

Closure properties of coalgebra automata

Clemens Kupke
CWI Amsterdam &
Universiteit van Amsterdam, ILLC
Plantage Muidergracht 24
1018TV Amsterdam, Netherlands
ckupke@science.uva.nl

Yde Venema
Universiteit van Amsterdam, ILLC
Plantage Muidergracht 24
1018TV Amsterdam, Netherlands
yde@science.uva.nl

Abstract

We generalize some of the central results in automata theory to the abstraction level of coalgebras.

In particular, we show that for any standard, weak pullback preserving functor F , the class of recognizable languages of F -coalgebras is closed under taking unions, intersections and projections. Our main technical result concerns a construction which transforms a given alternating F -automaton into an equivalent non-deterministic one.

1. Introduction

An important branch of automata theory, itself one of the classical subdisciplines of computer science, concerns the study of finite automata as devices for classifying infinite, or possibly infinite, objects. This perspective on finite automata has found important applications in areas of computer science where one investigates the ongoing behavior of non-terminating programs such as operating systems. As an example we mention the automata-based verification method of *model checking* [4]. This research also has a long and strong theoretical tradition, in which an extensive body of knowledge has been developed, with a number of landmark results. Many of these link the field to neighboring areas such as logic and game theory, see [6] for an overview. The outstanding example here is of course Rabin's decidability theorem [10] for the monadic second order logic of trees; to mention a more recent result, Janin & Walukiewicz [7] identified the modal μ -calculus

as the bisimulation invariant fragment of the monadic second order logic of labelled transition systems.

An interesting phenomenon in automata theory is that most (but not all) key results hold for word and tree automata alike, and that many can even be formulated and proved for automata that operate on yet other objects such as trees of unbounded branching degree, or labelled transition systems. This applies for instance to various closure properties of the class of recognizable languages, and to the fact that alternating automata can be transformed into equivalent non-deterministic ones. These observations naturally raise the question, whether these results can perhaps be formulated at a more general level of abstraction. Of course, such a universal approach towards automata theory would first of all require the introduction of an abstract notion that generalizes structures like words, trees and transition systems. Fortunately, such an abstract notion already exists in the form of *coalgebra*.

The theory of universal coalgebra (see [11] for an overview) seeks to provide a general framework for the study of notions related to (possibly infinite) behavior, such as invariance and observational indistinguishability (bisimilarity, in most cases). Intuitively, coalgebras (as objects) are simple but fundamental mathematical structures that capture the essence of dynamics. In this paper we will restrict our attention to *systems*; these are state-based coalgebras consisting of a set S and a map $S \rightarrow FS$, where F is some set functor determining the *signature* of the coalgebra. The general theory of coalgebra has already developed some general tools for the specification of properties of coalgebras. In particular, start-

ing with Moss’ coalgebraic logic [8], several logical languages have been proposed, usually with a strong modal flavor. Most of these languages are not designed for talking about *ongoing* behaviour, but in [13], the second author introduced a coalgebraic fixed point logic that does enable specifications of this kind.

The same paper also introduces, for coalgebras over a standard set functor F that preserves weak pullbacks, the notion of an *F-automaton* — we will recall the definition in section 2. These automata provide a common generalization of the familiar automata that operate on specific coalgebras such as words, trees or graphs. They also come in various shapes and kinds, the most important distinction being between alternating, non-deterministic, and deterministic ones, respectively.

Basically, F -automata are meant to either accept or reject pointed coalgebras (that is, pairs (\mathbb{S}, s) consisting of an F -coalgebra \mathbb{S} together with a selected state s in the carrier S of \mathbb{S}), and thus *express properties* of states in F -coalgebras. This makes them very similar to formulas, and explains the close connection with coalgebraic (fixed point) logic. Another important aspect of F -automata involves game theory: the criterion under which an F -automaton accepts or rejects a pointed coalgebra is formulated in terms of an infinite two-player parity game which enjoys history-free determinacy.

The aim of developing this coalgebraic framework is not so much to introduce new ideas in automata theory, as to provide a common generalization for existing notions that are well known from the theory of more specific automata. Apart from its general mathematical interest, this abstract approach may be motivated from various sources. To start with, the abstract perspective may be of help to find the right notion of automaton for other kinds of coalgebras, besides the well known kinds like words and trees. It may also be used to prove interesting results on coalgebraic logics — we will briefly come back to this in section 5.

It is the aim of the present paper to provide further motivation for taking a coalgebraic perspective on automata, by showing that some of the key results in automata theory can in fact be lifted to this more abstract level. In particular, this allows for *uniform* proofs of these results, which in its turn may lead to a better understanding of automata theory as such. The

concrete results that we prove concern the relation between alternating and non-deterministic automata, and some of the closure properties that one may associate with automata. For a proper formulation, we need to develop some terminology.

A class of pointed F -coalgebras will be referred to as an *F-language*. Such a language L is *recognized by an F-automaton* \mathbb{A} if a pointed F -coalgebra belongs to L if and only if it is accepted by \mathbb{A} , and (*non-deterministically*) *recognizable* if it is recognized by some (non-deterministic) F -automaton. Our main technical result can now be formulated as follows.

Theorem 1 *Let F be some standard set functor that preserves weak pullbacks. Then every F -automaton has a non-deterministic equivalent. Hence, an F -language is recognizable iff it is non-deterministically recognizable.*

In order to discuss closure properties, let \mathcal{O} be some operation on F -languages, then we say that a class of languages is closed under \mathcal{O} if we obtain a language from this class whenever we apply \mathcal{O} to a family of languages from the class. For example, one may easily prove that recognizable F -languages are closed under taking intersection and union; with some more effort we will show that the class of non-deterministically recognizable F -languages is closed under projection. Theorem 1 allows us to strengthen the above list of closure properties as follows.

Theorem 2 *Let F be some standard set functor that preserves weak pullbacks. Then the class of recognizable F -languages is closed under union, projections and intersection.*

Conspicuously absent in this list is closure under complementation — we will come back to this in section 5.

Finally, our *proofs* for these results are of course built on generalizations, to the coalgebraic level, of (well) known ideas from the theory of specific automata. This applies in particular to results on graph automata [7] and the abstract universal algebraic approach of [2].

2. Preliminaries

We presuppose familiarity with the basic concepts of universal coalgebra [11] and automata theory [6]. Here we fix some notation and terminology.

2.1. Set-functors

We let Set denote the category of sets and functions and Rel the category of sets and binary relations. A *relator*, that is, a functor $Q : \text{Rel} \rightarrow \text{Rel}$, *extends* a functor $F : \text{Set} \rightarrow \text{Set}$ if $QS = FS$ for any object (set) S and $Q(\text{Gr}(f)) = \text{Gr}(Ff)$ for any arrow (function) f ; here $\text{Gr}(f)$ denotes the graph of f . As shown by Trnková [12], functors that preserve weak pullbacks can always be extended to a relator. A second restriction on our functor F is *standardness*, i.e., F is assumed to preserve inclusions: if $A' \subseteq A$ then $FA' \subseteq FA$. For proofs of the following facts we refer to [3] and references therein.

Fact 2.1 *Let $F : \text{Set} \rightarrow \text{Set}$ be a standard, weak pullback preserving functor. Then F has a unique extension $\bar{F} : \text{Rel} \rightarrow \text{Rel}$; this relator satisfies, in addition:*

- (1) $\bar{F}(R^\circ) = (\bar{F}R)^\circ$ ($(\cdot)^\circ$ denotes converse),
- (2) \bar{F} is monotone: if $R \subseteq S$ then $\bar{F}(R) \subseteq \bar{F}(S)$,
- (3) If $A \subseteq A'$, $B \subseteq B'$ and $R \subseteq A' \times B'$, then $\bar{F}(R|_{A \times B}) = \bar{F}(R)|_{FA \times FB}$.

Example 2.2 The power set functor \mathcal{P} is weak pullback preserving and standard. The filter functor \mathcal{F} , mapping a set S to the set of all filters over S and a function f to the function $(f^{-1})^{-1}$, preserves weak pullbacks but it is not standard. The functor, mapping a set S to the set of all upward closed subsets of $(\mathcal{P}S, \subseteq)$ and a function to $(f^{-1})^{-1}$, is neither weak pullback preserving nor standard.

From now on all functors appearing in the paper will be standard and preserve weak pullbacks.

2.2. Coalgebra automata

Let $F : \text{Set} \rightarrow \text{Set}$ be a standard weak pullback preserving functor. We will briefly recall the definition of an F-automaton as introduced in [13].

Definition 2.3 An (alternating) F-*automaton* is a quadruple $\mathbb{A} = (A, a_I, \Delta, \Omega)$ with A some finite set of objects called *states*, $a_I \in A$ the *initial state*, $\Delta : A \rightarrow \mathcal{P}PFA$ the *transition function* and $\Omega : A \rightarrow \omega$ a parity map. An F-automaton is called *non-deterministic* if all members of each $\Delta(a)$ are singleton sets. \triangleleft

F-automata are designed to accept or reject pointed F-coalgebras. The acceptance condition is formulated in terms of a parity game [6].

Definition 2.4 Let $\mathbb{A} = (A, a_I, \Delta, \Omega)$ be an F-automaton, and let $\mathbb{S} = (S, \sigma)$ be an F-coalgebra. The *acceptance game* $\mathcal{G}(\mathbb{A}, \mathbb{S})$ associated with \mathbb{A} and \mathbb{S} is the parity graph game $(B_\exists, B_\forall, E, \bar{\Omega})$ with

$$\begin{aligned} B_\exists &:= S \times A \quad \cup \quad S \times FA \\ B_\forall &:= S \times PFA \quad \cup \quad \mathcal{P}(S \times A), \end{aligned}$$

while E and $\bar{\Omega}$ are given in Table 1. The set of winning positions for \exists in this game is denoted as $\text{Win}_\exists(\mathcal{G}(\mathbb{A}, \mathbb{S}))$, or Win_\exists if no confusion is likely. \mathbb{A} *accepts* the pointed F-coalgebra (\mathbb{S}, s) if $(s, a_I) \in \text{Win}_\exists$. \triangleleft

Remark 2.5 It is clear from the definition of $\bar{\Omega}$ that only the *basic* positions of a match, i.e., positions of the form $(s, a) \in S \times A$, are relevant to determine the winner of the match. Accordingly, in the sequel we will frequently represent a match of the game with the sequence of basic positions visited during the match.

Parity games are known to enjoy a strong form of determinacy: in any position of the game board either \exists or \forall has a history-free winning strategy. Therefore we can focus on \exists 's history-free strategies.

Definition 2.6 Given an F-coalgebra (S, σ) and an F-automaton \mathbb{A} a *strategy* of \exists is a pair of functions

$$(\Phi : S \times A \rightarrow PFA, Z : S \times FA \rightarrow \mathcal{P}(S \times A)).$$

Such a strategy is *legitimate* at a position if it maps the position to an admissible next position. A partial strategy of the kind $\Phi : S \times A \rightarrow PFA$ will often be represented as a map $\Phi : S \rightarrow (PFA)^A$; values of this map will be denoted as Φ_s , etc. \triangleleft

In the paper we will also need a slight variation of an F-automaton. For any set C , a C -coloring of a coalgebra \mathbb{S} is a map $\gamma : S \rightarrow C$; the C -colored F-coalgebra $\mathbb{S} \oplus \gamma := \langle S, \gamma, \sigma \rangle$ can be identified with the $C \times$ F-coalgebra $\langle S, \langle \gamma, \sigma \rangle \rangle$.

Definition 2.7 Let C be a finite set. A C -*chromatic* F-*automaton* is a quintuple $\mathbb{A} = (A, a_I, C, \Delta, \Omega)$ such that $\Delta : A \times C \rightarrow \mathcal{P}PFA$ (and A, a_I , and Ω are as before). Given such an automaton and a

Position: b	$P(b)$	Admissible moves: $E[b]$	$\overline{\Omega}(b)$
$(s, a) \in S \times A$	\exists	$\{(s, \Phi) \in S \times \mathcal{P}(FA) \mid \Phi \in \Delta(a)\}$	$\Omega(a)$
$(s, \Phi) \in S \times \mathcal{P}(FA)$	\forall	$\{(s, \varphi) \in S \times FA \mid \varphi \in \Phi\}$	0
$(s, \varphi) \in S \times FA$	\exists	$\{Z \in \mathcal{P}(S \times A) \mid (\sigma(s), \varphi) \in \overline{F}Z\}$	0
$Z \in \mathcal{P}(S \times A)$	\forall	Z	0

Table 1. Acceptance game for an F-automaton

$C \times F$ -coalgebra $\mathbb{S} = (S, \gamma, \sigma)$, the acceptance game $\mathcal{G}_C(\mathbb{A}, \mathbb{S})$ is defined as the acceptance game for F-automata with the only difference that \exists has to move from a position (s, a) to a position (s, Φ) such that $\Phi \in \Delta(a, \gamma(s))$. \triangleleft

Example 2.8 The well-known word, tree and graph automata are instantiations of this notion. Infinite words over an alphabet C can be seen as C -colored coalgebras over the identity functor Id , so parity automata for infinite C -words correspond to C -chromatic Id -automata. Likewise, C -chromatic $\text{Id} \times \text{Id}$ -automata correspond to automata on infinite binary trees, and C -chromatic \mathcal{P} -automata correspond to amorphous tree or graph automata, with C in all cases denoting the *alphabet*.

For each functor F mentioned above, our notion may also take care of the automata operating on the corresponding *finite* objects. For instance, finite words over C may be taken as C -colored coalgebras for the functor $\{\downarrow\} + \text{Id}$, with \downarrow denoting termination.

We need the following fact from [13].

Fact 2.9 *With any $C \times F$ -automaton \mathbb{A} we may associate a C -chromatic F-automaton \mathbb{A}_C , the C -chromatic F-companion of \mathbb{A} , such that \mathbb{A} and \mathbb{A}_C accept the same $C \times F$ -coalgebras.*

3. Closure properties

In this section we show that the class of non-deterministically recognizable languages is closed under taking union and projection, whereas the class of recognizable languages is shown to be closed under union and intersection. Combined with Theorem 1, this suffices to prove Theorem 2.

3.1. Closure under union and intersection

In this subsection we define the sum and product of two F-automata, and prove that they recognize, re-

spectively, the union and the intersection of the languages associated with the original automata.

Definition 3.1 Let $\mathbb{A}_1 = (A_1, a_I^1, \Delta_1, \Omega_1)$ and $\mathbb{A}_2 = (A_2, a_I^2, \Delta_2, \Omega_2)$ be two F-automata. We will define their *sum* \mathbb{A}_\cup and *product* \mathbb{A}_\cap .

Both of these automata will have the *disjoint union* $A_{12} := \{*\} \uplus A_1 \uplus A_2$ as their collection of states. Also, the parity function Ω will be the same for both automata:

$$\Omega(a) := \begin{cases} 0 & \text{if } a = *, \\ \Omega_i(a) & \text{if } a \in A_i. \end{cases}$$

The only difference between the automata lies in the transition functions, which are defined as follows:

$$\Delta_\cup(a) := \begin{cases} \Delta_1(a_I^1) \cup \Delta_2(a_I^2) & \text{if } a = * \\ \Delta_i(a) & \text{if } a \in A_i, \end{cases}$$

$$\Delta_\cap(a) := \begin{cases} \{\Phi_1 \cup \Phi_2 \mid \Phi_i \in \Delta_i(a_I^i)\} & \text{if } a = * \\ \Delta_i(a) & \text{if } a \in A_i. \end{cases}$$

Finally, we put $\mathbb{A}_\cup := (A_{12}, a_I, \Delta_\cup, \Omega)$ and $\mathbb{A}_\cap := (A_{12}, a_I, \Delta_\cap, \Omega)$. \triangleleft

Proposition 3.2 *Let \mathbb{A}_1 and \mathbb{A}_2 be two F-automata. Then for any pointed F-coalgebra (\mathbb{S}, s) we have:*

1. \mathbb{A}_\cup accepts (\mathbb{S}, s) iff \mathbb{A}_1 or \mathbb{A}_2 accepts (\mathbb{S}, s) ,
2. \mathbb{A}_\cap accepts (\mathbb{S}, s) iff both \mathbb{A}_1 and \mathbb{A}_2 accept (\mathbb{S}, s) .
3. \mathbb{A}_\cup is non-deterministic if \mathbb{A}_1 and \mathbb{A}_2 are so.

Proof. First suppose that the automaton \mathbb{A}_\cup accepts (\mathbb{S}, s) . Hence by definition, \exists has a winning strategy f in the game $\mathcal{G} := \mathcal{G}(\mathbb{A}_\cup, \mathbb{S})$ starting from position $(s, *)$. Let i be such that $f(*, s) \in \Delta(a_I^i)$. It is then straightforward to verify that f , restricted to \exists 's positions in $\mathcal{G}(\mathbb{A}_i, \mathbb{S})$, is a winning strategy for \exists from position (s, a_I^i) . From this it is immediate that \mathbb{A}_i accepts (\mathbb{S}, s) . The other statements of the proof admit similarly straightforward proofs. \square

3.2. Closure under projection

In this subsection all F-automata are assumed to be non-deterministic. To facilitate the presentation we will think of the transition function Δ as a map $A \rightarrow \mathcal{P}FA$ and the first component Φ of a strategy (Φ, Y) for \exists in an acceptance game $\mathcal{G}(\mathbb{A}, \mathbb{S})$ will be regarded as a function of type $A \times S \rightarrow FA$ (that is, we identify singleton sets with their unique elements).

Definition 3.3 Let C be a set, $\mathbb{A} = (A, a_I, \Delta, \Omega)$ be a $(C \times F)$ -coalgebra automaton and $\mathbb{A}_C = (A, a_I, C, \Delta_C, \Omega)$ its C -chromatic F-companion, see Fact 2.9. Then we define the C -projection $\pi_C \mathbb{A} := (A, a_I, \Delta_\pi, \Omega)$ where $\Delta_\pi(a) := \bigcup_{c \in C} \Delta_C(c, a)$. \triangleleft

Lemma 3.4 If \mathbb{A} accepts the $(C \times F)$ -coalgebra $(\mathbb{S}, s) := (S, \gamma, \sigma, s)$ then $\pi_C \mathbb{A}$ accepts $(\mathbb{S}^\pi, s) := (S, \sigma, s)$.

Proof. The proof is easy. One has to realize that all the moves of \exists in the game for \mathbb{A}_C are still legitimate moves of \exists in the $\pi_C \mathbb{A}$ acceptance game. QED

The converse of this lemma however fails in general. Let \mathbb{A} be some $C \times F$ -automaton and let (S, σ, r) be a pointed F-coalgebra that is accepted by $\pi_C \mathbb{A}$. Then we know that \exists has a winning strategy (Φ, Y) in $\mathcal{G}(\pi_C \mathbb{A}, \mathbb{S})$ from position (r, a_I) . We would like to ensure that (Φ, Y) is also a winning strategy in $\mathcal{G}(\mathbb{A}_C, \mathbb{S})$ by defining a coloring $\gamma : S \rightarrow C$ as follows: $\gamma(s) := c$ if there is a match of $\mathcal{G}(\pi_C \mathbb{A}, \mathbb{S})$, starting from position (r, a_I) and conform \exists 's strategy, in which a position (s, a) occurs and $\Phi_{s,a} \in \Delta_C(c, a)$. In general, however, there may be *distinct* positions (s, a_1) and (s, a_2) that \forall may force the match to pass through, and it may not be possible to find a single $c \in C$ such that both $\Phi_{s,a_1} \in \Delta(c, a_1)$ and $\Phi_{s,a_2} \in \Delta(c, a_2)$. To avoid this problem we introduce now the notion of *strong* acceptance.

Definition 3.5 Let \mathbb{A} be an F-automaton and (\mathbb{S}, r) a pointed F-coalgebra. A history free strategy (Φ, Y) for \exists in the game $\mathcal{G}(\mathbb{A}, \mathbb{S})$ initialized at (r, a_I) is called *scattered* if the relation

$$\{(r, a_I)\} \cup \bigcup \{Y_{s,\varphi} \subseteq S \times A \mid (s, \varphi) \in \text{Win}_\exists\}$$

is the graph of some possibly partial function. Furthermore we say that \mathbb{A} *strongly accepts* the pointed coalgebra (\mathbb{S}, r) if \exists has a scattered winning strategy in the game $\mathcal{G}(\mathbb{A}, \mathbb{S})$ initialized at position (r, a_I) . \triangleleft

Lemma 3.6 Let \mathbb{A} be a $C \times F$ -automaton, and let (\mathbb{S}, r) be a pointed F-coalgebra that is strongly accepted by $\pi_C \mathbb{A}$. Then there is a C -colouring $\gamma : S \rightarrow C$ of \mathbb{S} such that \mathbb{A} accepts (S, γ, σ, r) .

Proof. Let (Φ, Y) be a scattered winning strategy for \exists in $\mathcal{G}(\pi_C \mathbb{A}, \mathbb{S})$. According to the definition of scatteredness we can assign to every $s \in S$ a state $a_s \in A$ such that $a_r = a_I$, and if $(s, a) \in Y_{s,\varphi}$ for some winning position (s, φ) , then $a = a_s$. Then we define a function $\gamma : S \rightarrow C$ as follows. If there is a $c \in C$ such that $\Phi_{s,a_s} \in \Delta_C(c, a)$, then we pick such a c and put $\gamma(s) := c$; if there is no such c , then we define $\gamma(s) := d$ for some arbitrary $d \in C$. It follows from these definitions that (Φ, Y) is a strategy for \exists in $\mathcal{G}(\mathbb{A}_C, \mathbb{S} \oplus \gamma)$ that guarantees her winning every match starting from (r, a_I) . From this it is immediate that \mathbb{A} accepts (S, γ, σ, r) . QED

The next lemma shows that if a pointed coalgebra is accepted by some automaton, but not strongly so, then we can always find a bisimilar pointed coalgebra that is strongly accepted.

Lemma 3.7 Let \mathbb{A} be an F-automaton, and let $(\mathbb{S}, r) = (S, \sigma, r)$ be a pointed F-coalgebra that is accepted by \mathbb{A} . Then there is a pointed F-coalgebra $(\bar{\mathbb{S}}, \bar{\sigma}, \bar{r})$ such that (\mathbb{S}, r) is bisimilar to $(\bar{\mathbb{S}}, \bar{\sigma}, \bar{r})$ and \mathbb{A} strongly accepts $(\bar{\mathbb{S}}, \bar{\sigma}, \bar{r})$.

Proof. The coalgebra $\bar{\mathbb{S}}$ will be based on the set $\bar{S} := S \times A$, and as the selected state \bar{r} of $\bar{\mathbb{S}}$ we take the pair (r, a_I) . For the definition of the coalgebra structure $\bar{\sigma}$, we need some auxiliary definitions.

First, it is not hard to see that we may endow the set \bar{S} with a coalgebra map $\bar{\sigma}$ which turns the structure $\bar{\mathbb{S}} := (\bar{S}, \bar{\sigma})$ into the A -fold coproduct $\coprod_{a \in A} \mathbb{S}$, in such a way that the first projection map $\pi_S : S \times A \rightarrow S$ is a coalgebra homomorphism from $\bar{\mathbb{S}}$ to \mathbb{S} . Second, given a relation $R \subseteq S \times A$, define the relation $\hat{R} \subseteq \bar{S} \times A$ by putting

$$\hat{R} := \{(s, a), a \mid (s, a) \in R\}.$$

Then clearly we have that $R = Gr(\pi_S)^\smile \circ \hat{R}$, and hence,

$$\bar{F}R = Gr(F\pi_S)^\smile \circ \bar{F}\hat{R}. \quad (1)$$

Now, for the definition of $\bar{\sigma} : \bar{S} \rightarrow F\bar{S}$, consider an arbitrary element $(s, a) \in \bar{S}$, and distinguish cases. If (s, a) is a *winning* position in the game $\mathcal{G}(\mathbb{A}, \mathbb{S})$, then using (1), it follows from $(\sigma(s), \varphi) \in$

$\bar{F}Y$, that we may define $\bar{\sigma}(s, a)$ to be some object in $F\bar{S}$ satisfying $(\sigma(s), \bar{\sigma}(s, a)) \in Gr(F\pi_S)^\smile$ and $(\bar{\sigma}(s, a), \varphi) \in \bar{F}\hat{Y}$. If, on the other hand, $(s, a) \notin Win_{\exists}$, then we simply put $\bar{\sigma}(s, a) := \bar{\sigma}(s, a)$.

It is completely straightforward to check that the map π_S is in fact an F-coalgebraic homomorphism from \bar{S} onto S . From this, the first statement of the proposition follows immediately.

For the second statement, define the strategy $(\bar{\Phi}, \bar{Y})$ with $\bar{\Phi} : \bar{S} \times A \rightarrow A$ and $\bar{Y} : \bar{S} \times FA \rightarrow \mathcal{P}(\bar{S} \times A)$ as follows:

$$\begin{aligned}\bar{\Phi} : ((s, a), b) &\mapsto \Phi_{s,b} \\ \bar{Y} : ((s, a), \varphi) &\mapsto \hat{Y}_{s,\varphi}.\end{aligned}$$

Since all relations chosen by \exists are of the form \hat{R} , and all elements of such relations are of the form $((s, a), b)$ with $a = b$, it is obvious that the set $\{((s, a_I), a_I)\} \cup \bigcup \{\hat{Y}_{s,\varphi} \mid (s, \varphi) \in Win_{\exists}\}$ is functional. In other words, the strategy is scattered.

Thus it is left to prove that $(\bar{\Phi}, \bar{Y})$ guarantees \exists to win any match of $\mathcal{G}(\mathbb{A}, \bar{S})$ starting from (\bar{r}, a_I) . To see why this is the case, consider an arbitrary position $((s, a), a)$ with $(s, a) \in Win_{\exists}(\mathcal{G}(\mathbb{A}, \bar{S}))$, and abbreviate $\varphi := \Phi_{s,a}$. Then by definition, $\bar{\Phi}((s, a), a) = \varphi$ and $\bar{Y}((s, a), \varphi) = \hat{Y}_{s,\varphi} = \{((t, b), b) \mid (t, b) \in Y_{s,\varphi}\}$. From this observation it is easy to derive that for any $\mathcal{G}(\mathbb{A}, \bar{S})$ match $(\bar{r}, a_I)((s_1, a_1), a_1)((s_2, a_2), a_2) \dots$ that is conform the strategy $(\bar{\Phi}, \bar{Y})$, the corresponding $\mathcal{G}(\mathbb{A}, S)$ match $(r, a_I)(s_1, a_1)(s_2, a_2) \dots$ is conform (Φ, Y) . And since this strategy was supposed to be winning for \exists from (r, a_I) , it follows that the $\mathcal{G}(\mathbb{A}, \bar{S})$ match is, indeed, a win for \exists . This proves the second statement of the proposition. QED

Proposition 3.8 *Let \mathbb{A} be some $(C \times F)$ -automaton. Then the following are equivalent, for every pointed F-coalgebra (S, s) :*

1. $\pi_{\mathbb{A}}$ accepts (S, s) ,
2. \mathbb{A} accepts a $(C \times F)$ -coalgebra $(S', \gamma, \sigma', s')$ such that (S', σ', s') and (S, s) are bisimilar.

Proof. The implication $(1 \Rightarrow 2)$ is immediate by the Lemmas 3.7 and 3.6. The other implication follows from Lemma 3.4 and the observation [13] that F-automata do not distinguish between bisimilar pointed F-coalgebras. QED

4. From alternating automata to non-deterministic ones

In this section we prove the main technical result of the paper, Theorem 1. That is, we will construct, for an arbitrary, alternating F-automaton an equivalent *non-deterministic* F-automaton. Throughout this section we will be working with a fixed (but arbitrary) F-automaton $\mathbb{A} = \langle A, a_I, \Delta, \Omega \rangle$.

Before going into the technical details of the construction, let us first provide some of the intuitions behind our approach. These intuitions ultimately go back to ideas of Muller and Schupp, see for instance [9], but in particular, our proof generalizes work by Janin and Walukiewicz [7], using the approach of Arnold and Niwiński [2]. (In fact, with some effort, it would be possible to prove our result here as a *corollary* of the work mentioned, but that would be to miss our point that a uniform, *coalgebraic* proof is possible.)

The main idea is to bring the players' interaction pattern $\exists\forall\exists\forall$ in one round of the acceptance games for \mathbb{A} , into the 'strategic form' $\exists\forall$ (or more precisely: $\exists\exists\forall$). Concretely, consider a basic position $(s, a) \in S \times A$ in the acceptance game $\mathcal{G}(\mathbb{A}, S)$ for some F-coalgebra S . From this position, play proceeds as follows:

- \exists picks $\Phi \in \Delta(a)$, moving to position (s, Φ) ;
- \forall picks $\varphi \in \Phi$, moving to position (s, φ) ;
- \exists picks $Y_\varphi \subseteq S \times A$ with $(\sigma(s), \varphi) \in \bar{F}Y_\varphi$ — this Y_φ is the new position;
- \forall picks $(t, b) \in Y_\varphi$ as the next basic position.

Now the crucial point is that \exists may *gather* her family $\{Y_\varphi \subseteq S \times A \mid \varphi \in \Phi\}$ into *one single* relation $Y_\Phi \subseteq S \times \mathcal{P}A$, and that we may modify the game in such a way that this is an appropriate answer for \exists . This approach would suggest to take (representations of) subsets of A as the states of the new automaton \mathbb{A}^d .

However, as is well-known from the literature, such a straightforward subset construction may work for automata that operate on finite objects, in the case of automata for (possibly) infinite objects this approach fails to make some subtle but crucial distinctions. The remedy, which brings us to the second fundamental idea underlying our construction,

is to use *binary relations on A*, rather than subsets of A , to bring the acceptance game into some kind of ‘layered-strategic’ form. Then, using the notion of a *trace* through a sequence of such relations, we have an established tool at our disposal for bringing the interaction pattern of the acceptance game into the required format. Our contribution here is to show that all of this can be done in the abstract context of coalgebras for an arbitrary standard, weak pullback preserving functor.

Now we are ready for the technical details of the construction.

Definition 4.1 Given a finite word $\rho = R_1 R_2 \dots R_n$ over the set $Rel(A)$ of binary relations over A , a *trace* through ρ is an A -word $\alpha = a_0 a_1 a_2 \dots a_k$ with $k \leq n$ such that $a_0 = a_I$ is the initial state a_I of the automaton, and $a_i R_{i+1} a_{i+1}$ for all $i < k$. Similar definitions apply to (finite or infinite) traces on infinite $Rel(A)$ -words.

A trace α is a *trap* for \exists if $\Delta(a_i) = \emptyset$ for some state a_i on α ; a trace α is *bad* if it is a trap for \exists or, in case α is infinite, if $\max\{\Omega(a_i) \mid i \in \mathbb{N}\}$ is odd. \triangleleft

As we will see, traces may be associated with matches of the acceptance game for \mathbb{A} , bad traces with the ones that are lost by \exists . Let us look at this in a bit more detail. As a consequence of the great generality that we aim for, there are two different ways in which \exists may loose a match. She may either get stuck at some finite stage of the match (either at a basic position or at a position of the form $(s, \varphi) \in S \times FA$), or survive for infinitely many rounds but fail to establish the winning condition. Now the traces that are traps for \exists will correspond to matches in which she gets stuck in a *basic* position, whereas the other kind of badness will turn out to be an encoding of \exists failing to win an infinite match. For finite matches that \exists looses because of getting stuck in a non-basic position, we do not need a corresponding notion for traces.

The first proposition that we need is a variation on well-known results. It concerns the existence of a deterministic *word* automaton that accepts those words over $Rel(A)$ which contain no bad traces. Since there are two kinds of bad traces, this automaton needs to perform a double task: it needs to recognize traps for \exists , and it needs to take proper care of the infinite words. It will be convenient to have the automaton perform these two jobs more or less separately.

That is, the automaton will have a special state m_\forall signalling that \exists has been trapped. In order to formulate the proposition we need some notation: given a deterministic automaton $\mathbb{D} = \langle D, d_I, \Sigma, \delta, \Omega_{\mathbb{D}} \rangle$ with alphabet Σ and transition function δ , we let $\hat{\delta} : D \times \Sigma^* \rightarrow D$ denote the *iterated* transition function, inductively defined by $\hat{\delta}(d, \epsilon) = d$ and $\hat{\delta}(d, \alpha a) = \delta(\hat{\delta}(d, \alpha), a)$.

Proposition 4.2 *There is a deterministic word automaton $\mathbb{M}_0 = \langle M, m_I, Rel(A), \mu_0, \Omega_0 \rangle$, operating on $Rel(A)$ -words, and containing a special state m_\forall , such that:*

1. $\mu_0(m_\forall, R) = m_\forall$ for all $R \in Rel(A)$,
2. for any finite $Rel(A)$ -word ρ : $\hat{\mu}(\rho) = m_\forall$ iff ρ contains a trap for \exists ,
3. for any infinite $Rel(A)$ -word ρ : \mathbb{M}_0 accepts ρ iff ρ contains no bad traces.

In the remainder of this section we *fix* the automaton $\mathbb{M}_0 = \langle M, m_I, Rel(A), \mu_0, \Omega_0 \rangle$ and state m_\forall as given in Proposition 4.2. We leave the proof of Proposition 4.2 as an exercise for the reader, and move on to the main construction of the proof. Below we define a non-deterministic automaton \mathbb{M}_1 which operates on $(PFA)^A$ -colored F-coalgebras, $\mathbb{S} \oplus \Phi$, that is, F-coalgebras $\mathbb{S} = \langle S, \sigma \rangle$ that are colored by the map $\Phi : S \rightarrow (PFA)^A$. Such a map Φ represents a (potential) *strategy* of \exists in the game $\mathcal{G}(\mathbb{A}, \mathbb{S})$ that is *partial* in the sense of dealing with basic positions only. More precisely, for any position $(s, a) \in S \times A$, we let the value $\Phi_{s,a} \in PFA$ encode the move $(s, \Phi_{s,a}) \in S \times PFA$. Our aim with the automaton \mathbb{M}_1 is that it will recognize precisely those pointed $(PFA)^A$ -colored F-coalgebras $\langle S, \sigma, \Phi, s \rangle$ of which Φ forms the basic part of a winning strategy in the game $\mathcal{G}(\mathbb{A}, \mathbb{S})$. Towards the end of this section we will see that this suffices to prove Theorem 1.

For the definition of \mathbb{M}_1 we need some preliminary definitions.

Definition 4.3 An object $\Xi \in FP(A)$ is called an *F-redistribution* of a subset $\Phi \subseteq FA$ if $(\varphi, \Xi) \in \bar{F}(\in_A)$ for all $\varphi \in \Phi$.

An object $\Pi \in FRel(A)$ is called an *F-redistributive relational representation* of an element $\Phi \in (PFA)^A$, or shortly: an *F-relation* for Φ , if $(Fev_a)(\Pi)$ is a redistribution of $\Phi(a)$ for all $a \in A$.

Here $ev_a : Rel(A) \rightarrow \mathcal{P}(A)$ is the map given by $ev_a : R \mapsto R[a]$. The collection of F-relations for $\Phi \in (\mathcal{PFA})^A$ is denoted as $\mathcal{R}_F(\Phi)$. \triangleleft

The intuitions on these notions are as follows. Concerning redistributions, the point is that for any F-coalgebra $\mathbb{S} = \langle S, \sigma \rangle$, any point $s \in S$ and any set $\Phi \in \mathcal{PFA}$, there is a 1-1 correspondence between:

- families $\{Y_\varphi \subseteq S \times A \mid \varphi \in \Phi\}$ of relations such that $(\sigma(s), \varphi) \in \overline{FY}_\varphi$ for all $\varphi \in \Phi$, and
- pairs (Y_Φ, Ξ) consisting of a relation $Y_\Phi \subseteq S \times \mathcal{P}A$, and an F-redistribution $\Xi \in F\mathcal{P}A$ of Φ , such that $(\sigma(s), \Xi) \in \overline{FY}_\Phi$.

In brief, redistributions enable us to gather the information of a family $\{Y_\varphi \subseteq S \times A \mid \varphi \in \Phi\}$ of relation moves of \exists into one single relation $Y_\Phi \subseteq S \times \mathcal{P}A$.

However, this regrouping of information on \exists 's strategy in terms of redistributions has one shortcoming: it is based on *subsets* of A whereas we already pointed out that such an encoding will not suffice to encode the full flow of information when transforming alternating automata into non-deterministic ones. This is where the notion of an F-relation for Φ comes in. The important observation is that any element Π of the set $FRel(A)$ has the right shape to represent a family $\{\Phi_a \in \mathcal{PFA} \mid a \in A\}$: the point is that we may use, for every $a \in A$, the map $Fev_a : FRel(A) \rightarrow F\mathcal{P}A$ to provide an element $(Fev_a)(\Pi)$ in the right set $F\mathcal{P}A$ of (potential) F-redistributions of Φ . Thus, the definition of a $\Pi \in FRel(A)$ being an F-redistributive relational representation of $\Phi \in (\mathcal{PFA})^A$ forms, at least potentially, an adequate formalization of the requirement that Π and Φ 'fit well together'. As we will see below, it also forms the key to lead the flow of information in acceptance games for alternating automata into a non-deterministic channel.

Definition 4.4 Let \mathbb{M}_1 be the non-deterministic $(\mathcal{PFA})^A$ -chromatic F-automaton $\langle M, m_I, (\mathcal{PFA})^A, \mu, \Omega_0 \rangle$, where $\mu : M \times (\mathcal{PFA})^A \rightarrow \mathcal{PPFM}$ is the map defined by

$$\mu(m, \Phi) := \begin{cases} \{(F\mu_m)(\Pi) \mid \Pi \in \mathcal{R}_F(\Phi)\} & \text{if } m \neq m_\forall, \\ \emptyset & \text{if } m = m_\forall. \end{cases}$$

Here $\mu_m : Rel(A) \rightarrow M$ is given by $\mu_m(R) := \mu_0(m, R)$. \triangleleft

Remark 4.5 Let \mathbb{M}_1 be as above, and $\mathbb{S} = \langle S, \sigma, \Phi \rangle$ some $(\mathcal{PFA})^A$ -colored F-coalgebra. Note that the acceptance game $\mathcal{G}(\mathbb{M}_1, \mathbb{S})$ is summarized in Table 2.

Given the definition of μ , it is not hard to see that, from a position $(s, m) \in S \times M$, with subsequent moves of \exists , say, $(s, \{K\}) \in S \times \mathcal{PFM}$ and $Y \subseteq S \times M$, we may associate an element $\Pi \in FRel(A)$ and a relation $Z \subseteq S \times Rel(A)$ such that Π is an F-relation for Φ_s , $(F\mu_m)(\Pi) = K$ and $(\sigma(s), \Pi) \in \overline{FZ}$.

To start with, it is obvious from the definitions that there is some $\Pi \in \mathcal{R}_F(\Phi_s)$, such that $(F\mu_m)(\Pi) = K$. Now define the relation $Z := \{(t, R) \in S \times Rel(A) \mid (t, \mu_m(R)) \in Y\}$. Clearly, this relation is the composition of Y with the converse relation $Gr(\mu_m)^\smile$ of the graph of the function μ_m . From this it follows that $\overline{FZ} = \overline{FY} \circ \overline{F(Gr(\mu_m)^\smile)}$. Also, rewriting $(F\mu_m)(\Pi) = K$, we obtain that $(\Pi, K) \in Gr(F\mu_m) = \overline{F}Gr(\mu_m)$, so that $(K, \Pi) \in (\overline{F}Gr(\mu_m))^\smile = \overline{F}(Gr(\mu_m)^\smile)$. Hence, from $(\sigma(s), K) \in \overline{FY}$ it is immediate that $(\sigma(s), \Pi) \in \overline{FZ}$.

Proposition 4.6 For any pointed F-coalgebra $\langle \mathbb{S}, roots_{\mathbb{S}} \rangle$ and any $(\mathcal{PFA})^A$ -coloring Φ of S , the following are equivalent:

1. Φ is part of a winning strategy for \exists in $\mathcal{G}(\mathbb{A}, \mathbb{S})$ initialized at $roots_{\mathbb{S}}$;
2. \mathbb{M}_1 accepts $(\mathbb{S} \oplus \Phi, roots_{\mathbb{S}})$.

Proof. Recall that every infinite game may be represented as a tree, and that strategies of either player, limiting the possible course of actions, can be represented as *subtrees* of this game tree. Thus, both with a Φ -extending strategy of \exists in $\mathcal{G} = \mathcal{G}(\mathbb{A}, \mathbb{S})$, and with a strategy of \exists in the acceptance game $\mathcal{G}' = \mathcal{G}(\mathbb{M}_1, \mathbb{S} \oplus \Phi)$, we may associate such subtrees of the game trees of \mathcal{G} and \mathcal{G}' , respectively. As it turns out, these two trees turn out to be rather similar, and in fact, may be coded up into one and the same structure. This observation forms the basis of our proof of the proposition.

More specifically, we will show the equivalence of both (1) and (2) to the statement (3) below.

3. There is a labelled tree

$$\mathbb{X} = \langle X, root_{\mathbb{X}}, \xi, u, \Pi, Q \rangle,$$

where $root_{\mathbb{X}} \in X$ and $\xi : X \rightarrow \mathcal{P}(X)$ denote, respectively, the root and the successor function

Position: b	Type	$P(b)$	Admissible moves: $E[b]$	$\overline{\Omega_0}(b)$
(s, m)	$S \times M$	\exists	$\{(s, \{K\}) \in S \times \mathcal{PFM} \mid K \in \mu(m, \Phi(s))\}$	$\Omega_0(m)$
$(s, \{K\})$	$S \times \mathcal{PFM}$	\forall	$\{(s, K)\}$	0
(s, K)	$S \times FM$	\exists	$\{Z \in \mathcal{P}(S \times M) \mid (\sigma(s), K) \in \overline{F}Z\}$	0
Z	$\mathcal{P}(S \times M)$	\forall	Z	0

Table 2. Acceptance game for \mathbb{M}_1

of the tree, and $u : X \rightarrow S$, $\Pi : X \rightarrow FRel(A)$, and $Q : X \rightarrow Rel(A)$ are labellings.

This tree is supposed to satisfy the conditions 3a–3d below. Here, and in the sequel, we abbreviate Φ_{u_x} as Φ_x , and define $W_x := \{(u_y, Q_y) \mid y \in \xi(x)\}$. Branches of the tree start at the root, and thus induce (finite or infinite) words over $Rel(A)$.

- (a) $u_{root_x} = root_S$ and $Q_{root_x} = \{(a_I, a_I)\}$,
- (b) for all $x \in X$, Π_x is an F-relation for Φ_x ,
- (c) for all $x \in X$, $(\sigma(u_x), \Pi_x) \in \overline{F}W_x$.
- (d) \mathbb{X} has no bad traces (that is, no branch of \mathbb{X} induces a $Rel(A)$ -word containing a bad trace).

As hinted at above, our intuition about \mathbb{X} is that it represents a winning strategy for \exists both in \mathcal{G} and in \mathcal{G}' (in the case of \mathcal{G} , of course, a strategy completing the partial strategy Φ). Counterstrategies of \forall in \mathcal{G}' correspond to *branches* of \mathbb{X} , while his strategies in \mathcal{G} appear as *traces* on \mathbb{X} . Further details of the proof are left for the full version of this paper. QED

In the final step of the construction we have to transform \mathbb{M}_1 into a non-deterministic F-automaton \mathbb{A}^d that is equivalent to \mathbb{A} . This last transformation is in fact easy — relatively that is: we need an application of the closure under projection of non-deterministically recognizable languages.

Definition 4.7 Let \mathbb{A}^d be the F-automaton $\langle M, m_I, \mu^d, \Omega_0 \rangle$ where M , m_I and Ω_0 are as in Definition 4.4, while

$$\mu^d(m) := \bigcup_{e \in (\mathcal{P}FA)^A} \mu(m, e)$$

defines the transition map $\mu^d : M \rightarrow \mathcal{P}(FM)$. \triangleleft

It is easy to check that \mathbb{A}^d is indeed non-deterministic, so clearly, the following proposition, which is a straightforward corollary of the Propositions 3.8 and 4.6, suffices to prove Theorem 1.

Proposition 4.8 *The automata \mathbb{A} and \mathbb{A}^d accept exactly the same pointed F-coalgebras.*

Remark 4.9 Although we do not go into the algorithmic details of our construction, we want to stress here that complexity theoretically, our results match known results in automata theory. If we define the size of an automaton as its number of states, the main observation is that the size of \mathbb{A}^d is equal to the size of \mathbb{M}_0 and, in particular, does not depend on the functor F. In fact, combining well known results about word automata, one may show that basically, the size of \mathbb{M}_0 is exponential in the size of \mathbb{A} .

5. Conclusions & Questions

There is a long list of issues that need some further discussion, but for reasons of space limitations we confine ourselves to a very brief discussion of the following four items.

To start with, we believe that this paper provides evidence for the claim that universal coalgebra forms an appropriate *abstraction level* for studying automata theory. Our results show that important automata-theoretic phenomena have a natural existence at the coalgebraic level of abstraction.

Second, although we have hardly mentioned logic at all, the results in the paper have in fact significant logical corollaries. For instance, generalizing results in [5], we can show that the coalgebraic fixed point logics of [13] all have *uniform interpolation*. We hope to report on this in future work.

Probably the most important issue to be addressed concerns the closure of the class of recognizable languages under *complementation*. For our coalgebraic

automata it is not so easy to prove a complementation lemma, even for alternating or deterministic automata. The reason for this is that the acceptance game for coalgebraic automata has some crucial *non-symmetric* interaction between the two players, with \exists choosing relations and \forall picking elements of such relations. The fact that for many well-known functors (including the ones that yield simple coalgebras such as trees and transition systems), this game can be brought into a symmetric form, simply reveals the existence of an interesting *property* that some functors have, and others may not. We have to leave this matter as an intriguing area for further research, however. Should there be a strong need for closure of recognizable languages under complementation, one may always consider to move to a different notion of coalgebra automaton that is tailored towards a more symmetric acceptance game. This is also a matter that we leave for future investigations.

In any case, closure under complementation may be a less important property than it appears to be at first sight. Explained in logical terms, the point is that coalgebraic logics (with or without fixed points) *without negation* already have considerable expressive power. For instance, A. Baltag (private communication) has shown that any state in a *finite* coalgebra can be completely characterized (modulo bisimilarity) by a negation free coalgebraic fixed point formula.

Finally, we are quite interested to see whether the conditions on the functor are really needed. It seems that the condition of standardness can be lifted without too much difficulty — note also that every set functor is ‘almost’ standard [1, Theorem III.4.5]. However, we believe that our main result crucially depends on the fact the functor preserves weak pullbacks. This is in line with results by Trnková [1] indicating that for a related class of functorial automata, nondeterministic and deterministic recognizability coincide if *and only if* the functor preserves weak pullbacks. The precise connection with these results clearly needs to be investigated.

References

- [1] J. Adámek and V. Trnková. *Automata and Algebras in Categories*. Kluwer Academic Publishers, 1990.
- [2] A. Arnold and D. Niwiński. *Rudiments of μ -calculus*, volume 146 of *Studies in Logic*. North-Holland, 2001.
- [3] A. Baltag. A logic for coalgebraic simulation. In H. Reichel, editor, *Coalgebraic Methods in Computer Science (CMCS’00)*, volume 33 of *Electronic Notes in Theoretical Computer Science*, pages 41–60, 2000.
- [4] E. Clarke, O. Grumberg, and D. Peled, editors. *Model Checking*. The MIT Press, 2000.
- [5] G. d’Agostino and M. Hollenberg. Logical questions concerning the μ -calculus. *Journal of Symbolic Logic*, 65:310–332, 2000.
- [6] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logic, and Infinite Games*, volume 2500 of *LNCS*. Springer, 2002.
- [7] D. Janin and I. Walukiewicz. Automata for the modal μ -calculus and related results. In *Proc. MFCS’95*, pages 552–562. Springer, 1995. LNCS 969.
- [8] L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999. (Erratum published *APAL* 99:241–259, 1999.
- [9] D. Muller and P. Schupp. Simulating alternating tree automata by nondeterministic ones. *Theoretical Computer Science*, 141:69–107, 1995.
- [10] M.O. Rabin. Decidability of second-order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141:1–35, 1969.
- [11] J. Rutten. Universal coalgebra: A theory of systems. *Theoretical Computer Science*, 249:3–80, 2000.
- [12] V. Trnková. Relational automata in a category and theory of languages. In *Proc. FCT 1977*, pages 340–355. Springer, 1977. LNCS 56.
- [13] Y. Venema. Automata and fixed point logics for coalgebras. *Electronic Notes in Theoretical Computer Science*, 106:355–375, 2004.