds.

In order to prove the proposition in this simplified case, we make a case distinction as to whether $c \in B$ or $c \in D$. The case where $c \in D$ is in fact a special case of Proposition 6.7(3), and easier than the case where $c \in B$, and so we focus on the latter one. Assuming that $c \in B$, observe that $\Theta_{BD}(c) = \Theta_B(c)[\Theta_D(d_I)/p]$. We now make a further case distinction.

If $\Gamma \notin [\![\Theta_D(d_I)]\!]^1_c$, then consider the map $F : \mathcal{Q} \to (B \cup D)^\sharp$ given by

$$F(Q) := \{(a, a') \in Q \mid a \neq c \text{ or } a' \in B\},$$

and set $\mathcal{Q}' := \mathsf{Ran}F$ and $\Gamma' := (\mathsf{T}_\mathtt{X} F)\Gamma$. Clearly F is a surjective one-step frame homomorphism, $F : (\mathcal{Q}, \Gamma) \to (\mathcal{Q}', \Gamma')$, satisfying $F(Q) \subseteq Q$, for all $Q \in \mathcal{Q}$. From this it is immediate by Proposition 3.13 that $\subseteq : (\mathcal{Q}', \Gamma') \leftrightarroweq^1_{\Lambda,f} (\mathcal{Q}, \Gamma)$. We now show that

$$\mathcal{Q}', \Gamma', n_a \Vdash^1 \Theta_{BD}(a), \text{ for all } a \in \mathsf{Ran}R. \tag{29}$$

This is trivial in case $a \neq c$, and so we focus on the case where $a = c$. In this case (29) follows by the following observation, which can be proved by a straightforward induction on the complexity of $\alpha \in \mathtt{1ML}^+_\Lambda(\mathtt{X}, B)$:

$$\mathcal{Q}, \Gamma, n_c \Vdash^1 \alpha[\Theta_D(d_I)/p] \text{ implies } \mathcal{Q}', \Gamma', n_c \Vdash^1 \alpha[\Theta_D(d_I)/p].$$

If $\Gamma \in [\![\Theta_D(d_I)]\!]^1_c$, then by disjunctivity of $\mathbb{D}$, the proposition follows by a variation of the proof of Proposition 6.7(3). □

We are now ready to prove Proposition 8.6; our proof generalizes the proof of the analogous proposition in [12]. We do provide all details here, since there are some subtle differences with the mentioned proof, due to the fact that here we work with a slightly modified definition of the consequence game.

Proof of Proposition 8.6. Starting with notation, let $\mathbb{A} = (A, \Theta_A, \Omega_A, a_I)$, $\mathbb{B} = (B, \Theta_B, \Omega_B, b_I)$ and $\mathbb{D} = (D, \Theta_D, \Omega_D, d_I)$, and let $G : D \to A^\sharp$ be the map witnessing that $\mathbb{D}$ is a disjunctive companion of $\mathbb{A}$.

Our goal is to provide player II with a winning strategy χ in the consequence game $\mathcal{C}$ between $\mathbb{B}[\mathbb{D}/p]$ and $\mathbb{B}[\mathbb{A}/p]$. It will be convenient to make some simplifying assumptions on player I's moves in the game.

Claim 1. *Without loss of generality we may assume that at any position (R, R'), player I always plays a move $(\mathcal{Q}, \Gamma)$ such that*
(Ass1) $\mathsf{Dom}(Q) \subseteq \mathsf{Ran}(R)$ *for all* $Q \in \mathcal{Q}$*;*
(Ass2) $Q \cap (D \times B) = \varnothing$*, for all* $Q \in \mathcal{Q}$*;*
(Ass3) $|Q[c] \cap D| \leq 1$ *for all* $c \in B \cup D$ *and all* $Q \in \mathcal{Q}$*.*

Proof of Claim. Immediate by the Propositions 6.17, 6.7, and 8.7. ◀

Consider an arbitrary partial play

$$\Sigma = (R_0, R'_0), \ldots, (R_k, R'_k),$$

with $R_0 = R'_0 = \{(b_I, b_I)\}$. It follows by Claim 1 that we may assume each element $c \in \mathsf{Ran}R_k$ to lie on some trace through $R_0, \ldots, R_k$, and that every trace through $R_0, \ldots, R_k$ is either a $\mathbb{B}$-trace, or else it consists of an initial, non-empty $\mathbb{B}$-trace, followed by a non-empty $\mathbb{D}$-trace. By the second and third assumption of the claim, traces are *D-functional*, in the sense that if $d \in D \cap \mathsf{Ran} R_n$ for some $n < k$, then d has at most one R_{n+1}-successor, that we will denote as d^+, if it exists. As a consequence, every trace τ on $R_0, \ldots, R_n$ ending at d has at most one continuation through $R_{n+1}, \ldots, R_k$.

A key role in our proof is played by a Σ-induced total order on $\mathsf{Ran}_D R_k$ that we will introduce now. Intuitively, we say, for $d, d' \in \mathsf{Ran}_D R_k$, that d is Σ-older than d' if d lies on a trace τ that entered D at an earlier stage than any trace arriving at d'.

For a formal definition of this ordering, we need to assume some arbitrary but fixed total order on D, given by an injective map $\mathsf{mb} : D \to \omega$; we call $\mathsf{mb}(d)$ the *birth minute* of d. The reason is that there may be "ties", i.e. situations where the longest D-trace leading to two different states in D are of the same length. Following the analogy: we can have cases where two states have the same "birth date", and we then refer to the birth minute to decide which is the oldest.

Given a state $d \in \mathsf{Ran}_D R_k$, by Claim 1(1) there is a trace τ through $R_0, \ldots, R_k$ such that $\tau(k) = d$. By Claim 1(2), all such traces start in B and at some moment j move to the $\mathbb{D}$-part of the automaton. We let $\mathsf{tb}_\Sigma(d)$ be the smallest pair of natural numbers (j, l) in the lexicographic order on $\omega \times \omega$ such that there is some $e \in \mathsf{Ran}_D R_j$ with $\mathsf{mb}(e) = l$ and such that the unique trace on $R_j \ldots R_k$ beginning with e ends with d (this trace is unique because of trace functionality in D). The pair $\mathsf{tb}_\Sigma(d) = (j, l)$ is called the *time of birth* of d relative to the play Σ; we simply write $\mathsf{tb}(d)$ if Σ is clear from context.

Note that tb_Σ is always an injective map. To see this, suppose that $\mathsf{tb}_\Sigma(d) = \mathsf{tb}_\Sigma(d') = (j, l)$. Then there are $e, e' \in \mathsf{Ran}_D R_j$ such that the unique trace on $R_j, \ldots, R_k$ beginning with e ends with d, and the unique trace beginning with e' ends with d', and such that $\mathsf{mb}(e) = \mathsf{mb}(e') = l$. By injectivity of mb, we get $e = e'$, and so we get $d = d'$ by uniqueness of traces in the $\mathbb{D}$-part of $R_0, \ldots, R_k$.

Finally, we define a strict total ordering on $\mathsf{Ran}_D R_k$ relative to Σ by saying that d is Σ-*older* than d' if $\mathsf{tb}(d)$ is smaller than $\mathsf{tb}(d')$ (in the lexicographic order). We leave it for the reader to verify that, for $d \in \mathsf{Ran} R_n$ with $n < k$, it holds that $\mathsf{tb}(d^+) \leq \mathsf{tb}(d)$.

We now turn to the definition of player II's winning strategy χ. By a simultaneous induction on the length of a partial χ-play

$$\Sigma = (R_0, R'_0), \ldots, (R_n, R'_n),$$

with $R_0 = R'_0 = \{(b_I, b_I)\}$, we will define maps

$$F_n : (B \cup D)^\sharp \to (B \cup A)^\sharp$$

and

$$g_n : \mathsf{Ran}_A R'_n \to \mathsf{Ran}_D R_n.$$

We let the F-maps determine player II's strategy in the following sense. Suppose that in the mentioned partial play Σ, player I legitimately picks an element $(\mathcal{R}, \Gamma)$. Then player II's response will be the map $F_{n+1}\restriction_{\mathcal{R}}$, that is, the map F_{n+1}, restricted to the set $\mathcal{R} \subseteq (B \cup D)^\sharp$, together with the one-step frame $(F_{n+1}[\mathcal{R}], \mathsf{T}_\mathsf{x}(F_{n+1}\restriction_{\mathcal{R}})\Gamma)$.
Inductively we will ensure that the following conditions are maintained:

(*) $F_n R_n = R'_n$,
(†0) $R'_n = \mathsf{Res}_B R'_n \cup (R'_n \cap (B \times A)) \cup \mathsf{Res}_A R'_n$,
(†1) $\mathsf{Res}_B R'_n = \mathsf{Res}_B R_n$,
(†2) $R'_n \cap (B \times A) \subseteq \bigcup_{d \in D} \{(b, a) \mid (b, d) \in R_n \cap (B \times D) \ \& \ (a_I, a) \in G(d)\}$,
(†3) $\mathsf{Res}_A R'_n \subseteq \bigcup \{G(d) \mid d \in \mathsf{Ran}_D R_n\}$,
(‡) $a \in \mathsf{Ran} G(g_n a)$, for all $a \in \mathsf{Ran}_A R_n$.

For some explanation and motivation of these conditions, observe that (*) indicates that Σ itself is indeed χ-guided. For condition (†), first observe that while by Claim 1, all $\mathbb{B}[\mathbb{D}/p]$-traces consist of an initial $\mathbb{B}$-part followed by an $\mathbb{D}$-tail, condition (†0) states that similarly, all $\mathbb{B}[\mathbb{A}/p]$-traces consist of an initial $\mathbb{B}$-part followed by an $\mathbb{A}$-tail. Condition (†1) then states that the $\mathbb{B}$-part on the left and right side of a $\mathcal{C}(\mathbb{B}[\mathbb{D}/p], \mathbb{B}[\mathbb{A}/p]$-play is the same, and condition (†3) states that every pair $(a, b) \in \mathsf{Res}_A \mathsf{Ran} R'_n$ is 'covered' or 'implied' by some $d \in \mathsf{Ran}_D R_n$. Finally, (‡) states that, for every $a \in \mathsf{Ran} R'_n$, the map g_n picks a specific element $d \in \mathsf{Ran}_D R_n$ such that $a \in \mathsf{Ran}(Gd)$. As we will see in Claim 4 below, it will be this condition, together with the condition on the reflection of traces in Definition 8.4 and the actual definition of the maps g_n, that is pivotal in proving that player II wins all infinite plays.

Setting up the induction, observe that $R_0 = R'_0 = \{(b_I, b_I)\}$. Defining F_0 as the map $R \mapsto \mathsf{Res}_B R$ and g_0 as the empty map, we can easily check that (*), (†) and (‡) hold.

In the inductive case we will define the maps F_{n+1} and g_{n+1} for a partial play Σ as above. For the definition of $F_{n+1} : (B \cup D)^\sharp \to (B \cup A)^\sharp$, first observe that by (†0) we are only interested in relations $R \in (B \cup D)^\sharp$ that are of the form $R = \mathsf{Res}_B R \cup (R \cap (B \times D)) \cup \mathsf{Res}_D R$. We will define F_{n+1} by treating these three parts of R separately, using, respectively, the identity map on $B^\sharp$ and two auxiliary maps that we define now.

For the D-part of R, we define an auxiliary map $H_{n+1} : D \times D \to A^\sharp$:

$$H_{n+1}(d, d') := G(d') \cap (g_n^{-1}(d) \times A),$$

that is, $H_{n+1}(d, d')$ consists of those pairs $(a, a') \in G(d')$ for which $g_n(a) = d$. For the $B \times D$-part of R, we need a second auxiliary map $L : B \times D \to \mathsf{P}(B \times A)$, given by

$$L(b, d) := \{(b, a) \in B \times A \mid (a_I, a) \in G(d)\}.$$

Now we define $F_{n+1} : (B \cup D)^\sharp \to (B \cup A)^\sharp$ as follows:

$$\begin{aligned} F_{n+1}(R) := \ & \mathsf{Res}_B R \\ & \cup \bigcup \{L(b, d) \mid (b, d) \in R \cap (B \times D)\} \\ & \cup \bigcup \{H_{n+1}(d, d^+) \mid (d, d^+) \in \mathsf{Res}_D R\}. \end{aligned}$$

That is, we define $F_{n+1}(R)$ as the union of three disjoint parts: a $B \times B$-part, a $B \times A$-part and an $A \times A$-part.

For the definition of g_{n+1}, let (R_{n+1}, R'_{n+1}) be an arbitrary next basic position following the partial play Σ. Note that we may assume that R_{n+1} satisfies the assumptions formulated in Claim 1, and that we have $R'_{n+1} = F_{n+1}(R_{n+1})$ by the fact that player II's strategy is given by the map F_{n+1}. Given $a \in \mathsf{Ran}_A R'_{n+1}$, distinguish cases:

Case 1 If a has no R'_{n+1}-predecessor in A, then by definition of F_{n+1} and L, the set of states $d \in D$ for which there is a $b \in B$ with $(b, d) \in R_{n+1}$ and $(a_I, a) \in G(d)$ is non-empty. We define $g_{n+1}a$ to be the *oldest* element of this set, that is, in this case, the element with the earliest birth minute.

Case 2 If a does have an R'_{n+1}-predecessor in A, that is, the set $\{b \in A \mid (b, a) \in R'_{n+1}\}$ is non-empty, then we can define $g_{n+1}a$ to be the *oldest* element (with respect to the play $\Sigma \cdot (R_{n+1}, R'_{n+1})$) of the set $\{(g_n b)^+ \mid (b, a) \in R'_{n+1}\} \subseteq D$. Note that this set is indeed non-empty, by definition of F_{n+1}.

To gain some intuitions concerning this definition, observe that in the first case, we cannot define $g_{n+1}a$ inductively on the basis of the map g_n applied to an R'_{n+1}-predecessor of a: we have to start from scratch. This case only applies, however, in a situation where a does have an R'_{n+1}-successor $b \in B$ such that in R_{n+1},

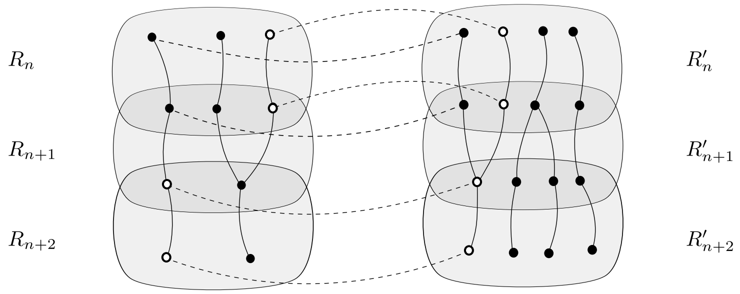

Fig. 1. A trace merge results in a trace jump.

this same b has a R_{n+1}-successor $d \in D$ such that $(a_I, a) \in Gd$. In this case we simply define $g_{n+1}a := d$, and if there are more such pairs (b, d), then for $g_{n+1}a$ we may pick any of these d's, for instance the one with the earliest birth minute.

We now turn to the second clause of the definition of g_{n+1} — here lies, in fact, the heart of the proof of Proposition 8.6. Consider a situation where a_0 and a_1, both in A, are the two R_{n+1}-predecessors of $a \in A$. Both $g_n a_0$ and $g_n a_1$ are states in D, and therefore their R_{n+1}-successors in D, if existing, are unique, and will be denoted by $(g_n a_0)^+$ and $(g_n a_1)^+$, respectively. We want to define $g_{n+1}a$ as either $(g_n a_0)^+$ or $(g_n a_1)^+$, but then we are facing a *choice* between these two states of D in case they are *distinct*. It is here that our play-dependent ordering of states in D comes in: we will define $g_{n+1}a$ as the *oldest* element of the two, relative to the (extended) play $\Sigma \cdot (R_{n+1}, R'_{n+1})$. Suppose (without loss of generality) it holds that $(g_n a_0)^+$ is older than $(g_n a_1)^+$, so that we put $g_{n+1}a := (g_n a_0)^+$. In this case we say that the trace through $g_n a_0$ is *continued*, while there is also a *trace jump* witnessed by the fact that $(a_1, a) \in R'_{n+1}$ but $(g_n a_1, g_{n+1}a) \notin R_{n+1}$ (see Fig. 1, where the dashed lines represent the g-maps, and the partial trace of white points on the right is not mapped to a partial trace on the left, due to a trace jump).

Claim 2. *By playing according to the strategy χ, player II indeed maintains the conditions (*), (†) and (‡).*

Proof of Claim.[6] Let Σ be a partial χ-play satisfying the conditions (*), (†) and (‡), and let $(R_{n+1}, R'_{n+1}) \in \mathsf{Gr}(F_{n+1})$ be any possible next position. It suffices to show that (R_{n+1}, R'_{n+1}) also satisfies (*), (†) and (‡).

The conditions (*), (†0), (†1) and (†2) are direct consequences of the definition of F_{n+1}, while (†3) is immediate by the fact that

$$(b, a) \in F_{n+1}R_{n+1} \iff (b, a) \in G((g_n b)^+) \tag{30}$$

for all $b, a \in A$. To prove (30), consider the following chain of equivalences, which hold for all $b, a \in A$:

$$\begin{aligned}
(b, a) \in F_{n+1}R_{n+1} &\iff (b, a) \in H_{n+1}(d, d^+), \text{ some } (d, d^+) \in \mathsf{Res}_D R_n && (\text{Def. } F_{n+1}) \\
&\iff (b, a) \in G(d^+), \text{ some } (d, d^+) \in \mathsf{Res}_D R_n \text{ with } d = g_n b && (\text{Def. } H_{n+1}) \\
&\iff (b, a) \in G((g_n b)^+). && (\text{obvious})
\end{aligned}$$

Finally, for condition (‡), let $a \in \mathsf{Ran}_A R'_{n+1}$ be arbitrary. If a has an R'_{n+1}-predecessor in A, then we are in case 2 of the definition of $g_{n+1}a$, where $g_{n+1}a$ is of the form $(g_n b)^+$ for some b with $(b, a) \in \mathsf{Res}_A R'_{n+1}$. But then $(b, a) \in G((g_n b)^+)$ by (30), so that indeed we find $a \in \mathsf{Ran} G(g_{n+1}a)$. If, on the other hand, a has no R_{n+1}-predecessor in A, then we are in case 1 of the definition of $g_{n+1}a$. In this case, $g_{n+1}a$ is an element of a set, each of whose elements d satisfies $a \in \mathsf{Ran} G(d)$; so we certainly have $a \in \mathsf{Ran} G(g_{n+1}a)$. ◄

[6] The proof of this Claim is verbatim the same as that of Claim 2 in the proof of Proposition 7.4 in [12].

Claim 3. *The moves for player II prescribed by the strategy χ are legitimate.*

Proof of Claim. Let Θ_{BD} and Θ_{BA} denote the transition maps of the automata $\mathbb{B}[\mathbb{D}/p]$ and $\mathbb{B}[\mathbb{A}/p]$, respectively. Consider a partial play Σ ending with the position (R_n, R'_n) and a subsequent move $(\mathcal{R}, \Gamma) \in \mathsf{P}((B \cup D)^\sharp) \times \mathsf{T_X}(B \cup D)^\sharp$ by player I such that

$$(B \cup D)^\sharp, \Gamma, n_e^{B\cup D} \Vdash^1 \Theta_{BD}(e), \tag{31}$$

for all $e \in \mathsf{Ran} R_n$. By Proposition 3.13, in order to prove the claim it suffices to show that, for an arbitrary element $c \in \mathsf{Ran} R'_n = \mathsf{Ran}(F_{n+1} R_n)$, we have

$$(B \cup A)^\sharp, (\mathsf{T_X} F_{n+1})\Gamma, n_c^{B\cup A} \Vdash^1 \Theta_{BA}(c). \tag{32}$$

But since $c \in B \cup A$ by definition of $\mathbb{B}[\mathbb{A}/p]$, one of the following two cases applies:

Case 1: $c \in A$. Then by (‡) we find $c \in \mathsf{Ran}(G(d))$, where $d := g_n c$ belongs to $\mathsf{Ran}_D R_n$. As an immediate consequence of (31) and the fact that $\Theta_{BD}(d) = \Theta_D(d)$, we find

$$(B \cup D)^\sharp, \Gamma, n_d^{B\cup D} \Vdash^1 \Theta_D(d), \tag{33}$$

from which it follows by naturality that

$$\mathcal{R}, \Gamma, n_d^{B\cup D}{\restriction_{\mathcal{R}}} \Vdash^1 \Theta_D(d). \tag{34}$$

Let the map $\mathsf{succ}_d : \mathcal{R} \to D$ be given by

$$\mathsf{succ}_d(Q) := \begin{cases} e & \text{if } Q[d] = \{e\}, \\ d_\top & \text{if } Q[d] = \varnothing. \end{cases}$$

Observe that this provides a well-defined (total) map by (Ass3) in Claim 1, and an easy calculation reveals that the diagram below commutes:

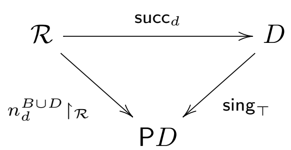

so that we may conclude that succ_d is a one-step model morphism:

$$\mathsf{succ}_d : (\mathcal{R}, \Gamma, n_d^{B\cup D}{\restriction_{\mathcal{R}}}) \to (\mathsf{succ}_d[\mathcal{R}], (\mathsf{T_X}\mathsf{succ}_d)\Gamma, \mathsf{sing}_\top).$$

From this, (34), the fact that $\Theta_D(d)$ is a one-step formula in D, and Corollary 3.9 we conclude that

$$D, (\mathsf{T_X}\mathsf{succ}_d)\Gamma, \mathsf{sing}_\top \Vdash^1 \Theta_D(d). \tag{35}$$

Now we may use the assumption that $(\mathbb{D}, d)$ is a disjunctive companion of $(\mathbb{A}, a)$, obtaining from clause (DC2) that

$$A^\sharp, (\mathsf{T_X} G)(\mathsf{T_X}\mathsf{succ}_d)\Gamma, n_c^A \Vdash^1 \Theta_A(c). \tag{36}$$

By functoriality of $\mathsf{T}_{\mathtt{X}}$ and the fact that $\Theta_A(c) = \Theta_{BA}(c)$, this is equivalent to

$$A^\sharp, (\mathsf{T}_{\mathtt{X}}(G \circ \mathsf{succ}_d))\Gamma, n_c^A \Vdash^1 \Theta_{BA}(c), \tag{37}$$

and so by Corollary 3.9 and Proposition 3.10 we obtain

$$(B \cup A)^\sharp, (\mathsf{T}_{\mathtt{X}}(G \circ \mathsf{succ}_d))\Gamma, n_c^{B\cup A} \Vdash^1 \Theta_{BA}(c). \tag{38}$$

From here on for conciseness we will write n_c for $n_c^{B\cup A}$. We now claim that, comparing the two A-markings $n_c \circ (G \circ \mathsf{succ}_d)$ and $n_c \circ F$, we have

$$\big(n_c \circ (G \circ \mathsf{succ}_d)\big)(Q) \subseteq \big(n_c \circ F_{n+1}\big)(Q) \tag{39}$$

for all $Q \in \mathcal{R}$. To see this, assume that $a \in \big(n_c \circ (G \circ \mathsf{succ}_d)\big)(Q)$, that is, $(c, a) \in G(\mathsf{succ}_d(Q))$. Observe that since $G(d_\top) = \varnothing$ by (DC1), by definition of the map succ_d it must be the case that $\mathsf{succ}_d(Q) = e$ for some unique $e = d_Q^+ \in D$ such that $Q[d] = \{d_Q^+\}$. Then (c, a) belongs to $H_{n+1}(d, d_Q^+)$ by definition of H_{n+1}, and to $F_{n+1}Q$ by definition of F_{n+1}. But from $(c, a) \in F_{n+1}(Q)$ we immediately obtain $a \in \big(n_c \circ F_{n+1}\big)(Q)$. This proves (39).

We use this observation in the following line of reasoning, where the key observation is that in fact both maps $G \circ \mathsf{succ}_d$ and F_{n+1} are one-step model morphisms.

$$\begin{aligned}
&(B\cup A)^\sharp, (\mathsf{T}_{\mathtt{X}}(G \circ \mathsf{succ}_d))\Gamma, n_c \Vdash^1 \Theta_{BA}(c) \\
&\quad\Leftrightarrow (B \cup A)^\sharp, \Gamma, n_c \circ (G \circ \mathsf{succ}_d) \Vdash^1 \Theta_{BA}(c) && \text{(Proposition 3.8)}\\
&\quad\Leftrightarrow \mathcal{R}, \Gamma, \big(n_c \circ (G \circ \mathsf{succ}_d)\big)\restriction_{\mathcal{R}} \Vdash^1 \Theta_{BA}(c) && \text{(Corollary 3.9)}\\
&\quad\Rightarrow \mathcal{R}, \Gamma, \big(n_c \circ F_{n+1}\big)\restriction_{\mathcal{R}} \Vdash^1 \Theta_{BA}(c) && \text{((39), Proposition 3.11)}\\
&\quad\Leftrightarrow (B \cup A)^\sharp, \Gamma, n_c \circ F_{n+1} \Vdash^1 \Theta_{BA}(c) && \text{(Corollary 3.9)}\\
&\quad\Leftrightarrow (B \cup A)^\sharp, (\mathsf{T}_{\mathtt{X}}F_{n+1})\Gamma, n_c \Vdash^1 \Theta_{BA}(c). && \text{(Proposition 3.8)}
\end{aligned}$$

This proves (32), as required.

Case 2 $c \in B$. Note that in this case we have $\Theta_{BA}(c) = \Theta_B(c)[\Theta_A(a_I)/p]$ and $\Theta_{BD}(c) = \Theta_B(c)[\Theta_D(d_I)/p]$. Thus by assumption we know that $(B \cup D)^\sharp, \Gamma, n_c^{B\cup D} \Vdash^1 \Theta_B(c)[\Theta_D(d_I)/p]$, while we need to establish that $(B \cup A)^\sharp, (\mathsf{T}_{\mathtt{X}}F_{n+1})\Gamma, n_c^{B\cup A} \Vdash^1 \Theta_B(c)[\Theta_A(a_I)/p]$. To achieve this it clearly suffices to show that

$$(B \cup D)^\sharp, \Gamma, n_c^{B\cup D} \Vdash^1 \alpha[\Theta_D(d_I)/p] \text{ implies } (B \cup A)^\sharp, (\mathsf{T}_{\mathtt{X}}F_{n+1})\Gamma, n_c^{B\cup A} \Vdash^1 \alpha[\Theta_A(a_I)/p] \tag{40}$$

for all $\alpha \in \mathtt{1ML}_\Lambda^+(\mathtt{X}, B)$. We will prove (40) by induction on the one-step formula α, taken as a lattice term over the set $\{p\} \cup \mathtt{1ML}_\Lambda^+(\mathtt{X} \setminus \{p\}, B))$. This perspective allows us to distinguish the following two cases in the induction base.

Base Case a: $\alpha = p$. Here we find $\alpha[\Theta_D(d_I)/p] = \Theta_D(d_I)$ and $\alpha[\Theta_A(a_I)/p] = \Theta_A(a_I)$. In other words, in order to prove (40) we assume that

$$(B \cup D)^\sharp, \Gamma, n_c^{B\cup D} \Vdash^1 \Theta_D(d_I), \tag{41}$$

and we need to show that

$$(B \cup A)^\sharp, (\mathsf{T}_{\mathtt{X}}F_{n+1})\Gamma, n_c^{B\cup A} \Vdash^1 \Theta_A(a_I). \tag{42}$$

Our line of reasoning here will be close to that in Case 1, and for this reason we are a bit more sketchy. By (Ass3) we may define a map $\mathsf{succ}_c : \mathcal{R} \to D$ by setting $\mathsf{succ}_c(Q)$ to be the unique Q-successor of c if $Q[c]$ is nonempty, and the true state $d_\top$ otherwise. As in Case 1 this map is a one-step morphism of models:

$$\mathsf{succ}_c : (\mathcal{R}, \Gamma, n_c^{B\cup D}{\restriction_{\mathcal{R}}}) \to (D, (\mathsf{T}_{\mathtt{X}}\mathsf{succ}_c)\Gamma, \mathsf{sing}_\top). \tag{43}$$

We also claim that our definition of the map F_{n+1} has been tailored towards the following inclusion:

$$\big(n_{a_I}^{B\cup A} \circ (G \circ \mathsf{succ}_c)\big)(Q) \subseteq \big(n_c^{B\cup A} \circ F_{n+1}\big)(Q) \tag{44}$$

for all $Q \in \mathcal{R}$. For a proof of (44), assume that $a \in \big(n_{a_I}^{B\cup A} \circ (G \circ \mathsf{succ}_c)\big)(Q)$ for some $Q \in \mathsf{Base}(\Gamma)$. In other words, we have $(a_I, a) \in G(\mathsf{succ}_c(Q))$, and so by definition of succ_c there is a unique $d \neq d_\top \in D$ such that $(c, d) \in Q$. But then we obtain $(a_I, a) \in L(b, d)$ by definition of the map L, and since $(c, d) \in Q \cap (B \times D)$ this gives $(c, a) \in F_{n+1}Q$ by definition of F_{n+1}. But from $(c, a) \in F_{n+1}Q$ we directly see that $a \in n_c^{B\cup A}(F_{n+1}Q)$, as required. This proves (44).

We can now show how to prove (42) from (41):

$$\begin{aligned}
(B\cup D)^\sharp, \Gamma, n_c^{B\cup D} \Vdash^1 \Theta_D(d_I) &\Leftrightarrow \mathcal{R}, \Gamma, n_c^{B\cup D}{\restriction_{\mathcal{R}}} \Vdash^1 \Theta_D(d_I) && \text{(Corollary 3.9)}\\
&\Leftrightarrow D, (\mathsf{T}_{\mathtt{X}}\mathsf{succ}_c)\Gamma, \mathsf{sing}_\top \Vdash^1 \Theta_D(d_I) && \text{(Proposition 3.8, (43))}\\
&\Rightarrow A^\sharp, (\mathsf{T}_{\mathtt{X}}G)((\mathsf{T}_{\mathtt{X}}\mathsf{succ}_c)\Gamma), n_{a_I}^{A} \Vdash^1 \Theta_A(a_I) && \text{(DC1,DC2)}\\
&\Leftrightarrow A^\sharp, (\mathsf{T}_{\mathtt{X}}(G \circ \mathsf{succ}_c))\Gamma, n_{a_I}^{A} \Vdash^1 \Theta_A(a_I) && \text{(functoriality)}\\
&\Leftrightarrow (B\cup A)^\sharp, (\mathsf{T}_{\mathtt{X}}(G \circ \mathsf{succ}_c))\Gamma, n_{a_I}^{B\cup A} \Vdash^1 \Theta_A(a_I) && \text{(as in Case 1)}\\
&\Leftrightarrow (B\cup A)^\sharp, (\mathsf{T}_{\mathtt{X}}F_{n+1})\Gamma, n_c^{B\cup A} \Vdash^1 \Theta_A(a_I). && \text{(as in Case 1, by (44))}
\end{aligned}$$

Base Case b: $\alpha \in \mathtt{1ML}_\Lambda^+(\mathtt{X} \setminus \{p\}, B)$, that is, α is a p-free one-step formula over B. In this case the proof of (40) is straightforward: clearly the substitutions in (40) have no effect, so what we have to prove is that

$$(B\cup D)^\sharp, \Gamma, n_c^{B\cup D} \Vdash^1 \alpha \text{ implies } (B\cup A)^\sharp, (\mathsf{T}_{\mathtt{X}}F_{n+1})\Gamma, n_c^{B\cup A} \Vdash^1 \alpha. \tag{45}$$

But intuitively this is clear, since α only uses variables from B, and 'when restricted to B', the two models in (45) are the same.

Formally, our proof of (45) proceeds as follows:

$$\begin{aligned}
(B\cup D)^\sharp, \Gamma, n_c^{B\cup D} \Vdash^1 \alpha &\Leftrightarrow (B\cup D)^\sharp, \Gamma, n_c^{B} \circ \mathsf{Res}_B \Vdash^1 \alpha && \text{(Proposition 3.10)}\\
&\Leftrightarrow B^\sharp, (\mathsf{T}_{\mathtt{X}}\mathsf{Res}_B)\Gamma, n_c^{B} \Vdash^1 \alpha && \text{(Proposition 3.8)}\\
&\Leftrightarrow B^\sharp, (\mathsf{T}_{\mathtt{X}}(\mathsf{Res}_B \circ F_{n+1}))\Gamma, n_c^{B} \Vdash^1 \alpha && (\dagger 1)\\
&\Leftrightarrow B^\sharp, (\mathsf{T}_{\mathtt{X}}\mathsf{Res}_B)((\mathsf{T}_{\mathtt{X}}F_{n+1})\Gamma), n_c^{B} \Vdash^1 \alpha && \text{(functoriality)}\\
&\Leftrightarrow (B\cup A)^\sharp, (\mathsf{T}_{\mathtt{X}}F_{n+1})\Gamma, n_c^{B} \circ \mathsf{Res}_B \Vdash^1 \alpha && \text{(Proposition 3.8)}\\
&\Leftrightarrow (B\cup A)^\sharp, (\mathsf{T}_{\mathtt{X}}F_{n+1})\Gamma, n_c^{B\cup A} \Vdash^1 \alpha && \text{(Proposition 3.10)}
\end{aligned}$$

Inductive case: The inductive cases in the proof of (40), where α is of the form $\alpha_0 \vee \alpha_1$ or $\alpha_0 \wedge \alpha_1$, are trivial.

This finishes the proof of Claim 3. ◀

Claim 4. *Suppose Σ is an infinite χ-guided play with basic positions*

$$(R_0, R'_0)(R_1, R'_1)(R_2, R'_2)\dots$$

such that the stream $R'_0R'_1R'_2\dots$ contains a bad trace. Then there is a bad trace on $R_0R_1R_2\dots$ as well.

Proof of Claim.[7] Fix a χ-guided play $\Sigma = (R_i, R'_i)_{i\geq 0}$ and a bad trace τ on $(R'_i)_{i\geq 0}$, as above. We will show that there is a bad trace on the stream $(R_i)_{i\geq 0}$ as well.

There are two possibilities for τ. In case τ stays entirely in B, then by (†1), τ is also a trace on $R_0R_1R_2\dots$, and so we are done. Hence we may focus on the second case, where from some finite stage onwards, τ stays entirely in A. So suppose τ is an infinite trace of the form

$$\tau = b_0b_1\dots b_na_{n+1}a_{n+2}a_{n+3}\dots,$$

where the b_j are all in B, and the a_i are all in A. Our key claim is the following:

$$\text{there exists an index } k > n \text{ such that } g_{j+1}a_{j+1} = (g_ja_j)^+ \text{ for all } j \geq k. \tag{46}$$

In order to prove (46), recall that a *trace jump* occurs at the index $j > n$ if we have $g_{j+1}a_{j+1} \neq (g_ja_j)^+$. We want to show that there can only be finitely many j at which a trace jump occurs. If no trace jump occurs at j, then we have

$$\mathsf{tb}(g_ja_j) \geq \mathsf{tb}((g_ja_j)^+) = \mathsf{tb}(g_{j+1}a_{j+1}).$$

Hence, it suffices to prove that if a trace jump occurs at j then $\mathsf{tb}(g_{j+1}a_{j+1})$ is strictly smaller than $\mathsf{tb}(g_ja_j)$ in the lexicographic order. It then follows that the stream

$$\mathsf{tb}(g_ka_k), \mathsf{tb}(g_{k+1}a_{k+1}), \mathsf{tb}(g_{k+2}a_{k+2}), \dots$$

is a stream of pairs of natural numbers that never increases, and strictly decreases at each j at which a trace jump occurs. By well-foundedness of the lexicographic order on $\omega \times \omega$ this can therefore only happen finitely many times, as required.

So we are left with the task of proving that tb is strictly decreasing at each index j for which a trace jump occurs. To see that this is indeed so, suppose that $g_{j+1}a_{j+1} \neq (g_ja_j)^+$. Recall that we defined $g_{j+1}a_{j+1}$ to be the oldest element of the set

$$\{(g_jc)^+ \mid (c, a_{j+1}) \in R'_{j+1}\}.$$

But since $(a_j, a_{j+1}) \in R'_{j+1}$, it follows that $g_{j+1}a_{j+1}$ must be older than $(g_ja_j)^+$, with respect to the age relation induced by the play $(R_0, R'_0), \dots, (R_{j+1}, R'_{j+1})$, and so $\mathsf{tb}(g_{j+1}a_{j+1})$ must be strictly smaller than $\mathsf{tb}((g_ja_j)^+) \leq \mathsf{tb}(g_ja_j)$, as required. This completes the proof of (46).

Let us finally see how (46) entails Claim 4. Suppose there exists an index k as in (46), and consider $g_ka_k \in \mathsf{Ran}_D R_k$. Pick an arbitrary initial trace $b_0\dots b_nd_{n+1}\dots d_k$ of $R_0\dots R_k$ leading up to $g_ka_k = d_k$ (as mentioned already after Claim 1, the existence of such a trace follows from our assumptions on player I's strategy). Then the stream

$$b_0, \dots, d_{k-1}, g_ka_k, g_{k+1}a_{k+1}, g_{k+2}a_{k+2}, \dots$$

[7] The proof of this Claim is verbatim the same as that of Claim 4 in the proof of Proposition 7.4 in [12].

is a trace of $R_0R_1R_2\dots$ by the property of the index k described in (46). Furthermore, it follows that $a_ka_{k+1}a_{k+2}\dots$ is a trace of the stream

$$G(g_ka_k), G(g_{k+1}a_{k+1}), G(g_{k+2}a_{k+2}), \dots$$

To see why, consider the pair (a_j, a_{j+1}) where $j \geq k$. Then $(a_j, a_{j+1}) \in R'_{j+1} = F_k(R_{j+1})$, so there is some $(d, d') \in R_{j+1}$ with $(a_j, a_{j+1}) \in H_{j+1}(d, d')$. Hence $d = g_ja_j$ and $(a_j, a_{j+1}) \in G(d')$.

But $d' = d^+$ by functionality of traces on D (which follows from the third assumption in Claim 1), and so we find $d' = d^+ = (g_ja_j)^+ = g_{j+1}a_{j+1}$. From this we get $(a_j, a_{j+1}) \in G(g_{j+1}a_{j+1})$ as required. Note too that $a_ka_{k+1}a_{k+2}\dots$ has the same tail as τ, and hence it is a *bad* trace too. It now follows from the trace reflection clause of Definition 8.4 that $g_ka_k, g_{k+1}a_{k+1}, g_{k+2}a_{k+2}, \dots$ is itself a bad trace, and so we have found a bad trace on $R_0R_1R_2\dots$ as required. ◀

Finally, the proof of the Proposition is immediate by the last two claims: it follows from Claim 3 that player II never gets stuck, so that we need not worry about finite plays. But Claim 4 states that II wins all infinite plays of $\mathcal{C}(\mathbb{B}[\mathbb{D}/p], \mathbb{B}[\mathbb{A}/p])$ as well. □

9. A generic completeness theorem for coalgebraic μ-calculi

We now set out to prove our generic completeness result, Theorem 1.1. Throughout this section we will fix a set functor T, a monotone signature Λ for T, and a one-step sound and complete axiomatization $\mathbf{H}$.

After our preparatory work in the previous sections, we have almost all pieces in place; the one result that is missing is the analogue, in our setting, of Kozen's completeness result for the aconjunctive fragment of the (standard) modal μ-calculus [22]. Here, and in the remainder of this section, we will freely apply proof-theoretic terminology and notation to Λ-automata, see Remark 5.20.

Proposition 9.1 *(Kozen's Lemma). If the Λ-automaton $\mathbb{A}$ is consistent, then $\exists$ has a winning strategy in $\mathcal{S}_{thin}(\mathbb{A})$ starting at $\{(a_I, a_I)\}$.*

Proof. The proof of this Proposition is almost verbatim a copy of the proof of the analogous result, viz., Theorem 5, in [12]: the only difference is that here we need the Consistency Reduction Lemma, Proposition 4.6. □

As an immediate consequence of this result and Proposition 7.5, we find that for semi-disjunctive automata, consistency implies satisfiability (this is why we think of Proposition 9.1 as our analog of Kozen's partial completeness result). So, since disjunctive automata are semi-disjunctive, we have left to prove the following theorem, which is the main technical result of this section.

Theorem 9.2. *For every formula φ, there exists a semantically equivalent disjunctive automaton $\mathbb{D}$ such that $\varphi \vdash_{\mathbf{H}} \mathbb{D}$.*

As an auxiliary result, we first prove the following proposition.

Proposition 9.3. *Let $\mathbb{A}$ be any semi-disjunctive modal automaton. Then $\mathbb{A} \vdash_{\mathbf{H}} \mathsf{sim}(\mathbb{A})$.*

Proof. It is clear from Theorem 8.2 that there is a winning strategy for Player II in the consequence game $\mathcal{C}(\mathbb{A}, \mathsf{sim}(\mathbb{A}))$. Since $\mathbb{A}$ is semi-disjunctive it follows by Lemma 7.7 that $\forall$ has a winning strategy in the thin satisfiability game for $\mathbb{A} \wedge \neg\mathsf{sim}(\mathbb{A})$. But then by Kozen's Lemma (Proposition 9.1), the automaton $\mathbb{A} \wedge \neg\mathsf{sim}(\mathbb{A})$ is inconsistent. From this and the clauses 1 and 2 of Proposition 5.18, it is immediate that $\mathbb{A} \vdash_{\mathbf{H}} \mathsf{sim}(\mathbb{A})$. □

Proof of Theorem 9.2. Since any fixpoint formula is provably equivalent to a formula in negation normal form, without loss of generality we may prove the theorem for formulas in this shape, and proceed by an induction on the complexity of such formulas. That is, the base cases of the induction are the literals, and we need to consider induction steps for conjunctions, disjunctions, both modal operators and both fixpoint operators.

The base case for literals follows immediately since it is easy to see that the modal automaton $\mathbb{A}_\varphi$ corresponding to a literal φ is already disjunctive. Disjunctions are easy since the operation $\vee$ on automata preserves the property of being disjunctive. For conjunctions: given formulas φ, φ' we have semantically equivalent disjunctive automata $\mathbb{D}, \mathbb{D}'$ such that $\varphi \vdash_{\mathbf{H}} \mathbb{D}$ and $\varphi' \vdash_{\mathbf{H}} \mathbb{D}'$. By the first clause of Proposition 5.18 we get $\varphi \wedge \varphi' \vdash_{\mathbf{H}} \mathbb{D} \wedge \mathbb{D}'$. But by Proposition 7.6(4) the automaton $\mathbb{D} \wedge \mathbb{D}'$ is semi-disjunctive modulo provable equivalence, and we can apply Proposition 9.3 to obtain the desired conclusion. The cases for the modalities are easy since these operations on automata preserve the property of being disjunctive.

For the greatest fixpoint operator, consider the formula $\varphi = \nu x.\alpha(x)$, and assume inductively that there is a disjunctive automaton $\mathbb{A}$ for α such that $\alpha \equiv \mathbb{A}$ and $\alpha \vdash_{\mathbf{H}} \mathbb{A}$. It follows by Proposition 5.18(4) that $\varphi = \nu x.\alpha \vdash_{\mathbf{H}} \nu x.\mathbb{A}$, and since $\nu x.\mathbb{A}$ is semidisjunctive modulo provable equivalence by Proposition 7.6(6), by Proposition 9.3 we are done.

Finally, we cover the crucial case for $\varphi = \mu x.\alpha(x)$. By the induction hypothesis there is a semantically equivalent disjunctive automaton $\mathbb{A}$ for α such that $\alpha \vdash_{\mathbf{H}} \mathbb{A}$. Let $\mathbb{D} := \mathsf{sim}(\mu x.\mathbb{A})$. This automaton is clearly semantically equivalent to φ. We want to show that

$$\mu x.\mathbb{A} \vdash_{\mathbf{H}} \mathbb{D}, \tag{47}$$

from which the result follows since $\varphi = \mu x.\alpha \vdash_{\mathbf{H}} \mu x.\mathbb{A}$ by Proposition 5.18(4) and the induction hypothesis.

In order to prove (47) we will work with the automaton $\mathbb{A}^x$ (see Definition 5.10). First observe that

$$\mathbb{A}^x[\mathbb{D}/x] \vDash_{\mathsf{G}} \mathbb{A}^x[\mu x.\mathbb{A}/x],$$

by Theorem 8.2, and that

$$\mathbb{A}^x[\mu x.\mathbb{A}/x] \vDash_{\mathsf{G}} \mu x.\mathbb{A},$$

as a straightforward argument shows (see Proposition 5.19 in [12]). But since

$$\mu x.\mathbb{A} \vDash_{\mathsf{G}} \mathsf{sim}(\mu x.\mathbb{A}) = \mathbb{D}$$

by Theorem 8.2 again, we find by transitivity of the game consequence relation (Proposition 6.15) that

$$\mathbb{A}^x[\mathbb{D}/x] \vDash_{\mathsf{G}} \mathbb{D}.$$

By Proposition 7.6(5) the automaton $\mathbb{A}^x[\mathbb{D}/x]$ is semi-disjunctive modulo provable equivalence, and so by Proposition 7.7 the automaton $\mathbb{A}^x[\mathbb{D}/x] \wedge \neg\mathbb{D}$ has a thin refutation, whence by Kozen's Lemma (Proposition 9.1) and Proposition 5.18 this automaton is inconsistent. In other words, we have

$$\mathbb{A}^x[\mathbb{D}/x] \vdash_{\mathbf{H}} \mathbb{D}.$$

Then by Proposition 5.18(5) we obtain that

$$\mathtt{tr}(\mathbb{A}^x[\mathtt{tr}(\mathbb{D})/x]) \vdash_{\mathbf{H}} \mathtt{tr}(\mathbb{D}),$$

so that one application of the fixpoint rule yields

$$\mu x.\mathtt{tr}(\mathbb{A}^x) \vdash_{\mathbf{H}} \mathbb{D}.$$

By Proposition 5.18(6) this suffices to prove (47). □

Finally we see how Theorem 9.2 implies completeness.

Proof of Theorem 1.1. Given a consistent formula φ, by Theorem 9.2 there exists a semantically equivalent disjunctive automaton $\mathbb{D}$ such that $\varphi \vdash_{\mathbf{H}} \mathbb{D}$. Clearly then, $\mathbb{D}$ is consistent too, whence by Proposition 9.1, $\exists$ has a winning strategy in the thin satisfiability game for $\mathbb{D}$. But $\mathbb{D}$ is disjunctive and hence semi-disjunctive, and so by Proposition 7.5 $\exists$ also has a winning strategy in $\mathcal{S}(\mathbb{D})$. It then follows by the adequacy of the satisfiability game (Proposition 6.9) that $\mathbb{D}$ is satisfiable, and so φ, being semantically equivalent to $\mathbb{D}$, is satisfiable as well. □

10. Applications

As an immediate consequence of Theorem 1.1, we get a number of completeness results:

Theorem 10.1. *The proof system $\mu\mathbf{H}$ is sound and complete for validity over T-models, where:*

1. $\mathsf{T} = \mathsf{Id}$ *and* $\mathbf{H} = \mathbf{I}$,
2. $\mathsf{T} = \mathsf{Id}^k$ *and* $\mathbf{H} = \mathbf{I}^k$,
3. $\mathsf{T} = \mathsf{P}^L$ *and* $\mathbf{H} = \mathbf{K}^L$,
4. $\mathsf{T} = \mathsf{B}$ *and* $\mathbf{H} = \mathbf{B}$.

The third item on this list is Walukiewicz' completeness theorem for the modal μ-calculus. The first item is a completeness result for the linear-time μ-calculus, and thus places Kaivola's theorem [21] under a common roof with Walukiewicz' result. The second item, $\mathsf{T} = \mathsf{Id}^k$, extends this to a completeness result for μ-calculi on trees of a fixed branching degree. The fourth item is Theorem 1.3, our completeness result for the graded μ-calculus, by an extension of known axioms for graded modal logic with the fixpoint axiom and Kozen–Park induction rule. Note that the proof of this result requires an application of Theorem 3.23.

10.1. Completeness for the monotone μ-calculus

For our next application, we will prove Theorem 1.4, stating the completeness for the monotone μ-calculus of the axiomatization below. In this section we let $\Sigma = \{\Diamond, \Box\}$ denote the signature of monotone modal logic.

Definition 10.2. Let $\mathbf{M}$ be the axiomatization for Σ consisting of the *empty* set of axioms. ◁

Recall from Definition 4.1 that *every* one-step axiomatization contains the monotonicity and dual axioms, for all its operators. Consequently, the proof system $\mu\mathbf{M}$ induced by $\mathbf{M}$ basically consists of the axioms (Du) and (Mon) for the two modalities of the signature Σ. Just as for $\mu\mathbf{B}$, the completeness of $\mu\mathbf{M}$ appears to be a new result.

Theorem 1.1 does not apply directly in this case, since in fact one can show that the monotone neighborhood functor M does *not* admit a disjunctive basis (a proof of this can be found in [13]). However, the problem is really more apparent than substantial: using a trick from [10], we can easily "repair" the monotone neighborhood functor to obtain a *companion* functor $\underline{\mathsf{M}}$ of M that preserves weak pullbacks. The signature Σ can then be extended to a signature for $\underline{\Sigma}$ that does have a disjunctive basis, so that Theorem 1.1 applies to this extended signature. The completeness result for the monotone μ-calculus can then be obtained via a relatively straightforward conservative extension theorem.

Definition 10.3. We define the *supported companion functor* $\underline{\mathsf{M}}$ of M as the subfunctor of $\mathsf{P} \times \mathsf{M}$, given, on objects, by $\underline{\mathsf{M}}S := \{(S_0, \gamma) \in \mathsf{P}S \times \mathsf{M}S \mid S_0 \text{ supports } \gamma\}$. Here we say that a subset $S_0 \subseteq S$ *supports* an object $\gamma \in \mathsf{M}S$ whenever $T \in \gamma$ iff $T \cap S_0 \in \gamma$, for all $T \in \mathsf{P}S$. ◁

Note that an $\underline{\mathsf{M}}$-model can be taken as a structure $\mathbb{S} = (S, R, \sigma, V)$, where $R \subseteq S \times S$ and $U(\mathbb{S}) := (S, \sigma, V)$ is a neighborhood model, such that $R[s]$ supports $\sigma(s)$, for all $s \in S$. We will call the structure $U(\mathbb{S})$ the *underlying neighborhood model* of $\mathbb{S}$, and (S, R, V) its *supporting Kripke model*.

The point of introducing this auxiliary functor is explained by the following result:

Proposition 10.4. *The functor* $\underline{\mathsf{M}}$ *preserves weak pullbacks.*

Proof. We first establish the following claim, where L is the relation lifting defined by $(\gamma, \gamma') \in LR$ for $R \subseteq X \times Y$ and $\gamma \in \mathsf{M}X$, $\gamma' \in \mathsf{M}Y$, iff:

$$\forall Z \in \gamma \exists Z' \in \gamma' : Z' \subseteq R[Z] \;\&\; \forall Z' \in \gamma' \exists Z \in \gamma : Z \subseteq R^{\circ}[Z']$$

In the proof we will make use of some basic laws for this relation lifting, see [25] for more details.

Claim 1. *Let $R \subseteq X \times Y$ be any binary relation that is full on both X and Y, and let $\gamma_X \in \mathsf{M}X$ and $\gamma_Y \in \mathsf{M}Y$ be such that $(\gamma_X, \gamma_Y) \in LR$. Then there exists a $\gamma_R \in \mathsf{M}R$ such that $\mathsf{M}\pi_X(\gamma_R) = \gamma_X$ and $\mathsf{M}\pi_Y(\gamma_R) = \gamma_Y$, where $\pi_X : R \to X$ and $\pi_Y : R \to Y$ are the two projection maps.*

Proof of Claim. We set $Z \in \gamma_R$ iff either there exists $Z' \in \gamma_X$ such that $\pi_X^{-1}[Z'] \subseteq Z$ or there exists $Z' \in \gamma_Y$ such that $\pi_Y^{-1}[Z'] \subseteq Z$. We prove that $\mathsf{M}\pi_X(\gamma_R) = \gamma_X$, and leave out the completely analogous argument for γ_Y.

We need to show that $Z \in \gamma_X$ iff $\pi_X^{-1}[Z] \in \gamma_R$. The direction from left to right is clear, so suppose $\pi_X^{-1}[Z] \in \gamma_R$. We make a case division:

(Case 1:) $\pi_X^{-1}[Z'] \subseteq \pi_X^{-1}[Z]$ for some $Z' \in \gamma_X$. Since R was full on X the projection π_X is surjective onto X, which means that in fact $Z' \subseteq Z$ (since if $y \in Z' \setminus Z$ then there is some $y' \in R$ with $\pi_X(y') = y$, hence $y' \in \pi_X^{-1}[Z'] \setminus \pi_X^{-1}[Z]$). Since $Z' \in \gamma_X$ we now get $Z \in \gamma_X$ too by monotonicity.

(Case 2:) $\pi_Y^{-1}[Z'] \subseteq \pi_X^{-1}[Z]$ for some $Z' \in \gamma_Y$. Since $(\gamma_X, \gamma_Y) \in LR$, there is some $Z'' \in \gamma_X$ such that for all $z'' \in Z''$ there exists a $z' \in Z'$ such that $z'' R z'$. So let $z'' \in Z''$. Pick some $z' \in Z'$ with $z'' R z'$. That is, $(z'', z') \in R$. Furthermore, $\pi_Y(z'', z') \in Z'$, so $(z'', z') \in \pi_Y^{-1}[Z']$. But we assumed that $\pi_Y^{-1}[Z'] \subseteq \pi_X^{-1}[Z]$ so we get $(z'', z') \in \pi_X^{-1}[Z]$, which means that $\pi_X(z'', z') = z'' \in Z$. We have shown that $Z'' \subseteq Z$, and since $Z'' \in \gamma_X$ we get $Z \in \gamma_X$ by monotonicity as required. ◀

Now, let $f : X \to W$ and $g : Y \to W$ be a span in **Set**, let $(S_X, \gamma_X) \in \underline{\mathsf{M}}X$ and $(S_Y, \gamma_Y) \in \underline{\mathsf{M}}Y$ be such that $\underline{\mathsf{M}}f(S_X, \gamma_X) = \underline{\mathsf{M}}g(S_Y, \gamma_Y)$. Let the relation $R \subseteq X \times Y$ together with projection maps π_X, π_Y be the pullback of f, g as standardly constructed in **Set**, i.e. we take $R = \{(x, y) \in X \times Y \mid f(x) = g(y)\}$. Consider the inclusion maps $\iota_X : S_X \to X$ and $\iota_Y : S_Y \to Y$. Then we have

$$\underline{\mathsf{M}}(f \circ \iota_X)(\gamma_X|_{S_X}) = \underline{\mathsf{M}}(g \circ \iota_Y)(\gamma_Y|_{S_Y})$$

Hence we get:

$$\begin{aligned}(\gamma_X|_{S_X}, \gamma_Y|_{S_Y}) &\in \mathsf{M}(f \circ \iota_X); \underline{\mathsf{M}}(g \circ \iota_Y)^\dagger \\ &\subseteq L(f \circ \iota_X); L(g \circ \iota_Y)^\dagger \\ &= L(\iota_X; f); L(g^\dagger; \iota_Y^\dagger) \\ &\subseteq L(\iota_X; f; g^\dagger; \iota_Y^\dagger) \\ &= LR'\end{aligned}$$

where we have used the identity:

$$\iota_X; f; g^\dagger; \iota_Y^\dagger = R'$$

which follows from the assumption that R was the pullback of W, f, g and the definition of R'. Here, we have denoted by $\gamma|_S$ for $\gamma \in \mathsf{M}X$ and $S \subseteq X$ the unique $\gamma' \in \mathsf{M}S$ with $\mathsf{M}\iota_{S,X}(\gamma') = \gamma$, which is concretely given by $\gamma|_S = \{Z \cap S \mid Z \in \gamma\}$.

Let $R' = R \cap (S_X \times S_Y)$. Then R' is full on both S_X and S_Y since $f[S_X] = g[S_Y]$ and R is the pullback of f and g. So by the Claim, we find some $\gamma_{R'} \in \mathsf{M}R'$ such that $\mathsf{M}(\pi_X{\restriction_{R'}})(\gamma_{R'}) = \gamma_X$ and $\mathsf{M}(\pi_Y{\restriction_{R'}})(\gamma_{R'}) = \gamma_Y$. Now, let $\iota_{R'} : R' \to R$ be the inclusion map and set $\gamma_R = \mathsf{M}\iota_{R'}(\gamma_{R'})$. We have $(R', \gamma_R) \in \underline{\mathsf{M}}R$. We also get:

$$\begin{aligned}\mathsf{M}\pi_X(\gamma_R) &= \mathsf{M}\pi_X \circ \mathsf{M}\iota_{R'}(\gamma_{R'}) \\ &= \mathsf{M}(\pi_X \circ \iota_{R'})(\gamma_{R'}) \\ &= \mathsf{M}(\iota_X \circ (\pi_X{\restriction_{R'}}))(\gamma_{R'}) \\ &= \mathsf{M}\iota_X \circ \mathsf{M}(\pi_X{\restriction_{R'}}))(\gamma_{R'}) \\ &= \mathsf{M}\iota_X(\gamma_X|_{S_X}) \\ &= \gamma_X\end{aligned}$$

where we have used the obvious identity $\pi_X \circ \iota_{R'} = \iota_X \circ (\pi_X{\restriction_{R'}})$. Furthermore we have $\mathsf{P}\pi_X(R') = S_X$ since R' was full on S_X. So we get $\underline{\mathsf{M}}\pi_X(R', \gamma_R) = (S_X, \gamma_X)$. Similarly, we get $\underline{\mathsf{M}}\pi_Y(R', \gamma_R) = (S_Y, \gamma_Y)$. It now follows, by the usual characterization of weak pullback squares in **Set** (see the appendix), that $\underline{\mathsf{M}}$ weakly preserves the pullback R, π_X, π_Y. □

Definition 10.5. The signature $\underline{\Sigma}$ is an expansion of the language Σ with two modalities $\underline{\Diamond}, \underline{\Box}$ that are interpreted as the standard diamond and box operators in the supporting Kripke models of a $\underline{\mathsf{M}}$-model. ◁

Definition 10.6. Let $\underline{\mathbf{M}}$ be the axiomatization for $\underline{\Sigma}$ consisting of the following axioms:
a) $\underline{\Box}(a \wedge b) \leftrightarrow (\underline{\Box}a \wedge \underline{\Box}b)$
b) $\underline{\Box}\top$
c) $(\underline{\Box}a \wedge \Box b) \to \Box(a \wedge b)$ ◁

We have previously proved that $\underline{\Sigma}$ is expressively complete for $\underline{\mathsf{M}}$ in [10]. Since $\underline{\mathsf{M}}$ restricts to finite sets, and one-step completeness of $\underline{\mathbf{M}}$ is straightforward, completeness of $\mu\underline{\mathbf{M}}$ is a direct consequence of Theorem 1.1, Proposition 10.4 and the fact that expressively complete signatures for weak-pullback preserving functors always admit a disjunctive basis.

Theorem 10.7. *The system $\mu\underline{\mathbf{M}}$ is sound and complete for validity over $\underline{\mathsf{M}}$-models.*

It turns out that completeness for $\mathbf{M}$ follows from this via a relatively easy argument. First, note that every pointed $\underline{\mathsf{M}}$-model $(\mathbb{S}, s)$ satisfies precisely the same $\mu\mathtt{ML}_\Sigma$-formulas as the underlying pointed M-model

$(U(\mathbb{S}), s)$. Furthermore, since it is easy to see that *every* M-model is of the form $U(\mathbb{S})$ for some $\underline{\mathsf{M}}$-model $\mathbb{S}$, it follows that a formula φ of $\mu\mathtt{ML}_\Sigma$ is valid over M-models if and only if it is valid over $\underline{\mathsf{M}}$-models. So by Theorem 10.7, to axiomatize the valid formulas of $\mu\mathtt{ML}_\Sigma$ over M-models, it suffices to show the following.

Proposition 10.8. *The logic* $\mu\underline{\mathbf{M}}$ *is a conservative extension of* $\mu\mathbf{M}$*.*

The proof of this result will make use of algebras for μ-calculi, which are called *μ-algebras* and have been thoroughly studied by Santocanale, see, e.g., [32]. Since our argument takes places in a fully Boolean context, we simplify our notation somewhat by working with the box modalities only.

Definition 10.9. An algebra $\mathbb{A} = (A, 0, 1, \wedge, \neg, \Box)$ is a *monotone modal algebra* if its Boolean reduct $(A, 0, 1, \wedge, \neg)$ is a Boolean algebra, and $\Box$ is a monotone (i.e., order preserving) operation on A. An algebra $\mathbb{A} = (A, 0, 1, \wedge, \neg, \Box, \underline{\Box})$ is a *supported monotone modal algebra* if its Σ-reduct $(A, 0, 1, \wedge, \neg, \Box)$ is a monotone modal algebra, and $\mathbb{A}$ satisfies the (equational versions of the) $\underline{\mathbf{M}}$-axioms a)–c) of Definition 10.6. ◁

Definition 10.10. A monotone modal algebra $\mathbb{A}$ is said to be a *monotone modal μ-algebra* if every map $v : \mathtt{X} \to A$ uniquely extends to a map $v^* : \mu\mathtt{ML}_\Sigma \to A$ which is a homomorphism with respect to all connectives, and which respects the least fixpoint operator in the following sense. Let $\varphi^v_p : A \to A$ denote the map defined by $\varphi^v_p(a) = v[a/p]^*(\varphi)$ where $v[a/p]$ is like v except it maps p to a. Then the map φ^v_p has a smallest pre-fixpoint m, and $v^*(\mu p.\varphi) = m$.

The notion of a *supported monotone modal algebra* is defined in a completely analogous way. ◁

We shall need the following simple little observation about fixpoints in lattices, the proof of which is a straightforward exercise:

Proposition 10.11. *Let L be any lattice, and let $f : L \times L \to L$ be a two-place map that is monotone in both its arguments. Suppose that, for all $b \in L$, the least prefixpoint l_b of the map $\lambda x.f(x, b)$ exists. Suppose furthermore that the meet*

$$m = \bigwedge\{l_b \mid f(b, b) \leq b,\ b \in L\}$$

exists. Then m is the least prefixpoint of the map $\lambda z.f(z, z)$.

Proof. First, note that $m \leq p$ for any prefixpoint p of the map $\lambda z.f(z, z)$: if $f(p, p) \leq p$ then $m \leq l_p$, but $l_p \leq p$ since l_p was the smallest prefixpoint of the map $\lambda x.f(x, p)$.

It now suffices to prove that m is a prefixpoint of $\lambda z.f(z, z)$, since it will then follow that it is the least prefixpoint. Let $X := \{b \in L \mid f(b, b) \leq b\}$. If $b \in X$ then $m \leq l_b$ by definition of m, and furthermore $l_b \leq b$ since l_b was the least prefixpoint of the map $\lambda x.f(x, b)$. So we get, for all $b \in X$:

$$\begin{array}{rcll} f(m, m) & \leq & f(l_b, l_b) & \text{(Monotonicity)} \\ & \leq & f(l_b, b) & \text{(Monotonicity \& } l_b \leq b\text{)} \\ & \leq & l_b & \text{(Prefixpoint property of } l_b\text{)} \end{array}$$

But since $m = \bigwedge\{l_b \mid b \in X\}$, we now get $f(m, m) \leq m$, as required. □

Using a standard Lindenbaum–Tarski algebra construction, one can prove the following algebraic completeness results.

Proposition 10.12. *Let φ be any formula of $\mu\mathtt{ML}_\Sigma$. Then $\mu\mathbf{M} \vdash \varphi$ if, and only if, $v^*(\varphi) = 1$ for every monotone modal μ-algebra $\mathbb{A}$ and every valuation $v : \mathtt{X} \to A$.*

Proposition 10.13. *Let φ be any formula of $\mu\mathtt{ML}_{\underline{\Sigma}}$. Then $\mu\underline{\mathbf{M}} \vdash \varphi$ if, and only if, $v^*(\varphi) = 1$ for every supported monotone modal μ-algebra $\mathbb{A}$ and every valuation $v : \mathtt{X} \to A$.*

With these completeness results in place, the conservative extension theorem we want to prove boils down to the following statement:

Proposition 10.14. *Every monotone modal μ-algebra is a reduct of some supported monotone modal μ-algebra.*

Before we turn to the proof of this proposition, we show how it entails that $\mu\underline{\mathbf{M}}$ is a conservative extension of $\mu\mathbf{M}$: it is clear that every formula of $\mu\mathtt{ML}_\Sigma$ provable in $\mu\mathbf{M}$ is provable in $\mu\underline{\mathbf{M}}$ also. Conversely, suppose that $\varphi \in \mu\mathtt{ML}_\Sigma$ is *not* provable in $\mu\mathbf{M}$. Then by Proposition 10.12, there is a monotone modal μ-algebra $\mathbb{A}$ and a valuation $v : \mathtt{X} \to A$ such that $v^*(\varphi) \neq 1$. By Proposition 10.14 there is a supported monotone modal μ-algebra $\mathbb{A}'$ over the same carrier, whose reduct is equal to $\mathbb{A}$. So the map v extends uniquely to the map $v^\dagger$ witnessing that $\mathbb{A}'$ is a supported monotone modal μ-algebra, and clearly $v^\dagger(\varphi) = v^*(\varphi) \neq 1$. So by Proposition 10.13, $\mu\underline{\mathbf{M}} \nvdash \varphi$ as required.

We now turn to the proof of Proposition 10.14:

Proof of Proposition 10.14. Let $\mathbb{A} = (A, 0, 1, \wedge, \neg, \Box)$ be a monotone modal μ-algebra. We want to define an operation $\underline{\Box} : A \to A$ that makes $\mathbb{A}' = (A, 0, 1, \wedge, \neg, \Box, \underline{\Box})$ a supported monotone modal μ-algebra. The construction is straightforward: for each $a \in A$, set $\underline{\Box} a = 1$ if $a = 1$, and $\underline{\Box} a = 0$ otherwise. It is a purely mechanical task to check that this is in fact a supported monotone modal algebra. The argument showing that $\mathbb{A}'$ is, in addition, a supported monotone modal μ-algebra is based on finding, for each $\mu\mathtt{ML}_{\underline{\Sigma}}(\mathtt{X})$-formula φ with a positive variable p, and every map $v : \mathtt{X} \to A$, a formula $t(\varphi, v)$ of $\mu\mathtt{ML}_\Sigma(\mathtt{X})$ such that $v^*(\mu p.t(\varphi, v))$ is a least prefixpoint of the map φ^v_p. More precisely, the proof is based on the following claim, which is proved by induction on the complexity of formulas in $\mu\mathtt{ML}_{\underline{\Sigma}}$:

Claim 1. *There exists an assignment $t(-, -)$ mapping every formula φ of $\mu\mathtt{ML}_{\underline{\Sigma}}$, and every valuation $v : \mathtt{X} \to A$, to a formula $t(\varphi, v)$ in $\mu\mathtt{ML}_\Sigma$, such that:*

- *if φ is positive (negative) in p then so is $t(\varphi, v)$,*
- *if φ is positive in p and $a \leq b$ then $w^*(t(\varphi, v[a/p])) \leq w^*(t(\varphi, v[b/p]))$ for every valuation w,*
- *if φ is negative in p and $a \leq b$ then $w^*(t(\varphi, v[b/p])) \leq w^*(t(\varphi, v[a/p]))$ for every valuation w,*
- *$t(\underline{\Box}\varphi, v) = \top$ if $v^*(t(\varphi, v)) = 1$, $t(\underline{\Box}\varphi, v) = \bot$ otherwise,*
- *$t(\varphi \wedge \psi, v) = t(\varphi, v) \wedge t(\psi, v)$,*
- *$t(\Box\varphi, v) = \Box t(\varphi, v)$ and $t(\neg\varphi, v) = \neg t(\varphi, v)$,*
- *for every formula φ the set $\{t(\varphi, v) \mid v : \mathtt{X} \to A\}$ is finite, and if φ is positive in p then:*

$$t(\mu p.\varphi, v) = \bigwedge \{\mu p.t(\varphi, v[a/p]) \mid v^*(t(\varphi, v[a/p])) \leq a\}$$

where the big conjunction on the right-hand side is defined since the set $\{t(\varphi, v[a/p]) \mid a \in A\}$ is finite.

With this claim in place, we can define the map $v^\dagger : \mu\mathtt{ML}_{\underline{\Sigma}} \to A$ by setting:

$$v^\dagger(\varphi) := v^*(t(\varphi, v))$$

Note that this map commutes with all the connectives, including the operator $\boxminus$: to see why, we make a case distinction. If $v^*(t(\varphi, v)) = 1$ then $t(\boxminus\varphi, v) = \top$, so $v^\dagger(\boxminus\varphi) = v^*(\top) = 1$. But $v^*(t(\varphi, v)) = 1$ means that $v^\dagger(\varphi) = 1$, hence $\boxminus v^\dagger(\varphi) = 1$ by definition of the operation $\boxminus$ in $\mathbb{A}'$. On the other hand, if $v^*(t(\varphi, v)) \neq 1$ then $t(\boxminus\varphi, v) = \bot$, so $v^\dagger(\boxminus\varphi) = v^*(\bot) = 0$. But $v^*(t(\varphi, v)) \neq 1$ means that $v^\dagger(\varphi) \neq 1$, hence $\boxminus v^\dagger(\varphi) = 0$, again by definition of the operation $\boxminus$ in $\mathbb{A}'$.

Finally, we find that the value $v^\dagger(\mu p.\varphi)$ is indeed a least pre-fixpoint of the map φ_p^v, as an instance of Proposition 10.11: just put $f(a, b) := v[a/p]^*(t(\varphi, v[b/p]))$ and note that $f(a, a) = v[a/p]^\dagger(\varphi)$. □

This concludes our proof of Theorem 1.4.

10.2. Transferring completeness from coalgebraic modal logics

As a final application, we prove Corollary 1.2 which allows one to transfer any previously established completeness result for a coalgebraic modal logic to a completeness result for its fixpoint extension. Formally, given a monotone modal signature Λ for a functor T, the formulas of the coalgebraic modal logic $\mathtt{ML}_\Lambda$ are defined by the following grammar:

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

Semantics of these formulas in a T-model are as before, and we say that a formula $\varphi \in \mathtt{ML}_\Lambda$ is *valid* if it is true in every pointed T-model. We denote this by $\vDash \varphi$ as before.

We take a Hilbert-style *axiom system* $\mathbf{L}$ for $\mathtt{ML}_\Lambda$ to be a set of formulas in $\mathtt{ML}_\Lambda$, and we say that a formula φ is derivable in the system, $\vdash_{\mathbf{L}} \varphi$, if it is provable from axioms in $\mathbf{L}$, substitution instances of propositional tautologies, (Mon) and (Du) using the rules of modus ponens, uniform substitution and the congruence rule. We define the derivation system $\mu\mathbf{L}$ for the extended language $\mu\mathtt{ML}_\Lambda$ by simply adding the fixpoint axiom and Kozen–Park induction rule, i.e. we say that φ is derivable in $\mu\mathbf{L}$, $\vdash_{\mu\mathbf{L}} \varphi$, if it is derivable using axioms in $\mathbf{L}$ and the fixpoint axiom using the rules of $\mathbf{L}$ and the Kozen–Park induction rule.

Definition 10.15. The system $\mathbf{L}$ ($\mu\mathbf{L}$) is said to be sound and complete if, for any formula $\varphi \in \mathtt{ML}_\Lambda$ ($\varphi \in \mu\mathtt{ML}_\Lambda$), we have $\vDash \varphi$ iff $\vdash_{\mathbf{L}} \varphi$ ($\vdash_{\mu\mathbf{L}} \varphi$). ◁

We can now prove our transfer result:

Proof of Corollary 1.2. Soundness clearly transfers from $\mathbf{L}$ to $\mu\mathbf{L}$. To prove completeness, we define a one-step axiom system $\mathbf{H}$ by setting, for a one-step formula $\alpha \in \mathtt{1ML}_\Lambda(\mathtt{Var})$, $\alpha \in \mathbf{H}$ if $\vdash_{\mathbf{L}} \alpha$. To prove that the system $\mu\mathbf{H}$ is sound, we need the following claim, the proof of which we leave to the reader:

$$\vDash \alpha \text{ iff } \vDash^1 \alpha, \tag{48}$$

for any $\alpha \in \mathtt{1ML}_\Lambda(\mathtt{Var})$. The proof of (48) basically consists of noting that every pointed T-model $(\mathbb{S}, s)$ induces a one-step model by simply applying the map $\sigma : S \to \mathsf{T}S$ to s, and conversely every one-step model (S, σ, m) can be viewed as a pointed T-model by simply adding a new point u and mapping this to σ (and mapping elements of S to arbitrary elements of $\mathsf{T}S$).

It clearly follows from (48) that the one-step derivation system $\mathbf{H}^1$ is one-step complete, so by Theorem 1.1 the system $\mu\mathbf{H}$ is complete.

It now suffices to prove that every formula φ that is provable in $\mu\mathbf{H}$ is also provable in $\mu\mathbf{L}$. But this is in fact an easy consequence of the definition of $\mathbf{H}$: given any axiom of $\mu\mathbf{H}$ of the form $\alpha[\tau]$ where $\alpha \in \mathbf{H}$ and $\tau : \mathtt{Var} \to \mu\mathtt{ML}_\Lambda$, we have $\vdash_\mathbf{L} \alpha$ by definition of $\mathbf{H}$ and so $\vdash_{\mu\mathbf{L}} \alpha$, hence $\vdash_{\mu\mathbf{L}} \alpha[\tau]$ by an application of the uniform substitution rule. All other axioms of $\mu\mathbf{H}$ are axioms of $\mu\mathbf{L}$ too, and all rules in $\mu\mathbf{H}$ are in $\mu\mathbf{L}$. Hence the system $\mu\mathbf{L}$ indeed proves all theorems of $\mu\mathbf{H}$, and so is complete. $\square$

11. Concluding remarks

We conclude the paper by mentioning a few topics for future investigation. Clearly, a question that is high on the priority list is whether the existence of a disjunctive basis is necessary to obtain completeness for modal μ-calculi, or if only one-step completeness suffices. It could be that one can prove a stronger generic completeness result, to the effect that one-step completeness *always* lifts to completeness for the full modal μ-calculus corresponding to some modal signature. This would provide us with a very powerful tool to prove completeness for modal fixpoint logics, which would make it a completely routine task in many cases. A possible path to such a result may have opened up due to very recent work by Afshari & Leigh [2], who proved completeness for the modal μ-calculus in a way that avoids the detour via disjunctive normal forms. Our hope, and belief, is that much of our technical machinery can be merged with their proof theoretic approach. In particular one of the first tasks would be to incorporate into their setting games for coalgebraic modal automata and the connection we set up between automata and proofs. This promises not only to put Afshari & Leigh's result to use to obtain a stronger coalgebraic completeness result for μ-calculi, but it could also bring in conceptually interesting automata- and game-theoretic perspectives on their proof. It is even possible that this approach could streamline and simplify parts of the proof, as we previously tried to do for Walukiewicz's original proof in [12].

But we want to stress that even if the approach via disjunctive bases and the simulation theorem for automata can ultimately be avoided, this does not mean that it is no longer interesting. Disjunctive normal forms and non-deterministic automata for the modal μ-calculus were originally invented for the purpose of proving completeness for the μ-calculus, but they have later gained a much wider significance, lying at the heart of much (or even most) of the metatheory that has been developed for the μ-calculus since then. So it certainly seems a significant fact that some of the most central features of Walukiewicz's completeness proof turn out to be essentially coalgebraic, and have natural coalgebraic generalizations. Indeed, the first and third authors have recently explored some further applications of disjunctive bases in [13], where we used it to generalize the uniform interpolation and Lyndon theorems for the modal μ-calculus [8] to arbitrary coalgebraic modal μ-calculi. Furthermore, in our completeness proof we have come across a number of coalgebraic and automata-theoretic concepts that deserve a deeper further investigation. Especially, we think the consequence game for coalgebraic modal automata should be given special attention. As a first observation, it is easy to see that every modal automaton has a winning "identity strategy" for Player II, and that winning strategies for Player II can be composed. This suggests that we may identify an interesting category of coalgebraic modal automata, where the arrows are player II's winning strategies (possibly modulo some suitable equivalence relation over strategies). Such a category would have a much richer structure than the poset category given by semantic consequence between automata. If so, this would open a range of questions, such as: Can disjunctive automata be characterized as objects in this category by some purely categorical property?

Finally, what is the proof-theoretic significance of the consequence game, the satisfiability game, the concept of semi-disjunctive automata etc.? Can these concepts be used to study the structure of proofs in the μ-calculus by automata theoretic methods? Or vice versa, perhaps the connection can be used to apply proof theoretic methods to address problems about the μ-calculus that up to now have been handled mainly using automata, like interpolation.

Appendix A. Basic definitions

A.1. Basic mathematical concepts and notation

Definition A.1. Let A be some set. We denote its size as $|A|$, and its power set as $\mathsf{P}A$. ◁

Since binary relations play an important role in our work, we will frequently use the following notation.

Definition A.2. The collection of binary relations over a set A is denoted as $A^\sharp$. Given a relation $R \subseteq A \times A'$, we let $\mathsf{Dom}R$ and $\mathsf{Ran}R$ denote its domain and range, respectively; for a subset $B' \subseteq A'$, we define $\mathsf{Ran}_{B'}R := \mathsf{Ran}R \cap B'$. Furthermore, we denote the converse relation of R as $R^\circ := \{(a', a) \in A' \times A \mid (a, a') \in R\}$, and we set $R[a] := \{a' \in A' \mid Raa'\}$. The composition of two relations R and S is denoted as $R ; S$, and the diagonal relation on a set S is denoted as Id_S. Given a relation $R \subseteq A \times A$ and a subset $B \subseteq A$, we let $\mathsf{Res}_B R := R \cap (B \times B)$ denote the *restriction* of R to B. ◁

Definition A.3. Given a relation $R \subseteq A \times A'$, we define the following relations between $\mathsf{P}A$ and $\mathsf{P}A'$:

$$\overrightarrow{\mathsf{P}}R := \{(B, B') \in \mathsf{P}A \times \mathsf{P}A' \mid \text{ for all } b \in B \text{ there is a } b' \in B' \text{ with } Rbb'\}$$
$$\overleftarrow{\mathsf{P}}R := \{(B, B') \in \mathsf{P}A \times \mathsf{P}A' \mid \text{ for all } b' \in B' \text{ there is a } b \in B \text{ with } Rbb'\}$$
$$\overline{\mathsf{P}}R := \overrightarrow{\mathsf{P}}R \cap \overleftarrow{\mathsf{P}}R.$$

The relation $\overline{\mathsf{P}}R$ is called the *Egli–Milner* lifting of R. ◁

Definition A.4. We write $f : A \to B$ to denote that f is a map from A to B, and we will frequently identify f with its *graph* $\mathsf{Gr}f := \{(a, fa) \mid a \in A\}$. The composition of two functions $f : A \to B$ and $g : B \to C$ is denoted as $g \circ f : A \to C$. ◁

Definition A.5. Given a set A, we let A^* and A^ω denote, respectively, the set of *words* (finite sequences) and *streams* (infinite sequences) over A. We will write both ww' and $w \cdot w'$ to denote the concatenation of the words w and w', and similar for the concatenation of a word and a stream. The last symbol of a word w is denoted as $\mathsf{last}(w)$.

Two A-streams σ and τ are *eventually equal*, denoted as $\sigma =_\infty \tau$, if there is a $k \in \omega$ such that $\sigma(j) = \tau(j)$ for all $j \geq k$. ◁

A.2. Set functors

As mentioned in section 2, we let Set denote the category with sets as objects and functions as arrows. An endofunctor on Set will simply be called a *set functor*. In this section we briefly define and review some pertinent categorical notions regarding set functors.

Convention A.6. Throughout this paper we shall assume that T is a set functor that preserves injections. For convenience we will in fact assume that T preserves inclusions; that is, with $\iota_B^A : A \hookrightarrow B$ denoting the inclusion map from a subset A of B to B, we have

$$\mathsf{T}X \subseteq \mathsf{T}Y \text{ and } \mathsf{T}(\iota_Y^X) = \iota_{\mathsf{T}Y}^{\mathsf{T}X}$$

for all pairs X, Y of sets such that $X \subseteq Y$.

For completeness we recall some definitions related to the notion of a (weak) pullback.

Definition A.7. Recall that a set P together with functions $p_1 : P \to X_1$ and $p_2 : P \to X_2$ is a *pullback* of two functions $f_1 : X_1 \to X$ and $f_2 : X_2 \to X$ if $f_1 \circ p_1 = f_2 \circ p_2$ and for all sets P' and all functions $p'_1 : P' \to X_1$, $p'_2 : P' \to X_2$ such that $f_1 \circ p'_1 = f_2 \circ p'_2$ there exists a *unique* function $e : P' \to P$ such that $p_i \circ e = p'_i$ for $i = 1, 2$:

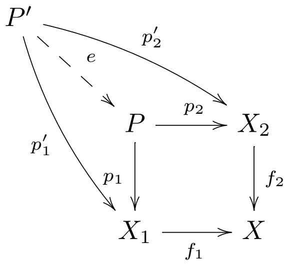

If the function e is not necessarily unique we call (P, p_1, p_2) a *weak pullback*. Furthermore we call a relation $R \subseteq X_1 \times X_2$ a (weak) pullback of f_1 and f_2 if R together with the projection maps π^R_1 and π^R_2 is a (weak) pullback of f_1 and f_2. ◁

In the category of sets, (weak) pullbacks have a straightforward characterization.

Fact A.8. [17]. *Given two functions $f_1 : X_1 \to X_3$ and $f_2 : X_2 \to X_3$, let*

$$\mathsf{pb}(f_1, f_2) := \{(x_1, x_2) \mid f_1(x_1) = f_2(x_2)\}.$$

Furthermore, given a set P with functions $p_1 : P \to X_1$ and $p_2 : P \to X_2$, let

$$e : y \mapsto (p_1(y), p_2(y))$$

define a function $e : P \to \mathsf{pb}(f_1, f_2)$. Then

(1) *(P, p_1, p_2) is a pullback of f_1 and f_2 iff $f_1 \circ p_1 = f_2 \circ p_2$ and e is an isomorphism.*
(2) *(P, p_1, p_2) is a weak pullback of f_1 and f_2 iff $f_1 \circ p_1 = f_2 \circ p_2$ and e is surjective.*

Definition A.9. A functor T *preserves weak pullbacks* if it transforms every weak pullback (P, p_1, p_2) for f_1 and f_2 into a weak pullback $(\mathsf{T}P, \mathsf{T}p_1, \mathsf{T}p_2)$ for $\mathsf{T}f_1$ and $\mathsf{T}f_2$. ◁

An equivalent characterization is to require T to *weakly preserve pullbacks*, that is, to turn pullbacks into weak pullbacks. In Fact 2.10 we give another, and probably more motivating, characterization of this property.

Proposition A.10. *Let $f : S \to S'$ be some map, and let $X \subseteq S$ be a subset of S. Then for any $\xi \in \mathsf{T}X$ we have $(\mathsf{T}f)\xi \in \mathsf{T}(f[X])$.*

Proof. Since $X \subseteq S$ and $f[X] \subseteq S'$ we have that $\mathsf{T} \subseteq \mathsf{T}S$ and $\mathsf{T}(f[X]) \subseteq \mathsf{T}S'$. Now consider the following diagram:

$$\begin{array}{ccc} \mathsf{T}X & \xrightarrow{\mathsf{T}f\restriction_X} & \mathsf{T}(f[X]) \\ \downarrow & & \downarrow \\ \mathsf{T}S & \xrightarrow{\mathsf{T}f} & \mathsf{T}S' \end{array}$$

Chasing ξ in this diagram yields the statement of the proposition. □

Definition A.11. Given a finite set S we let

$$\mathsf{Base}_S : \sigma \mapsto \bigcap \{X \subseteq S \mid \sigma \in \mathsf{T}X\}$$

define a map $\mathsf{Base}_S : \mathsf{T}S \to \mathsf{P}S$. ◁

Fact A.12. *Let $f : S \to S'$ be some map between finite sets S, S', and let $\sigma \in \mathsf{T}S$.*

(1) $\mathsf{Base}_S(\sigma)$ is the smallest set X such that $\sigma \in \mathsf{T}X$.

(2) $\mathsf{Base}_{S'}((\mathsf{T}f)\sigma) \subseteq (\mathsf{P}f)(\mathsf{Base}_x(\alpha))$.

(3) $\mathsf{Base}_{S'}((\mathsf{T}f)\sigma) = (\mathsf{P}f)(\mathsf{Base}_x(\alpha))$ if T is weak pullback preserving; hence in this case Base is a natural transformation, $\mathsf{Base} : \mathsf{T}_\omega \to \mathsf{P}_\omega$:

$$\begin{array}{ccccc} S & \quad & \mathsf{T}S & \xrightarrow{\mathsf{Base}_S} & \mathsf{P}S \\ {\scriptstyle f}\downarrow & & {\scriptstyle \mathsf{T}f}\downarrow & & \downarrow{\scriptstyle \mathsf{P}f} \\ S' & & \mathsf{T}S' & \xrightarrow[\mathsf{Base}_{S'}]{} & \mathsf{P}S' \end{array}$$

A.3. Graph games

Definition A.13. 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 an *initial* position p, usually denoted as $\mathbb{G}@p$. ◁

Definition A.14. A *play* 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 play 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 plays ρ ending in a position $\mathsf{last}(\rho) \in G_\Pi$, and define $\mathrm{PM}_\Pi@p$ as the set of partial plays in PM_Π starting at position p.

The *winner* of a full play 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. ◁

Definition A.15. 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 play, 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 play $\rho = (p_i)_{i<\kappa}$ 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 $p_0 \dots p_{n-1} \in \mathrm{PM}_\Pi$ (that is, $p_{n-1} \in G_\Pi$). A Π-strategy χ is *legitimate* in $\mathbb{G}@p$ if the moves that it prescribes to χ-guided partial plays 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 plays 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 plays that are themselves guided by χ.

The following fact, independently due to Emerson & Jutla [9] and Mostowski [27], will be quite useful to us.

Fact A.16 *(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.*

References

[1] P. Aczel, Non-Well-Founded Sets, CSLI Publications, 1988.
[2] B. Afshari, G.E. Leigh, Cut-free completeness for modal mu-calculus, in: Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, IEEE Computer Society, 2017, pp. 1–12.
[3] A. Arnold, D. Niwiński, Rudiments of μ-calculus, in: Studies in Logic and the Foundations of Mathematics, vol. 146, North-Holland Publishing Co., Amsterdam, 2001.
[4] H. Barringer, R. Kuiper, A. Pnueli, A really abstract concurrent model and its temporal logic, in: Conference Record of the Thirteenth Annual ACM Symposium on Principles of Programming Languages, POPL 13, 1986, pp. 173–183.
[5] P. Blackburn, M. de Rijke, Y. Venema, Modal logic, in: Cambridge Tracts in Theoretical Computer Science, vol. 53, Cambridge University Press, 2001.
[6] C. Cîrstea, C. Kupke, D. Pattinson, EXPTIME tableaux for the coalgebraic μ-calculus, in: E. Grädel, R. Kahle (Eds.), Computer Science Logic 2009, in: Lecture Notes in Computer Science, vol. 5771, Springer, 2009, pp. 179–193.
[7] C. Cîrstea, D. Pattinson, Modular construction of modal logics, in: Ph. Gardner, N. Yoshida (Eds.), Proceedings of the 15th International Conference on Concurrency Theory, CONCUR 2004, in: LNCS, vol. 3170, Springer, 2004, pp. 258–275.
[8] G. D'Agostino, M. Hollenberg, Logical questions concerning the μ-calculus, J. Symbolic Logic 65 (2000) 310–332.
[9] E.A. Emerson, C.S. Jutla, Tree automata, mu-calculus and determinacy (extended abstract), in: Proceedings of the 32nd Symposium on the Foundations of Computer Science, IEEE Computer Society Press, 1991, pp. 368–377.
[10] S. Enqvist, F. Seifan, 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, 2015, pp. 353–365.
[11] S. Enqvist, F. Seifan, Y. Venema, Completeness for coalgebraic fixpoint logic, in: Proceedings of the 25th EACSL Annual Conference on Computer Science Logic, CSL 2016, in: LIPIcs, vol. 62, 2016, pp. 7:1–7:19.
[12] S. Enqvist, F. Seifan, Y. Venema, Completeness for the modal μ-calculus: separating the combinatorics from the dynamics, Theoret. Comput. Sci. 737 (2018) 37–100.
[13] S. Enqvist, Y. Venema, Disjunctive bases: normal forms for modal logics, in: F. Bonchi, B. König (Eds.), 7th Conference on Algebra and Coalgebra in Computer Science, CALCO 2017, Dagstuhl, Germany, in: LIPIcs, vol. 72, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017, pp. 11:1–11:16. An expanded and updated version is available at arXiv:1710.10706.
[14] G. Fontaine, R. Leal, Y. Venema, Automata for coalgebras: an approach using predicate liftings, in: Automata, Languages and Programming: 37th International Colloquium, ICALP'10, in: LNCS, vol. 6199, Springer, 2010, pp. 381–392.
[15] D. Gorín, L. Schröder, Simulations and bisimulations for coalgebraic modal logics, in: Proceedings of the 5th International Conference on Algebra and Coalgebra in Computer Science, CALCO 2013, in: LNCS, vol. 8089, Springer, 2013, pp. 253–266.
[16] E. Grädel, W. Thomas, T. Wilke (Eds.), Automata, Logic, and Infinite Games, LNCS, vol. 2500, Springer, 2002.
[17] H.P. Gumm, Functors for coalgebras, Algebra Universalis 45 (2–3) (2001) 135–147.
[18] B. Jacobs, Introduction to Coalgebra, Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, 2016.
[19] D. Janin, A Contribution to Formal Methods: Games, Logic and Automata, Habilitation thesis, Computer Science, Université Sciences et Technologies – Bordeaux I, 2005.
[20] D. Janin, I. Walukiewicz, Automata for the modal μ-calculus and related results, in: J. Wiedermann, P. Hájek (Eds.), Mathematical Foundations of Computer Science 1995, 20th International Symposium, MFCS'95, in: LNCS, vol. 969, Springer, 1995, pp. 552–562.
[21] R. Kaivola, Axiomatising linear time mu-calculus, in: I. Lee, S.A. Smolka (Eds.), Proceedings of the 6th International Conference on Concurrency Theory, CONCUR '95, in: LNCS, vol. 962, Springer, 1995, pp. 423–437.
[22] D. Kozen, Results on the propositional μ-calculus, Theoret. Comput. Sci. 27 (1983) 333–354.
[23] O. Kupferman, U. Sattler, M.Y. Vardi, The complexity of the graded μ-calculus, in: Andrei Voronkov (Ed.), Proceedings of the 18th International Conference on Automated Deduction, CADE 18, in: LNCS, vol. 2392, Springer, 2002, pp. 423–437.
[24] C. Kupke, A. Kurz, Y. Venema, Completeness for the coalgebraic cover modality, Log. Methods Comput. Sci. 8 (3) (2012).
[25] J. Marti, Y. Venema, Lax extensions of coalgebra functors and their logic, J. Comput. System Sci. 81 (5) (2015) 880–900.
[26] L. Moss, Coalgebraic logic, Ann. Pure Appl. Logic 96 (1999) 277–317 (Erratum published: Ann. Pure Appl. Logic 99 (1999) 241–259).
[27] A.M. Mostowski, Games with Forbidden Positions, Technical Report 78, Instytut Matematyki, Uniwersytet Gdański, Poland, 1991.

[28] D.E. Muller, P.E. Schupp, Simulating alternating tree automata by nondeterministic automata, Theoret. Comput. Sci. 141 (1995) 69–107.
[29] E. Pacuit, S. Salame, Majority logic, in: Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning, KR2004, 2004, pp. 598–605.
[30] D. Pattinson, Coalgebraic modal logic: soundness, completeness and decidability of local consequence, Theoret. Comput. Sci. 309 (1–3) (2003) 177–193.
[31] J. Rutten, Universal coalgebra: a theory of systems, Theoret. Comput. Sci. 249 (2000) 3–80.
[32] L. Santocanale, Completions of μ-algebras, Ann. Pure Appl. Logic 154 (1) (2008) 27–50.
[33] L. Santocanale, Y. Venema, Completeness for flat modal fixpoint logics, Ann. Pure Appl. Logic 162 (2010) 55–82.
[34] L. Schröder, A finite model construction for coalgebraic modal logic, J. Log. Algebr. Program. 73 (2007) 97–110.
[35] L. Schröder, D. Pattinson, PSPACE bounds for rank-1 modal logics, ACM Trans. Comput. Log. 10 (2) (2009) 13.
[36] L. Schröder, D. Pattinson, Rank-1 modal logics are coalgebraic, J. Logic Comput. 20 (5) (2010) 1113–1147.
[37] L. Schröder, Y. Venema, Flat coalgebraic fixed point logics, in: Proceedings of the 21st International Conference on Concurrency Theory, CONCUR 2010, 2010, pp. 524–538.
[38] Y. Venema, Lectures on the Modal μ-Calculus, Lecture Notes, ILLC, University of Amsterdam, 2012.
[39] I. Walukiewicz, Completeness of Kozen's axiomatisation of the propositional μ-calculus, Inform. and Comput. 157 (2000) 142–182.
[40] T. Wilke, Alternating tree automata, parity games, and modal μ-calculus, Bull. Belg. Math. Soc. 8 (2001) 359–391.