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 this 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 Cîrstea 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 lifitings, 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 an 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 [14] provides the notion of a *coalgebra* as the natural mathematical generalization of state-based evolving systems such as streams,

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

S. Abramsky et al. (Eds.): ICALP 2010, Part II, LNCS 6199, pp. 381–392, 2010.

(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 space of the coalgebra, and $\sigma : S \longrightarrow \mathcal{T}S$ 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 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 [11] 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 [12] 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 ML_Λ associated with a set Λ 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 [15].

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 [9]. A coalgebraic fixpoint language on the basis of Moss' modality was introduced by Venema [16]. Recently, Cîrstea, Kupke and Pattinson [5] introduced the *coalgebraic μ-calculus* $\mu\mathrm{ML}_\Lambda$ parametrized by a set Λ 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 [16]. Kupke & Venema [10] 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 $\mu\mathrm{ML}_\Lambda$-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 $\mu\mathrm{ML}_\Lambda$-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 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* $\mathcal{T} : \mathsf{Set} \longrightarrow \mathsf{Set}$ *be a functor. A* $\mathcal{T}$-coalgebra *is a pair* (S, σ) *where* S *is a set and* σ *is a function* $\sigma : S \longrightarrow \mathcal{T}S$*. A* morphism of $\mathcal{T}$-coalgebras *from* $\mathbb{S}$ *to* $\mathbb{S}'$*, written* $f : \mathbb{S} \longrightarrow \mathbb{S}'$*, is a function* $f : S \longrightarrow S'$ *such that* $T(f)\sigma = \sigma' f$*. The* size of a coalgebra $\mathbb{S}$ *is the cardinality of the set* S*.*

1. We write $\mathcal{Q} : \mathsf{Set}^{op} \longrightarrow \mathsf{Set}$ 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 $\mathcal{D}(X) = \{\mu : X \longrightarrow [0,1] \mid \sum_{x \in X} \mu(x) = 1\}$ and a function f to the function $\mathcal{D}(f) : \mathcal{D}(X) \longrightarrow \mathcal{D}(Y)$ 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 [17].

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 $Inf(\pi)$ for the set of elements in A that appear infinitely often in π.*

(2) A graph game *is a tuple $\mathbb{G} = (G_\exists, G_\forall, E, Win)$ where $G_\exists$ and $G_\forall$ are disjoint sets, and (with $G := G_\exists + G_\forall$) we have $E \subseteq G^2$, and $Win \subseteq G^\omega$. In case $\mathbb{G}$ is a parity game, that is, Win is given by a parity function $\Omega : G \longrightarrow \mathbb{N}$, we write $\mathbb{G} = (G_\exists, G_\forall, E, \Omega)$.*

(3) A strategy *for a player P in a game $\mathbb{G} = (G_\exists, G_\forall, E, Win)$ is a map $\alpha : G^* \longrightarrow G$. A $\mathbb{G}$-match $\pi = v_0 v_1 \ldots$ is α-*conform *if $v_{i+1} = \alpha(v_0 \ldots v_i)$ for all $i \geq 0$ 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 is a finite* memory set *M, an element $m_I \in M$ and a map $(\alpha_1, \alpha_2) : G \times M \longrightarrow G \times M$ such that for all pairs of sequences $v_0 \ldots v_k \in V^*$ and $m_0 \ldots m_k \in M^*$ 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 $\alpha(v_0 \ldots v_k) = \alpha_1(v_k, m_k)$.*

(5) A game $\mathbb{G} = (G_\exists, G_\forall, E, Win)$ is called regular *if there exists an ω-regular language L over a finite alphabet C, and a map $col : G \longrightarrow C$, such that $Win = \{v_0 v_1 \ldots \in G^\omega : col(v_0) col(v_1) \ldots \in L\}$.*

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

Fact 1. *Let $\mathbb{G} = (G_\exists, G_\forall, E, Win)$ be a regular game, let $col : G \longrightarrow C$ be a coloring of G, and let $\mathbb{B}$ be a deterministic parity stream automaton such that $Win = \{v_0 v_1 \ldots \in G^\omega \mid col(v_0) col(v_1) \ldots \in L(\mathbb{B})\}$. 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 $\mathcal{O}\left(d \cdot m \cdot b \cdot \left(\frac{n \cdot b}{\lfloor d/2 \rfloor}\right)^{\lfloor d/2 \rfloor}\right)$.*

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*

$$\lambda : \mathcal{Q}^n \longrightarrow \mathcal{Q}\mathcal{T}.$$

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

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* $\heartsuit_\lambda$. Its semantics in a coalgebra $\mathbb{S}$ is given by the following:

$$[\![\heartsuit_\lambda(\phi_1, \ldots, \phi_n)]\!]_{\mathbb{S}} = \sigma^{-1}\lambda_S([\![\phi_1]\!]_{\mathbb{S}}, \ldots, [\![\phi_n]\!]_{\mathbb{S}}) \tag{1}$$

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

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. $[\![\heartsuit_\lambda \phi]\!]^\sigma_V = [\![\Box\phi]\!]^\sigma_V$, on Kripke Frames.

(2) Consider the monotone neighborhood functor. We can obtain the standard modalities as predicate liftings. The universal modality is given by $\lambda_S(U) = \{N \in \mathcal{M}(S) \mid U \in N\}$.

(3) Let k be a natural number. A graded modality can be seen as a predicate lifting for the multiset functor; $\lambda^k_S(U) = \{B : S \longrightarrow \mathbb{N} \mid \sum_{x \in U} B(x) \geq k\}$. In this case $\mathbb{S}, V, s \Vdash \heartsuit^k_\lambda \phi$ 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 $\lambda^p_S(U) = \{\mu : S \longrightarrow [0, 1] \mid \sum_{x \in U} \mu(x) \geq p\}$. In this case $\mathbb{S}, V, s \Vdash \heartsuit^p_\lambda \phi$ 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 $\mathcal{P}(\mathsf{P}) \times \mathcal{T}$, where P is a fixed set of proposition letters. The semantics of the proposition letter $p \in \mathsf{P}$ is given by the predicate liftings $\lambda^p_S(U) = \{(X, t) \in \mathcal{P}(\mathsf{P}) \times \mathcal{T}(S) \mid p \in X\}$, and $\lambda^{\neg p}_S(U) = \{(X, t) \in \mathcal{P}(\mathsf{P}) \times \mathcal{T}(S) \mid p \notin X\}$.

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 $\mu\mathrm{ML}_\Lambda$ of fixpoint formulas ϕ, ϕ_i as follows:

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

where $\lambda \in \Lambda$. Syntactic notions pertaining to formulas, such as alternation depth, are defined as usual. The size of a formula is defined as its length (with the proviso that if Λ is infinite, to each occurrence of a modality we add a weight associated with its index, as in [15]).

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

$$[\![x]\!]_{\mathbb{S},V} := V(x), \quad [\![\mu x.\phi]\!]_{\mathbb{S},V} := \mathsf{LFP}.\phi_x^{\mathbb{S},V}, \quad [\![\nu x.\phi]\!]_{\mathbb{S},V} := \mathsf{GFP}.\phi_x^{\mathbb{S},V}.$$

Here $\mathsf{LFP}.\phi_x^{\mathbb{S},V}$ and $\mathsf{GFP}.\phi_x^{\mathbb{S},V}$ are the least and greatest fixpoint, respectively, of the monotone map $\phi_x^{\mathbb{S},V} : \mathcal{P}(S) \longrightarrow \mathcal{P}(S)$ given by $\phi_x^{\mathbb{S},V}(A) := [\![\phi]\!]_{\mathbb{S},V[x\mapsto A]}$ (with $V[x \mapsto A](x) = A$ and $V[x \mapsto A](y) = V(y)$ for $y \neq x$). For *sentences*, that is, formulas without free variables, the valuation does not matter; we write $\mathbb{S}, s \Vdash \phi$ iff $s \in [\![\phi]\!]_{\mathbb{S},V}$ for some/any valuation V.

By Convention 2, we may assume that the language $\mu\mathrm{ML}_\Lambda$ 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 $\mathcal{L}_0(X)$:

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

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

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

$$[\![\heartsuit_\lambda(\phi_1, \ldots, \phi_n)]\!]_V^1 = \lambda_S([\![\phi_1]\!]_V, \ldots, [\![\phi_n]\!]_V) \qquad (2)$$

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 $\tau \in [\![\phi]\!]_V^1$, 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, $\delta : A \longrightarrow \mathcal{L}_0\Lambda(A)$ is the transition map, and $\Omega : A \longrightarrow \mathbb{N}$ 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 $V : A \longrightarrow \mathcal{P}S$ 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 $Acc(\mathbb{A}, \mathbb{S})$ is the parity game given by the table below.*

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

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 is routine, and deferred to the Appendix. Here we say that a Λ-automaton $\mathbb{A}$ is *equivalent* to a sentence $\phi \in \mu\mathrm{ML}_\Lambda$ 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 sentence ϕ in $\mu\mathrm{ML}_\Lambda$ into an equivalent Λ-automaton $\mathbb{A}_\phi$ 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 $\mu\mathrm{ML}_\Lambda$ 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 $Tr(R_0 \dots R_k) := \{a_0 \dots a_{k+1} \in A^* \mid (a_i, a_{i+1}) \in R_i$ for all $i \leq k.\}$. Similarly $Tr(R_0 R_1 \dots) \subseteq A^\omega$ denotes the set of (infinite) traces through $R_0 R_1 \dots$. With $NBT(A, \Omega)$ 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* $\max\{\Omega(a) \mid a \in Inf(a_0 a_1 \dots)\}$ *is odd.*

Definition 7 (Satisfiability game). *The satisfiability game $Sat(\mathbb{A})$ 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 $\mathcal{R} \subseteq \mathcal{P}(A \times A)$, $\varsigma^a : A \longrightarrow A \times A$ maps b to (a, b) and $U_\mathcal{R} : A \times A \longrightarrow \mathcal{P}(\mathcal{R})$ denotes the valuation given by such that $U_\mathcal{R}(a, b) = \{R \in \mathcal{R} \mid (a, b) \in R\}$. The range of a relation R is denoted by $Ran(R)$.*

Position	*Player*	*Admissible moves*
$R \subseteq A \times A$	$\exists$	$\{\mathcal{R} \in \mathcal{PP}(A \times A) \mid [\![\bigwedge\{\varsigma^a\delta(a) \mid a \in Ran(R)\}]\!]^1_{U_\mathcal{R}} \neq \emptyset\}$
$\mathcal{R} \in \mathcal{PP}(A \times A)$	$\forall$	$\{R \mid R \in \mathcal{R}\}$

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 [13], 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 Λ-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 $\alpha_1 : G_\exists \times M \longrightarrow G$ and $\alpha_2 : G_\exists \times M \longrightarrow M$ 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 $m_0m_1\dots$ with $R_0 = R, m_0 = m$ 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 $\xi : G_\exists \times M \longrightarrow \mathcal{T}(G_\exists \times M)$. 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 $g(R, m) \in \mathcal{T}\mathcal{R}$ such that for every $a \in Ran(R)$, the formula $\varsigma^a\delta(a)$ 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 $g : W_\exists \longrightarrow \mathcal{T}G_\exists$. 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 $\mathcal{T}G_\exists$. Second, consider the map $add_m : G_\exists \longrightarrow G_\exists \times M$, given by $add_m(R) = (R, m)$. Based on this map we define the function $h : \mathcal{T}(G_\exists) \times M \longrightarrow \mathcal{T}(G_\exists \times M)$ such that $h(\tau, m) = \mathcal{T}(add_m)(\tau)$.

We let $\mathbb{S}$ be the coalgebra $(G_\exists \times M, \xi)$, where $\xi : G_\exists \times M \longrightarrow \mathcal{T}(G_\exists \times M)$ 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 $(a_0, (R_0, m_0)), \ldots, (a_k, (R_k, m_k))$ is a partial match of the acceptance game (during which $\exists$ plays the inductively defined strategy), then $a_I a_0 \ldots a_k$ is a trace through $R_0 \ldots R_k$ (and so in particular, a_k belongs to $Ran(R_k)$), and for all $i \in \{0, \ldots, k-1\}$, $R_{i+1} \in \alpha_1(R_i, m_i)$ 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 $a_k \in Ran(R_k)$ 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 $V : A \longrightarrow \mathcal{P}(G_\exists \times M)$ 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

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

for all formulas ϕ in $\mathcal{L}_0(\Lambda(A))$. 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 (recall that $n = \alpha_2(R_k, m_k)$):

$$\begin{aligned}
\mathbb{S}, V, (R_k, m_k) \Vdash^1 \heartsuit_\lambda b &\iff \xi(R_k, m_k) \in \lambda_{G_\exists \times M}([\![b]\!]_V). && \text{(definition of } \Vdash^1\text{)}\\
&\iff (\mathcal{T}add_n)(\gamma) \in \lambda_{G_\exists \times M}([\![b]\!]_V). && \text{(definition of } \xi\text{)}\\
&\iff \gamma \in (\mathcal{T}add_n)^{-1}(\lambda_{G_\exists \times M}[\![b]\!]_V) && \text{(definition of } (\cdot)^{-1}\text{)}\\
&\iff \gamma \in \lambda_{G_\exists}(add_n^{-1}([\![b]\!]_V)) && \text{(naturality of } \lambda\text{)}\\
&\iff \gamma \in \lambda_\mathcal{R}([\![(a_k, b)]\!]_U) && (\ddagger)\\
&\iff \mathcal{TR}, U, \gamma \Vdash^1 \heartsuit_\lambda(a_k, b) && \text{(definition of } \Vdash^1\text{)}\\
&\iff \mathcal{TR}, U, \gamma \Vdash^1 \varsigma^{a_k}\heartsuit_\lambda b && \text{(definition of } \varsigma^{a_k}\text{)}
\end{aligned}$$

For $(\ddagger)$, consider the following valuation $U' : A \times A \longrightarrow \mathcal{P}(G_\exists)$ such that $U'(a', b') := U(a', b') \cap \mathcal{R}$. It follows from $\mathcal{R} \subseteq G_\exists$ and standardness that $\lambda_\mathcal{R}[\![a]\!]_U = \lambda_{G_\exists}[\![a]\!]_{U'}$. But then $(\ddagger)$ follows because $add_n^{-1}([\![b]\!]_V) = [\![(a, b)]\!]_{U'}$, which holds by a relatively routine proof. *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 $Acc(\mathbb{A}, \mathbb{S})$. Let $\pi = (a_0, (R_0, m_0))(a_1, (R_1, m_1))\ldots$ be such a match, then it follows from (*) that $a_I a_0 a_1 \ldots$ is a trace through $R_0 R_1 \ldots$, and so we may infer from the assumption that (α_1, α_2) is a winning strategy for $\exists$ in $Sat(\mathbb{A})$, that $a_I a_0 a_1 \ldots$ is not bad. This means that 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 $\phi \in \mu\mathrm{ML}_\Lambda$ 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 $U_\mathcal{X} : A \longrightarrow \mathcal{P}\mathcal{X}$ 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 $\mu\mathrm{ML}_\Lambda$ over $\mathcal{T}$-coalgebras is decidable in 2EXPTIME.*

Proof. Let ϕ be a given sentence in $\mu\mathrm{ML}_\Lambda$ of size n, and let $\mathbb{A}_\phi$ 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 $Sat(\mathbb{A}_\phi)$ is a regular game of size doubly exponential in n, and with a winning condition that is recognizable by a deterministic parity stream automaton of size exponential in n and index polynomial in n. Hence by Fact 1 the problem of determining the winner of this game can be solved in doubly exponential time.

However, the game $Sat(\mathbb{A}_\phi)$ 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 $\mathbb{A}_\phi$ may be exponential in n.

5 One-Step Tableau Completeness

In this section we show how our satisfiability game relates to the work of Cîrstea, 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* $\gamma_1, \ldots, \gamma_n \subseteq_\omega X$*, 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* $Conc(\mathsf{d})$ *for the set* Γ_0 *and* $Prem(\mathsf{d})$ *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* $\tau : Y \longrightarrow \Lambda(X)$ *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* D *be a set of one-step rules for* Λ*. The game* $Tab(\mathbb{A}, \mathsf{D})$ *is the two-player graph game given by the table below.*

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

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

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 $Tab(\mathbb{A}, \mathsf{D})$ is (in some natural sense) *equivalent* to the satisfiability game for $\mathbb{A}$ if we assume the set D to be *one-step complete* with respect to $\mathcal{T}$.

Definition 11. *A set* 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* $V : Y \longrightarrow \mathcal{P}(S)$ *the following are equivalent:*
(a) $[\![\bigwedge \Gamma]\!]^1_V \neq \emptyset$
(b) for all rules $\mathsf{d} \in \mathsf{D}$ *and all substitutions* $\tau : X \longrightarrow Y$ *with* $\tau[Conc(d)] \subseteq_\omega \Gamma$*, there exists* $\gamma_i \in Prem(\mathsf{d})$ *such that* $[\![\bigwedge \tau[\gamma_i]]\!]^1_V \neq \emptyset$*.*

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

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

For the purpose of obtaining good complexity results for the coalgebraic μ-calculus, in case we have a nice set 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 $\phi \in \mu\mathrm{ML}_\Lambda$, the *size* of the game board of $Tab(\mathbb{A}, \mathsf{D})$ is not necessarily *doubly* exponential in the size of ϕ. If we follow exactly the ideas of [5], with a suitable restriction of 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 $\mu\mathrm{ML}_\Lambda$-sentences over $\mathcal{T}$-coalgebras may be solved in exponential time.

References

1. Aczel, P.: Non-Well-Founded Sets. CSLI Publications, Stanford (1988)
2. Adámek, J., Trnková, V.: Automata and Algebras in Categories. Kluwer Academic Publishers, Dordrecht (1990)
3. Bartels, F., Sokolova, A., de Vink, E.: A hierarchy of probabilistic system types. Theoretical Computer Science 327, 3–22 (2004)
4. Büchi, J.R., Landweber, L.H.: Solving sequential conditions by finite state strategies. Transactions of the American Mathematical Society 138, 295–311 (1969)
5. Cîrstea, C., Kupke, C., Pattinson, D.: EXPTIME tableaux for the coalgebraic μ-calculus. In: Grädel, E., Kahle, R. (eds.) CSL 2009. LNCS, vol. 5771, pp. 179–193. Springer, Heidelberg (2009)
6. Grädel, E., Thomas, W., Wilke, T. (eds.): Automata, Logics, and Infinite Games. LNCS, vol. 2500. Springer, Heidelberg (2002)
7. Hansen, H., Kupke, C.: 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. Jurdziński, M.: Small progress measures for solving parity games. In: Reichel, H., Tison, S. (eds.) STACS 2000. LNCS, vol. 1770, pp. 290–301. Springer, Heidelberg (2000)
9. Kozen, D.: Results on the propositional μ-calculus. Theoretical Computer Science 27(3), 333–354 (1983)
10. Kupke, C., Venema, Y.: Coalgebraic automata theory: basic results. Logical Methods in Computer Science 4, 1–43 (2008)
11. Moss, L.: Coalgebraic logic. Annals of Pure and Applied Logic 96, 277–317 (1999); Erratum published APAL 99, 241–259, 1999
12. Pattinson, D.: Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. Theoretical Computer Science 309, 177–193 (2003)
13. Piterman, N.: From nondeterministic Büchi and Streett automata to deterministic parity automata. In: Proceedings of the Twentyfirst Annual IEEE Symposium on Logic in Computer Science (LICS 2006), pp. 255–264. IEEE Computer Society, Los Alamitos (2006)
14. Rutten, J.: Universal coalgebra: A theory of systems. Theoretical Computer Science 249, 3–80 (2000)
15. Schröder, L., Pattinson, D.: PSPACE bounds for rank–1 modal logics. ACM Transactions on Computational Logic 10, 1–32 (2009)
16. Venema, Y.: Automata and fixed point logic: a coalgebraic perspective. Information and Computation 204, 637–678 (2006)
17. Visser, A., D'Agostino, G.: Finality regained: A coalgebraic study of Scott-sets and multisets. Archive for Mathematical Logic 41, 267–298 (2002)