

Automata for Coalgebras: an approach using predicate liftings^{*}

Gaëlle Fontaine, Raul Leal, and Yde Venema

Institute for Logic, Language and Computation, Universiteit van Amsterdam,
Science Park 904, 1098XH Amsterdam, the Netherlands.

Abstract. Universal Coalgebra provides the notion of a *coalgebra* as the natural mathematical generalization of state-based evolving systems such as (infinite) words, trees, and transition systems. We lift the theory of parity automata to a coalgebraic level of abstraction by introducing, for a set Λ of predicate liftings associated with a set functor \mathcal{T} , the notion of a Λ -automata operating on coalgebras of type \mathcal{T} . In a familiar way these automata correspond to extensions of coalgebraic modal logics with least and greatest fixpoint operators.

Our main technical contribution is a general bounded model property result: We provide a construction that transforms an arbitrary Λ -automaton \mathbb{A} with nonempty language into a small pointed coalgebra (\mathbb{S}, s) of type \mathcal{T} that is recognized by \mathbb{A} , and of size exponential in that of \mathbb{A} . \mathbb{S} is obtained in a uniform manner, on the basis of the winning strategy in our satisfiability game associated with \mathbb{A} . On the basis of our proof we obtain a general upper bound for the complexity of the non-emptiness problem, under some mild conditions on Λ and \mathcal{T} . Finally, relating our automata-theoretic approach to the tableaux-based one of Cirstea et alii, we indicate how to obtain their results, based on the existence of a complete tableau calculus, in our framework.

Keywords: coalgebra, modal logic, parity automata, predicate liftings, fixpoint logic

1 Introduction

The theory of finite automata, seen as devices for classifying (possibly) infinite structures [6], combines a rich mathematical theory, dating back to the seminal work of Büchi and Rabin, with a wide range of applications, particularly in areas related to program verification and synthesis. The main purpose of our paper is to contribute to this theory by showing that some of its fundamental ideas can be lifted to a coalgebraic level of generality.

Universal Coalgebra [?] provides the notion of a *coalgebra* as the natural mathematical generalization of state-based evolving systems such as streams, (infinite) trees, Kripke models, (probabilistic) transition systems, and many others. Formally, a coalgebra is a pair $\mathbb{S} = (S, \sigma)$, where S is the carrier or state

^{*} The research of the authors has been made possible by VICI grant 639.073.501 of the Netherlands Organization for Scientific Research (NWO).

space of the coalgebra, and $\sigma : S \rightarrow \mathcal{T}S$ is its unfolding or transition map. This approach combines simplicity with generality and wide applicability: many features, including input, output, nondeterminism, probability, and interaction, can easily be encoded in the coalgebra type \mathcal{T} (formally an endofunctor on the category \mathbf{Set} of sets as objects with functions as arrows).

Logic enters the picture if one wants to specify and reason about *behavior*, one of the most fundamental notions admitting a coalgebraic formalization. With Kripke structures constituting key examples of coalgebras, it should come as no surprise that most coalgebraic logics are some kind of modification or generalization of *modal logic*.

Moss [?] introduced a modality $\nabla_{\mathcal{T}}$ generalizing the so-called ‘cover modality’ from Kripke structures to coalgebras of arbitrary type. This approach is uniform in the functor \mathcal{T} , but as a drawback only works properly if \mathcal{T} satisfies a certain category-theoretic property (viz., it should preserve weak pullbacks); also the nabla modality is syntactically rather nonstandard. As an alternative, Pattinson [?] and others developed coalgebraic modal formalisms, based on a completely standard syntax, that work for coalgebras of arbitrary type. In this approach, the semantics of each modality is determined by a so-called *predicate lifting* (see Definition 3 below). Many well-known variations of modal logic in fact arise as the coalgebraic logic $\text{ML}_{\mathcal{A}}$ associated with a set \mathcal{A} of such predicate liftings; examples include both standard and (monotone) neighborhood modal logic, graded and probabilistic modal logic, coalition logic, and conditional logic. The theory of coalgebraic modal logic has developed rather rapidly; to mention just one example, it presently includes generic PSPACE upper bounds for the satisfiability problem [?].

The fact that ordinary modal formulas have a finite *depth* severely restricts the expressive power of plain coalgebraic modal logic, and thus limits its usefulness as a language for specifying *ongoing* behavior. For the latter purpose one needs to extend the language with *fixpoint operators*, generalizing the modal μ -calculus [?]. A coalgebraic fixpoint language on the basis of Moss’ modality was introduced by Venema [10]. Recently, Cirstea, Kupke and Pattinson [5] introduced the *coalgebraic μ -calculus* $\mu\text{ML}_{\mathcal{A}}$ parametrized by a set \mathcal{A} of predicate liftings for a functor \mathcal{T} .

Given the success of automata-theoretic approaches towards fixpoint logics, one may expect a rich and elegant *universal automata theory* that generalizes the theory of specific devices for streams, trees or graphs, by dealing with automata that operate on coalgebras. A first step in this direction was the introduction of so-called *coalgebra automata* by Venema [10]. Kupke & Venema [?] generalized many results in automata theory, such as closure properties of recognizable languages, to this class of automata. However, coalgebra automata are related to fixpoint languages based on Moss’ modality ∇ , and do not correspond directly to coalgebraic modal languages associated with predicate liftings (such as the graded modal μ -calculus). In addition, the theory of coalgebra automata needs the *type* of the coalgebras to be a functor that preserves weak pullbacks, and hence cannot be applied as generally as possible.

This paper introduces automata for coalgebras of *arbitrary* type (Definition 4). More precisely, given a set Λ of monotone predicate liftings, we introduce Λ -automata as devices that accept or reject pointed \mathcal{T} -coalgebras (that is, coalgebras with an explicitly specified starting point) on the basis of so-called *acceptance games*. Λ -automata provide the counterpart to the coalgebraic μ -calculus for Λ . In particular, there is a construction transforming a μML_Λ -formula into an equivalent Λ -automaton (of size quadratic in the length of the formula). Hence we may use the theory of Λ -automata in order to obtain results about coalgebraic modal fixpoint logic.

The main technical contribution of this paper concerns a *small model property* for Λ -automata (Theorem 3). We show that any Λ -automaton \mathbb{A} with a non-empty language recognizes a pointed coalgebra (\mathbb{S}, s) that can be obtained from \mathbb{A} via some uniform construction involving a satisfiability game (Definition 7) that we associate with \mathbb{A} . The size of \mathbb{S} is exponential in the size of \mathbb{A} . On the basis of our proof, in Theorem 4 we give a doubly exponential bound on the complexity of the satisfiability problem of μML_Λ -formulas in \mathcal{T} -coalgebras (provided that the one-step satisfiability problem of Λ over \mathcal{T} has a reasonable complexity).

Compared to the work of Cîrstea, Kupke and Pattinson [5], our results are more general in the sense that they do not depend on the existence of a complete tableau calculus. On the other hand, the cited authors obtain a much better complexity result: Under some mild conditions on the efficiency of their complete tableau calculus (conditions that are met by e.g. the modal μ -calculus and the graded μ -calculus), they establish an EXPTIME upper bound for the satisfiability problem of the μ -calculus for Λ . However, in Section 5 below we shall make a connection between our satisfiability game and their tableau game, and on the basis of this connection one may obtain the same complexity bound as in [5] (if one assumes the same conditions on the existence and nature of the tableau system).

2 Preliminaries

We assume familiarity with basic notions from category theory such as categories, functors, natural transformations. We let \mathbf{Set} denote the category with sets as objects and functions as arrows. For convenience, and without loss of generality [2], we assume our functors to be standard i.e. to preserve set inclusions.

Definition 1. *Let $\mathcal{T} : \mathbf{Set} \rightarrow \mathbf{Set}$ be a functor. A \mathcal{T} -coalgebra is a pair (S, σ) where S is a set and σ is a function $\sigma : S \rightarrow \mathcal{T}S$. A morphism of \mathcal{T} -coalgebras from \mathbb{S} to \mathbb{S}' , written $f : \mathbb{S} \rightarrow \mathbb{S}'$, is a function $f : S \rightarrow S'$ such that $T(f)\sigma = \sigma'f$. The size of a coalgebra \mathbb{S} is the cardinality of the set S .*

1. We write $\mathcal{Q} : \mathbf{Set}^{op} \rightarrow \mathbf{Set}$ for the contravariant power set functor, and \mathcal{P} for the covariant power set functor. Coalgebras for \mathcal{P} are Kripke frames [1].
2. The monotone neighborhood functor \mathcal{M} maps a set X to $\mathcal{M}(X) = \{U \in \mathcal{Q}\mathcal{Q}(X) \mid U \text{ is upwards closed}\}$, and a function f to $\mathcal{M}(f) = \mathcal{Q}\mathcal{Q}(f) = (f^{-1})^{-1}$. Coalgebras for this functor are monotone neighborhood frames [7].

3. We write \mathcal{D} for the distribution functor which maps a set X to $\mathcal{D}(X) = \{\mu : X \rightarrow [0, 1] \mid \sum_{x \in X} \mu(x) = 1\}$ and a function f to the function $\mathcal{D}(f) : \mathcal{D}(X) \rightarrow \mathcal{D}(Y)$ which maps a probability distribution μ to $\mathcal{D}(f)(\mu)(y) = \sum_{f(x)=y} \mu(x)$. In this case coalgebras correspond to Markov chains [3].
4. We write \mathcal{B} for the bags, or multiset, functor which maps a set X to $\overline{\mathbb{N}}^X$, where $\overline{\mathbb{N}} = \mathbb{N} + \{\infty\}$, the action on arrows is similar to that of \mathcal{D} . Coalgebras for \mathcal{B} are often referred to as multigraphs [11].

We assume familiarity with the basic notions of the theory of automata and infinite games [6]. Here we fix some notation and terminology.

- Definition 2.** 1. Given a set A , we let A^* and A^ω denote, respectively, the set of words (finite sequences) and streams (infinite sequences) over A . Automata operating on streams will be called stream automata (rather than ω -automata). Given $\pi \in A^* + A^\omega$ we write $\text{Inf}(\pi)$ for the set of elements in A that appear infinitely often in π .
2. A graph game is a tuple $\mathbb{G} = (G_\exists, G_\forall, E, \text{Win})$ where G_\exists and G_\forall are disjoint sets, and (with $G := G_\exists + G_\forall$) E is a subset of G^2 , and Win is a subset of G^ω . In case \mathbb{G} is a parity game, that is, Win is given by a parity function $\Omega : G \rightarrow \mathbb{N}$, we write $\mathbb{G} = (G_\exists, G_\forall, E, \Omega)$.
 3. A strategy for a player P in a game $\mathbb{G} = (G_\exists, G_\forall, E, \text{Win})$ is a map $\alpha : G^* \rightarrow G$. A \mathbb{G} -match $\pi = v_0 v_1 \dots$ is α -conform if $v_i = \alpha(v_0 \dots v_{i-1})$ for all $i > 0$ such that $v_i \in G_\exists$. A strategy α is winning for a player P if all α -conform matches are winning for P .
 4. A strategy α is a finite memory strategy if there exists a finite set M , called the memory set, an element $m_I \in M$ and a map $(\alpha_1, \alpha_2) : G \times M \rightarrow G \times M$ such that for all pairs of sequences $v_0 \dots v_k \in V^*$ and $m_0 \dots m_k \in M^*$ if $m_0 = m_I$, $v_k \in G_\exists$ and $m_{i+1} = \alpha_2(v_i, m_i)$ (for all $i < k$), then $\alpha(v_0 \dots v_k) = \alpha_1(v_k, m_k)$.
 5. A game $\mathbb{G} = (G_\exists, G_\forall, E, \text{Win})$ is called regular if there exists an ω -regular language L over a finite alphabet C , and a map $\text{col} : G \rightarrow C$, such that $\text{Win} = \{v_0 v_1 \dots \in G^\omega : \text{col}(v_0) \text{col}(v_1) \dots \in L\}$.

The following fact on regular games can be proved by putting together various known results from [4] and [8].

Fact 1 Let $\mathbb{G} = (G_\exists, G_\forall, E, \text{Win})$ be a regular game, let $\text{col} : G \rightarrow C$ be a polynomial coloring of G , and let \mathbb{B} be a deterministic parity stream automaton such that $\text{Win} = \{v_0 v_1 \dots \in G^\omega \mid \text{col}(v_0) \text{col}(v_1) \dots \in L(\mathbb{B})\}$. Let n, m , and b be the size of G , E , and \mathbb{B} , respectively, and let d be the index of \mathbb{B} . Then for each player P we may assume winning strategies for P to be finite memory ones, with memory of size b . In addition, the problem, whether a given position $v \in G$ is winning for P , is decidable in time $\mathcal{O}\left(d \cdot m \cdot b \cdot \left(\frac{b \cdot d}{\lfloor d/2 \rfloor}\right)^{\lfloor d/2 \rfloor}\right)$.

3 Automata for the coalgebraic μ -calculus

As mentioned in the introduction, the following notion is fundamental in the development of coalgebraic modal logic.

Definition 3. An n -ary predicate lifting for \mathcal{T} is a natural transformation

$$\lambda : \mathcal{Q}^n \rightarrow \mathcal{QT}.$$

Such a predicate lifting is monotone if for each set S , the operation $\lambda_S : \mathcal{Q}(S)^n \rightarrow \mathcal{Q}(S)$ preserves the (subset) order in each coordinate. The (Boolean) dual of a predicate lifting $\lambda : \mathcal{Q}^n \rightarrow \mathcal{QT}$ is the lifting $\bar{\lambda} : \mathcal{Q}^n \rightarrow \mathcal{QT}$ given by $\bar{\lambda}_S(A_1, \dots, A_n) = S \setminus \lambda(S \setminus A_1, \dots, S \setminus A_n)$.

Predicate liftings allow one to see coalgebras as (polyadic) neighborhood frames. Accordingly, with each n -ary predicate lifting we will associate an n -ary modality \heartsuit_λ . Its semantics in a coalgebra \mathbb{S} is given by the following:

$$\llbracket \heartsuit_\lambda(\phi_1, \dots, \phi_n) \rrbracket_{\mathbb{S}} = \sigma^{-1} \lambda_S(\llbracket \phi_1 \rrbracket_{\mathbb{S}}, \dots, \llbracket \phi_n \rrbracket_{\mathbb{S}}) \quad (1)$$

where we inductively assume that $\llbracket \phi_i \rrbracket_{\mathbb{S}} \subseteq S$ is the meaning of the formula ϕ_i . In words, $\heartsuit_\lambda(\phi_1, \dots, \phi_n)$ is true at a state s iff the unfolding $\sigma(s)$ belongs to the set $\lambda_S(\llbracket \phi_1 \rrbracket_{\mathbb{S}}, \dots, \llbracket \phi_n \rrbracket_{\mathbb{S}})$.

Example 1. (1) In case of the covariant power set functor the predicate lifting given by $\lambda_S(U) = \{V \in \mathcal{P}S \mid V \subseteq U\}$ induces the usual universal modality \Box , i.e. $\llbracket \heartsuit_\lambda \phi \rrbracket_{\mathcal{V}} = \llbracket \Box \phi \rrbracket_{\mathcal{V}}$, on Kripke Frames.

(2) Consider the monotone neighborhood functor. We can obtain the standard modalities as predicate liftings. The universal modality is given by $\lambda_S(U) = \{N \in \mathcal{M}(S) \mid U \in N\}$. In the case we have a coalgebra $\sigma : S \rightarrow \mathcal{M}(S)$ which is given by some topology τ , i.e. $\sigma(s) = \{X \in \mathcal{P}(S) \mid (\exists V \in \tau)(s \in V \subseteq X)\}$, the interpretation of this predicate lifting is $\mathbb{S}, V, s \Vdash \heartsuit_\lambda \phi$ iff $s \in \text{Int}(\llbracket \phi \rrbracket_V)$ which is the usual interpretation of the universal modality on topological models.

(3) Let k be a natural number. A graded modality can be seen as a predicate lifting for the multiset functor; $\lambda_S^k(U) = \{B : S \rightarrow \mathbb{N} \mid \sum_{x \in U} B(x) \geq k\}$. In this case $\mathbb{S}, V, s \Vdash \heartsuit_\lambda^k \phi$ holds iff s has at least k many successors satisfying ϕ .

(4) Let p be an element in the closed interval $[0, 1]$. The following defines a predicate lifting for the distribution functor $\lambda_S^p(U) = \{\mu : S \rightarrow [0, 1] \mid \sum_{x \in U} \mu(x) \geq p\}$. In this case $\mathbb{S}, V, s \Vdash \heartsuit_\lambda^p \phi$ holds if the probability that s has a successor satisfying ϕ is at least p .

(5) Propositional information can be provided by predicate liftings for the functor $\mathcal{P}(\mathbb{P}) \times \mathcal{T}$, where \mathbb{P} is a fixed set of proposition letters. The semantics of the proposition letter $p \in \mathbb{P}$ is given by the predicate liftings $\lambda_S^p(U) = \{(X, t) \in \mathcal{P}(\mathbb{P}) \times \mathcal{T}(S) \mid p \in X\}$, and $\lambda_S^{\neg p}(U) = \{(X, t) \in \mathcal{P}(\mathbb{P}) \times \mathcal{T}(S) \mid p \notin X\}$.

Convention 2 In the remainder of this paper we fix a functor \mathcal{T} on Set , and a set Λ of monotone predicate liftings that we assume to be closed under taking Boolean duals. In case we are dealing with a language containing proposition letters, these are supposed to be encoded in appropriate liftings, as in Example 1(5).

We can now introduce coalgebraic modal fixpoint logic, or the coalgebraic μ -calculus. We fix a set X of variables, and define the set μML_Λ of fixpoint formulas ϕ, ϕ_i as follows:

$$\phi ::= x \in X \mid \perp \mid \top \mid \phi_0 \wedge \phi_1 \mid \phi_0 \vee \phi_1 \mid \heartsuit_\lambda(\phi_0, \dots, \phi_n) \mid \mu x. \phi \mid \nu x. \phi$$

where $\lambda \in \Lambda$. Syntactic notions pertaining to formulas, such as size and alternation depth, are defined as usual.

The semantics of this language is completely standard. Let $\mathbb{S} = (S, \sigma)$ be a \mathcal{T} -coalgebra. Given a valuation $V : X \rightarrow \mathcal{P}(S)$, we define the *meaning* $\llbracket \phi \rrbracket_{\mathbb{S}, V}$ of a formula ϕ by a standard induction which includes the following clauses:

$$\llbracket x \rrbracket_{\mathbb{S}, V} := V(x), \quad \llbracket \mu x. \phi \rrbracket_{\mathbb{S}, V} := \text{LFP}. \phi_x^{\mathbb{S}, V}, \quad \llbracket \nu x. \phi \rrbracket_{\mathbb{S}, V} := \text{GFP}. \phi_x^{\mathbb{S}, V}.$$

Here $\text{LFP}. \phi_x^{\mathbb{S}, V}$ and $\text{GFP}. \phi_x^{\mathbb{S}, V}$ are the least and greatest fixpoint, respectively, of the monotone map $\phi_x^{\mathbb{S}, V} : \mathcal{P}(S) \rightarrow \mathcal{P}(S)$ given by $\phi_x^{\mathbb{S}, V}(A) := \llbracket \phi \rrbracket_{\mathbb{S}, V[x \mapsto A]}$ (with $V[x \mapsto A](x) = A$ and $V[x \mapsto A](y) = V(y)$ for $y \neq x$).

By Convention 2, we may assume that the language μML_Λ contains proposition letters and their negations, and we may see negation itself as a definable connective.

Before we can turn to the definition of our automata we need some preliminary notions. Given a set X , we denote the set of positive propositional formulas, or lattice terms, over X , by $\mathcal{L}_0(X)$:

$$\phi ::= x \in X \mid \perp \mid \top \mid \phi_0 \wedge \phi_1 \mid \phi_0 \vee \phi_1,$$

and we let $\Lambda(X)$ denote the set $\{\heartsuit_\lambda(x_1, \dots, x_n) \mid \lambda \in \Lambda, x_i \in X\}$. Elements of the set $\mathcal{L}_0 \Lambda \mathcal{L}_0(X)$ will be called *depth-one* formulas over X .

Any valuation $V : X \rightarrow \mathcal{P}(S)$ can be extended to a meaning function $\llbracket - \rrbracket_V : \mathcal{L}_0 X \rightarrow \mathcal{P}(S)$ in the usual manner. We write $S, V, s \Vdash \phi$ to indicate $s \in \llbracket \phi \rrbracket_V$. The meaning function $\llbracket - \rrbracket_V$ naturally induces a map $\llbracket - \rrbracket_V^1 : \mathcal{L}_0 \Lambda \mathcal{L}_0(X) \rightarrow \mathcal{P}(\mathcal{T}S)$ interpreting depth-one formulas as subsets of $\mathcal{T}S$. This map is defined inductively, with

$$\llbracket \heartsuit_\lambda(\phi_1, \dots, \phi_n) \rrbracket_V^1 = \lambda_S(\llbracket \phi_1 \rrbracket_V, \dots, \llbracket \phi_n \rrbracket_V) \quad (2)$$

being the clause for the modalities, and with the standard clauses for the boolean connectives. We write $\mathcal{T}S, V, \tau \Vdash^1 \phi$ to indicate $\tau \in \llbracket \phi \rrbracket_V^1$, and refer to this relation as the *one-step semantics*.

We are now ready for the definition of the key structures of this paper, viz., Λ -automata, and their semantics.

Definition 4 (Λ -automata). *A Λ -automaton \mathbb{A} is a quadruple $\mathbb{A} = (A, a_I, \delta, \Omega)$, where A is a finite set of states, $a_I \in A$ is the initial state, $\delta : A \rightarrow \mathcal{L}_0 \Lambda(A)$ is the transition map, and $\Omega : A \rightarrow \mathbb{N}$ is a parity map. The size of \mathbb{A} is defined as its number of states, and its index as the size of the range of Ω .*

The acceptance game of Λ -automata proceeds in *rounds* moving from one basic position in $A \times S$ to another. In each round, at position (a, s) first \exists picks

a valuation V that makes the depth-one formula $\delta(a)$ true at $\sigma(s)$. Looking at this $V : A \rightarrow \mathcal{P}S$ as a binary relation $\{(a', s') \mid s' \in V(a')\}$ between A and S , \forall closes the round by picking an element of this relation.

Definition 5 (Acceptance game). Let $\mathbb{S} = (S, \sigma)$ be a \mathcal{T} -coalgebra and let $\mathbb{A} = (A, a_I, \delta, \Omega)$ be a Λ -automaton. The associated acceptance game $\text{Acc}(\mathbb{A}, \mathbb{S})$ is the parity game given by the table below.

Position	Player	Admissible moves	Priority
$(a, s) \in A \times S$	\exists	$\{V : A \rightarrow \mathcal{P}(S) \mid \mathbb{S}, V, \sigma(s) \Vdash^1 \delta(a)\}$	$\Omega(a)$
$V \in \mathcal{P}(S)^A$	\forall	$\{(a', s') \mid s' \in V(a')\}$	0

A pointed coalgebra (\mathbb{S}, s_0) is accepted by the automaton \mathbb{A} if the pair (a_I, s_0) is a winning position for player \exists in $\text{Acc}(\mathbb{A}, \mathbb{S})$.

As expected, and generalizing the automata-theoretic perspective on the modal μ -calculus as in [6], Λ -automata are the counterpart of the coalgebraic μ -calculus associated with Λ . As a formalization of this we need the following Proposition, the proof of which can be found in the Appendix. Here we say that a Λ -automaton \mathbb{A} is *equivalent* to a formula $\phi \in \mu\text{ML}_\Lambda$ if any pointed \mathcal{T} -coalgebra (\mathbb{S}, s) is accepted by \mathbb{A} iff $\mathbb{S}, s \Vdash \phi$.

Proposition 1. There is an effective procedure transforming a formula ϕ in μML_Λ into an equivalent Λ -automaton \mathbb{A}_ϕ of size dn and index d , where n is the size and d is the alternation depth of ϕ .

4 Finite model property

In this section we show that μML_Λ has the small model property. The key tool in our proof is a satisfiability game that characterizes whether the class of pointed coalgebras accepted by a given Λ -automaton, is empty or not.

Definition 6. Let A be a finite set and Ω a map from A to \mathbb{N} . Given a sequence $R_0 \dots R_k$ in $(\mathcal{P}(A \times A))^*$ the set of traces through $R_0 \dots R_k$ is defined as $\text{Tr}(R_0 \dots R_k) := \{a_0 \dots a_{k+1} \in A^* \mid (a_i, a_{i+1}) \in R_i \text{ for all } i \leq k\}$. Similarly $\text{Tr}(R_0 R_1 \dots) \subseteq A^\omega$ denotes the set of (infinite) traces through $R_0 R_1 \dots$. With $\text{NBT}(A, \Omega)$ we denote the set of $R_0 R_1 \dots \in (\mathcal{P}(A \times A))^\omega$ that contain no bad trace, that is, no trace $a_0 a_1 \dots$ such that $\max\{\Omega(a) \mid a \in \text{Inf}(a_0 a_1 \dots)\}$ is odd.

Definition 7 (Satisfiability game). The satisfiability game $\text{Sat}(\mathbb{A})$ associated with an automaton $\mathbb{A} = (A, a_I, \delta, \Omega)$ is the graph game given by the rules of the tableau below. Here for an element $a \in A$ and for a collection $\mathcal{R} \subseteq \mathcal{P}(A \times A)$, $\zeta^a : A \rightarrow A \times A$ maps b to (a, b) and $U_{\mathcal{R}} : A \times A \rightarrow \mathcal{P}(\mathcal{R})$ denotes the valuation given by $U_{\mathcal{R}}(a, b) = \{R \in \mathcal{R} \mid (a, b) \in R\}$. The range of a relation R is denoted by $\text{Ran}(R)$.

Position	Player	Admissible moves
$R \subseteq A \times A$	\exists	$\{\mathcal{R} \subseteq \mathcal{P}(A \times A) \mid \llbracket \bigwedge \{\zeta^a \delta(a) \mid a \in \text{Ran}(R)\} \rrbracket_{U_{\mathcal{R}}}^1 \neq \emptyset\}$
$\mathcal{R} \subseteq \mathcal{P}(A \times A)$	\forall	$\{R \mid R \in \mathcal{R}\}$

Unless specified otherwise, we assume $\{(a_I, a_I)\}$ to be the starting position of $Sat(\mathbb{A})$. An infinite match $R_0\mathcal{R}_0R_1\dots$ is winning for \exists if $R_0R_1\dots \in NBT(A, \Omega)$.

We leave it for the reader to verify that $Sat(\mathbb{A})$ is a regular game, and that its winning condition is an ω -regular language L of which the complement is recognized by a nondeterministic parity stream automaton of size $|A|$ and index $|Ran(\Omega)|$. So by [9], L is recognized by a deterministic parity stream automaton of size exponential in $|A|$ and index polynomial in $|A|$.

We are now ready to state and prove our main result.

Theorem 3. *Given a A -automaton \mathbb{A} , the following are equivalent.*

- (1) $L(\mathbb{A})$ is not empty.
- (2) \exists has a winning strategy in the game $Sat(\mathbb{A})$.
- (3) $L(\mathbb{A})$ contains a finite pointed coalgebra of size exponential in the size of \mathbb{A} .

Proof. Details for the implication $(1 \Rightarrow 2)$ are in the appendix, and $(3 \Rightarrow 1)$ is immediate. We focus on the hardest implication $(2 \Rightarrow 3)$. Suppose that \exists has a winning strategy in the game $Sat(\mathbb{A}) = (G_\exists, G_\forall, E, Win)$. By the remark following Definition 7 and by Fact 1, we may assume this strategy to use finite memory only: there is a finite set M , $m_I \in M$ and maps $\alpha_1 : G_\exists \times M \rightarrow G$ and $\alpha_2 : G_\exists \times M \rightarrow M$ which satisfy the conditions of Definition 2(3). Moreover, the size of M is at most exponential in the size of \mathbb{A} . Without loss of generality, we may assume that for all $(R, m) \in G_\exists \times M$, $\alpha_2(R, m) = m$.

We denote by W_\exists the set of pairs $(R, m) \in G_\exists \times M$ satisfying the following: For all $Sat(\mathbb{A})$ -matches $R_0\mathcal{R}_0R_1\mathcal{R}_1\dots$ for which there exists a sequence $m_0m_1\dots$ such that $R_0 = R$, $m_0 = m$ and for all $i \in \mathbb{N}$, $\mathcal{R}_i = \alpha_1(R_i, m_i)$, $m_{i+1} = \alpha_2(R_i, m_i)$, we have that $R_0\mathcal{R}_0R_1\mathcal{R}_1\dots$ is won by \exists .

The finite coalgebra in $L(\mathbb{A})$ that we are looking for will have the set $G_\exists \times M$ as its carrier. Therefore we first define a coalgebra map $\xi : G_\exists \times M \rightarrow \mathcal{T}(G_\exists \times M)$. We base this construction on two observations.

First, let (R, m) be an element of W_\exists , and write $\mathcal{R} := \alpha_1(R, m)$; then by the rules of the satisfiability game, there is an object $g(R, m) \in \mathcal{TR}$ such that for every $a \in Ran(\mathcal{R})$, the formula $\zeta^a\delta(a)$ is true at $g(R, m)$ under the valuation $U_{\mathcal{R}}$. Note that $\mathcal{R} \subseteq G_\exists$, and thus we may think of the above as defining a function $g : W_\exists \rightarrow \mathcal{T}G_\exists$. Choosing some dummy values for elements $(R, m) \in (G_\exists \times M) \setminus W_\exists$, the domain of this function can be extended to the full set $G_\exists \times M$. To simplify our notation we will also let g denote the resulting map, with domain $G_\exists \times M$ and codomain $\mathcal{T}G_\exists$. Next, consider the map $add_m : G_\exists \rightarrow G_\exists \times M$, given by $add_m(R) = (R, m)$. Based on this map we define the function $h : \mathcal{T}G_\exists \times M \rightarrow \mathcal{T}(G_\exists \times M)$ such that $h(\tau, m) = \mathcal{T}(add_m)(\tau)$.

We let \mathbb{S} be the coalgebra $(G_\exists \times M, \xi)$, where $\xi : G_\exists \times M \rightarrow \mathcal{T}(G_\exists \times M)$ is the map $\xi := h \circ (g, \alpha_2)$. Observe that the size of \mathbb{S} is at most exponential in the size of \mathbb{A} , since G_\exists is the set $\mathcal{P}(A \times A)$ and M is at most exponential in the size of A . As the designated point of \mathbb{S} we take the pair (R_I, m_I) , where $R_I := \{(a_I, a_I)\}$.

It is left to prove that the pointed coalgebra $(\mathbb{S}, (R_I, m_I))$ is accepted by \mathbb{A} . That is, using \exists 's winning strategy α in the satisfiability game we need to find a

winning strategy for \exists in the acceptance game for the automaton \mathbb{A} with starting position $(a_I, (R_I, m_I))$. We will define this strategy by induction on the length of a partial match, simultaneously setting up a shadow match of the satisfiability game. Inductively we maintain the following relation between the two matches: (*) If $(a_0, (R_0, m_0)), \dots, (a_k, (R_k, m_k))$ is a partial match of the acceptance game (during which \exists plays the inductively defined strategy), then $a_I a_0 \dots a_k$ is a trace through $R_0 \dots R_k$ (and so in particular, a_k belongs to $\text{Ran}(R_k)$), and for all $i \in \{0, \dots, k-1\}$, $R_{i+1} \in \alpha_1(R_i, m_i)$ and $m_{i+1} = \alpha_2(R_i, m_i)$.

Setting up the induction, it is easy to see that the above condition is met at the start $(a_0, (R_0, m_0)) = (a_I, (R_I, m_I))$ of the acceptance match: $a_I a_I$ is the (unique) trace through the one element sequence R_I .

Inductively assume that, with \exists playing as prescribed, the play of the acceptance game has reached position $(a_k, (R_k, m_k))$. By the induction hypothesis, we have $a_k \in \text{Ran}(R_k)$ and the position (R_k, m_k) is a winning position for \exists in the acceptance game. Abbreviate $\mathcal{R} := \alpha_1(R_k, m_k)$ and $n := \alpha_2(R_k, m_k)$. As the next move for \exists we propose the valuation $V : A \rightarrow \mathcal{P}(G_\exists \times M)$ such that $V(a) := \{(R, n) \mid (a_k, a) \in R \text{ and } R \in \mathcal{R}\}$.

Claim. V is a legitimate move at position $(a_k, (R_k, m_k))$.

Proof of Claim. We need to show that $\mathbb{S}, V, (R_k, m_k) \Vdash^1 \delta(a_k)$. First, recall that (R_k, m_k) belongs to W_\exists . Hence, the element $\gamma := g(R_k, m_k)$ of \mathcal{TR} satisfies the formula $\varsigma^{a_k} \delta(a_k)$ under the valuation $U := U_{\mathcal{R}}$ (where $U_{\mathcal{R}}$ is defined as in Definition 7). That is $\mathcal{TR}, U_{\mathcal{R}}, \gamma \Vdash^1 \varsigma^{a_k} \delta(a_k)$. Thus in order to prove the claim it clearly suffices to show that

$$\mathbb{S}, V, (R_k, m_k) \Vdash^1 \phi \text{ iff } \mathcal{TR}, U, \gamma \Vdash^1 \varsigma^{a_k} \phi \quad (3)$$

for all formulas ϕ in $\mathcal{L}_0(A(A))$. The proof of (3) proceeds by induction on the complexity of ϕ . We only consider a simplified version of the base step, where ϕ is of the form $\heartsuit_\lambda a$. We can prove (3) as follows:

$$\begin{aligned} \mathbb{S}, V, (R_k, m_k) \Vdash^1 \heartsuit_\lambda b &\iff \xi(R_k, m_k) \in \lambda_{G_\exists \times M}(\llbracket b \rrbracket_V). && \text{(definition of } \Vdash) \\ &\iff (\mathcal{T} \text{ add}_n)(\gamma) \in \lambda_{G_\exists \times M}(\llbracket b \rrbracket_V). && \text{(definition of } \xi) \\ &\iff \gamma \in (\mathcal{T} \text{ add}_n)^{-1}(\lambda_{G_\exists \times M} \llbracket b \rrbracket_V) && \text{(definition of } (\cdot)^{-1}) \\ &\iff \gamma \in \lambda_{G_\exists}(\text{add}_n^{-1}(\llbracket b \rrbracket_V)) && \text{(naturality of } \lambda) \\ &\iff \gamma \in \lambda_{\mathcal{R}}(\llbracket (a_k, b) \rrbracket_U) && (\ddagger) \\ &\iff \mathcal{TR}, U, \gamma \Vdash^1 \heartsuit_\lambda(a_k, b) && \text{(definition of } \Vdash) \\ &\iff \mathcal{TR}, U, \gamma \Vdash^1 \varsigma_{a_k} \heartsuit_\lambda b && \text{(definition of } \varsigma_{a_k}) \end{aligned}$$

For (\ddagger) , consider the following valuation $U' : A \times A \rightarrow \mathcal{P}(G_\exists)$ such that $U'(a', b') := U(a', b') \cap \mathcal{R}$. It follows from $\mathcal{R} \subseteq G_\exists$ and standardness that $\lambda_{\mathcal{R}} \llbracket a \rrbracket_U = \lambda_{G_\exists} \llbracket a \rrbracket_{U'}$. But then (\ddagger) follows because $\text{add}_n^{-1}(\llbracket b \rrbracket_V) = \llbracket (a, b) \rrbracket_{U'}$; this is more or less routine. *This finishes the proof of the Claim.*

We leave it for the reader to verify that with this definition of a strategy for \exists , the inductive hypothesis (including the relation (*) between the two matches)

remains true. In particular this shows that \exists will never get stuck. Hence in order to verify that the strategy is winning for \exists , we may confine our attention to infinite matches of $\text{Acc}(\mathbb{A}, \mathbb{S})$. Let $\pi = (a_0, (R_0, m_0))(a_1, (R_1, m_1)) \dots$ be such a match, then it follows from (*) that $a_I a_0 a_1 \dots$ is a trace through $R_0 R_1 \dots$, and so we may infer from the assumption that (α_1, α_2) is a winning strategy for \exists in $\text{Sat}(\mathbb{A})$, that $a_I a_0 a_1 \dots$ is not bad. That is, the match π is won by \exists .

Putting this theorem together with Proposition 1, we obtain a small model property for the coalgebraic μ -calculus, for *every* set of predicate liftings.

Corollary 1. *If $\phi \in \mu\text{ML}_\Lambda$ is satisfiable in a \mathcal{T} -coalgebra, it is satisfiable in a \mathcal{T} -coalgebra of size exponential in the size of ϕ .*

Moreover, given some mild condition on Λ and \mathcal{T} , we obtain the following complexity result.

Definition 8. *Given sets A and $\mathcal{X} \subseteq \mathcal{P}A$, let $U_{\mathcal{X}} : A \rightarrow \mathcal{P}\mathcal{X}$ be the valuation given by $U_{\mathcal{X}} : a \mapsto \{B \in \mathcal{X} \mid a \in B\}$. The one-step satisfiability problem for Λ over \mathcal{T} is the problem whether, for fixed A and \mathcal{X} , a given formula ϕ is satisfiable in $\mathcal{T}\mathcal{X}$ under $U_{\mathcal{X}}$.*

Theorem 4. *If Λ has an EXPTIME one-step satisfiability problem over \mathcal{T} , then the satisfiability problem of μML_Λ over \mathcal{T} -coalgebras is decidable in 2EXPTIME.*

Proof. Let ϕ be a given formula in μML_Λ of size n , and let \mathbb{A}_ϕ be the Λ -automaton associated with ϕ , as in Proposition 1. On the basis of the remark following Definition 7, the reader may easily check that $\text{Sat}(\mathbb{A}_\phi)$ is a regular game of size doubly exponential in n , and with a winning condition that is recognizable by an deterministic parity stream automaton of size exponential in n and index polynomial in n . Hence by Theorem 1 the problem of determining the winner of this game can be solved in doubly exponential time.

However, the game $\text{Sat}(\mathbb{A}_\phi)$ has to be *constructed* in doubly exponential time as well. The problem here concerns the complexity of the problem whether a given pair (R, \mathcal{R}) is an edge of the game graph. Under the assumption of the Theorem, this can be done in time doubly exponential in n — note that the *length* of the one-step formulas in the range of the transition function of \mathbb{A}_ϕ may be exponential in n .

5 One-step tableau completeness

In this section we show how our satisfiability game relates to the work of Cirstea, Kupke and Pattinson [5]. We need some definitions — for reasons of space limitations we omit proofs and refer to *opus cit.* for motivation and examples.

Definition 9. *A one-step rule d for Λ is of the form*

$$\frac{\Gamma_0}{\gamma_1 \cdots \gamma_n}$$

where $\Gamma_0 \subseteq_{\omega} \Lambda(X)$ and $\gamma_1, \dots, \gamma_n \subseteq_{\omega} X$, every propositional variable occurs at most once in Γ_0 and all variables occurring in each of the γ_i 's ($i > 0$) also occur in Γ_0 . We write $\text{Conc}(\mathbf{d})$ for the set Γ_0 and $\text{Prem}(\mathbf{d})$ for the set $\{\gamma_i \mid 1 \leq i \leq n\}$.

Given $\Gamma, \{\phi\} \subseteq_{\omega} \mathcal{L}_0\Lambda(X)$, we say that Γ propositionally entails ϕ , notation: $\Gamma \vdash_{PL} \phi$, if there are $\Gamma', \{\phi'\} \subseteq_{\omega} \mathcal{L}_0(Y)$ and a substitution $\tau : Y \rightarrow \Lambda(X)$ such that $\tau[\Gamma'] = \Gamma$, $\tau(\phi') = \phi$ and $\Gamma' \vdash \phi'$ in propositional logic,.

For a set of such rules, with an automaton \mathbb{A} we associate a so-called *tableau game*, in which the rules themselves are part of the game board.

Definition 10. Let $\mathbb{A} = (A, a_I, \delta, \Omega)$ be a Λ -automaton and let \mathbf{D} be a set of one-step rules for Λ . The game $\text{Tab}(\mathbb{A}, \mathbf{D})$ is the two-player graph game given by the table below.

Position	Player	Admissible moves
$R \in \mathcal{P}(A \times A)$	\exists	$\{\Gamma \subseteq_{\omega} \Lambda(A \times A) \mid (\forall a \in \text{ran}(R))(\Gamma \vdash_{PL} \zeta^a \delta(a))\}$
$\Gamma \subseteq_{\omega} \Lambda(A \times A)$	\forall	$\{(\mathbf{d}, \tau) \in \mathbf{D} \times (A \times A)^X \mid \tau[\text{Conc}(\mathbf{d})] \subseteq \Gamma\}$
$(\mathbf{d}, \tau) \in \mathbf{D} \times (A \times A)^X$	\exists	$\{\tau[\gamma] \mid \tau : X \rightarrow A \times A, \gamma \in \text{Prem}(\mathbf{d})\}$

Unless specified differently, the starting position is $\{(a_I, a_I)\}$. An infinite match $R_0\Gamma_0(\mathbf{d}_0, \tau_0)R_1\Gamma_1(\mathbf{d}_1, \tau_1) \dots$ is won by \exists if $R_0R_1 \dots$ belongs to $\text{NBT}(A, \Omega)$.

Given the connection of Proposition 1 between formulas and automata, one may show that our tableau games are virtually the same as the ones in [5]. Our tableau game $\text{Tab}(\mathbb{A}, \mathbf{D})$ is (in some natural sense) *equivalent* to the satisfiability game for \mathbb{A} if we assume the set \mathbf{D} to be *one-step complete* with respect to \mathcal{T} .

Definition 11. A set \mathbf{D} of one-step rules is one-step complete for \mathcal{T} if for any set Y of variables, any set S , $\Gamma \subseteq_{\omega} \Lambda(Y)$ and valuation $V : Y \rightarrow \mathcal{P}(S)$ the following are equivalent:

- (a) $\llbracket \bigwedge \Gamma \rrbracket_V^1 \neq \emptyset$
- (b) for all rules $\mathbf{d} \in \mathbf{D}$ and all substitutions $\tau : X \rightarrow Y$ with $\tau[\text{Conc}(\mathbf{d})] \subseteq_{\omega} \Gamma$, there exists $\gamma_i \in \text{Prem}(\mathbf{d})$ such that $\llbracket \bigwedge \tau[\gamma_i] \rrbracket_V^1 \neq \emptyset$.

The proof of the following equivalence is deferred to the appendix.

Theorem 5. Let \mathbb{A} be a Λ -automaton and let \mathbf{D} be a set of one-step rules for Λ . If \mathbf{D} is one-step complete with respect to \mathcal{T} , then \exists has a winning strategy in $\text{Sat}(\mathbb{A})$ iff \exists has a winning strategy in $\text{Tab}(\mathbb{A}, \mathbf{D})$.

For the purpose of obtaining good complexity results for the coalgebraic μ -calculus, in case we have a nice set \mathbf{D} of derivation rules at our disposal, then the tableau game has considerable advantages over the satisfiability game. The point is that starting from a sentence $\phi \in \mu\text{ML}_A$, the *size* of the game board of $\text{Tab}(\mathbb{A}, \mathbf{D})$ is not necessarily *doubly* exponential in the size of ϕ . If we follow exactly the ideas of [5], with a suitable restriction of \mathbf{D} , some further manipulations may in fact yield a *single exponential* size game board, which may also be constructed in single exponential time (in the size of the original sentence).

More specifically, in our framework of Λ -automata we may prove the main result of [5] stating that if Λ admits a so-called *exponentially tractable, contraction closed* one-step complete set D of rules, then the satisfiability problem for μML_Λ -sentences over \mathcal{T} -coalgebras may be solved in exponential time.

6 Conclusions

In this paper we have introduced Λ -automata which are automata using predicate liftings (Definition 4). We generalize [10] in that our presentation works for any type of coalgebra i.e. no restriction on the functor.

We introduced an acceptance game (Definition 5) for Λ -automata, and established a finite model property (Theorem 3) using a satisfiability game (Definition 7) for Λ -automata. We used these games to establish a 2EXPTIME bound on the satisfiability problem of μML_Λ (Theorem 4).

We showed how our approach relates to the work in [5] by means of a game based on tableau rules (Definition 10, Theorem 5).

There are still some unsolved issues. For instance, it seems worth to investigate the relation between coalgebra automata in [10] and the automata introduced here.

References

1. P. Aczel. *Non-Well-Founded Sets*. CSLI Publications, 1988.
2. J. Adámek and V. Trnková. *Automata and Algebras in Categories*. Kluwer Academic Publishers, 1990.
3. F. Bartels, A. Sokolova, and E. de Vink. A hierarchy of probabilistic system types. *Theoretical Computer Science*, 327:3–22, 2004.
4. J.R. Büchi and L.H. Landweber. Solving sequential conditions by finite state strategies. *Transactions of the American Mathematical Society*, 138:295–311, 1969.
5. C. Cirstea, C. Kupke, and D. Pattinson. EXPTIME tableaux for the coalgebraic μ -calculus. In E. Grädel and R. Kahle, editors, *Computer Science Logic 2009*, volume LNCS 5771, pages 179–193. Springer, 2009.
6. E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logic, and Infinite Games*, volume 2500 of *LNCS*. Springer, 2002.
7. H. Hansen and C. Kupke. A coalgebraic perspective on monotone modal logic. *Electronic Notes in Theoretical Computer Science*, 106:121 – 143, 2004. Proceedings of the Workshop on Coalgebraic Methods in Computer Science (CMCS).
8. M. Jurdziński. Small progress measures for solving parity games. In *Proceedings of the 17th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume LNCS 1770, pages 290–301, 2000.
9. D. Kozen. Results on the propositional μ -calculus. *Theoretical Computer Science*, 27(3):333–354, 1983.
10. C. Kupke and Y. Venema. Coalgebraic automata theory: basic results. *Logical Methods in Computer Science*, 4:1–43, 2008.
11. L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999. (Erratum published *APAL* 99:241–259, 1999.)

12. D. Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoretical Computer Science*, 309:177–193, 2003.
13. N. Piterman. From nondeterministic Büchi and Street automata to deterministic parity automata. In *Proceedings of the Twentyfirst Annual IEEE Symposium on Logic in Computer Science (LICS 2006)*, pages 255–264, 2006.
14. J. Rutten. Universal coalgebra: A theory of systems. *Theoretical Computer Science*, 249:3–80, 2000.
15. L. Schröder and D. Pattinson. Pspace bounds for rank-1 modal logics. *ACM Transactions on Computational Logic*, 10:1–32, 2009.
16. Y. Venema. Automata and fixed point logic: a coalgebraic perspective. *Information and Computation*, 204:637–678, 2006.
17. A. Visser and G. D’Agostino. Finality regained: A coalgebraic study of Scott-sets and multisets. *Archive for Mathematical Logic*, 41:267–298, 2002.

Appendix

Proof (of Proposition 1). The proof of this proposition is more or less routine. The basic idea is to define the intermediate concept of an *unguarded* Λ -automaton $\mathbb{A} = (A, a_I, \delta, \Omega)$ with $\delta : A \rightarrow \mathcal{L}_0(A \cup \Lambda(A))$ admitting unguarded occurrences of A in the unfolding of states. It is completely standard [6] to transform ϕ into such an automaton \mathbb{B}_ξ , with size $\leq n$ and index $\leq d$. Then we use the approach of [10] to transform \mathbb{B}_ξ into an equivalent guarded automaton of size $\leq dn$ and index d .

Proof (Implication (1 \Rightarrow 2) in Theorem 3). Let $\mathbb{S} = (S, \sigma, s_0)$ be a pointed coalgebra accepted by \mathbb{A} . We will show that \exists has a winning strategy in the non-emptiness game $Sat(\mathbb{A})$.

Before we go into the details of this definition, we need some terminology and notation. By assumption, player \exists has a winning strategy α in the acceptance game for the automaton \mathbb{A} with starting position (a_I, s_0) . Since the acceptance game is a parity game, we may assume this strategy to be positional. Given two finite sequences $\mathbf{s} = s_0 \dots s_k \in S^*$ and $\mathbf{a} = a_0 \dots a_k \in A^*$, we say that \mathbf{a} α -corresponds to \mathbf{s} if there is an α -conform partial match which has basic positions $(a_0, s_0) \dots (a_k, s_k)$. The set of all sequences in A^* that α -correspond to \mathbf{s} is denoted as $Corr_\alpha(\mathbf{s})$. Intuitively, this set represents the collection of all α -conform matches passing through \mathbf{s} .

The definition of the winning strategy for \exists in the non-emptiness game $Sat(\mathbb{A})$ will be given by induction on the length of partial matches. Simultaneously we will select, through the coalgebra \mathbb{S} , a path $s_1 s_2 \dots$, which is related to the $Sat(\mathbb{A})$ -match π as follows: At each finite stage $R_0 \mathcal{R}_0 R_1 \dots R_k$ of π , $R_0 = R_I$,

$$\begin{aligned} Tr(R_1 \dots R_k) &\subseteq Corr_\alpha(s_0 \dots s_k) \\ \text{and each } a \in Ran(R_k) &\text{ occurs in some trace through } R_0 \dots R_k \end{aligned} \quad (*)$$

This implies in particular that for each element $a \in Ran(R_k)$, the pair (a, s_k) is a winning position for \exists in the acceptance game.

First, we check that when the satisfiability game starts, condition (*) is satisfied. In this case, we have $R_0 = \{(a_I, a_I)\}$ and \mathbf{s} is the one element sequence s_0 . It is routine to verify that (*) holds.

For the induction step, assume that in the satisfiability game, the partial match $\mathbf{R} = R_0 \mathcal{R}_0 R_1 \dots R_k$ has been played. First we will provide \exists with an appropriate response $\mathcal{R} \subseteq \mathcal{P}(A \times A)$.

Inductively, we have selected a sequence $\mathbf{s} = s_0 s_1 \dots s_k$ satisfying condition (*). Since α is by assumption a *winning* strategy for \exists in the acceptance game, the pair (a, s_k) is a winning position for \exists for each $a \in Ran(R_k)$. This means that \exists 's strategy α will provide her with a collection of valuations $\{V_a : A \rightarrow \mathcal{P}(S) \mid a \in ran(R_k)\}$ such that

$$\mathcal{TS}, V_a, \sigma s_k \Vdash^1 \delta(a). \quad (4)$$

for all $a \in Ran(R_k)$. The collection $\{V_a \mid a \in Ran(R_k)\}$ induces a map

$$f_V : S \rightarrow \mathcal{P}(A \times A)$$

given by

$$f_V(s) := \{(a, b) \in A \times A \mid a \in \text{Ran}(R_k) \text{ and } s \in V_a(b)\}.$$

Define \mathcal{R}_k as the image of S under f_V , that is,

$$\mathcal{R}_k := f_V[S].$$

Thus we may and will see f_V as a surjective map from S to \mathcal{R}_k .

Claim. \mathcal{R}_k is a legitimate move for \exists in $\text{Sat}(\mathbb{A})$ at position R_k .

Proof. To see this, first observe that we have

$$\mathcal{T}f_V : \mathcal{T}S \rightarrow \mathcal{T}\mathcal{R}_k,$$

and so the object $(\mathcal{T}f_V)\sigma s_k$ is indeed a member of the set $\mathcal{T}\mathcal{R}_k$.

Now, in order to prove that \exists may legitimately play \mathcal{R}_k at R_k , it suffices to prove that, for all $a \in \text{Ran}(R_k)$:

$$\mathcal{T}\mathcal{R}_k, U_{\mathcal{R}_k}, (\mathcal{T}h_V)\sigma s_k \Vdash^1 \varsigma_a \delta(a). \quad (5)$$

Fix $a \in \text{Ran}(R_k)$, and abbreviate $U := U_{\mathcal{R}_k}$, where $U_{\mathcal{R}_k}$ is defined as in Definition 7. Given (4), it clearly suffices to prove that

$$\mathcal{T}\mathcal{R}_k, U, (\mathcal{T}f_V)\sigma s_k \Vdash^1 \varsigma_a \phi \text{ iff } \mathcal{T}S, V_a, \sigma s_k \Vdash^1 \phi \quad (6)$$

for all formulas ϕ in $\mathcal{L}_0(\Lambda(A))$. We will prove (6) by induction on the complexity of ϕ .

In the base case we are dealing with a formula $\phi = \heartsuit_\lambda(b_1, \dots, b_n)$. For simplicity however we confine ourselves to the (representative) special case where $n = 1$, and write $b = b_1$. In this setting, (6) follows from the following chain of equivalences:

$$\begin{aligned} \mathcal{T}\mathcal{R}_k, U, (\mathcal{T}f_V)\sigma s_k \Vdash^1 \varsigma_a \phi, &\iff \mathcal{T}\mathcal{R}_k, U, (\mathcal{T}f_V)\sigma s_k \Vdash^1 \heartsuit_\lambda(a, b) \\ &\quad \text{(definition of } \varsigma_a \text{ and } \phi) \\ &\iff (\mathcal{T}f_V)(\sigma s_k) \in \lambda_{\mathcal{R}_k} \llbracket (a, b) \rrbracket_U \quad \text{(definition of } \Vdash) \\ &\iff \sigma s_k \in (\mathcal{T}f_V)^{-1}(\lambda_{\mathcal{R}_k} \llbracket (a, b) \rrbracket_U) \\ &\quad \text{(definition of } (\cdot)^{-1}) \\ &\iff \sigma s_k \in \lambda_S f_V^{-1}(\llbracket (a, b) \rrbracket_U) \quad (\dagger) \\ &\iff \sigma s_k \in \lambda_S(\llbracket b \rrbracket_{V_a}) \quad (\ddagger) \\ &\iff \mathcal{T}S, V_a, \sigma s_k \Vdash^1 \heartsuit_\lambda b \quad \text{(definition of } \Vdash) \end{aligned}$$

Here the step marked (\dagger) follows from λ being a natural transformation, which implies that the following diagram commutes:

$$\square < 1' - 1' - 1'1; 700'400 > [\mathcal{Q}S' \mathcal{Q}\mathcal{T}(S)' \mathcal{Q}\mathcal{R}_k' \mathcal{Q}\mathcal{T}(\mathcal{R}_k); \lambda_S' f_V^{-1'} (\mathcal{T}f_V)^{-1'} \lambda_{\mathcal{R}_k}]$$

The step marked (‡) follows from the identity $\llbracket b \rrbracket_{V_a} = f_V^{-1}(\llbracket (a, b) \rrbracket_U)$, which follows from the following chain of equivalences, all applying to an arbitrary $s \in S$:

$$\begin{aligned}
s \in \llbracket b \rrbracket_{V_a} &\iff s \in V_a(b) && \text{(definition of } \llbracket \cdot \rrbracket \text{)} \\
&\iff (a, b) \in f_V(s) && \text{(definition of } f_V \text{)} \\
&\iff b \in U(f_V(s)) && \text{(definition of } U = U_{\mathcal{R}_k} \text{)} \\
&\iff f_V(s) \in \llbracket (a, b) \rrbracket_U && \text{(definition of } \llbracket \cdot \rrbracket \text{)} \\
&\iff s \in f_V^{-1}(\llbracket (a, b) \rrbracket_U)
\end{aligned}$$

Since the inductive steps in the proof of (6) are completely routine, and therefore, omitted, this finishes the proof of (6), and thus also the proof of the claim.

Given the legitimacy of \mathcal{R}_k as a move for \exists at position R_k , we may propose it as her move in the satisfiability game. Note that this yields the definition of a strategy.

Playing this strategy enables \exists to maintain the inductive condition (*). Indeed, by definition of \mathcal{R}_k , for every $R \in \mathcal{R}_k$ there is an $s_R \in S$ such that $R = f_V(s_R)$. Hence if \forall picks such a relation R , that is putting $R_{k+1} := R$, \exists adds state s_R to her sequence \mathbf{s} , putting $s_{k+1} := s_R$.

To verify that the sequences $R_0 \dots R_{k+1}$ and $s_0 \dots s_{k+1}$ satisfy (*), let $a_0 \dots a_{k+1}$ be a trace through $R_1 \dots R_{k+1}$. Since $R_0 \dots R_k$ and $s_0 \dots s_k$ satisfy (*), there is an α -conform match of the form $(a_0, s_0) \dots (a_k, s_k)$. In this match, when the position (a_k, s_k) is reached, \exists choose a marking $V_{a_k} : A \rightarrow \mathcal{P}(S)$ such that $\mathbb{S}, V_{a_k}, s_k \Vdash^1 \delta(a_k)$. Then, \forall picks a pair (s, a) such that $s \in V_{a_k}(a)$. So in order to show that there is a partial α -conform match of the form $(a_0, s_0) \dots (a_{k+1}, s_{k+1})$, it suffices to prove that $s_{k+1} \in V_{a_k}(a_{k+1})$. Recall that $(a_k, a_{k+1}) \in R_{k+1}$. Since $R_{k+1} = f_V(s_{k+1})$, a_{k+1} belongs to $\text{Ran}(R_k)$ and $a_{k+1} \in V_a(s_{k+1})$, which finishes the proof that the first part of (*) holds for $R_0 \dots R_{k+1}$ and $s_0 \dots s_{k+1}$.

It remains to show that for all $a \in \text{Ran}(R_{k+1})$, a occurs in a trace through $R_0 \dots R_{k+1}$. Fix $a \in \text{ran}(R_{k+1})$. So there exists $a_k \in A$ such that (a_k, a) belongs to R_{k+1} . Since $R_{k+1} = f_V(s_{k+1})$, a_k belongs to $\text{Ran}(R_k)$. Moreover, it follows from the induction hypothesis that if $a \in \text{Ran}(R_k)$, there is a sequence $a_{-1} \dots a_{k-1}$ such that $a_{-1} R_0 a_0 \dots R_k a_k$. Putting this together with $(a_k, a) \in R_{k+1}$, this finishes to prove that a occurs in a trace through $R_0 \dots R_{k+1}$.

Finally we show why this strategy is winning for her in the game $\text{Sat}(\mathbb{A})$, initiated at $\{(a_I, a_I)\}$. Consider an arbitrary match of this game, where \exists plays the strategy as defined above. First, suppose that this match is finite. It should be clear from our definition of \exists 's strategy in $\text{Sat}(\mathbb{A})$ that she never gets stuck. So if the match is finite, \forall got stuck and \exists wins.

In case the match is infinite, \exists has constructed an infinite sequence $\mathbf{s} = s_0 s_1 s_2 \dots$ corresponding to the infinite sequence $\mathbf{R} = R_0 R_1 R_2 \dots$ induced by the $\text{Sat}(\mathbb{A})$ -match. It is easy to see that since the relation (*) holds at each finite level, for every infinite trace $a_0 a_1 a_2 \dots$ through \mathbf{R} there is an α -conform infinite

match of the acceptance game on \mathbb{S} with basic positions $(a_0, s_0)(a_1, s_1) \dots$. Since α was assumed to be a winning strategy, none of these traces is bad. In other words, the sequence \mathbf{R} satisfies the winning condition of $Sat(\mathbb{A})$ for \exists , and thus she is declared to be the winner of the $Sat(\mathbb{A})$ -match. Since we considered an arbitrary match in which she is playing the given strategy, this shows that this strategy is winning, and thus finishes the proof of the implication $(1 \Rightarrow 2)$.

Proof (Of Theorem 5). For the direction from left to right, suppose \exists has a winning strategy α in $Sat(\mathbb{A})$; we define a winning strategy β for \exists in $Tab(\mathbb{A}, \mathbb{D})$. The idea is that during a β -conform match $R_0\Gamma_0(\mathbf{d}_0, \tau_0) \dots R_{k-1}\Gamma_{k-1}(\mathbf{d}_{k-1}, \tau_{k-1})R_k$, \exists will maintain an α -conform shadow match $R'_0\mathcal{R}_0 \dots R'_{k-1}\mathcal{R}_{k-1}R'_k$ such that for all $i \leq k$, $R_i \subseteq R'_i$.

The first position of any β -conform match is $R_0 = \{(a_I, a_I)\}$. The first position of its α -conform shadow match is $R'_0 = \{(a_I, a_I)\}$, hence $R_0 \subseteq R'_0$.

For the induction step, let $\pi = R_0\Gamma_0(\mathbf{d}_0, \tau_0) \dots R_{k-1}\Gamma_{k-1}(\mathbf{d}_{k-1}, \tau_{k-1})R_k$ be a β -conform match and let $\pi' = R'_0\mathcal{R}_0 \dots R'_{k-1}\mathcal{R}_{k-1}R'_k$ be its α -conform shadow match such that for all $i \leq k$, $R_i \subseteq R'_i$. At position R'_k in the α -conform match, suppose that \exists 's choice is the set $\mathcal{R}_k \subseteq \mathcal{P}(A \times A)$. It follows from the rules of $Tab(\mathbb{A}, \mathbb{D})$ that $\llbracket \bigwedge \{\zeta^a \delta(a) \mid a \in \text{Ran}(R_k)\} \rrbracket_{U_{\mathcal{R}_k}}^1 \neq \emptyset$.

Recall that for each $a \in A$ the formula $\zeta^a \delta(a)$ belongs to $\mathcal{L}_0(\mathcal{A}(A))$. Hence the formula $\bigwedge \{\zeta_a \delta(a) \mid a \in \text{Ran}(R_k)\}$ can be rewritten in disjunctive normal form, i.e. as a disjunction of conjunctions of formulas of the form $\heartsuit_\lambda(a, b)$. So there is a conjunct $\bigwedge \Gamma_k$ of this disjunctive normal form such that $\llbracket \bigwedge \Gamma_k \rrbracket_{U_{\mathcal{R}_k}}^1 \neq \emptyset$.

We define the strategy β such that \exists 's next move in π is the set $\Gamma_k \subseteq \text{Up}(A \times A)$. Using basic propositional logic, it follows that for all $a \in \text{ran}(R_k)$, $\Gamma_k \vdash_{PL} \zeta_a \delta(a)$.

Next, in the $Tab(\mathbb{A}, \mathbb{D})$ -match, suppose that \forall plays and chooses a pair (\mathbf{d}_k, τ_k) , where $\mathbf{d}_k \in \mathbb{D}$ and $\tau_k : X \rightarrow A \times A$ are such that $\tau_k[\text{Conc}(\mathbf{d}_k)] \subseteq \Gamma_k$. Now \exists has to find a set $\gamma \in \text{Prem}(\mathbf{d}_k)$; since $\llbracket \bigwedge \Gamma_k \rrbracket_{U_{\mathcal{R}_k}}^1 \neq \emptyset$ and \mathbb{D} is one-step tableau-complete for \mathcal{T} , there exists $\gamma \in \text{Prem}(\mathbf{d}_k)$ such that $\llbracket \bigwedge \tau_k[\gamma] \rrbracket_{U_{\mathcal{R}_k}} \neq \emptyset$. We continue the definition of β by letting \exists choose the set $\tau_k[\gamma]$ as the next position in the $Tab(\mathbb{A}, \mathbb{D})$ -match, i.e. $R_{k+1} = \tau_k[\gamma] \subseteq A \times A$.

Since $\llbracket \bigwedge \tau_k[\gamma] \rrbracket_{U_{\mathcal{R}_k}} \neq \emptyset$, there exists $R'_{k+1} \in \mathcal{R}_k$ such that $\mathcal{R}_k, U_{\mathcal{R}_k}, R'_{k+1} \Vdash \bigwedge \tau_k[\gamma]$. We define the next move for \forall in the α -conform match π' as the relation R'_{k+1} .

Now we check that the induction hypothesis remains true, i.e. $R_{k+1} \subseteq R'_{k+1}$. Fix $(a, b) \in R_{k+1}$. By definition of R_{k+1} , we get that (a, b) belongs to $\tau_k[\gamma]$. Since $\mathcal{R}_k, U_{\mathcal{R}_k}, R'_{k+1} \Vdash \bigwedge \tau_k[\gamma]$, we have that $R'_{k+1} \in U_{\mathcal{R}_k}(a, b)$. Recalling the definition of $U_{\mathcal{R}_k}$ (see Definition 7), we see that (a, b) belongs to R'_{k+1} .

It remains to show that such a strategy β is winning for \exists in $Tab(\mathbb{A}, \mathbb{D})$. From the definition of β , it is straightforward to check that \exists will never get stuck. Hence we may confine our attention to infinite β -conform matches. Let $\pi = R_0\Gamma_0(\mathbf{d}_0, \tau_0)R_1\Gamma_1(\mathbf{d}_1, \tau_1) \dots$ be such a match. By definition of β , there exists an α -conform shadow match $\pi' = R'_0\mathcal{R}_0R'_1\mathcal{R}_1 \dots$ such that for all $i \in \mathbb{N}$, $R_i \subseteq R'_i$. Since α is a winning strategy for \exists , the stream $R'_0R'_1 \dots$ does not contain any bad trace. Putting this together with the fact that $R_i \subseteq R'_i$ (for all

$i \in \mathbb{N}$), we obtain that $R_0R_1 \dots$ does not contain any bad trace; that is, π is won by \exists .

For the direction from right to left, suppose \exists has a winning strategy β in $Tab(\mathbb{A}, \mathbb{D})$. We need to provide a winning strategy α for \exists in $Sat(\mathbb{A})$. During an α -conform match $R_0\mathcal{R}_0 \dots R_{k-1}\mathcal{R}_{k-1}R_k$, \exists will maintain a β -conform shadow match $R_0\Gamma_0(\mathbf{d}_0, \tau_0) \dots R_{k-1}\Gamma_{k-1}(\mathbf{d}_{k-1}, \tau_{k-1})R_k$. The first position of any α -conform match is $R_0 = \{(a_I, a_I)\}$ and the first position of its β -conform shadow match is also R_0 . For the induction step, let $\pi = R_0\mathcal{R}_0 \dots R_{k-1}\mathcal{R}_{k-1}R_k$ be an α -conform match and let $\pi' = R_0\Gamma_0(\mathbf{d}_0, \tau_0) \dots R_{k-1}\Gamma_{k-1}(\mathbf{d}_{k-1}, \tau_{k-1})R_k$ be its β -conform shadow match. In the α -conform match, \exists has to define a set $\mathcal{R}_k \subseteq \mathcal{P}(A \times A)$ such that

$$\llbracket \bigwedge \{ \zeta^a \delta(a) \mid a \in \text{Ran}(R_k) \} \rrbracket_{U_{\mathcal{R}_k}}^1 \neq \emptyset. \quad (7)$$

At position R_k in the β -conform shadow match, \exists chooses a set $\Gamma_k \subseteq Up(A \times A)$ such that $\Gamma_k \vdash_{PL} \bigwedge \{ \zeta^a \delta(a) \mid a \in \text{Ran}(R_k) \}$.

We say that a set $R \subseteq A \times A$ is β -reachable from R_k if there is a β -conform match of the form $R_0\Gamma_0(\mathbf{d}_0, \tau_0) \dots R_{k-1}\Gamma_{k-1}(\mathbf{d}_{k-1}, \tau_{k-1})R_k\Gamma_k(\mathbf{d}, \tau)R$. We define α such that the next move for \exists in the match π is the set $\mathcal{R}_k = \{ R \subseteq A \times A \mid R \text{ is } \beta\text{-reachable from } R_k \}$. To show that such a move is a valid move for \exists in $Sat(\mathbb{A})$, we have to check that (7) holds. Since $\Gamma_k \vdash_{PL} \bigwedge \{ \zeta^a \delta(a) \mid a \in \text{Ran}(R_k) \}$, it is enough to prove that $\llbracket \bigwedge \Gamma_k \rrbracket_{U_{\mathcal{R}_k}}^1 \neq \emptyset$. As \mathbb{D} is one-step tableau-complete for \mathcal{T} , it suffices to verify that for all rules \mathbf{d} in \mathbb{D} and for all maps $\tau : X \rightarrow A \times A$ such that $\tau[\text{Conc}(\mathbf{d})] \subseteq \Gamma_k$, there exists $\gamma \in \text{Prem}(\mathbf{d})$ such that $\llbracket \bigwedge \tau[\gamma] \rrbracket_{U_{\mathcal{R}_k}} \neq \emptyset$.

Fix a rule \mathbf{d} in \mathbb{D} and a map $\tau : X \rightarrow A \times A$ such that $\tau[\text{Conc}(\mathbf{d})] \subseteq \Gamma_k$. Now consider the β -conform match $R_0\Gamma_0(\mathbf{d}_0, \tau_0) \dots R_{k-1}\Gamma_{k-1}(\mathbf{d}_{k-1}, \tau_{k-1})R_k\Gamma_k(\mathbf{d}, \tau)$. In this match, it is \exists 's turn and according to β , she chooses a set $\gamma \in \text{Prem}(\mathbf{d})$, making the relation $\tau[\gamma]$ as the new position. So the relation $\tau[\gamma]$ is β -reachable from R_k ; therefore it belongs to \mathcal{R}_k . To show $\llbracket \bigwedge \tau[\gamma] \rrbracket_{U_{\mathcal{R}_k}} \neq \emptyset$, it is enough to prove that $\mathcal{R}_k, U_{\mathcal{R}_k}, \tau[\gamma] \Vdash \bigwedge \tau\gamma$; this follows from the definition of $U_{\mathcal{R}_k}$.

Next, in the α -conform match, \forall picks a relation $R_{k+1} \in \mathcal{R}_k$. By definition of \mathcal{R}_k , R_{k+1} is β -reachable from R_k . So there is a β -conform shadow match of the form $R_0\Gamma_0(\mathbf{d}_0, \tau_0) \dots R_{k-1}\Gamma_{k-1}(\mathbf{d}_{k-1}, \tau_{k-1})R_k\Gamma_k(\mathbf{d}_k, \tau_k)R_{k+1}$. This finishes the definition of α .

We still have to check that the strategy α is winning for \exists in $Sat(\mathbb{A})$. For this we observe that using α as a strategy \exists will never get stuck. Hence we assume that we are in an infinite α -conform match $\pi = R_0\mathcal{R}_0R_1\mathcal{R}_1 \dots$. By construction of α , there exists a β -conform shadow match $\pi' = R_0\Gamma_0(\mathbf{d}_0, \tau_0)R_1\Gamma_1(\mathbf{d}_1, \tau_1) \dots$. Since π' is won by \exists , $R_0R_1 \dots$ does not contain any bad trace, therefore π is won by \exists .