

COALGEBRAIC AUTOMATA THEORY: BASIC RESULTS

CLEMENS KUPKE AND YDE VENEMA

CWI Amsterdam, Kruislaan 413, 1098 SJ Amsterdam, Netherlands
e-mail address: kupke@cwi.nl

Universiteit van Amsterdam, ILLC, Plantage Muidergracht 24, 1018 TV Amsterdam, Netherlands
e-mail address: yde@science.uva.nl

ABSTRACT. We generalize some of the central results in automata theory to the abstraction level of coalgebras and thus lay out the foundations of a universal theory of automata operating on infinite objects.

Let F be any set functor that preserves weak pullbacks. We show that the class of recognizable languages of F -coalgebras is closed under taking unions, intersections, and projections. We also prove that if an F -automaton accepts some coalgebra it accepts a finite one of the size of the automaton. Our main technical result concerns an explicit construction which transforms a given alternating F -automaton into an equivalent nondeterministic one, of bounded size.

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 nonterminating programs such as operating systems. As an example we mention the automata-based verification method of *model checking* [5]. 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 [9] for an overview. The outstanding example here is of course Rabin's decidability theorem [24] for the monadic second order logic of trees; to mention a more recent result, Janin & Walukiewicz [12] identified the modal μ -calculus as the bisimulation invariant fragment of the monadic second order logic of labelled transition systems.

2000 ACM Subject Classification: F.1.1; F.4.3; F.3.2.

Key words and phrases: Automata, coalgebra, fixed-point logics, parity games.

The paper is a completely revised and extended version of [15].

Supported by the Netherlands Organisation for Scientific Research (NWO) under FOCUS/BRICKS grant number 642.000.502.

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 nondeterministic ones. These observations naturally raise the question, whether these results can perhaps be formulated and proved at a more general level of abstraction. This is by analogy to algebra, where it is more convenient, for instance, to formulate and prove the Homomorphism Theorems at the level of universal algebra, than to treat them separately for every algebraic signature. 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 [25] 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, starting with Moss' coalgebraic logic [19], 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 [29], the second author introduced a coalgebraic fixed-point logic that does enable specifications of this kind (see [30] for a more detailed exposition).

The same paper [29] 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, nondeterministic, and deterministic ones, respectively.

Basically, F -automata are meant to either accept or reject pointed coalgebras (that is, pairs $\langle S, s \rangle$ consisting of an F -coalgebra S together with a selected state s in the carrier S of 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. This connection generalizes the relation between the modal μ -calculus and the μ -automata [12] to the abstraction level of coalgebra. A key feature of the coalgebraic framework is that one restricts attention to observable, or *bisimulation-invariant properties*. 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 graph game.

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 7.

It is the aim of the present paper, which is a completely revised and extended version of [15], 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 nondeterministic automata, some of the closure properties that one may associate with automata, and the nonemptiness problem. 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 (*nondeterministically*) *recognizable* if it is recognized by some (nondeterministic) parity F -automaton. Our main technical result can now be formulated as follows.

Theorem 1. *Let F be some set functor that preserves weak pullbacks.¹ Then every alternating parity F -automaton \mathbb{A} has a nondeterministic equivalent \mathbb{A}^\bullet . Hence, an F -language is recognizable iff it is nondeterministically recognizable.*

More specifically, we will give a construction which is both *uniform*, in the sense that it takes the type F of \mathbb{A} as a parameter, and *concrete*, in that we give an explicit definition of \mathbb{A}^\bullet in terms of \mathbb{A} .

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 nondeterministically recognizable F -languages is closed under (existential) projection. Theorem 1 allows us to strengthen the above list of closure properties as follows.

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

Conspicuously absent in this list is closure under complementation — we will come back to this in section 7.

The third result that we want to mention here concerns the nonemptiness problem for coalgebra automata. As we will show, if a nondeterministic parity automaton \mathbb{A} accepts an F -coalgebra at all, then it accepts an F -coalgebra that somehow ‘lives inside \mathbb{A} ’. From this and Theorem 1 the following result is immediate.

Theorem 3. *Let F be some set functor that preserves weak pullbacks. Then the F -language recognized by a parity automaton \mathbb{A} is nonempty iff \mathbb{A} accepts a finite F -coalgebra of bounded size.*

We prove these results by generalizing, to the coalgebraic level, (well-)known ideas from the theory of specific automata. This applies in particular to the results on graph automata by Janin & Walukiewicz [12], whereas our approach is similar to the abstract universal algebraic

¹The meaning and importance of this side condition will be explained in section 2.

approach of Nawiński and Arnold [22, 2]. Just as in the literature on specific kinds of automata, our proofs crucially depend on the “import” of two fundamental results from the theory of automata and infinite games. The proof of our main result, Theorem 1 contains an essential instance of the determinization of ω -automata (a result originally due to MacNaughton [18], whereas the construction we use goes back to Safra [26]). Theorem 3, our solution to the nonemptiness problem for coalgebra automata, can be seen as an application of the history-free determinacy of parity games (see Fact 2.13).

Finally, let us stress again that our aim here is not to prove new results for well-known structures. Rather, the point that we try to make is that many of the well-known theorems in automata theory in fact belong to the field of Universal Coalgebra, in the same way that the Homomorphism Theorems are results in Universal Algebra. In our opinion, the abstract, coalgebraic perspective not only generalizes existing results, the uniformity of the coalgebraic presentation has helped us to obtain a better understanding of automata theory itself. To mention one example, for various automata-theoretic constructions, our coalgebraic proofs show that generally, *size issues* do not depend on the type of the automata (that is, on the functor F), because most of these issues only depend on constructions for ω -automata (cf. also Remark 5.3 below).

The paper is organized as follows: In Section 2 we equip the reader with the necessary background material on coalgebra automata. After that, we show in Section 3 how to transform an alternating parity F -automaton into an equivalent nondeterministic F -automaton that has a so-called regular acceptance condition. These ‘regular automata’ will be discussed in detail in Section 4, where we prove that any regular F -automaton can be transformed into an equivalent F -automaton with a parity acceptance condition. In Section 5 we combine the results from Section 3 and Section 4 in order to obtain Theorem 1. We then discuss other closure properties of recognizable languages and prove Theorem 2. Finally, in Section 6, we demonstrate that the nonemptiness problem of F -automata can be effectively solved, proving Theorem 3. Section 7 concludes the paper with a short summary of our results and an outlook on future research.

Acknowledgements. We are grateful to the anonymous referees for providing a number of very helpful suggestions concerning the presentation of our results.

2. PRELIMINARIES

Since coalgebra automata bring together notions from two research areas (coalgebra and automata theory), we have made an effort to make this paper accessible to both communities. As a consequence, we have included a fairly long section containing the background material that is needed for understanding the main definitions and proofs of the paper. Readers who want to acquire more knowledge could consult RUTTEN [25] and GRÄDEL, THOMAS & WILKE [9] for further details on universal coalgebra and automata theory, respectively. For a more gentle introduction to coalgebra automata, the reader is referred to VENEMA [30].

First we fix some mathematical notation and terminology.

Convention 2.1. Let $f : S \rightarrow T$ be a function. Then the *graph* of f is the relation

$$Gr(f) := \{(s, f(s)) \in S \times T \mid s \in S\}$$

Given a relation $R \subseteq S \times T$, we denote the *domain* and *range* of R by $\text{dom}(R)$ and $\text{rng}(R)$, respectively. Given subsets $S' \subseteq S$, $T' \subseteq T$, the *restriction* of R to S' and T' is given as

$$R \upharpoonright_{S' \times T'} := R \cap (S' \times T').$$

The composition of two relations $R \subseteq S \times T$ and $R' \subseteq T \times U$ is denoted by $R \circ R'$, whereas the composition of two functions $f : S \rightarrow T$ and $f' : T \rightarrow U$ is denoted by $f' \circ f$. Thus, we have $\text{Gr}(f' \circ f) = \text{Gr}(f) \circ \text{Gr}(f')$. Finally, the set of all functions from S to T is denoted both by T^S and by $S \rightarrow T$.

2.1. Coalgebras, bisimulations and relation lifting. Informally, a coalgebra \mathbb{S} consists of a set S of *states*, together with a *transition structure* σ mapping a state $s \in S$ to an object $\sigma(s)$ in some set FS , where F is the *type* of the structure \mathbb{S} . Technically, the type of a coalgebra is formalized as a *set functor*, i.e. an endofunctor on the category Set that has sets as objects and functions as arrows. Such a functor maps a set S to a set FS , and a function $f : S \rightarrow S'$ to a function $\text{F}f : \text{FS} \rightarrow \text{FS}'$.

Definition 2.2. Given a set functor F , an F -coalgebra is a pair $\mathbb{S} = \langle S, \sigma \rangle$ with $\sigma : S \rightarrow \text{FS}$. A *pointed coalgebra* is a pair consisting of a coalgebra together with an element of (the carrier set of) that coalgebra. \triangleleft

Remark 2.3. The action of F on functions is needed to define the *morphisms* between two coalgebras. Formally, given two F -coalgebras $\mathbb{S} = \langle S, \sigma \rangle$ and $\mathbb{S}' = \langle S', \sigma' \rangle$, a function $f : S \rightarrow S'$ is a coalgebra morphism if $\sigma' \circ f = (\text{F}f) \circ \sigma$. Since these morphisms do not play a large role in this paper, the reader that has no familiarity with coalgebras may safely ignore this aspect. In the sequel, we will often even introduce set functors just by defining their action on sets, trusting that the initiated reader will be able to supplement the action on functions.

Example 2.4. (1) For instance, consider the functor B which associates with a set S the cartesian product $S \times S$, and with a map $f : S \rightarrow S'$, the map $\text{B}f : S \times S \rightarrow S' \times S'$ given by $(\text{B}f)(s, s') := (f(s), f(s'))$. Thus every state in an B -coalgebra has both a left and a right successor. As a special example of a B -coalgebra, consider the *binary tree* given as the set 2^* of finite words over the alphabet $2 = \{0, 1\}$, with the coalgebra map given by $s \mapsto (s0, s1)$.

- (2) Directed graphs can be seen as coalgebras of the power set functor P . The functor P maps a set to its collection of subsets, and a function $f : S \rightarrow S'$ to its direct image function $\text{P}f : \mathcal{P}S \rightarrow \mathcal{P}S'$ given by $(\text{P}f)(X) := \{f(x) \in S' \mid x \in X\}$. A graph (G, E) is then modelled as the coalgebra $(G, \lambda x.E[x])$, that is, the relation E is given by the function mapping a point x to the collection $E[x]$ of its (direct) successors.
- (3) The labelled transition systems from [12] can be represented as coalgebras for the functor $\mathcal{P}\Phi \times \text{P}^A$, i.e. for the functor that maps a set S to the set $\mathcal{P}\Phi \times (\mathcal{P}S)^A$. Here Φ is a set of propositional variables and A is the set of actions. Hence, in the coalgebraic presentation, a point s is mapped to the pair consisting of the set of proposition letters in Φ that are true at s , and, for each action a , the set of ‘ a -successors’ of s .
- (4) 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 F_C -coalgebra $\langle S, \langle \gamma, \sigma \rangle \rangle$. Here F_C is the functor that takes a set S to the set $C \times S$ (and that takes a map $f : S \rightarrow S'$

to the function $F_C f : C \times S \rightarrow C \times S'$ given by $(F_C f)(c, s) = (c, f(s))$. So B_C and P_C are, respectively, the functors that we may associate with C -labelled binary trees and C -labelled graphs, respectively.

- (5) Let D_ω be the functor that maps a set S to the set

$$D_\omega S := \{\rho : S \rightarrow [0, 1] \mid \rho \text{ has finite support and } \sum_{s \in S} \rho(s) = 1\}$$

where we say that ρ has finite support if $\rho(s) \neq 0$ for only finitely many elements s of S . Then coalgebras for the functor $1 + D_\omega$ correspond to the probabilistic transition systems by Larsen and Skou in [16]. Here $1 + D_\omega$ denotes the functor that maps a set S to the disjoint union of the one-element set and the set $D_\omega S$. Further details about this example can be found in [7].

- (6) Transition systems in which every state has a *multiset* of successors can be modeled as coalgebras for the functor M_ω that maps a set S to the set $M_\omega S$ of all functions $\mu : S \rightarrow \mathbb{N}$ with finite support, i.e. $\mu(s) \neq 0$ for only finitely many elements s of S . More information can be found in [10].

As mentioned in the introduction, the theory of coalgebra aims to provide a simple but general framework for formalizing and studying the concept of *behavior*. For this purpose it is of importance to develop a notion of *behavioral equivalence* between states. For almost all important coalgebraic types, this notion can be naturally expressed using *bisimulations*. Intuitively, these are relations between the state sets of two coalgebras that witness the observational indistinguishability of the pairs that they relate. For our purposes it will be convenient to use a definition of bisimilarity in terms of *relation lifting*.

Definition 2.5. Let F be a set functor. Given two sets S and S' , and a binary relation Z between $S \times S'$, we define the *lifted relation* $\bar{F}(Z) \subseteq FS \times FS'$ as follows:

$$\bar{F}(Z) := \{((F\pi)(\varphi), (F\pi')(\varphi)) \mid \varphi \in FZ\},$$

where $\pi : Z \rightarrow S$ and $\pi' : Z \rightarrow S'$ are the projection functions given by $\pi(s, s') = s$ and $\pi'(s, s') = s'$.

$$\begin{array}{ccc} S & \xleftarrow{\pi} R \xrightarrow{\pi'} & S' \\ & \Downarrow \text{lifting} & \\ FS & \xleftarrow{F\pi} FR \xrightarrow{F\pi'} & FS' \end{array}$$

Now let $\mathbb{S} = \langle S, \sigma \rangle$ and $\mathbb{S}' = \langle S', \sigma' \rangle$ be two F -coalgebras. Then a relation $Z \subseteq S \times S'$ satisfying

$$(\sigma(s), \sigma'(s')) \in \bar{F}(Z) \text{ for all } (s, s') \in Z. \quad (2.1)$$

is an (F -)bisimulation between \mathbb{S} and \mathbb{S}' . Two states s and s' in such coalgebras are *bisimilar*, notation: $\mathbb{S}, s \Leftrightarrow \mathbb{S}', s'$, iff they are linked by some bisimulation. \triangleleft

Intuitively, for $Z \subseteq S \times S'$ to be a bisimulation we require that whenever s and s' are linked by Z , then $\sigma(s)$ and $\sigma'(s')$ are linked by the lifted version $\bar{F}(Z)$ of Z .

Remark 2.6. Strictly speaking, the definition of the relation lifting of a given relation R depends on the type of the relation, i.e. given sets S', S, T', T such that $R \subseteq S' \times T'$ and $R \subseteq S \times T$, it matters whether we look at R as a relation from S' to T' or as a relation from S to T . This possible source of ambiguity can be avoided if we require the functor F to be

standard. We come back to a detailed discussion of the notion of a standardness at the end of this section.

Example 2.7. Let us see how these definitions apply to coalgebras for the functors \mathbf{B}_C and \mathbf{P} of Example 2.4. For this purpose, fix two sets S and S' , and a relation $Z \subseteq S \times S'$.

For the definition of the relation $\overline{\mathbf{B}_C}(Z)$, it is easy to see that (c, s_0, s_1) and (c', s'_0, s'_1) are related iff $c = c'$ and both (s_0, s'_0) and (s_1, s'_1) belong to Z . Hence the relation Z is a bisimulation between two \mathbf{B}_C -coalgebras if Z -related points have the same color and both their left successors and their right successors are Z -related. From this it easily follows that two labelled binary trees are bisimilar iff they are identical. However, the notion becomes less trivial if we consider other coalgebras for the functor \mathbf{B}_C . In fact, bisimilarity can be used to formulate some well-known notions in the theory of tree automata; for instance, a labelled binary tree is *regular* iff it is bisimilar to a finite \mathbf{B}_C -coalgebra.

Concerning the Kripke functor \mathbf{P} , observe that $(X, X') \in \overline{\mathbf{P}}(Z)$ iff for all $x \in X$ there is an $x' \in X'$, and for all $x' \in X'$ there is an $x \in X$ such that $(x, x') \in Z$. That is, $\overline{\mathbf{P}}(Z)$ is the *Egli-Milner* lifting of Z . Thus, a relation relating nodes of one graph \mathbb{S} to those of another graph \mathbb{S}' is a *bisimulation* if for every linked pair (s, s') , every successor t of s is related to some successor t' of s' , and, vice versa, every successor t' of s' is related to some successor t of s .

Finally, the connection between $\overline{\mathbf{F}}$ and $\overline{\mathbf{F}_C}$ is given by the following: $((c, \varphi), (c', \varphi')) \in \overline{\mathbf{F}_C}(Z)$ iff $c = c'$ and $(\varphi, \varphi') \in \overline{\mathbf{F}}(Z)$.

Given the key role that relation lifting plays in this paper, we need some of its properties. It can be shown that relation lifting interacts well with the operation of taking the graph of a function $f : S \rightarrow S'$, and with most operations on binary relations. In fact, the properties listed in Fact 2.12 are *all* the information on relation lifting that is needed in this paper. Readers that have no interest in categorical details may safely skip some material and move on to Fact 2.12.

Unfortunately, for two properties of $\overline{\mathbf{F}}$ which are crucial for our results, we need the functor to satisfy a certain categorical property, namely the preservation of weak pullbacks (to be defined in Remark 2.11 below). Fortunately, many set functors in fact do preserve weak pullbacks, which guarantees a wide scope for the results in this paper.

Example 2.8. All functors from Example 2.4 preserve weak pullbacks. For the functors \mathbf{B} and \mathbf{P} this is not difficult to check. A proof for the fact that \mathbf{D}_ω is weak pullback preserving can be found in [7]. That \mathbf{M}_ω is weak pullback preserving is a direct consequence of the results in [10].

Furthermore it can be shown that the class of weak pullback preserving functors contains all constant functors, and is closed under composition, taking products, coproducts (disjoint unions), and under exponentiation. We use these observations in order to define a class of functors that are all weak pullback preserving.

Definition 2.9. The class KPF of *Kripke polynomial functors* is inductively defined by putting

$$\mathbf{F} ::= A \in \mathbf{Set} \mid \text{Id} \mid \mathcal{P} \mid \mathbf{F} + \mathbf{F} \mid \mathbf{F} \times \mathbf{F} \mid \mathbf{F}^D, D \in \mathbf{Set} \mid \mathbf{F} \circ \mathbf{F},$$

where Id denotes the identity functor on \mathbf{Set} . Enlarging the basis of this inductive definition with the probabilistic and multiset functor of Example 2.4:

$$\mathbf{F} ::= A \in \mathbf{Set} \mid \text{Id} \mid \mathcal{P} \mid \mathbf{D}_\omega \mid \mathbf{M}_\omega \mid \mathbf{F} + \mathbf{F} \mid \mathbf{F} \times \mathbf{F} \mid \mathbf{F}^D, D \in \mathbf{Set} \mid \mathbf{F} \circ \mathbf{F}.$$


we arrive at the definition of the *extended Kripke polynomial functors*. \triangleleft

As explained in the example one can prove the following fact.

Fact 2.10. *All extended Kripke polynomial functors preserve weak pullbacks.*

Although the precise definition of weak pullback preservation is not relevant in order to understand this paper, for the interested reader we provide some details in Remark 2.11 below.

Remark 2.11. Given two functions $f_0 : S_0 \rightarrow S$, $f_1 : S_1 \rightarrow S$, a *weak pullback* is a set P , together with two functions $p_i : P \rightarrow S_i$ such that $f_0 \circ p_0 = f_1 \circ p_1$, and in addition, for every triple (Q, q_0, q_1) also satisfying $f_0 \circ q_0 = f_1 \circ q_1$, there is an arrow $h : Q \rightarrow P$ such that $q_0 = h \circ p_0$ and $q_1 = h \circ p_1$.



A functor F *preserves weak pullbacks* if it transforms every weak pullback (P, p_0, p_1) for f_0 and f_1 into a weak pullback for Ff_0 and Ff_1 . (The difference with *pullbacks* is that in the definition of a weak pullback, the arrow h is not required to be unique.)

A category-theoretically nicer way of formulating this property involves the category Rel , i.e. the category of sets (as objects) and binary relations (as arrows). A functor Q on Rel is called a *relator* if for all binary relations R, S such that $R \subseteq S$ we have $Q(R) \subseteq Q(S)$. A relator Q *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 . Then one may prove that a set functor F can be extended to a relator iff F preserves weak pullbacks, and that this extension is unique if it exists. For a sketch of the proof of this fact note that it is easy to see (cf. e.g. [4, Chap. 5]) that any relator Q extending a set functor F satisfies $Q(R) = \bar{F}(R)$ for all binary relations R . Furthermore Trnková proved in [28] that \bar{F} is a relator iff F preserves weak pullbacks.

Finally, for an example of a functor that does not preserve weak pullbacks, consider the functor that takes a set S to the set of upward closed subsets of $(\mathcal{P}S, \subseteq)$, and a function f to its double inverse $(f^{-1})^{-1}$.

As mentioned already, in the following fact we list all the properties concerning relation lifting that we need in this paper.

Fact 2.12. *Let F be a set functor. Then the relation lifting \bar{F} satisfies the following properties, for all functions $f : S \rightarrow S'$, all relations $R, Q \subseteq S \times S'$, and all subsets $T \subseteq S$, $T' \subseteq S'$:*

- (1) \bar{F} extends F : $\bar{F}(\text{Gr}(f)) = \text{Gr}(Ff)$;
- (2) \bar{F} preserves the diagonal: $\bar{F}(\text{Id}_S) = \text{Id}_{FS}$;
- (3) \bar{F} commutes with relation converse: $\bar{F}(R^\smile) = (\bar{F}R)^\smile$;
- (4) \bar{F} is monotone: if $R \subseteq Q$ then $\bar{F}(R) \subseteq \bar{F}(Q)$;
- (5) \bar{F} distributes over composition: $\bar{F}(R \circ Q) = \bar{F}(R) \circ \bar{F}(Q)$, if F preserves weak pullbacks.

For proofs we refer to [19, 3], and references therein. The proof that Fact 2.12(5) depends on the property of weak pullback preservation goes back to Trnková [28]. In the remainder of the paper we will usually assume that all functors that we consider preserve weak pullbacks (but we will always mention this explicitly when we formulate important results).

2.2. Graph games. Two-player infinite graph games, or *graph games* for short, are defined as follows. For a more comprehensive account of these games, the reader is referred to GRÄDEL, THOMAS & WILKE [9].

First some preliminaries on sequences. Given a set A , let A^* , A^ω and A^\star denote the collections of finite, infinite, and all, sequences over A , respectively. (Thus, $A^\star = A^* \cup A^\omega$.) Given $\alpha \in A^*$ and $\beta \in A^\star$ we define the *concatenation* of α and β in the obvious way, and we denote this element of A^\star simply by juxtaposition: $\alpha\beta$. Given an infinite sequence $\alpha \in A^\omega$, let $\text{Inf}(\alpha)$ denote the set of elements $a \in A$ that occur infinitely often in α .

A graph game is played on a *board* B , that is, a set of *positions*. Each position $b \in B$ belongs to one of the two *players*, \exists (Éloise) and \forall (Abélard). Formally we write $B = B_\exists \cup B_\forall$, and for each position b we use $P(b)$ to denote the player i such that $b \in B_i$. Furthermore, the board is endowed with a binary relation E , so that each position $b \in B$ comes with a set $E[b] \subseteq B$ of *successors*. Formally, we say that the *arena* of the game consists of a directed two-sorted graph $\mathbb{B} = (B_\exists, B_\forall, E)$.

A *match* or *play* of the game consists of the two players moving a pebble around the board, starting from some *initial position* b_0 . When the pebble arrives at a position $b \in B$, it is player $P(b)$'s turn to move; (s)he can move the pebble to a new position of their liking, but the choice is restricted to a successor of b . Should $E[b]$ be empty then we say that player $P(b)$ *got stuck* at the position. A *match* or *play* of the game thus constitutes a (finite or infinite) sequence of positions $b_0b_1b_2\dots$ such that b_iEb_{i+1} (for each i such that b_i and b_{i+1} are defined). A *full play* is either (i) an infinite play or (ii) a finite play in which the last player got stuck. A non-full play is called a *partial play*.

The rules of the game associate a *winner* and (thus) a *loser* for each full play of the game. A finite full play is lost by the player who got stuck; the winning condition for infinite games is given by a subset Ref of B^ω (Ref is short for ‘referee’): our convention is that \exists is the winner of $\beta \in B^\omega$ precisely if $\beta \in \text{Ref}$. A *graph game* is thus formally defined as a structure $\mathcal{G} = (B_\exists, B_\forall, E, \text{Ref})$. Sometimes we want to restrict our attention to matches of a game with a certain initial position; in this case we will speak of a game that is *initialized* at this position.

Various kinds of winning conditions are known. In a *parity game*, the set Ref is defined in terms of a *parity function* on the board B , that is, a map $\Omega : B \rightarrow \omega$ with finite range. More specifically, the set Ref is defined by

$$B_\Omega^\omega := \{\beta \in B^\omega \mid \max(\text{Inf}(\Omega \circ \beta)) \text{ is even}\} \quad (2.2)$$

(where Inf was defined at the beginning of this subsection). In words, \exists wins a match if the highest parity encountered infinitely often during the match, is even.

A *strategy* for player i is a function mapping partial plays $\beta = b_0 \cdots b_n$ with $P(b_n) = i$ to admissible next positions, that is, to elements of $E[b_n]$. In such a way, a strategy tells i how to play: a play β is *conform* or *consistent with* strategy f for i if for every proper initial sequence $b_0 \cdots b_n$ of β with $P(b_n) = i$, we have that $b_{n+1} = f(b_0 \cdots b_n)$. A strategy is *history free* if it only depends on the current position of the match, that is, $f(\beta) = f(\beta')$ whenever

β and β' are partial plays with the same last element (which belongs to the appropriate player). Occasionally, it will be convenient to extend the name ‘strategy’ to arbitrary functions mapping partial plays to positions; in other words, we allow strategies enforcing *illegal* moves. In this context, the strategies proper, that is, the ones that dictate admissible moves only will be called *legitimate*.

A strategy is *winning for player i* from position $b \in B$ if it guarantees i to win any match with initial position b , no matter how the adversary plays — note that this definition also applies to positions b for which $P(b) \neq i$. A position $b \in B$ is called a *winning position* for player i , if i has a winning strategy from position b ; the set of winning positions for i in a game \mathcal{G} is denoted as $Win_i(\mathcal{G})$.

Parity games form an important game model because they have many attractive properties, such as *history-free determinacy*.

Fact 2.13. *Let $\mathcal{G} = (B_\exists, B_\forall, E, \Omega)$ be a parity graph game. Then*

(1) *\mathcal{G} is determined: $B = Win_\exists(\mathcal{G}) \cup Win_\forall(\mathcal{G})$.*

(2) *Each player i has a history-free strategy which is winning from any position in $Win_i(\mathcal{G})$.*

The determinacy of parity games follows from a far more general game-theoretic result concerning Borel games, due to MARTIN [17]. The fact that winning strategies in parity games can always be taken to be history free, was independently proved in MOSTOWSKI [20] and EMERSON & JUTLA [8].

2.3. Coalgebra automata and the acceptance game. For a detailed exposition of coalgebra automata, the reader is referred to VENEMA [30]. Here we confine ourselves to a self-contained survey of the definitions.

Probably the easiest introduction to coalgebra automata involves a reformulation of the notion of bisimilarity in game-theoretic terms. It follows from the characterization (2.1) that a bisimulation between two coalgebras $\mathbb{S} = \langle S, \sigma \rangle$ and $\mathbb{A} = \langle A, \alpha \rangle$ is nothing but a postfixpoint of the following operation on the set $\mathcal{P}(S \times A)$ of binary relations between S and A :

$$Z \mapsto \{(s, a) \in S \times A \mid (\sigma(s), \alpha(a)) \in \bar{F}(Z)\}.$$

By monotonicity of relation lifting, this operation is monotone, and thus it is an immediate consequence of standard fixpoint theory that there is a largest bisimulation between the two coalgebras, which is given as the union of all bisimulations between \mathbb{S} and \mathbb{A} . Furthermore, this largest bisimulation has a nice *game-theoretic* characterization, formulated in this generality for the first time in BALTAG [3].

For an informal description of this game, the admissible moves of the players are given as follows:

- in position (s, a) , \exists may choose a *local bisimulation* for s and a , i.e., a relation $Z \subseteq S \times A$ satisfying $(\sigma(s), \alpha(a)) \in \bar{F}Z$;
- in position $Z \subseteq S \times A$, \forall may choose any element (s', a') of Z .

Finally, the winning conditions for infinite matches of this game are straightforward: if \exists manages to survive all finite stages of a match, she is declared the winner of the resulting infinite match.

This bisimilarity game can be formulated as a graph game with a very simple parity winning condition (namely, all positions have the same, even, priority).

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

Table 1: Bisimilarity game for F-coalgebras

Definition 2.14. Let \mathbb{F} be a set functor, and let $\mathbb{S} = \langle S, \sigma \rangle$ and $\mathbb{A} = \langle A, \alpha \rangle$ be two \mathbb{F} -coalgebras. The *bisimilarity game* $\mathcal{B}(\mathbb{A}, \mathbb{S})$ associated with \mathbb{S} and \mathbb{A} is the parity graph game $(B_{\exists}, B_{\forall}, E, \Omega)$ with

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

while E and Ω are given in Table 1. Positions of the form $(s, a) \in S \times A$ are called *basic*. The set of winning positions for \exists in this game is denoted as $Win_{\exists}(\mathcal{B}(\mathbb{A}, \mathbb{S}))$, or Win_{\exists} if no confusion is likely. \triangleleft

Remark 2.15. We leave it for the reader to verify that

$$(s, a) \in Win_{\exists}(\mathcal{B}) \text{ iff } \mathbb{S}, s \Leftrightarrow \mathbb{A}, a.$$

The key observation here is that the relation $\{(s, a) \in S \times A \mid (s, a) \in Win_{\exists}(\mathcal{B})\}$ is the largest bisimulation between \mathbb{S} and \mathbb{A} .

In order to make the key observation in understanding coalgebra automata, we take a slightly *different perspective* on the bisimilarity game. The point is to think of one structure (\mathbb{A}) as *classifying* the other one (\mathbb{S}). This conceptual removal of the symmetry between the two structures, enables us to think of \mathbb{A} as an *automaton* operating on \mathbb{S} , and of the game as an *acceptance game* rather than as a comparison game. But from this perspective it is natural to impose some modifications on \mathbb{A} , making it resemble the standard concept of an automaton more closely:

To start with, the state space A of \mathbb{A} is required to be finite, and \mathbb{A} will have a fixed *initial state* $a_I \in A$.

Second, some infinite matches may be won by \forall . This can easily be implemented as follows. With each infinite match of the game we may associate an infinite stream of *basic positions* of the form $(s, a) \in S \times A$, and thus an infinite stream of states from A . Hence, the winning conditions in the game can be formulated using a subset $Acc \subseteq A^{\omega}$. And since A is finite we may formulate more specific conditions; for instance, a parity condition can be formulated by a map $\Omega : A \rightarrow \omega$.

And finally, we may introduce *nondeterminism* or even *alternation* on \mathbb{A} . Here the idea is that, whereas in a coalgebra $\mathbb{A} = \langle A, \alpha \rangle$, the ‘successor object’ of a state $a \in A$ is fixed as $\alpha(a) \in \mathbb{F}A$, we may now allow \exists to dynamically *choose* this object from some *set* $\Delta(a) \in \mathcal{P}(\mathbb{F}A)$ of objects in $\mathbb{F}A$. Or, the ‘successor object’ of $a \in A$ may be dynamically determined via some game-theoretic interaction between the *two* players.

Putting these observations together, we arrive at the following definition of VENEMA [29].

Convention 2.16. In the sequel we will frequently denote the power set $\mathcal{P}S$ of a set S by either $\mathcal{P}_{\exists}S$ or $\mathcal{P}_{\forall}S$. This notation indicates that we are in a game-theoretic context, where $X \in \mathcal{P}_iS$ means that X represents a collection of possible moves, and that i is the player who may choose an element from X .

Definition 2.17. Let F be some set functor. An (alternating) F -automaton is a quadruple $\mathbb{A} = \langle A, a_I, \Delta, Acc \rangle$ with A some finite set of objects called *states*, $a_I \in A$ the *initial state*, $\Delta : A \rightarrow \mathcal{P}_{\exists} \mathcal{P}_{\forall} F A$ the *transition function* and $Acc \subseteq A^\omega$ the *acceptance condition*. An F -automaton is called *nondeterministic* if each member of each $\Delta(a)$ is a singleton set. The *size* of \mathbb{A} is defined as the number of elements of A .

A *parity* F -automaton is an F -automaton $\mathbb{A} = \langle A, a_I, \Delta, Acc \rangle$ where Acc is given by some parity condition $\Omega : A \rightarrow \omega$; such a structure will usually be denoted as $\mathbb{A} = \langle A, a_I, \Delta, \Omega \rangle$. The *index* of a parity automaton is defined as the size of the range of Ω . \triangleleft

F -automata are designed to accept or reject pointed F -coalgebras. The acceptance condition is formulated in terms of a graph game. For an informal description of this game, the first observation is that matches of this game proceed in *rounds* that start and end in a *basic position* of the form $(s, a) \in S \times A$. From such a basic position (s, a) , a round of the match proceeds along the moves (a) - (d) below (of course, unless one of the players gets stuck):

- (a) \exists picks an element $\Phi \in \Delta(a)$, making $(s, \Phi) \in S \times \mathcal{P}_{\forall} F A$ the next position.
- (b) \forall picks $\varphi \in \Phi$, moving to position $(s, \varphi) \in S \times F A$.

Note that this interaction between \exists and \forall has fixed the ‘successor object’ $\varphi \in F A$ of a , whereas the ‘successor object’ $\sigma(s)$ of s was determined from the outset of the match.

- (c) \exists picks a ‘local bisimulation’ Y for φ and $\sigma(s)$, that is, a binary relation $Y \subseteq S \times A$ such that $(\sigma(s), \varphi) \in \overline{F} Y$. This relation $Y_{s, \varphi}$ is itself the new position.
- (d) \forall chooses a pair $(t, b) \in Y$ as the next basic position.

In the sequel we will refer to the first two moves in the round of the game as the *static* part of the round (static because the match does not pass to another state in the coalgebra), and to the last two moves as the *dynamic* or *coalgebraic* stage of the round.

For the winning conditions, recall that finite matches are lost by the player who gets stuck. For infinite matches, consider an arbitrary such match:

$$\mu = (s_0, a_0)(s_0, \Phi_0)(s_0, \varphi_0)Y_0(s_1, a_1)(s_1, \Phi_1)(s_1, \varphi_1)Y_1(s_2, a_2) \dots$$

Clearly, μ induces an infinite sequence of basic positions

$$(s_0, a_0)(s_1, a_1)(s_2, a_2) \dots$$

and, thus, an infinite sequence of states in A :

$$\mu \upharpoonright_A := a_0 a_1 a_2 \dots$$

Now the winner of the match is determined by whether $\mu \upharpoonright_A$ belongs to the set Acc or not.

Definition 2.18. Let $\mathbb{A} = \langle A, a_I, \Delta, Acc \rangle$ be an F -automaton, and let $\mathbb{S} = \langle S, \sigma \rangle$ be an F -coalgebra. The *acceptance game* $\mathcal{G}(\mathbb{A}, \mathbb{S})$ associated with \mathbb{A} and \mathbb{S} is the graph game $(B_{\exists}, B_{\forall}, E, \underline{Acc})$ with

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

where E is given in Table 2. Positions of the form $(s, a) \in S \times A$ are called *basic*.

For the winning conditions of $\mathcal{G}(\mathbb{A}, \mathbb{S})$, observe that every infinite match μ induces an infinite sequence $\mu \upharpoonright_A := a_0 a_1 a_2 \dots \in A^\omega$. We put $\mu \in \underline{Acc}$ if $\mu \upharpoonright_A \in Acc$, i.e. the winner of μ is \exists if $\mu \upharpoonright_A \in Acc$, and \forall otherwise.

The set of winning positions for \exists in this game is denoted as $Win_{\exists}(\mathcal{G}(\mathbb{A}, \mathbb{S}))$, or Win_{\exists} if no confusion is likely. \mathbb{A} *accepts* the pointed F -coalgebra (\mathbb{S}, s) if $(s, a_I) \in Win_{\exists}$. \triangleleft

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

Table 2: Acceptance game for an F-automaton

Remark 2.19. It is clear from the definition of \underline{Acc} 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 by the sequence of basic positions visited during the match.

It is easy to see that the acceptance games associated with *parity* automata, are parity games. (Simply define the priority of a basic position (s, a) as $\underline{\Omega}(s, a) := \Omega(a)$, putting $\underline{\Omega}(p) := 0$ for all other positions.) But 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 2.20. Given an F-coalgebra (S, σ) and a parity F-automaton \mathbb{A} a *positional* or *history-free* strategy for \exists is a pair of functions

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

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

Remark 2.21. In the case of a nondeterministic F-automaton \mathbb{A} we shall usually simplify our notation a little. Recall that the transition map of such an automaton is of the form $\Delta : A \rightarrow \mathcal{P}_{\exists}\mathcal{P}_{\forall}FA$, with each element of each $\Delta(a)$ a singleton. As a consequence, the move of \forall in the static part of the game is completely determined — he has nothing to choose. Consequently, we may eliminate these vacuous moves from the game by simplifying the presentation of the automaton.

Identifying singleton sets with their unique elements, we think of the transition function Δ as a map of type $A \rightarrow \mathcal{P}_{\exists}FA$. Accordingly then, we present the first component Φ of a positional strategy (Φ, Y) for \exists as a function of type $S \times A \rightarrow FA$.

It should be stressed that, in the case of automata operating on well-known infinite objects such as labelled binary trees, we have not really introduced a new kind of device, but rather, given a slightly different presentation of the more standard automata.

Example 2.22. Consider the case of binary tree automata over an alphabet C . In our presentation, C -labelled binary trees are coalgebras for the functor \mathbb{B}_C , with $\mathbb{B}_C(S) = C \times S \times S$, see Example 2.4.

The transition map of nondeterministic tree automata is usually presented in the form

$$\Delta : A \times C \rightarrow \mathcal{P}(A \times A),$$

whereas in our presentation, following Remark 2.21, the transition map is of the form

$$\Delta : A \rightarrow \mathcal{P}_{\exists}(C \times A \times A).$$

It is not difficult to see that these two presentations are in fact equivalent. Using the principle of currying $(P \times Q) \rightarrow R \cong P \rightarrow (Q \rightarrow R)$, and the notion of characteristic function $\mathcal{P}(Q) \cong Q \rightarrow 2$, we obtain

$$\begin{aligned} (A \times C) \rightarrow \mathcal{P}(A \times A) &\cong (A \times C) \rightarrow ((A \times A) \rightarrow 2) \\ &\cong (A \times C \times A \times A) \rightarrow 2 \\ &\cong A \rightarrow ((C \times A \times A) \rightarrow 2) \\ &\cong A \rightarrow \mathcal{P}(C \times A \times A). \end{aligned}$$

In [30] the second author explains the equivalence between the two presentations in detail.

In fact, in the case the functor is of the form F_C for some functor F and color set C — that is, if we are investigating C -colored F -coalgebras, we could have defined coalgebra automata in a different way, which is more in line with the standard usage. This alternative definition leads to the notion of *chromatic F-automata* [29], which is needed (only) in subsection 5.3.

Definition 2.23. Let C be a finite set. A C -chromatic F -automaton is a quintuple $\mathbb{A} = \langle A, a_I, C, \Delta, \Omega \rangle$ such that $\Delta : A \times C \rightarrow \mathcal{P}_{\exists} \mathcal{P}_{\forall} F A$ (and A , a_I , and Ω are as before). Given such an automaton and a F_C -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

It was shown in [29] that C -chromatic F -automata and F_C -automata have the same recognizing power. We need the following fact.

Fact 2.24. *With any parity F_C -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 F_C -coalgebras.*

2.4. Standardization. We already mentioned that we will work with a weak pullback preserving functor F throughout the paper. Moreover, for a smooth presentation, it will sometimes be useful to require the functor F to be *standard*. Of course we will always clearly state when exactly we assume the property of standardness. The purpose of this section is to convince the reader that the restriction to standard set functors is not essential, in that every set functor is ‘almost standard’. Let us start by formally defining the notion of a standard set functor.

Definition 2.25. Given two sets S and T such that $S \subseteq T$, let $\iota_{S,T}$ denote the inclusion map from S into T . A set functor F is *standard* if $F \iota_{S,T} = \iota_{FS,FT}$ for every inclusion map $\iota_{S,T}$. \triangleleft

Many but not all set functors have this property. For instance, all Kripke polynomial functors of Definition 2.9 are standard, but not the multiset functor of Example 2.4.

In words, a set functor is standard iff it turns inclusions into inclusions. This means that in particular, $S \subseteq T$ implies $FS \subseteq FT$. An immediate observation is that standardness ensures that the definition of the lifting of a relation R is independent of its type (cf. Remark 2.6).

Proposition 2.26. *Let F be a standard set functor, let S', S, T', T be sets and let $R' \subseteq S' \times T'$ be a relation. Furthermore let $R \subseteq S \times T$ be the relation R' - but now seen as a relation between S and T . Then the relations $\overline{F}R$ and $\overline{F}R'$ are equal.*

Proof. In order to prove the proposition, let $R' \subseteq S' \times T'$ and $R \subseteq S \times T$ represent the same relation, but with different type information. We prove the proposition under the additional assumption that $S' \subseteq S$ and $T' \subseteq T$. The case in which this is not true can be reduced to this special case by considering the relation $R'' := R \cap ((S' \cap S) \times (T' \cap T))$, which has to be equal to both R and R' . It then follows from our simpler claim that $\bar{F}R = \bar{F}R'' = \bar{F}R'$.

We now turn to the proof for the case that $S' \subseteq S$ and $T' \subseteq T$. The situation can be summarized by the following commuting diagram:

$$\begin{array}{ccccc} S' & \xleftarrow{\pi'_1} & R' & \xrightarrow{\pi'_2} & T' \\ \iota_{S',S} \downarrow & & \parallel & & \downarrow \iota_{T',T} \\ S & \xleftarrow{\pi_1} & R & \xrightarrow{\pi_2} & T \end{array}$$

Here the ι 's denote the inclusion maps. Moreover one has to keep in mind that R' and R denote the same set - we only use two distinct letters in order to be able to distinguish between the two representations of the relation R . If we apply the standard functor F to this diagram we get

$$\begin{array}{ccccc} FS' & \xleftarrow{F\pi'_1} & FR' & \xrightarrow{F\pi'_2} & FT' \\ F\iota_{S',S} = \iota_{FS',FS} \downarrow & & \parallel & & \downarrow F\iota_{T',T} = \iota_{FT',FT} \\ FS & \xleftarrow{F\pi_1} & FR & \xrightarrow{F\pi_2} & FT \end{array}$$

Therefore we can calculate that

$$\begin{aligned} (x, y) \in \bar{F}R' & \text{ iff there is a } z \in FR' \text{ with } F\pi'_1(z) = x \text{ and } F\pi'_2(z) = y \\ & \text{ iff there is a } z \in FR' \text{ with } (\iota_{FS',FS} \circ F\pi'_1)(z) = F\pi'_1(z) = x \\ & \quad \text{and } (\iota_{FT',FT} \circ F\pi'_2)(z) = F\pi'_2(z) = y \\ & \text{ iff there is a } z \in FR \text{ with } F\pi_1(z) = x \text{ and } F\pi_2(z) = y \\ & \text{ iff } (x, y) \in \bar{F}R \end{aligned}$$

□

As already mentioned every weak pullback preserving set functor is 'almost' standard. This statement is made formal using the notion of a natural isomorphism between functors.

Definition 2.27. Let F and G be two set functors, and suppose that for every set S there is a bijection $\lambda_S : FS \rightarrow GS$. This collection λ is a *natural isomorphism* between F and G , if $(Gf) \circ \lambda_S = \lambda_T \circ (Ff)$, for every $f : S \rightarrow T$:

$$\begin{array}{ccc} S & & FS \xrightarrow{\lambda_S} GS \\ f \downarrow & & \downarrow Ff \quad \downarrow Gf \\ T & & FT \xrightarrow{\lambda_T} GT \end{array}$$

In this situation, we say that F is *naturally isomorphic* to G via λ , notation: $\lambda : F \cong G$. If F is naturally isomorphic to a standard functor G , we call G a *standardization* of F . ◁

Naturally isomorphic functors are ‘almost the same’. For instance, it is not hard to show that if F and G are naturally isomorphic, then the categories of F -coalgebras is isomorphic to that of G -coalgebras. The following fact shows that the requirement of the functor F to be standard is not essential at all. A proof of this fact can be based on the construction in part (a) of the proof of Theorem III.4.5 in [1].²

Fact 2.28. *Every weak pullback preserving functor has a standardization.*

3. FROM ALTERNATION TO NONDETERMINISM

In this section we construct, for an arbitrary alternating parity F -automaton, an equivalent *nondeterministic* 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 briefly mention the intuitions behind our approach.

Remark 3.1. These intuitions ultimately go back to ideas of Muller and Schupp, see for instance [21], but in particular, our proof generalizes work by Janin and Walukiewicz [12], 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. That, however, would be to miss our point that a *coalgebraic* proof is possible, which is both *uniform*, in the sense that it is parametrical in the functor F , and *concrete* in the sense that we give an explicit definition which constructs the nondeterministic equivalent.

Consider a single round of the acceptance game $\mathcal{G}(\mathbb{A}, \mathbb{S})$ for some F -coalgebra \mathbb{S} , starting at a basic position $(s, a) \in S \times A$, with \exists employing some positional strategy (Φ, Z) :

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

Our proof is based on the following four ideas:

strategic normal form: First, we may bring the players’ interaction pattern $\exists\forall\exists\forall$ in each round of the acceptance game for \mathbb{A} , into the *strategic form* $\exists\forall$ (or more precisely: $\exists\exists\forall\forall$). Concretely, instead of choosing a separate $Y_{s,\varphi}$ for each $\varphi \in \Phi$, we will show that \exists may in fact choose the same relation $\bigcup_{\varphi \in \Phi} Y_{s,\varphi}$ in response to each $\varphi \in \Phi$ picked by \forall .

relations as states: The crucial part of our proof involves a natural refinement of the classical power set construction which is used for the determinization of automata operating on finite words. We will define a nondeterministic automaton \mathbb{A}^\sharp based on the set $A^\sharp := \mathcal{P}(A \times A)$ of *binary relations* over the state space A of \mathbb{A} . The acceptance condition of \mathbb{A}^\sharp is phrased in terms of *traces* through infinite sequences of such relations.

regular automata: The nondeterministic automaton \mathbb{A}^\sharp is nonstandard in the sense that its acceptance condition is expressed as an ω -regular language Acc over the set A . In the next section we will show that any such automaton is equivalent to a

²Note that the construction in *loc.cit.* requires the functor to preserve arbitrary monomorphisms. It is not difficult to see that weak pullback preserving functors meet this requirement.

standard nondeterministic automaton which is obtained as a kind of wreath product of \mathbb{A}^\sharp with the deterministic word automaton recognizing the set of infinite A -words in Acc .

coalgebraic perspective: While none of the above ideas in itself is essentially new, we believe that our coalgebraic perspective simplifies matters. It enables us to carry out the entire construction uniformly in the functor, with *relation lifting* (see Definition 2.5) being the novel, unifying concept.

Let us now look at the construction in more detail. Before arriving at the actual definition of the regular, nondeterministic automaton \mathbb{A}^\sharp , we discuss and motivate the ideas mentioned above.

3.1. Relations and traces. We start with motivating the use of *binary relations* on A as the states of \mathbb{A}^\sharp .

First, the construction is based on the principle that \exists should be prepared to counter many of \forall 's moves *simultaneously*. Intuitively, then, it would be a natural move to construct an automaton \mathbb{A}^* taking subsets of A as its states. Such a *macro-state* would represent the set of states that \mathbb{A} could be in and that \exists should be able to somehow handle simultaneously. Building on this intuition we could proceed to give a precise definition of the automaton \mathbb{A}^* , generalizing the subset construction for automata over finite words.

Continuing along these lines, we might establish a tight link between the basic positions

$$(s_0, \{a_I\})(s_1, B_1) \dots (s_k, B_k) \quad (3.1)$$

of a partial match of $\mathcal{G}(\mathbb{A}^*, \mathbb{S})$, and a *collection* of partial matches

$$(s_0, a_I)(s_1, a_1) \dots (s_k, a_k) \quad (3.2)$$

in $\mathcal{G}(\mathbb{A}, \mathbb{S})$ such that every a_i is an element of B_i . This link would then naturally extend to infinite matches.

Unfortunately however, we encounter a difficulty when we try to formulate an adequate acceptance condition for \mathbb{A}^* . The problem is that, just on the basis of an infinite sequence of subsets of \mathbb{A} , we may fail to make some subtle but crucial distinctions. The point is that the acceptance condition for \mathbb{A}^* should declare \exists as the winner of the match (3.1) if and only if she is the winner of each associated match of the form (3.2). But we may mistakenly declare \forall as the winner of (3.1) on the basis of a sequence of the form (3.2), which satisfies $a_i \in B_i$ for each i , and meets the winning conditions for \forall , but *which did not come about as an actual match of $\mathcal{G}(\mathbb{A}, \mathbb{S})$ associated with (3.1)*. (As many readers will have recognized, this is exactly the problem one faces when transforming a nondeterministic word automaton into an equivalent deterministic one, and explains why the Safra construction is so much more involved than the power set construction.)

An elegant way to avoid this problem is to use *binary relations over A* rather than subsets. When considered statically, the relation R simply represents the macrostate $\text{rng}(R)$ (that is, the *range* of R). The additional structure of binary relations comes into play when we look at infinite sequences: The key notion of a *trace* through a sequence of binary relations allows us to make the required subtle distinctions referred to above.

Definition 3.2. Given an infinite word $\rho = R_1 R_2 R_3 \dots$ over the set A^\sharp of binary relations over a set A , a *trace* through ρ is a finite or infinite A -word $\alpha = a_0 a_1 a_2 \dots a_k$ or $\alpha = a_0 a_1 a_2 \dots$ such that $a_i R_{i+1} a_{i+1}$ for all i (respectively, for all $i < k$).

Relative to a priority map Ω on A , call a trace α *bad* if it is infinite and the maximum priority occurring infinitely often on α , is an odd number. Let NBT_Ω denote the set of infinite A^\sharp -words that contain no bad traces relative to Ω . \triangleleft

Since A^\sharp is the carrier set of \mathbb{A}^\sharp , infinite matches of the acceptance game $\mathcal{G}(\mathbb{A}^\sharp, \mathbb{S})$ induce A^\sharp -streams, whereas traces on such a sequence may be linked to associated matches of the acceptance game for \mathbb{A} . Thus it will be natural to declare \exists as the winner of a $\mathcal{G}(\mathbb{A}^\sharp, \mathbb{S})$ -match if there is no bad trace on the induced infinite A^\sharp -sequence, since bad traces correspond to $\mathcal{G}(\mathbb{A}, \mathbb{S})$ matches that are won by \forall . This explains the acceptance condition NBT_Ω of the automaton \mathbb{A}^\sharp .

Remark 3.3. Note, however, that this acceptance condition is not a parity condition. This means that the automaton \mathbb{A}^\sharp is not the automaton \mathbb{A}^\bullet referred to in the statement of Theorem 1. In order to transform \mathbb{A}^\sharp into a nondeterministic parity automaton, we need to prove that \mathbb{A}^\sharp is a regular automaton, that is, its acceptance condition NBT_Ω is an ω -regular language, recognized by a word automaton. This result will be proved in section 4.

3.2. Normalized strategies. In the sequel it will be convenient to work with so-called *normalized* strategies for \exists . Intuitively, these are positional strategies for \exists that provide her, in the “dynamic”, second half of each round of the game $\mathcal{G}(\mathbb{A}, \mathbb{S})$, with a relation between S and A that does not depend on \forall 's move in the first part of the round.

For more details, suppose that Φ , together with $Y : S \times FA \rightarrow \mathcal{P}(S \times A)$, is a positional strategy for \exists in \mathcal{G} , and consider a basic position (s, a) . Here first \exists chooses an element $\Phi_{s,a} \in \Delta(a)$, and then, for every choice $\varphi \in \Phi_{s,a}$, she can choose a separate relation $Y_{s,\varphi} \subseteq \mathcal{P}(S \times A)$. If she uses a normalized strategy however, then her choice of $Y_{s,\varphi}$ must be independent of φ ; it may however depend on the earlier basic position (s, a) . Formally then, we model the dynamic part of a normalized strategy as a map $Z : S \times A \rightarrow \mathcal{P}(S \times A)$. Intuitively, $Z_{s,a}$ consists of those elements (t, b) that \exists may expect as the next basic position after (s, a) .

For technical reasons it will be convenient to add one more condition to the definition of a normalized strategy: In the static part of the game, we require \exists to head for an immediate win if there is one. More precisely, consider a state $a \in A$ such that $\Delta(a)$ contains the empty set \emptyset as a choice. Clearly, for such an a at any position (s, a) \exists may choose (s, \emptyset) as the next position and win immediately, since \forall cannot choose an element from the empty set. A normalized strategy requires \exists to indeed choose \emptyset in such a position.

Let us now first give the formal definition of a normalized strategy.

Definition 3.4. Given an alternating F-automaton \mathbb{A} and an F-coalgebra \mathbb{S} , a *normalized* strategy for \exists in the game $\mathcal{G}(\mathbb{A}, \mathbb{S})$ is a pair (Ψ, Z) where $\Psi : S \times A \rightarrow \mathcal{P}_\forall FA$ and $Z : S \times A \rightarrow \mathcal{P}(S \times A)$ are such that $\Psi_{s,a} = \emptyset$ if $\emptyset \in \Delta(a)$. \triangleleft

If \exists uses a normalized strategy (Φ, Z) , we can present the interaction pattern of the players per round as follows, starting at a basic position $(s, a) \in S \times A$:

- \exists picks and plays $\Phi_{s,a} \in \Delta(a)$
- \exists chooses a relation $Z_{s,a} \subseteq S \times A$;
- \forall picks and plays a $\varphi \in \Phi_{s,a}$;
- \exists plays $Z_{s,a}$;
- \forall picks and plays a pair $(t, b) \in Z_{s,a}$ as the next basic position.

The resulting interaction pattern is indeed of the earlier announced *strategic normal form* ‘ $\exists\exists\forall\forall$ ’ rather than of the form ‘ $\exists\forall\exists\forall$ ’.

The following proposition states that without loss of generality we may always assume that \exists ’s winning strategies are normalized.

Proposition 3.5. *Fix an alternating parity F-automaton \mathbb{A} and an F-coalgebra \mathbb{S} . Then there is a normalized strategy for \exists which is winning from every position $(s, a) \in \text{Win}_\exists(\mathcal{G}(\mathbb{A}, \mathbb{S}))$.*

Proof. Let \mathbb{A} and \mathbb{S} be as in the proposition. By the historyfree determinacy of the parity game $\mathcal{G} = \mathcal{G}(\mathbb{A}, \mathbb{S})$ we may assume the existence of a positional strategy

$$(\Phi : S \times A \rightarrow \mathcal{P}A, Y : S \times \text{FA} \rightarrow \mathcal{P}(S \times A))$$

which is winning for \exists from every position $(s, a) \in \text{Win}_\exists(\mathcal{G}(\mathbb{A}, \mathbb{S}))$.

Define the strategy (Ψ, Z) as follows:

$$\begin{aligned} \Psi_{s,a} &:= \begin{cases} \emptyset & \text{if } \emptyset \in \Delta(a) \\ \Phi_{s,a} & \text{otherwise} \end{cases} \\ Z_{s,a} &:= \bigcup_{\varphi \in \Psi_{s,a}} Y_{s,\varphi}. \end{aligned}$$

Consider *one round* of the game \mathcal{G} starting at a winning position for \exists . We claim that either \exists wins already during this round, or else the match arrives at a new basic position that could also have been reached if \exists had played her original strategy (Φ, Y) . From this claim one may derive that the strategy (Ψ, Z) guarantees \exists to win any match of \mathcal{G} starting at a position in $\text{Win}_\exists(\mathcal{G}(\mathbb{A}, \mathbb{S}))$.

To prove our claim, take a position $(s, a) \in \text{Win}_\exists(\mathcal{G}(\mathbb{A}, \mathbb{S}))$. To start with, it is easy to check that $\Psi_{s,a}$ is a legitimate move for \exists . If $\Psi_{s,a} = \emptyset$, then \exists wins immediately. So suppose otherwise, and let \forall pick an element $\varphi \in \Psi_{s,a} = \Phi_{s,a}$. It follows from $(s, a) \in \text{Win}_\exists(\mathcal{G}(\mathbb{A}, \mathbb{S}))$ and the fact that (Φ, Y) is a winning strategy for \exists , that $(\sigma(s), \varphi) \in \overline{\text{F}}Y_{s,\varphi}$. Hence, by the monotonicity of $\overline{\text{F}}$ (see Fact 2.12) and the definition of $Z_{s,a}$, we find that $(\sigma(s), \varphi) \in \overline{\text{F}}Z_{s,a}$, so that $Z_{s,a}$ is a legitimate answer to \forall ’s move φ . If $Z_{s,a} = \emptyset$ then \exists wins immediately, otherwise \forall may finish the round by picking an element $(t, b) \in Z_{s,a}$.

It remains to be shown that such an element $(t, b) \in Z_{s,a}$ could also have been obtained if \exists had played her original strategy (Φ, Y) . But it follows by definition of $Z_{s,a}$ that $(t, b) \in Y_{s,\psi}$ for some $\psi \in \Psi_{s,a}$. Thus $\Psi_{s,a} \neq \emptyset$, and so $\Psi_{s,a} = \Phi_{s,a}$. Hence the position (t, b) could have been reached in the scenario where \exists had played $\Phi_{s,a}$, followed by \forall picking $\psi \in \Phi_{s,a}$, \exists choosing the move $Y_{s,\psi}$, and, finally, \forall playing the pair $(t, b) \in Y_{s,\psi}$. The only thing left to verify here is the legitimacy of the move $Y_{s,\psi}$, but this is immediate by the assumption that Y is part of a winning strategy for \exists . \square

3.3. Normalized strategies and binary relations. In order to see how the ideas of the previous two subsections fit together, consider again a round of the acceptance game in which \exists uses a normalized strategy (Φ, Z) :

- \exists plays $\Phi_{s,a} \in \Delta(a)$ (and chooses $Z_{s,a} \subseteq S \times A$)
- \forall plays $\varphi \in \Phi_{s,a}$;
- \exists plays $Z_{s,a} \subseteq S \times A$;
- \forall plays $(t, b) \in Z_{s,a}$ as the next basic position.

The point about normalized strategies is that in fact the first two moves of such a round are only of interest if they lead to an immediate end of the match in that one of the players gets stuck, i.e., if either $\Delta(a)$ or $\Phi_{s,a}$ is empty. For infinite matches, the only relevant interaction is between \exists choosing binary relations between S and A , and \forall choosing elements of those relations. Let us look at this interaction in a bit more detail.

Recall that $Z_{s,a}$ contains those elements (t, b) that \exists ‘expects’ as the next basic position after (s, a) . Thus the dynamic part Z of \exists ’s strategy induces a tree of basic positions for \forall to choose from. We will now reorganize this tree, as follows.

First observe that, given a relation $Z_{s,a}$, for a single $t \in S$, there may be *many* elements $b \in A$ such that $(t, b) \in Z_{s,a}$. These are the states that \exists should prepare for to meet ‘simultaneously’ at the point $t \in S$, and that may be grouped together in a ‘macro-state’, as discussed earlier on. But then inductively, at s , the state a might already have been one of many parallel states in some macro-state. Here it starts making a lot of sense to involve binary relations: Instead of having macro-states $\{b \in A \mid (t, b) \in Z_{s,a}\}$, for each $a \in A$, we consider the binary relation

$$\zeta_s(t) := \{(a, b) \in A \times A \mid (t, b) \in Z_{s,a}\}. \quad (3.3)$$

Formally, we may represent the dynamic part Z of a normalized strategy as a map $\zeta : S \rightarrow (S \rightarrow \mathcal{P}(A \times A))$, where ζ_s is a map assigning a binary relation on A to each $t \in S$. The connection between Z and ζ is given by

$$(t, b) \in Z_{s,a} \iff (a, b) \in \zeta_s(t). \quad (3.4)$$

It is not hard to show that (3.4) induces a natural bijection

$$S \times A \rightarrow \mathcal{P}(S \times A) \cong S \rightarrow (S \rightarrow \mathcal{P}(A \times A)). \quad (3.5)$$

In fact, using currying ($P \rightarrow (Q \rightarrow R) \cong (P \times Q) \rightarrow R$) and exponentiation ($\mathcal{P}(Q) \cong Q \rightarrow 2$), it is very easy to see why (3.5) must hold:

$$S \times A \rightarrow \mathcal{P}(S \times A) \cong (S \times S \times A \times A) \rightarrow 2 \cong S \rightarrow (S \rightarrow \mathcal{P}(A \times A)). \quad (3.6)$$

Conversely, an explicit way of obtaining Z from ζ is as follows. Let the map $ev_a : \mathcal{P}(A \times A) \rightarrow \mathcal{P}A$ be given by $ev_a : R \mapsto R[a]$, and recall that the graph $\{(x, fx) \mid x \in X\} \subseteq X \times Y$ of a function $f : X \rightarrow Y$ is denoted as $Gr(f)$. It is then easy to see that

$$Z_{s,a} = Gr(\zeta_s) \circ Gr(ev_a) \circ \exists_A. \quad (3.7)$$

Summarizing the above discussion we give the following proposition, of which we will make heavily use in the sequel.

Proposition 3.6. *For any pair of sets S and A , there is a natural bijection between maps $Z : S \times A \rightarrow \mathcal{P}(S \times A)$ and functions $\zeta : S \rightarrow (S \rightarrow \mathcal{P}(A \times A))$. This correspondence is explicitly given by (3.3) and (3.7) above.*

3.4. The definition of \mathbb{A}^\sharp . We have already announced that \mathbb{A}^\sharp will be a nondeterministic regular F-automaton based on the collection A^\sharp of binary relations on A , and with acceptance condition of the form NBT_Ω , where Ω is the parity condition of \mathbb{A} . Thus to complete the definition of the automaton \mathbb{A}^\sharp it suffices to give the transition structure $\Delta^\sharp : A^\sharp \rightarrow \mathcal{P}(FA^\sharp)$.

Roughly speaking, it works like this. Earlier on we already briefly mentioned that, intuitively, a relation $R \in A^\sharp$ represents the macrostate $\text{rng}(R) \subseteq A$. Now suppose we consider the

static part of a strategy for \exists at a certain point of the coalgebra (see the discussion following Definition 2.17 for a division of a round of the acceptance game into a static and a dynamic part). In order to handle *each* of the challenges $a \in \text{rng}(R)$, \exists needs to come up with a *family*

$$\{\Phi(a) \in \mathcal{P}(FA) \mid a \in \text{rng}(R)\} \text{ such that } \Phi(a) \in \Delta(a) \text{ for each } a \in \text{rng}(R). \quad (3.8)$$

Our definition of Δ^\sharp will be such that given $R \in A^\sharp$, the members of $\Delta^\sharp(R)$ are those elements $\Pi \in FA^\sharp$ that are in a natural correspondence with such a family.

The key to understanding this ‘natural correspondence’ is the notion of *F-redistribution*, which links sets of the form FPA and $\mathcal{P}FA$. For an introduction to this notion, first consider the membership relation \in_A on A . Since \in_A is a binary relation between A and $\mathcal{P}(A)$, we may lift it to a relation $\bar{F}\in_A$ between FA and FPA . The relation $\bar{F}\in_A$ is like the membership relation ‘behind an F-veil’. Now suppose that $\Phi \in \mathcal{P}FA$ and $\Xi \in FPA$ satisfy the condition that *each* element φ of Φ is such an ‘F-member’ of Ξ , i.e., $(\varphi, \Xi) \in \bar{F}\in_A$. In such a case we call Ξ an F-redistribution of Φ , and it makes sense to think of Ξ as a representation of Φ as a set of type FPA .

Definition 3.7. Given a set A with membership relation $\in_A \subseteq A \times \mathcal{P}(A)$, we call $\Xi \in FPA$ an *F-redistribution* of $\Phi \in \mathcal{P}FA$, or say that Ξ *redistributes* Φ , if $(\varphi, \Xi) \in \bar{F}\in_A$ for all $\varphi \in \Phi$. \triangleleft

Example 3.8. For the binary tree functor B_C , an element $(c, a^l, a^r) \in B_C A$ is an B_C -member of an object $(d, A^l, A^r) \in B_C(A)$ if $c = d$, $a^l \in A^l$ and $a^r \in A^r$. Hence $(d, A^l, A^r) \in B_C(A)$ is a B_C -redistribution of the set $\{(c_i, a_i^l, a_i^r) \mid i \in I\}$ iff $c_i = d$, $a_i^l \in A^l$ and $a_i^r \in A^r$, for each $i \in I$.

For the power set functor P , an object $X \in PA = \mathcal{P}(A)$ is a P -element of an object $B \in PP(A) = \mathcal{P}(\mathcal{P}(A))$ iff $X \subseteq \bigcup B$ and $X \cap B \neq \emptyset$ for all $B \in \mathcal{B}$. So $\mathcal{B} \in PP(A)$ is a P -redistribution of $\mathcal{X} \in \mathcal{P}(PA)$ iff $\bigcup \mathcal{X} \subseteq \bigcup \mathcal{B}$ and $X \cap B \neq \emptyset$ for all $X \in \mathcal{X}$ and all $B \in \mathcal{B}$.

Remark 3.9. In [11] Jacobs shows that for every weak pullback preserving functor F there is a so-called *distributive law* $\lambda : FP \Rightarrow PF$ of F over the power set monad, i.e. λ is a natural transformation that preserves the monad structure. This distributive law is defined using the relation lifting of the \in -relation. Therefore there is a close connection between Jacobs’s law and our *F*-redistributions: $\Xi \in FPA$ is an *F*-redistribution of $\Phi \in \mathcal{P}FA$ iff $\Phi \subseteq \lambda_A(\Xi)$.

We are now almost ready for the definition of Δ^\sharp . For the final step, recall that for any element $a \in A$ we may go from A^\sharp to $\mathcal{P}(A)$ using the *evaluation map*

$$ev_a : R \mapsto R[a].$$

Thus $Fev_a : FA^\sharp \rightarrow F\mathcal{P}(A)$. The function Fev_a enables us to link potential elements $\Pi \in \Delta^\sharp(R)$ to redistributions in $F\mathcal{P}(A)$ of objects $\Phi(a) \in \Delta(a)$.

Definition 3.10. Let F be a set functor that preserves weak pullbacks, and let $\mathbb{A} = \langle A, a_I, \Delta, \Omega \rangle$ be an alternating *F*-automaton. Then the automaton \mathbb{A}^\sharp is defined as the structure

$$\mathbb{A}^\sharp := \langle A^\sharp, R_I, \Delta^\sharp, \text{NBT}_\Omega \rangle,$$

where $A^\sharp := \mathcal{P}(A \times A)$ is the collection of binary relations over A , $R_I = \{(a_I, a_I)\}$, $\Delta^\sharp : A^\sharp \rightarrow \mathcal{P}_\exists FA^\sharp$ is given by

$$\Delta^\sharp(R) := \{ \Pi \in FA^\sharp \mid \forall a \in \text{rng}(R) \exists \Phi(a) \in \Delta(a) \\ (Fev_a)(\Pi) \text{ is an F-redistribution of } \Phi(a) \},$$

and $\text{NBT}_\Omega \subseteq (A^\sharp)^\omega$ is the set of those infinite sequences of binary relations that do not contain any bad trace. \triangleleft

It is obvious that \mathbb{A}^\sharp is a nondeterministic F-automaton. It remains to be shown that \mathbb{A}^\sharp is equivalent to \mathbb{A} .

3.5. Proof of equivalence.

Proposition 3.11. *Let F be a set functor that preserves weak pullbacks, and let $\mathbb{A} = \langle A, a_I, \Delta, \Omega \rangle$ be an alternating parity F-automaton. Then \mathbb{A}^\sharp is equivalent to \mathbb{A} .*

Proof. Fix a coalgebra \mathbb{S} and a point s_0 in \mathbb{S} . We will prove the following equivalence:

$$\mathbb{A} \text{ accepts } (\mathbb{S}, s_0) \iff \mathbb{A}^\sharp \text{ accepts } (\mathbb{S}, s_0). \quad (3.9)$$

Obviously, both directions of this equivalence will be proved via a comparison of the two acceptance games $\mathcal{G} := \mathcal{G}(\mathbb{A}, \mathbb{S})$ and $\mathcal{G}^\sharp := \mathcal{G}(\mathbb{A}^\sharp, \mathbb{S})$.

\Rightarrow For the direction from left to right, assume that \mathbb{A} accepts (\mathbb{S}, s_0) . Then by Proposition 3.5 we may assume that \exists has a normalized strategy (Φ, Z) which is winning for her in the game \mathcal{G} initialized at (s_0, a_I) . In the sequel we will also make use of the map $\zeta : S \times S \rightarrow A^\sharp$ that is associated with Z as in Proposition 3.6.

In order to prove that (s_0, R_I) is a winning position for \exists in the game \mathcal{G}^\sharp , we let \exists play according to the following positional strategy (Π, Q) . The *static* part $\Pi : S \times A^\sharp \rightarrow FA^\sharp$ of this strategy is given by

$$\Pi_{s,R} := (F\zeta_s)(\sigma(s)),$$

while the *dynamic* part $Q : S \times FA^\sharp \rightarrow \mathcal{P}(S \times A^\sharp)$ is defined as

$$Q_{s,\Sigma} := Gr(\zeta_s).$$

Since the function Q only depends on its first argument, in the sequel we will simply write Q_s instead of $Q_{s,\Sigma}$.

The claims 2 and 3 below state that, playing this strategy, \exists wins all finite, respectively, infinite matches. In the proof of these results we need the following additional claim which, roughly spoken, states that the strategy defined above is legitimate at any safe position (s, R) of \mathcal{G}^\sharp , and guarantees that the next basic position is safe as well. Here we call a basic position (s, R) of \mathcal{G}^\sharp *safe* if $(s, a) \in \text{Win}_\exists(\mathcal{G})$ for all $a \in \text{rng}(R)$.

Claim 1. Let (s, R) be a position of \mathcal{G}^\sharp such that $(s, a) \in \text{Win}_\exists(\mathcal{G})$ for all $a \in \text{rng}(R)$. Then

- (1) both Π and Q provide legitimate moves at (s, R) ,
- (2) $(t, b) \in Z_{s,a}$ for all $(t, R') \in Q_s$, all $a \in A$ and all $b \in R'[a]$.

PROOF OF CLAIM The main part of the proof consists in showing that $\Pi := \Pi_{s,R}$ is a legitimate move for \exists at position (s, R) .

In order to show that, indeed, $\Pi \in \Delta^\sharp(R)$, consider an arbitrary element $a \in \text{rng}(R)$. By assumption, $(s, a) \in \text{Win}_\exists(\mathcal{G})$. Recall that $\Phi_{s,a}$ and $Z_{s,a}$ are the moves of \exists in \mathcal{G} prescribed by \exists 's winning normalized strategy. Take an arbitrary element $\varphi \in \Phi_{s,a}$. It suffices to prove that

$$(\varphi, (Fev_a)(\Pi)) \in (\bar{F} \in_A), \quad (3.10)$$

since this implies that $(Fev_a)(\Pi)$ is an F-redistribution of $\Phi_{s,a} \in \Delta(a)$, and thus that $\Pi \in \Delta^\sharp(R)$, since a was arbitrary.

It follows from the fact that $Z_{s,a}$ is part of a winning, and thus legitimate strategy, that

$$(\sigma(s), \varphi) \in \bar{F}Z_{s,a}. \quad (3.11)$$

Now from $Z_{s,a} = Gr(\zeta_s) \circ Gr(ev_a) \circ \ni_A$, (see (3.7) and some elementary properties of relation lifting (cf. Fact 2.12)), it follows that

$$\bar{F}(Z_{s,a}) = Gr(F\zeta_s) \circ Gr(Fev_a) \circ \bar{F}\ni_A.$$

Thus from (3.11) and the fact that $\Pi = (F\zeta_s)(\sigma(s))$ is defined as the *unique* object such that $(\sigma(s), \Pi) \in Gr(F\zeta_s)$, it is immediate that

$$(\Pi, \varphi) \in Gr(Fev_a) \circ \bar{F}\ni_A,$$

which is easily seen to be equivalent to (3.10).

To finish the proof of part (1) of the claim, it then suffices to show that Q_s is a legitimate move at position (s, Π) (where still we write $\Pi = \Pi_{s,R}$). But this is immediate by the definitions. The point is that from $\Pi = (F\zeta_s)(\sigma(s))$ we may infer $(\sigma(s), \Pi) \in Gr(F\zeta_s) = \bar{F}Gr(\zeta_s) = \bar{F}Q_s$.

Part (2) of the claim is also straightforward. Let $a, b \in A$, $t \in S$ and $R' \in A^\sharp$ be such that $(t, R') \in Q$ and $(a, b) \in R'$. Recall that by definition of Q , $(t, R') \in Q$ implies that $R' = \zeta_s(t)$, so by (3.4) we have $(a, b) \in R'$ iff $(t, b) \in Z_{s,a}$. \triangleleft

Claim 2. As long as \exists plays her strategy (Π, Q) , she wins all finite matches starting at position (s_0, R_I) .

PROOF OF CLAIM A straightforward inductive proof using part (2) of Claim 1 shows that any partial \mathcal{G}^\sharp -match $(s_0, R_I)(s_1, R_1) \dots (s_n, R_n)$ of \mathcal{G}^\sharp in which \exists plays her strategy (Π, Q) has the property that

$$(s_n, b) \in Win_{\exists}(\mathcal{G}) \text{ for all } b \in \text{rng}(R_n).$$

Then by part (1) of Claim 1 it follows that Π and Q provide legitimate moves for \exists . In other words, she will not get stuck after position (s_n, R_n) . \triangleleft

Claim 3. As long as she plays her strategy (Π, Q) , \exists wins all infinite matches starting at position (s_0, R_I) .

PROOF OF CLAIM Consider an infinite match

$$(s_0, R_0)(s_1, R_1) \dots$$

of \mathcal{G}^\sharp in which \exists plays her strategy (Π, Q) (and with $R_I = R_0$). In order to show that this match is won by \exists , consider an arbitrary trace on the sequence $R_I R_1 R_2 \dots$. It suffices to show that this trace is even.

Clearly the trace is of the form $a_0 a_0 a_1 a_2 \dots$ with $a_I = a_0$, $a_0 R_0 a_0$ and $a_i R_{i+1} a_{i+1}$ for every i . A direct inductive proof, using part (2) of Claim 1, shows that $(s_{i+1}, a_{i+1}) \in Z_{s_i, a_i}$ for every i . From this it is easy to find a match

$$(s_0, a_I)(s_1, a_1) \dots$$

of \mathcal{G} in which \exists plays her strategy (Φ, Z) , cf. the proof of Proposition 3.5. But by assumption, this strategy is *winning* for \exists , so the trace $a_0 a_0 a_1 a_2 \dots$ is indeed *even*. \triangleleft

Finally, the direction \Rightarrow of (3.9) is a direct consequence of the Claims 2 and 3.

$\boxed{\Leftarrow}$ For the direction from right to left, assume that \mathbb{A}^\sharp accepts (\mathbb{S}, s_0) . In other words, we may assume that there is a strategy f which is winning for \exists in the game \mathcal{G}^\sharp starting at (s_0, R_I) . In order to show that \mathbb{A} accepts (\mathbb{S}, s_0) , we need to prove that (s_0, a_I) is a winning position for \exists in \mathcal{G} .

We will equip \exists with a strategy f' , in the game \mathcal{G} initialized at (s_0, a_I) , which has the following property. For any (possibly finite) f' -conform match $(s_0, a_0)(s_1, a_1) \dots$ of \mathcal{G} with $a_0 = a_I$, there is an f -conform match $(s_0, R_0)(s_1, R_1) \dots$ of \mathcal{G}^\sharp , with $R_0 = R_I$, satisfying the condition that

$$a_{i+1} \in R_{i+1}[a_i] \text{ for every stage } i. \quad (3.12)$$

Hence, the sequence of \mathbb{A} -states $a_0 a_1 a_2 \dots$ of such a match is a *trace* of the A^\sharp -sequence $R_0 R_1 R_2 \dots$ which we may associate with an f -conform match. Since f is by assumption winning for \exists , by definition of the winning condition NBT_Ω of \mathbb{A}^\sharp , the (maximum parity occurring infinitely often on) the trace must be *even*.

This guarantees that she wins all *infinite* matches of the game. Hence, it suffices to prove that at any finite stage of an f' -conform match, she either immediately, or else she can keep the above condition for one more round.

Suppose then that \exists has been able to keep this condition for k steps. That is, with the partial \mathcal{G} -match $(s_0, a_0) \dots (s_k, a_k)$ (where $a_I = a_0$) we may associate a partial, f -conform \mathcal{G}^\sharp -match $(s_0, R_0) \dots (s_k, R_k)$ such that $R_0 = R_I$ and

$$a_{i+1} \in R_{i+1}[a_i] \text{ for all } i < k. \quad (3.13)$$

For notational convenience, write $a = a_k$, $R = R_k$ and $s = s_k$, so we have $a \in \text{rng}(R)$. Let $\Pi \in \text{FA}^\sharp$ and $Q \subseteq S \times A^\sharp$, respectively, be the moves dictated by \exists 's winning strategy f in \mathcal{G}^\sharp . It follows from the fact that f is a winning strategy, that Π and Q are legitimate moves, that is, $\Pi \in \Delta^\sharp(R)$ and $(\sigma(s), \Pi) \in \bar{F}(Q)$. Then by definition of Δ^\sharp , and the fact that $a \in \text{rng}(R)$, there is some $\Phi \in \Delta(a)$ such that $(\text{Fev}_a)(\Pi)$ is an F-redistribution of Φ . This Φ is the next move of \exists in the game \mathcal{G} .

If $\Phi = \emptyset$ then \exists wins right away, in which case we are done immediately. So assume that $\Phi \neq \emptyset$, and suppose that \forall responds to \exists 's move with an object $\varphi \in \Phi$. Then \exists has to come up with a relation $Y \subseteq S \times A$ such that $(\sigma(s), \varphi) \in \bar{F}(Y)$. Our suggestion to \exists is to pick the relation given by

$$Y := Q \circ \text{Gr}(\text{ev}_a) \circ \ni_A,$$

or, spelled out,

$$Y = \{(t, b) \in S \times A \mid b \in R'[a] \text{ for some } R' \in A^\sharp \text{ with } (t, R') \in Q\}.$$

If this is a legitimate move for \exists , then we are done. For, distinguish the following cases. If $Y = \emptyset$ then \forall gets stuck so \exists wins immediately. But if $Y \neq \emptyset$, then with any $(s_{k+1}, a_{k+1}) \in Y$ that \forall chooses as his next move, by definition we may associate a relation $R_{k+1} \in A^\sharp$ such that $(a_k, a_{k+1}) \in R_{k+1}$ and $(s_{k+1}, R_{k+1}) \in Q$. In other words, we have showed that \exists can indeed maintain the above mentioned condition (3.13) for one more round of the game.

Thus it is left to show that Y is a legal move for \exists in \mathcal{G} ; that is, we must show that

$$(\sigma(s), \varphi) \in \bar{F}(Y). \quad (3.14)$$

For this purpose, first observe that the definition of Y and the properties of \bar{F} (cf. Fact 2.12) imply that

$$\bar{F}(Y) = \bar{F}(Q) \circ Gr(Fev_a) \circ \bar{F}(\exists_A). \quad (3.15)$$

Now it follows from the legitimacy of Π in the game \mathcal{G}^\sharp , that $(Fev_a)(\Pi)$ is an F-redistribution of Φ , i.e.,

$$(\Pi, \varphi) \in Gr(Fev_a) \circ \bar{F}(\exists_A). \quad (3.16)$$

From the legitimacy of Q it follows that

$$(\sigma(s), \Pi) \in \bar{F}Q. \quad (3.17)$$

But then (3.14) is immediate from (3.15), (3.16) and (3.17). \square

4. REGULAR AUTOMATA

In this section we look in detail at some of the acceptance conditions of coalgebra automata. Recall that in the case of an acceptance game $\mathcal{G}(\mathbb{A}, \mathbb{S})$, the winner of any infinite match is determined by the infinite sequence of \mathbb{A} -states

$$a_I a_1 a_2 \dots$$

that is induced by the match. More specifically, the acceptance condition of the automaton \mathbb{A} is of the form $L \subseteq L^\omega$, i.e., an ω -language over the set A of \mathbb{A} -states. In many cases, the set L has a fairly low complexity. For instance, in the case of parity automata, the criterion whether an A -stream α belongs to L or not is given in terms of the set $Inf(\alpha)$ of those states that occur infinitely often in α .

An interesting class of automata is given by those in which the acceptance condition is a so-called ω -regular language, that is, a subset $L \subseteq A^\omega$ that is itself recognized by some word automaton.

Remark 4.1. For readers that are not familiar with the theory of automata operating on infinite words, we summarize the definitions here. Fix an alphabet C .

A C -stream is an infinite C -word $\gamma = c_0 c_1 c_2 \dots$. A *nondeterministic C -automaton* is a quadruple $\mathbb{A} = \langle A, a_I, \Delta, Acc \rangle$, where A is a finite set, $a_I \in A$ is the *initial state* of \mathbb{A} , $\Delta : A \times C \rightarrow \mathcal{P}(A)$ its *transition function* of \mathbb{A} , and $Acc \subseteq A^\omega$ its *acceptance condition*. Such an automaton is *deterministic* if $\Delta(a, c)$ is a singleton for each $a \in A$ and $c \in C$.

A *run* of a deterministic automaton $\mathbb{A} = \langle A, a_I, \Delta, Acc \rangle$ on an C -stream $\gamma = c_0 c_1 c_2 \dots$ is an infinite A -sequence

$$\rho = a_0 a_1 a_2 \dots$$

such that $a_0 = a_I$ and $a_{i+1} \in \Delta(a_i, c_i)$ for every $i \in \omega$. Note that such a run is unique if \mathbb{A} is deterministic.

A nondeterministic C -automaton $\mathbb{A} = \langle A, a_I, \Delta, Acc \rangle$ *accepts* an C -stream γ if there is a successful run of \mathbb{A} on γ . The set of those streams is denoted by $L_\omega(\mathbb{A})$. A set $L \subseteq C^\omega$ is called ω -regular if $L = L_\omega(\mathbb{A})$ for some C -automaton \mathbb{A} with a parity acceptance condition.

A key result in the theory of stream automata states that the every nondeterministic parity automaton can be transformed into an equivalent deterministic parity automata. That is, every ω -regular language L is of the form $L_\omega(\mathbb{A})$ for some *deterministic* parity automaton \mathbb{A} .

Definition 4.2. An F -automaton $\mathbb{A} = \langle A, a_I, \Delta, L \rangle$ is called *regular* if $L \subseteq A^\omega$ is an ω -regular language. \triangleleft

It is sometimes attractive to use regular automata because the acceptance condition may be easier or more intuitive to formulate in the form of an ω -regular language than as a parity condition. An important example of this was given in the previous section where the nondeterministic automaton \mathbb{A}^\sharp was provided with a regular acceptance condition. Nevertheless the recognizing power of (nondeterministic) regular automata is not strictly greater than that of (nondeterministic) parity automata: Theorem 4.4, the main result of this section, states that every regular nondeterministic automaton can be replaced with an equivalent nondeterministic parity automaton. As a consequence of this result, we may use regular automata as a handy, auxiliary notion in the theory of coalgebra automata.

The key idea underlying the proof of Theorem 4.4 is the construction of a so-called *wreath product*. Given a regular nondeterministic F -automaton with state set A , and a parity word automaton \mathbb{W} operating on infinite A -words, we define the wreath product as some nondeterministic parity F -automaton. Informally this automaton $\mathbb{B} \odot \mathbb{W}$ runs the automaton \mathbb{B} on a given pointed F -coalgebra and feeds the resulting sequence of automata states into the automaton \mathbb{W} .

Definition 4.3. Let F be a set functor, and let $\mathbb{B} = \langle B, b_I, \Delta, L \rangle$ be a nondeterministic F -automaton, and let $\mathbb{W} = \langle W, w_I, \delta : W \times B \rightarrow W, \Omega \rangle$ be a deterministic parity word automaton.

Let, for $w \in W$, the map $\delta_w : B \rightarrow B \times W$ be defined by putting

$$\delta_w(b) := (b, \delta(w, b)).$$

Using this map, we define the element $a_I \in B \times W$ and the map $\Gamma : B \times W \rightarrow \mathcal{P}F(B \times W)$ be given by

$$\begin{aligned} a_I &:= (b_I, \delta(w_I, b_I)) \\ \Gamma(b, w) &:= \{(F\delta_w)(\varphi) \mid \varphi \in \Delta(b)\}. \end{aligned}$$

The map $\Psi : B \times W \rightarrow \omega$ is defined as $\Psi = \Omega \circ \pi_2$, that is,

$$\Psi(b, w) := \Omega(w).$$

Finally, the nondeterministic F -automaton

$$\mathbb{B} \odot \mathbb{W} = \langle B \times W, a_I, \Gamma, \Psi \rangle$$

is the *wreath product* of \mathbb{B} and \mathbb{W} . \triangleleft

As we will see now, if the acceptance condition of \mathbb{A} is actually the language recognized by \mathbb{W} , then the two automata \mathbb{A} and $\mathbb{A} \odot \mathbb{W}$ are equivalent.

Theorem 4.4. *Let F be a set functor that preserves weak pullbacks, let $\mathbb{B} = \langle B, b_I, \Delta, L \rangle$ be a nondeterministic regular F -automaton and let \mathbb{W} be a deterministic parity automaton such that L is the language accepted by \mathbb{W} . Then \mathbb{B} and $\mathbb{B} \odot \mathbb{W}$ are equivalent.*

Proof. Let \mathcal{G} and \mathcal{G}^\odot be the acceptance games $\mathcal{G} := \mathcal{G}(\mathbb{B}, \mathbb{S})$ and $\mathcal{G}^\odot := \mathcal{G}(\mathbb{B} \odot \mathbb{W}, \mathbb{S})$, respectively, and fix a pointed F -coalgebra (\mathbb{S}, s_0) . Our aim is to prove the following equivalence:

$$\mathbb{B} \text{ accepts } (\mathbb{S}, s_0) \text{ iff } \mathbb{B} \odot \mathbb{W} \text{ accepts } (\mathbb{S}, s_0). \quad (4.1)$$

In both cases our proof consists in showing that \exists can mimic the match in one game by a correlated match of the other game. Here we call a partial \mathcal{G}^\odot -match π *correlated* to a partial \mathcal{G} -match π' if (the respective sequences of basic positions in) π and π' are of the following form:

$$\begin{aligned}\pi &= (s_0, (b_0, w_1))(s_1, (b_1, w_2)) \dots (s_n, (b_n, w_{n+1})) \\ \pi' &= (s_0, b_0)(s_1, b_1) \dots (s_n, b_n)\end{aligned}$$

where $w_{i+1} = \delta(w_i, b_i)$ for each i . A similar definition applies to full matches.

It is not hard to see that infinite correlated matches have the same winner, i.e., if the infinite matches π and π' are correlated, then \exists wins π iff she wins π' . The key observation here is that if $\pi = (s_0, (b_0, w_1))(s_1, (b_1, w_2)) \dots$ is correlated to $\pi' = (s_0, b_0)(s_1, b_1) \dots$, then $w_0 w_1 \dots$ is the run of \mathbb{W} on the infinite B -word $b_0 b_1 \dots$.

We now turn to the proof of (4.1). For the direction from left to right, assume that \mathbb{B} accepts (\mathbb{S}, s_0) , that is, assume that \exists has a winning strategy in the game \mathcal{G} initiated at (s_0, b_0) . In order to show that $\mathbb{B} \odot \mathbb{W}$ accepts (\mathbb{S}, s_0) , we need to equip her with a winning strategy in the game \mathcal{G}^\odot starting from the position $(s_0, (b_0, w_1))$. We will show that with the running \mathcal{G}^\odot -match, \exists can maintain a correlated \mathcal{G} -match in which she plays her winning strategy. Then by the fact that infinite correlated matches are won by the same player, she is guaranteed to win all infinite matches. Hence it suffices to prove inductively that if she has maintained the shadow match for n rounds, she either directly wins the \mathcal{G}^\odot -match in the next round, or else she can maintain the shadow match for one more round.

Assume then that with the partial \mathcal{G}^\odot -match

$$\pi = (s_0, (b_0, w_1))(s_1, (b_1, w_2)) \dots (s_n, (b_n, w_{n+1}))$$

she has associated a correlated partial \mathcal{G} -match

$$\pi' = (s_0, b_0)(s_1, b_1) \dots (s_n, b_n)$$

which is conform her winning strategy in \mathcal{G} . Suppose that this winning strategy tells her to choose $\varphi \in \Delta(b_n)$, followed by the relation $Y \subseteq S \times B$ satisfying $(\sigma(s_n), \varphi) \in \overline{F}Y$. Then in the \mathcal{G}^\odot -match π she chooses $(F\delta_{w_{n+1}})(\varphi)$, followed by the relation $Z := Y \circ Gr(F\delta_{w_{n+1}})$. The legitimacy of these moves is immediate by the definitions.

Clearly if $Z = \emptyset$, \exists wins immediately, so assume otherwise, and suppose that \forall picks a pair $(s, (b, w))$ as the next basic position continuing π . Then by definition of Z we have $(s, b) \in Y$ and $w = \delta(w_{n+1}, b)$. Hence in \mathcal{G} we could have arrived at the position (s, b) if in π' , \exists had chosen φ and Y . And since $w = \delta(w_{n+1}, b)$, the two partial matches $\pi(s, (b, w))$ and $\pi'(s, b)$ are correlated. In other words, \exists has indeed maintained the required condition for one more round.

For the other direction of (4.1), assume that \exists has a winning strategy in the acceptance game \mathcal{G}^\odot initiated at $(s_0, (b_0, w_1))$. It suffices to show that in the game \mathcal{G} starting at (s_0, b_0) , \exists has a winning strategy. Analogously to the proof for the other direction, we will show that, round by round, \exists can maintain a shadow match in \mathcal{G}^\odot which is correlated to the running \mathcal{G} -match, and conform her supposed winning strategy.

More precisely, inductively assume that with the partial \mathcal{G} -match

$$\rho = (s_0, b_0)(s_1, b_1) \dots (s_n, b_n)$$

she has associated a correlated partial \mathcal{G}° -match conform her winning strategy:

$$\rho' = (s_0, (b_0, w_1))(s_1, (b_1, w_2)) \dots (s_n, (b_n, w_{n+1}))$$

Consider the moves suggested by the winning strategy in \mathcal{G}° when the partial match ρ' arrives at $(s_n, (b_n, w_{n+1}))$. \exists first picks an element from $\Gamma(s_n, (b_n, w_{n+1}))$, which by definition of Γ is of the form $(F\delta_{w_{n+1}})(\varphi)$ with $\varphi \in \Delta(b)$, followed by a relation $Z \subseteq S \times (B \times W)$ such that $(\sigma(s_n), (F\delta_{w_{n+1}})(\varphi)) \in \bar{F}Z$. Clearly then the pair $(\sigma(s_n), \varphi)$ belongs to the set $\bar{F}Z \circ (Gr(F\delta_{w_{n+1}}))^\smile$, which by properties of relation lifting (see Fact 2.12) is identical to the relation $\bar{F}(Z \circ (Gr(\delta_{w_{n+1}}))^\smile)$. This means that \exists may legitimately continue the \mathcal{G} -match ρ with the moves φ and $Y := Z \circ (Gr(\delta_{w_{n+1}}))^\smile$.

Suppose that \forall responds to these moves by playing a pair $(s, b) \in Y$. By definition of Y there must be a pair $(b', w) \in B \times W$ such that $(s, (b', w)) \in Z$ and $\delta_{w_{n+1}}(b) = (b', w)$. From this it is immediate that $b' = b$ and $w = \delta(w_{n+1}, b)$. But from $(s, (b, w)) \in Z$ it follows that in \mathcal{G}° , \forall may respond to \exists 's move Z by picking $(s, (b, w))$ as the next basic position. Then from $w = \delta(w_{n+1}, b)$ it follows that the partial matches $\rho(s, b)$ and $\rho'(s, (b, w))$ are correlated, and since the continuation of ρ' was conform \exists 's winning strategy, we are done. \square

Remark 4.5. Our construction does not use the fact that the word automaton \mathbb{W} is a *parity* word automaton. We could use any other acceptance condition on infinite words, such as a Büchi, a Muller or a Rabin condition. In these cases the resulting wreath product automaton would be an F-automaton with Büchi, Muller or Rabin condition respectively.

5. CLOSURE PROPERTIES

In this section we prove Theorem 1 and Theorem 2. First we combine the results of the previous two sections in order to show that every parity F-automaton can be transformed into an equivalent nondeterministic one. After that we will see that the class of nondeterministically 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.

5.1. The main theorem. In the previous sections we saw how to transform an arbitrary parity F-automaton \mathbb{A} into a nondeterministic regular F-automaton \mathbb{A}^\sharp . Furthermore we showed how to use the wreath product construction in order to transform a given regular F-automaton into an equivalent F-automaton with parity acceptance condition. We will combine these facts in order to prove that for every parity F-automaton we can effectively construct an equivalent nondeterministic parity F-automaton. To begin with we have a closer look at the size of the word automaton \mathbb{W} that witnesses the fact that the acceptance condition of \mathbb{A}^\sharp is regular.

Proposition 5.1. *Let F be a set functor, and let $\mathbb{A} = \langle A, a_I, \Delta, \Omega \rangle$ be a parity F-automaton with n states and index k . Then we can construct a deterministic parity A^\sharp -word automaton $\mathbb{W} = \langle W, w_I, \delta, \Omega' \rangle$, such that \mathbb{W} accepts $\alpha = R_0 R_1 R_2 \dots \in (A^\sharp)^\omega$ iff α contains no odd trace. This automaton has $2^{\mathcal{O}(nk \log(nk))}$ states and index $\mathcal{O}(nk)$.*

Proof. The construction of \mathbb{W} is done in four steps:

Step 1: We define a nondeterministic A^\sharp -word automaton $\mathbb{B}_1 := \langle A, a_I, \delta_1, \Omega^{+1} \rangle$ where we put $\delta_1(a, R) := R[a]$ for all $a \in A$ and all $R \in A^\sharp$, and $\Omega^{+1}(a) := \Omega(a) + 1$. A straightforward argument shows that \mathbb{B}_1 accepts a word $\alpha \in (A^\sharp)^\omega$ iff α contains an odd trace. The automaton \mathbb{B}_1 has n states and index k .

Step 2: Using a standard construction, see for instance [13]), we transform the automaton into an equivalent nondeterministic A^\sharp -word automaton \mathbb{B}_2 with Büchi acceptance condition. The size of this automaton is bounded by $\mathcal{O}(nk)$.

Step 3: Using the variant of the Safra construction described in [23], we transform the automaton \mathbb{B}_2 into an equivalent deterministic parity word automaton \mathbb{B}_3 . The size of \mathbb{B}_3 is bounded by $2^{\mathcal{O}(nk \log(nk))}$, and \mathbb{B}_3 has index $\mathcal{O}(nk)$.

Step 4: Finally we define \mathbb{W} to be the deterministic parity A^\sharp -word automaton that accepts the complement of the language of \mathbb{B}_3 , that is, \mathbb{W} accepts a word $\alpha \in (A^\sharp)^\omega$ iff α does *not* contain any odd trace. This automaton can be obtained by taking \mathbb{B}_3 and changing its parity function in the same way as in Step 1, i.e., we increase all the parities by 1. The size and index of \mathbb{W} are still bounded by $2^{\mathcal{O}(nk \log(nk))}$ and $\mathcal{O}(nk)$, respectively. □

Theorem 5.2. *Let F be a set functor that preserves weak pullbacks, and let $\mathbb{A} = \langle A, a_I, \Delta, \Omega \rangle$ be an alternating parity F -automaton with n states and index k . Then we can construct an equivalent nondeterministic parity automaton \mathbb{A}^\bullet of size $2^{\mathcal{O}(n^2 + nk \log(nk))}$ and with index $\mathcal{O}(nk)$.*

Proof. In Section 3 we showed how to transform \mathbb{A} into an equivalent *regular* nondeterministic F -automaton \mathbb{A}^\sharp . This automaton has at most size 2^{n^2} . Furthermore we can use Proposition 5.1 in order to construct a parity A^\sharp -word automaton \mathbb{W} of size $2^{\mathcal{O}(nk \log(nk))}$ and with index $\mathcal{O}(nk)$ that accepts exactly those words $\alpha \in (A^\sharp)^\omega$ that do not contain any odd trace. We define \mathbb{A}^\bullet to be the wreath product $\mathbb{A}^\sharp \odot \mathbb{W}$. From Theorem 4.4 we know that \mathbb{A}^\bullet is a nondeterministic F -automaton that is equivalent to \mathbb{A} . Furthermore, spelling out the definitions, one can easily check that \mathbb{A}^\bullet has size $2^{\mathcal{O}(n^2 + nk \log(nk))}$ and index $\mathcal{O}(nk)$. □

Remark 5.3. The complexity bound in our main theorem is the immediate consequence from known results in the literature on ω -automata. In particular we heavily rely on [23]. Our contribution is to show that known complexity bounds concerning ω -automata can be transferred to other types of structures essentially without increasing the complexity. In other words, we prove that the complexity of transforming a given alternating F -automaton into an equivalent nondeterministic one is bound by the complexity of the Safra construction on ω -automata. This observation has been further substantiated in [14] where it is shown in detail that both the lower and upper complexity bounds from the Safra construction on ω -automata yield the respective complexity bounds for transforming a given alternating tree automaton into an equivalent nondeterministic one.

5.2. Closure under union and intersection. In this subsection we prove that the class of recognizable languages is closed under taking unions and intersections. This is the content of the following proposition. Note that in this subsection we do require the functor to be standard. As we demonstrated in Section 2.4 this is not an essential condition. In fact, it

would be not difficult to modify the arguments in this subsection in order to show closure under union and intersection also for nonstandard functors. The fairly simple proofs would, however, look unnecessarily complicated.

Proposition 5.4. *Let F be some set functor. Given two parity F -automata \mathbb{A}_1 and \mathbb{A}_2 we can construct parity F -automata \mathbb{A}_\cup and \mathbb{A}_\cap such that $L(\mathbb{A}_\cup) = L(\mathbb{A}_1) \cup L(\mathbb{A}_2)$ and $L(\mathbb{A}_\cap) = L(\mathbb{A}_1) \cap L(\mathbb{A}_2)$. Both \mathbb{A}_\cup and \mathbb{A}_\cap have size $n_1 + n_2 + 1$, where n_i is the size of \mathbb{A}_i , and both automata have index $k := \max(k_1, k_2)$, where k_i is the index of \mathbb{A}_i . Moreover \mathbb{A}_\cup is nondeterministic if \mathbb{A}_1 and \mathbb{A}_2 are so.*

Before we prove the proposition we define the automata \mathbb{A}_\cup and \mathbb{A}_\cap .

Definition 5.5. Let $\mathbb{A}_1 = \langle A_1, a_I^1, \Delta_1, \Omega_1 \rangle$ and $\mathbb{A}_2 = \langle A_2, a_I^2, \Delta_2, \Omega_2 \rangle$ be two parity 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:

$$\begin{aligned} \Delta_\cup(a) &:= \begin{cases} \Delta_1(a_I^1) \cup \Delta_2(a_I^2) & \text{if } a = * \\ \Delta_i(a) & \text{if } a \in A_i, \end{cases} \\ \Delta_\cap(a) &:= \begin{cases} \{\Phi_1 \cup \Phi_2 \mid \Phi_i \in \Delta_i(a_I^i)\} & \text{if } a = * \\ \Delta_i(a) & \text{if } a \in A_i. \end{cases} \end{aligned}$$

Finally, we put $\mathbb{A}_\cup := \langle A_{12}, a_I, \Delta_\cup, \Omega \rangle$ and $\mathbb{A}_\cap := \langle A_{12}, a_I, \Delta_\cap, \Omega \rangle$. ◁

Let us now turn to the proof of Proposition 5.4.

Proof. The automata \mathbb{A}_\cup and \mathbb{A}_\cap are constructed as defined above in Definition 5.5. Clearly \mathbb{A}_\cup and \mathbb{A}_\cap meet the size requirements stated in the proposition. Furthermore it is easy to see that \mathbb{A}_\cup is nondeterministic if \mathbb{A}_1 and \mathbb{A}_2 are nondeterministic automata. We only show that $L(\mathbb{A}_\cup) = L(\mathbb{A}_1) \cup L(\mathbb{A}_2)$, the other statements of the proposition admit similarly straightforward proofs. It suffice to show that \mathbb{A}_\cup accepts an arbitrary pointed F -coalgebra (\mathbb{S}, s) iff \mathbb{A}_1 or \mathbb{A}_2 accepts (\mathbb{S}, s) . Let (\mathbb{S}, s) be a pointed F -coalgebra. and suppose first 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) . Conversely, suppose that \mathbb{A}_i accepts (\mathbb{S}, s) , and let g be a winning strategy for \exists in the game $\mathcal{G}(\mathbb{A}_i, \mathbb{S})$. Then in the game $\mathcal{G}(\mathbb{A}_\cup, \mathbb{S})$ starting at $(s, *)$, let \exists start with playing $g(s, a_I^i) \in \Delta_\cup(*)$, and from then on, play her strategy g . It is again straightforward to check that this constitutes a winning strategy for \exists . □

5.3. Closure under projection. In this subsection we shall prove that the recognizable F -languages are closed under *existential* projection. In order to make this notion precise, fix a set functor F and a set C of colours. Recall from Example 2.4 that we may identify F_C -coalgebras with C -coloured F -coalgebras: Given an F_C -coalgebra $\langle S, \sigma \rangle$, we may split the coalgebra map $\sigma : S \rightarrow F_C S$ into two parts, the colouring $\sigma_C : S \rightarrow C$ and the F -coalgebra map $\sigma_F : S \rightarrow F S$. Hence, with each F_C -coalgebra $\mathbb{S} = \langle S, \sigma \rangle$ we may associate its *F-projection* $\pi_F \mathbb{S} := \langle S, \sigma_F \rangle$, and likewise for pointed coalgebras.

A possibility for defining the existential F -projection $\pi_F L$ of an F_C -language L would be the class

$$\{\pi_F \mathbb{S} \mid \mathbb{S} \text{ in } L\}. \quad (5.1)$$

Note that we use pseudo-set notation here — recall that L may be a class rather than a set. It is, however, not difficult to see that this language will in general not be closed under bisimilarity. If we take the position that bisimilar states represent the *same* process, this means that (5.1) is not the right notion. This leads to the definition of the existential C -projection $\pi_F L$ of L as the closure of (5.1) under F -bisimilarity.

Definition 5.6. Let F be some set functor, let C be some set, and let L be an F_C -language. The *existential F-projection* of L , notation: $\pi_F L$, consists of all pointed F -coalgebras $(\langle S, \sigma \rangle, s)$ for which there is an F_C -coalgebra $(\langle S', \gamma, \sigma' \rangle, s')$ in L such that $\langle S, \sigma \rangle, s \Leftarrow_F \langle S', \sigma' \rangle, s'$. \triangleleft

In pseudo-set notation we could write

$$\pi_F L := \{(\langle S, \sigma \rangle, s) \mid \langle S, \sigma \rangle, s \Leftarrow_F \langle S', \sigma' \rangle, s' \text{ for some } (\langle S', \gamma, \sigma' \rangle, s') \text{ in } L\}.$$

Remark 5.7. This definition is in accordance with standard usage. In the case of binary trees, we are dealing with two alphabets C and D . Given a class K of $C \times D$ -labeled binary trees, one defines the *D-projection* of this class as the class of D -labeled binary trees $\langle 2^*, \tau_D : 2^* \rightarrow D \rangle$ for which there is a map (*C-colouring*) $\tau_C : 2^* \rightarrow C$ such that the $C \times D$ -labeled binary tree $\tau : 2^* \rightarrow C \times D$ given by $\tau(s) = (\tau_C(s), \tau_D(s))$ belongs to K . No reference to bisimilarity is needed here due to the fact that two labeled binary trees are bisimilar if and only if they are *identical*, see Example 2.7.

Second, for Kripke structures our notion of C -projection exactly corresponds to the usual interpretation of existential bisimulation quantifiers — a fact which can be used to prove that closure under projection of P_C -automata implies uniform interpolation of the modal μ -calculus. We refer to D'AGOSTINO & HOLLENBERG [6] for more details.

The main result of this section states that the class of recognizable languages is closed under this operation. We will show that, given an F_C -automaton \mathbb{A} , we will define an F -automaton $\pi_C \mathbb{A}$ that accepts a given pointed F -coalgebra $(\langle S, \sigma \rangle, s)$ iff there *exists* a bisimilar F -coalgebra $(\langle S', \sigma' \rangle, s')$ and a colouring $\gamma : S' \rightarrow C$ such that $(\langle S', \gamma, \sigma' \rangle, s')$ is accepted by \mathbb{A} . Before we start to prove this, let us say a word about *universal* projection.

Remark 5.8. The universal F -projection $\pi_F^\forall L$ of an F_C -language L is defined dually:

$$\pi_F^\forall L := \{(\langle S, \sigma \rangle, s) \mid (\langle S', \gamma, \sigma' \rangle, s') \text{ in } L \text{ whenever } \langle S, \sigma \rangle, s \Leftarrow_F \langle S', \sigma' \rangle, s'\}.$$

The question whether the class of recognizable languages is also closed under universal projection is still open and closely related to the question whether F -recognizable languages, in general, are closed under complementation.

We now turn to the proof that the recognizable languages are closed under (existential) projection. In the remainder of this section, all F-automata are assumed to be nondeterministic. To facilitate the presentation we will think of the transition function Δ as a map $A \rightarrow \mathcal{P}_{\exists}FA$ 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 $A \times S \rightarrow FA$, cf. Remark 2.21.

The main result of this subsection is stated in the following proposition.

Proposition 5.9. *Let F be a set functor that preserves weak pullbacks. For any nondeterministic parity F_C -automaton \mathbb{A} of size n and index k we can construct a nondeterministic parity F-automaton $\pi_F\mathbb{A}$ of size n and index k , such that for every pointed F-coalgebra (\mathbb{S}, s) the following are equivalent:*

- (1) $\pi_F\mathbb{A}$ accepts (\mathbb{S}, s) ,
- (2) \mathbb{A} accepts an F_C -coalgebra $(\langle S', \gamma, \sigma' \rangle, s')$ such that $(\langle S', \sigma' \rangle, s')$ and (\mathbb{S}, s) are bisimilar.

The remainder of this section is devoted to the proof of this proposition. First we define the automaton $\pi_C\mathbb{A}$ and then we show that it meets the requirements of the proposition.

Definition 5.10. Let C be a set, $\mathbb{A} = \langle A, a_I, \Delta, \Omega \rangle$ be a parity F_C -automaton and $\mathbb{A}_C = \langle A, a_I, C, \Delta_C, \Omega \rangle$ its C -chromatic F-companion, see Fact 2.24. Then we define the F-projection $\pi_F\mathbb{A} := \langle A, a_I, \Delta_\pi, \Omega \rangle$ where $\Delta_\pi(a) := \bigcup_{c \in C} \Delta_C(c, a)$. \triangleleft

Lemma 5.11. *If \mathbb{A} accepts the F_C -coalgebra $(\mathbb{S}, s) := (\langle S, \gamma, \sigma \rangle, s)$ then $\pi_F\mathbb{A}$ accepts $(\mathbb{S}^\pi, s) := (\langle S, \sigma \rangle, s)$.*

Proof. The proof is straightforward. 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 $\pi_F\mathbb{A}$ acceptance game. \square

The converse of this lemma however fails in general.

Let \mathbb{A} be some F_C -automaton and let $(\langle S, \sigma \rangle, r)$ be a pointed F-coalgebra that is accepted by $\pi_F\mathbb{A}$. Then we know that \exists has a winning strategy (Φ, Y) in $\mathcal{G}(\pi_F\mathbb{A}, \mathbb{S})$ 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 $\mathcal{G}(\pi_F\mathbb{A}, \mathbb{S})$, 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 5.12. Let \mathbb{A} be a parity 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 functional (that is, for every $s \in S$ there is at most one $a \in A$ such that the pair (s, a) belongs to the relation). 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

As we will see now, strong acceptance is the key to find colorings of pointed F-coalgebras.

Lemma 5.13. *Let \mathbb{A} be an F_C -automaton, and let (\mathbb{S}, r) be a pointed F -coalgebra that is strongly accepted by $\pi_F \mathbb{A}$. Then there is a C -colouring $\gamma : S \rightarrow C$ of \mathbb{S} such that \mathbb{A} accepts $(\langle S, \gamma, \sigma \rangle, r)$.*

Proof. Let (Φ, Y) be a scattered winning strategy for \exists in $\mathcal{G}(\pi_F \mathbb{A}, \mathbb{S})$. 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 $(\langle S, \gamma, \sigma \rangle, r)$. \square

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 5.14. *Let \mathbb{A} be an F -automaton, and let (\mathbb{S}, r) be a pointed F -coalgebra that is accepted by \mathbb{A} . Then \mathbb{A} strongly accepts some pointed F -coalgebra $(\tilde{\mathbb{S}}, \tilde{r})$ which is bisimilar to (\mathbb{S}, r) .*

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

First we endow the set \tilde{S} with a coalgebra map $\tilde{\sigma}$ such that the structure $\tilde{\mathbb{S}} := (\tilde{S}, \tilde{\sigma})$ is isomorphic to the A -fold coproduct ('disjoint union') $\coprod_{a \in A} \mathbb{S}$. For the exact definition of the coproduct of coalgebras the reader is referred to [25].

The canonical injections into $\tilde{\mathbb{S}}$ are given by the functions

$$\begin{aligned} \kappa_a : S &\rightarrow S \times A \\ s &\mapsto (s, a) \end{aligned}$$

for all $a \in A$.

Now consider the first projection map $\pi_S : S \times A \rightarrow S$. From $\pi_S(s, a) = s$ it follows that

$$\pi_S \circ \kappa_a = id_S \quad \text{for all } a \in A. \quad (5.2)$$

We are going to prove that

$$\pi_S : (\tilde{S}, \tilde{\sigma}) \rightarrow \mathbb{S} \text{ is a coalgebra morphism.} \quad (5.3)$$

That is, we will show that the following diagram commutes:

$$\begin{array}{ccc} S \times A & \xrightarrow{\pi_S} & S \\ \tilde{\sigma} \downarrow & & \sigma \downarrow \\ F(S \times A) & \xrightarrow{F\pi_S} & F\mathbb{S} \end{array}$$

In order to prove the commutativity of the diagram, take an arbitrary $(s, a) \in S \times A$. We

obtain the following sequence of identities:

$$\begin{aligned}
\mathbf{F}\pi_S(\bar{\sigma}(s, a)) &= \mathbf{F}\pi_S(\bar{\sigma}(\kappa_a(s))) \\
&\stackrel{\kappa_a \text{ coalg. morph.}}{=} \mathbf{F}\pi_S(\mathbf{F}\kappa_a(\sigma(s))) \\
&= \mathbf{F}(\pi_S \circ \kappa_a)(\sigma(s)) \\
&\stackrel{(5.2)}{=} \sigma(s) \\
&= \sigma(\pi_S(s, a)),
\end{aligned}$$

which proves (5.3).

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,

$$\bar{F}R = Gr(\mathbf{F}\pi_S)^\smile \circ \bar{F}\hat{R}. \quad (5.4)$$

We are now prepared to turn to the proof of the lemma. Assume that (Φ, Y) is a positional strategy for \exists in the acceptance game $\mathcal{G}(\mathbb{A}, \mathbb{S})$ which is winning for any match starting at the position (a_I, r) .

For the definition of $\bar{\sigma} : \bar{S} \rightarrow \mathbf{F}\bar{S}$, consider an arbitrary element $(s, a) \in \bar{S}$, and distinguish cases. If (s, a) is a *winning* position for \exists in the game $\mathcal{G}(\mathbb{A}, \mathbb{S})$, then using (5.4), it follows from $(\sigma(s), \varphi) \in \bar{F}Y$, that

$$(\sigma(s), \varphi) \in (Gr(\mathbf{F}\pi_S)^\smile) \circ \bar{F}\hat{Y}.$$

Hence we may take $\bar{\sigma}(s, a)$ to be some element $x \in \mathbf{F}\bar{S}$ such that $(\sigma(s), x) \in (Gr(\mathbf{F}\pi_S)^\smile)$, that is, $(\mathbf{F}\pi_S)(x) = \sigma(s)$, and $(x, \varphi) \in \bar{F}\hat{Y}$. If, on the other hand, $(s, a) \notin Win_\exists$, then we simply put $\bar{\sigma}(s, a) := \tilde{\sigma}(s, a)$.

We first check that π_S is indeed an F-coalgebra morphism from $\bar{\mathbb{S}}$ onto \mathbb{S} . Take an arbitrary element (s, a) in $\bar{\mathbb{S}}$, then we have to check that $(\mathbf{F}\pi_S)(\bar{\sigma}(s, a)) = \sigma(\pi_S(s, a))$. In case $(s, a) \notin Win_\exists$ this follows from the facts that $\bar{\sigma}(s, a) = \tilde{\sigma}(s, a)$ and the fact that $\tilde{\sigma}$ is a coalgebra morphism. In case $(s, a) \in Win_\exists$ the identity follows by definition of $\bar{\sigma}(s, a)$.

Thus we have proved the first statement of the proposition. For the second statement, define the strategy $(\bar{\Phi}, \bar{Y})$ with $\bar{\Phi} : \bar{S} \times A \rightarrow A$ and $\bar{Y} : \bar{S} \times \mathbf{F}A \rightarrow \mathcal{P}(\bar{S} \times A)$ 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 $\mathcal{G}(\mathbb{A}, \bar{\mathbb{S}})$ starting from (\bar{r}, a_I) . To see why this is the case, consider an arbitrary position $((s, a), a)$ with $(s, a) \in Win_\exists(\mathcal{G}(\mathbb{A}, \mathbb{S}))$, 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 $\mathcal{G}(\mathbb{A}, \bar{\mathbb{S}})$ 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 $\mathcal{G}(\mathbb{A}, \mathbb{S})$ 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 $\mathcal{G}(\mathbb{A}, \bar{\mathbb{S}})$ match is, indeed, a win for \exists . This proves the second statement of the proposition. \square

We are now ready to prove our main result.

Proof of Proposition 5.9. The implication $(1 \Rightarrow 2)$ is immediate by the Lemmas 5.14 and 5.13. The other implication follows from Lemma 5.11 and the observation [29] that F-automata do not distinguish between bisimilar pointed F-coalgebras. \square

Together with Theorem 5.2 the proposition entails what we call closure under (existential) projection.

Corollary 5.15. *Let F be a set functor that preserves weak pullbacks. Given an alternating parity F_C -automaton \mathbb{A} of size n and index k we can construct a nondeterministic parity F-automaton $\pi_F \mathbb{A}^\bullet$ such that the following are equivalent:*

- (1) $\pi_F \mathbb{A}^\bullet$ accepts a pointed F-coalgebra $(\langle S, \sigma \rangle, s_I)$,
- (2) \mathbb{A} accepts a pointed F_C -coalgebra $(\langle S', \gamma, \sigma' \rangle, s'_I)$ such that $(\langle S, \sigma \rangle, s_I)$ and $(\langle S', \sigma' \rangle, s'_I)$ are F-bisimilar.

The size of $\pi_F \mathbb{A}^\bullet$ is $2^{\mathcal{O}(n^2 + nk \log(nk))}$ and the index of $\pi_F \mathbb{A}^\bullet$ is $\mathcal{O}(nk)$.

6. SOLUTION OF THE NONEMPTINESS PROBLEM

In this section we prove that every parity F-automaton accepts a *finite* coalgebra, if it accepts a coalgebra at all. The key result leading to this observation is that any nondeterministic parity automaton with a nonempty language, actually accepts a coalgebra that ‘lives inside the automaton’, in the following sense.

Theorem 6.1. *Let F be some weak pullback preserving set functor, and let $\mathbb{A} = \langle A, a_I, \Delta, \Omega \rangle$ be a nondeterministic parity F-automaton. Then \mathbb{A} accepts some pointed F-coalgebra iff \mathbb{A} accepts a pointed F-coalgebra $(\langle S, \sigma \rangle, s_0)$ with $S \subseteq A$, $s_0 = a_I$ and $\sigma(s) \in \Delta(s)$ for all $s \in S$.*

As an immediate consequence of the above result, and of the fact that for every alternating F-coalgebra automaton we can effectively construct an equivalent nondeterministic automaton (Theorem 5.2), we obtain the following solution for the nonemptiness problem for parity F-automata.

Corollary 6.2. *Let F be some weak pullback preserving set functor, and let $\mathbb{A} = \langle A, a_I, \Delta, \Omega \rangle$ be a parity F-automaton of size n . Then $L(\mathbb{A}) \neq \emptyset$ iff \mathbb{A} accepts a pointed F-coalgebra (S, σ, s_0) with $|S| \leq 2^{\mathcal{O}(n^2 \log n)}$.*

Proof. Suppose $L(\mathbb{A}) \neq \emptyset$, i.e. there is some pointed F-coalgebra that is accepted by \mathbb{A} . The index of \mathbb{A} is smaller than or equal to its size n , and so it follows from Theorem 5.2 that we can transform \mathbb{A} into an equivalent nondeterministic parity automaton \mathbb{A}^\bullet of size $2^{\mathcal{O}(n^2 + n^2 \log(n^2))} = 2^{\mathcal{O}(n^2 \log n)}$. Because \mathbb{A}^\bullet is equivalent to \mathbb{A} we know that \mathbb{A}^\bullet accepts some pointed F-coalgebra. The claim follows now immediately from Theorem 6.1. \square

The remainder of this section is devoted to the proof of Theorem 6.1.

Remark 6.3. The heart of the proof of Theorem 6.1 will be the construction of the so-called nonemptiness game of a nondeterministic F-automaton. This nonemptiness game can be seen as a variant of the acceptance game of an F-automaton that is played *inside* the given automaton. Readers familiar with automata that operate on infinite objects will recognize an analogy to standard automata theoretic techniques: in order to solve the nonemptiness

problem of a given nondeterministic automaton one usually considers an *input-free* variant of the automaton and then decides whether the input-free automaton has a successful run (cf. e.g. [9, Chap. 8] concerning tree automata). Successful runs of such an input-free automaton correspond to winning strategies of \exists in our nonemptiness game.

6.1. Standardness and the nonemptiness game. For a smooth presentation of the proof of Theorem 6.1, we first consider the special case, where we impose the additional condition on the functor to be *standard* (cf. Section 2.4). This means that in particular, $S \subseteq T$ implies $FS \subseteq FT$. Furthermore we need the following properties of standard functors.

Proposition 6.4. *Let F be a standard, weak pullback preserving functor. Then, for all sets S, T, S' and T' , with $S' \subseteq S$ and $T' \subseteq T$, and for all relations $R \subseteq S \times T$:*

- (1) F commutes with intersections: $F(S \cap T) = FS \cap FT$.
- (2) \bar{F} commutes with restrictions: $\bar{F}(R \upharpoonright_{S' \times T'}) = (\bar{F}R) \upharpoonright_{FS' \times FT'}$;

Proof. By a result of Trnková (see [27]), property (1) holds for any standard functor F , provided that the intersection $S \cap T$ is *nonempty*. We use weak pullback preservation of F to show that the claim is true for arbitrary intersections. Let S and T be sets. The left diagram below is a pullback diagram which, by our assumption that F is standard, gets mapped to the lower right square:

$$\begin{array}{ccc}
 S \cap T & \xrightarrow{\iota_{S \cap T, S}} & S \\
 \downarrow \iota_{S \cap T, T} & & \downarrow \iota_{S, S \cup T} \\
 T & \xrightarrow{\iota_{T, S \cup T}} & S \cup T
 \end{array}
 \qquad
 \begin{array}{ccccc}
 FS \cap FT & & & & FS \\
 \downarrow \iota_{FS \cap FT, FS} & \searrow \iota_{FS \cap FT, FS} & & & \downarrow \iota_{FS, F(S \cup T)} \\
 & & F(S \cap T) & \xrightarrow{\iota_{F(S \cap T), FS}} & \\
 & & \downarrow \iota_{F(S \cap T), FT} & & \downarrow \iota_{FS, F(S \cup T)} \\
 & & FT & \xrightarrow{\iota_{FT, F(S \cup T)}} & F(S \cup T)
 \end{array}$$

From the fact that F preserves weak pullbacks, it follows that the square in the right diagram is a weak pullback diagram. Hence there exists some function h that makes both upper triangles commute as depicted in the diagram. But then a straightforward verification shows that h itself must be the inclusion from $FS \cap FT$ into $F(S \cap T)$, i.e. $FS \cap FT \subseteq F(S \cap T)$. The converse inclusion $F(S \cap T) \subseteq FS \cap FT$ is an immediate consequence of standardness of the functor F .

For (2), we first consider the inclusion \subseteq . By monotonicity of \bar{F} we have $\bar{F}(R \upharpoonright_{S' \times T'}) \subseteq \bar{F}(R)$.

From this it follows immediately that

$$\bar{F}(R \upharpoonright_{S' \times T'}) \subseteq (\bar{F}R) \upharpoonright_{FS' \times FT'}$$

A short proof of the opposite inclusion can be found in [30, Prop. 2.2]. \square

Position: b	Player	Admissible moves: $E[b]$	$\Omega'(b)$
$a \in A$	\exists	$\Delta(a)$	$\Omega(a)$
$\varphi \in FA$	\exists	$\{Base(\varphi)\}$	0
$B \subseteq A$	\forall	B	0

Table 3: Nonemptiness game of a nondeterministic F-automaton

The key concept in our proof of Theorem 6.1 is the so-called *nonemptiness game* $\mathcal{G}_{\neq\emptyset}(\mathbb{A})$ that we may associate with a nondeterministic automaton. Intuitively, one should think of this game as the simultaneous projection on \mathbb{A} of all acceptance games $\mathcal{G}(\mathbb{A}, \mathbb{S})$. For a formal definition we need the following notion.

Definition 6.5. Let F be a standard set functor. Given a finite set A and an element $\varphi \in FA$, the set

$$Base(\varphi) := \bigcap \{U \mid U \subseteq A \text{ and } \varphi \in FU\}.$$

is defined as the *base* of φ . \triangleleft

It follows from $\varphi \in FA$ that the set $\{U \mid U \subseteq A \text{ and } \varphi \in FU\}$ is nonempty, so that $Base(\varphi)$ is well defined.

Example 6.6. Fix some finite set A , and recall the definition of the functors B and P of Example 2.4. The base of an arbitrary element $(a_1, a_2) \in BA$ is the set $\{a_1, a_2\}$. An element of PX is a subset $B \subseteq A$; the base of such an element is the set B itself.

Intuitively the base of an element $\varphi \in FA$ consists exactly of those elements of A that we need to ‘construct’ φ . Bases have the following key property.

Proposition 6.7. *Let F be a standard set functor, and consider an object $\varphi \in FA$, where A is some finite set. Then $Base(\varphi)$ is the smallest set X such that $\varphi \in FX$.*

Proof. It is an easy consequence of Fact 6.4 and the finiteness of A that $\varphi \in F(Base(\varphi))$. Now suppose Y is a set such that $\varphi \in FY$. Then we have $\varphi \in FA \cap FY = F(A \cap Y)$, so that $Base(\varphi)$ is a subset of $A \cap Y$, and, hence, of Y . \square

In the remaining part of this subsection we assume a fixed standard functor F . We can now define the ‘nonemptiness game’ $\mathcal{G}_{\neq\emptyset}(\mathbb{A})$ associated with a given nondeterministic parity F-automaton \mathbb{A} .

Definition 6.8. Let $\mathbb{A} = \langle A, a_I, \Delta, \Omega \rangle$ be a nondeterministic parity F-automaton, where F is a standard set functor. The rules and the (parity) winning conditions of the *nonemptiness game* $\mathcal{G}_{\neq\emptyset}(\mathbb{A})$ of \mathbb{A} are given in Table 3. \triangleleft

For an informal description of this game, we first note that, just like the acceptance game, matches proceed in rounds. The basic positions of the game are now the states of the automaton A . At a basic position $a \in A$ player \exists has to move to some successor $\varphi \in \Delta(a)$. Then \exists moves further to the base of φ . Finally it is \forall ’s turn to chose some $a' \in Base(\varphi)$ as the next basic position.

Attentive readers may have noticed that the formulation of $\mathcal{G}_{\neq\emptyset}(\mathbb{A})$ looks unnecessarily complicated because \exists ’s second move (from $\varphi \in FA$ to $Base(\varphi) \in PA$) is entirely determined

by her first move. We keep this redundancy because it makes it easier to relate matches of the nonemptiness game to matches of the acceptance game of \mathbb{A} .

The name ‘nonemptiness game’ can be justified by the following two lemmas that, taken together, imply that \exists has a winning strategy in the nonemptiness game for \mathbb{A} iff the language recognized by \mathbb{A} is not empty. The first lemma takes care of the direction from left to right (recall that $\mathcal{G}_{\neq\emptyset}(\mathbb{A})$, being a parity game, satisfies history-free determinacy).

Lemma 6.9. *Let $\mathbb{A} = \langle A, a_I, \Delta, \Omega \rangle$ be a nondeterministic F-automaton. If $\varphi : A \rightarrow FA$ encodes a history-free winning strategy of \exists in $\mathcal{G}_{\neq\emptyset}(\mathbb{A})$ at position a_I then \mathbb{A} accepts the F-coalgebra (A, φ, a_I) .*

Proof. Let $\varphi : A \rightarrow FA$ be a winning strategy of \exists in $\mathcal{G}_{\neq\emptyset}(\mathbb{A})$ and define \mathbb{A}_φ to be the F-coalgebra $\langle A, \varphi \rangle$. In order to show that \mathbb{A} accepts the pointed coalgebra $(\mathbb{A}_\varphi, a_I)$ we have to equip \exists with a winning strategy in $\mathcal{G} = \mathcal{G}(\mathbb{A}, \mathbb{A}_\varphi)$. To this aim, we define:

$$\begin{array}{ccc} \Phi : \text{Id}_A & \rightarrow & FA \\ (a, a) & \mapsto & \varphi(a) \end{array} \qquad \begin{array}{ccc} Z : A \times FA & \rightarrow & \mathcal{P}(\text{Id}_A) \\ (a, \psi) & \mapsto & \text{Id}_{\text{Base}(\psi)} \end{array}$$

The functions Φ and Z encode a legitimate strategy for \exists in \mathcal{G} at position (a_I, a_I) : in order to see this, first note that any match that starts at (a_I, a_I) and in which \exists plays conform (Φ, Z) will only pass through basic positions of the form (a, a) . Hence, \exists 's strategy is defined on all positions that are possibly reached in such a match. Let us now see that at any position of the form (a, a) for some $a \in A$, the moves encoded by \exists 's strategy are legitimate. That $\varphi(a)$ is an element of $\Delta(a)$ is true by definition. In order to see that the move further to $\text{Id}_{\text{Base}(\varphi(a))}$ is also a legal move we use Fact 2.12(2) which yields that $\bar{F}(\text{Id}_{\text{Base}(\varphi(a))}) = \text{Id}_{F\text{Base}(\varphi(a))}$. Together with $\varphi(a) \in F\text{Base}(\varphi(a))$ this implies that $\varphi(a), \varphi(a) \in \bar{F}(\text{Id}_{\text{Base}(\varphi(a))})$.

It remains to show that (Φ, Z) is indeed a *winning* strategy for \exists . The key observation here is that the ‘projection’ of a \mathcal{G} -match in which \exists plays conform (Φ, Z) is a match of $\mathcal{G}_{\neq\emptyset}(\mathbb{A})$ in which \exists plays conform φ :

$$\begin{array}{ccccccc} \mathcal{G}\text{-match} & (a_I, a_I) & (\varphi(a), a_I) & \text{Id}_{\text{Base}(\varphi(a))} & (a_1, a_1) & \dots & (a_n, a_n) \\ & \downarrow & \downarrow & \downarrow & \downarrow & & \downarrow \\ \text{projection} & a_I & \varphi(a) & \text{Base}(\varphi(a)) & a_1 & \dots & a_n \end{array}$$

This suffices to prove that \exists wins all \mathcal{G} -matches in which she follows the strategy (Φ, Z) , because we assumed φ to be a winning strategy for \exists in $\mathcal{G}_{\neq\emptyset}(\mathbb{A})$. \square

The next lemma states that if the language recognized by \mathbb{A} is not empty, then \exists wins the nonemptiness game of \mathbb{A} indeed.

Lemma 6.10. *Let $\mathbb{A} = \langle A, a_I, \Delta, \Omega \rangle$ be a nondeterministic F-automaton. If \mathbb{A} accepts some pointed F-coalgebra (\mathbb{S}, s_0) then \exists has a winning strategy in $\mathcal{G}_{\neq\emptyset}(\mathbb{A})$ at position a_I .*

Proof. Suppose \mathbb{A} accepts the pointed F-coalgebra $(\mathbb{S}, s_0) = (\langle S, \sigma \rangle, s_0)$. Then \exists has a history-free winning strategy in the acceptance game $\mathcal{G} = \mathcal{G}(\mathbb{A}, \mathbb{S})$ starting at position (s_0, a_I) , that can be encoded as a pair of functions

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

Moreover, without loss of generality we may assume that

$$\text{rng}(Z_{s,\varphi}) = \text{Base}(\varphi) \tag{6.1}$$

for all positions (s, φ) that are winning for \exists . To see this, observe that for all $(s, \varphi) \in \text{Win}_{\exists}$ we have $(\sigma, \varphi) \in \bar{F}(Z_{s, \varphi})$ by legitimacy of the strategy. Thus it follows from $\varphi \in \text{FBase}(\varphi)$ that

$$(\sigma(s), \varphi) \in (\bar{F}Z_{(s, \varphi)}) \upharpoonright_{\text{FS} \times \text{FBase}(\varphi)} .$$

But Fact 2.12 yields

$$(\bar{F}Z_{(s, \varphi)}) \upharpoonright_{\text{FS} \times \text{FBase}(\varphi)} = \bar{F}(Z_{(s, \varphi)} \upharpoonright_{S \times \text{Base}(\varphi)}) ,$$

so that we may infer that

$$(\sigma(s), \varphi) \in \bar{F}(Z_{(s, \varphi)} \upharpoonright_{S \times \text{Base}(\varphi)}) .$$

From this it follows that instead of playing $Z_{s, \varphi}$, \exists could have played the relation $Z_{(s, \varphi)} \upharpoonright_{S \times \text{Base}(\varphi)}$ as well. Since the latter relation is *smaller* it decreases the choice of \forall , and so \exists will increase rather than decrease her chances of winning the game. This shows that indeed we may assume (6.1) without loss of generality.

We now turn to the nonemptiness game $\mathcal{G}_{\neq \emptyset}(\mathbb{A})$, and show that the strategy (Φ, Z) can be used to equip \exists with a winning strategy. Call a position (s, a) in a $\mathcal{G}(\mathbb{A}, \mathbb{S})$ -match *parallel* to a' if $a = a'$ and (s, a) is winning for \exists .

We first describe \exists 's strategy in one round of the game $\mathcal{G}_{\neq \emptyset}(\mathbb{A})$, and demonstrate how she constructs a parallel round of $\mathcal{G}(\mathbb{A}, \mathbb{S})$. Let a be a position in a $\mathcal{G}_{\neq \emptyset}(\mathbb{A})$ -match and let (s, a) be an (inductively defined) parallel position in $\mathcal{G}(\mathbb{A}, \mathbb{S})$. \exists 's strategy is to move from a to $\varphi := \Phi_{(s, a)}$ and further to $\text{Base}(\varphi)$. After that \forall chooses an element a' of $\text{Base}(\varphi)$. The corresponding round of the $\mathcal{G}(\mathbb{A}, \mathbb{S})$ -match is constructed as follows: \exists moves from (s, a) to (s, φ) and further to $Z_{(s, \varphi)}$. Now we use our assumption (6.1): from $\text{rng}(Z_{(s, \varphi)}) = \text{Base}(\varphi)$ and $a' \in \text{Base}(\varphi)$ we may infer the existence of some s' such that $(s', a') \in Z_{(s, \varphi)}$. Therefore in $\mathcal{G}(\mathbb{A}, \mathbb{S})$ \forall can move from $Z_{(s, \varphi)}$ to (s', a') . This position (s', a') is parallel to a' because \exists played according to her winning strategy in $\mathcal{G}(\mathbb{A}, \mathbb{S})$, and so (s', a') is winning for \exists .

It should then be obvious how this strategy leads to a victory for \exists in the nonemptiness game. The $\mathcal{G}_{\neq \emptyset}(\mathbb{A})$ -match π starts at position a_I and the $\mathcal{G}(\mathbb{A}, \mathbb{S})$ -match π' starts at the parallel position (s_0, a_I) . Now if \exists plays the strategy sketched above, then for any resulting $\mathcal{G}_{\neq \emptyset}(\mathbb{A})$ -match

$$\pi = a_I \dots a_1 \dots a_2 \dots a_3 \dots$$

there is a parallel $\mathcal{G}(\mathbb{A}, \mathbb{S})$ -match

$$\pi' = (s_0, a_I) \dots (s_1, a_1) \dots (s_2, a_2) \dots (s_3, a_3) \dots$$

which is conform her winning strategy (Φ, Z) . From this it follows immediately that \exists wins the match π . \square

In the case of a *standard* functor, Theorem 6.1 is an almost immediate consequence of the above two lemmas, together with the history-free determinacy of parity games. For the nontrivial direction of the theorem, suppose that \mathbb{A} accepts some pointed F-coalgebra. So, by Lemma 6.10, \exists has a winning strategy in the nonemptiness game $\mathcal{G}_{\neq \emptyset}(\mathbb{A})$ starting at position a_I . As parity games are history-free determined this implies \exists actually has a *positional* winning strategy φ from position a_I in $\mathcal{G}_{\neq \emptyset}(\mathbb{A})$. Then Lemma 6.9 implies that \mathbb{A} accepts (A, φ, a_I) which finishes the proof of the theorem. For the general case, we have to do a little more work.

6.2. Nonemptiness problem: the general case. The idea underlying the proof of Theorem 6.1 is simply a reduction of the general situation to the standard case. For this purpose we need to define the *standardization* of an F-automaton.

Definition 6.11. Let F be a weak pullback preserving set functor, and suppose that F_λ is a standardization of F , i.e. there is a natural isomorphism $\lambda : F \cong F_\lambda$ (see Definition 2.25). Given a nondeterministic F-automaton $\mathbb{A} = \langle A, a_I, \Delta, Acc \rangle$, we first define the map $\Delta_\lambda : A \rightarrow F_\lambda A$ by putting

$$\Delta_\lambda(a) := \lambda_A[\Delta(a)].$$

Then the λ -standardization of \mathbb{A} is the automaton $\mathbb{A}_\lambda := \langle A, a_I, \Delta_\lambda, Acc \rangle$. \triangleleft

In words, $\Delta_\lambda(a)$ is the direct image of $\Delta(a)$ under the bijection λ_A .

Proposition 6.12. *Let F be a weak pullback preserving set functor, suppose that $\lambda : F \cong F_\lambda$, and let $\mathbb{A} = \langle A, a_I, \Delta, Acc \rangle$ be a nondeterministic F-automaton. Then \mathbb{A} accepts a pointed F-coalgebra $\langle S, \sigma, s \rangle$ iff \mathbb{A}_λ accepts the pointed F_λ -coalgebra $\langle S, \lambda_S \circ \sigma, s \rangle$.*

Proof. Let $\mathbb{S} = \langle S, \sigma \rangle$ and let $\mathbb{S}_\lambda = \langle S, \sigma_\lambda \rangle$ be the F_λ -coalgebra given by $\sigma_\lambda = \lambda_S \circ \sigma$. We prove the implication from left to right, the converse direction can be proven analogously.

Suppose that \mathbb{S} is accepted by \mathbb{A} . We want to show that \mathbb{S}_λ is accepted by \mathbb{A}_λ . By assumption, \exists has a positional strategy $(\Phi : S \times A \rightarrow FA, Z : S \times FA \rightarrow S \times A)$ in $\mathcal{G} := \mathcal{G}(\mathbb{A}, \mathbb{S})$, which is winning for any match starting at position (s, a_I) .

Now define the following positional strategy $(\Phi_\lambda : S \times A \rightarrow F_\lambda A, Z : S \times F_\lambda A \rightarrow S \times A)$ for \exists in $\mathcal{G}_\lambda = \mathcal{G}(\mathbb{A}_\lambda, \mathbb{S}_\lambda)$:

$$\begin{aligned} \Phi_\lambda(s, a) &:= \lambda_A \circ \Phi(s, a), \\ Z_\lambda(s, \alpha) &:= Z(s, \lambda_A^{-1}(\alpha)). \end{aligned}$$

Here $\lambda_A^{-1} : F_\lambda A \rightarrow FA$ denotes the *inverse* of the bijection λ_A .

It is obvious, that for any basic position (s, a) , after two moves in \mathcal{G}_λ , \exists arrives at the *same* binary relation $Z(s, \Phi(s, a))$, as after two moves in \mathcal{G} . From this it follows immediately, that $(\Phi_\lambda, Z_\lambda)$ is a winning strategy for any \mathcal{G}_λ -match starting at position (s, a_I) .

The only thing that is left to prove is that $(\Phi_\lambda, Z_\lambda)$ only suggests *legitimate moves*, at least, when we start at a basic positions that is winning for \exists . The only case worth worrying about, is whether Z_λ is legitimate at position $(s, \lambda_A(\varphi)) \in S \times F_\lambda A$, if Z is legitimate at position $(s, \varphi) \in S \times FA$. That is, we assume that $(\sigma(s), \varphi) \in \bar{F}Z$, and need to show that $(\sigma_\lambda(s), \lambda_A(\varphi)) \in \bar{F}_\lambda Z_\lambda$.

By definition of relation lifting, it follows from $(\sigma(s), \varphi) \in \bar{F}Z$ that there is some $\zeta \in FZ$ such that

$$\begin{aligned} \sigma(s) &:= (F\pi_S)(\zeta), \\ \varphi &:= (F\pi_A)(\zeta), \end{aligned}$$

where $\pi_S : Z \rightarrow S$ and $\pi_A : Z \rightarrow A$ are the projections. Furthermore the following diagram commutes by naturality of λ :

$$\begin{array}{ccccc} FS & \xleftarrow{F\pi_S} & FZ & \xrightarrow{F\pi_A} & FA \\ \lambda_S \downarrow & & \lambda_Z \downarrow & & \lambda_A \downarrow \\ F_\lambda S & \xleftarrow{F_\lambda \pi_S} & F_\lambda Z & \xrightarrow{F_\lambda \pi_A} & F_\lambda A \end{array}$$

Hence, for $\zeta_\lambda := \lambda_Z(\zeta) \in F_\lambda Z$ we may compute:

$$\begin{aligned} (F_\lambda \pi_S)(\zeta_\lambda) &= ((F_\lambda \pi_S) \circ \lambda_Z)(\zeta) \\ &= (\lambda_S \circ F_\lambda \pi_S)(\zeta) \\ &= \lambda_S(\sigma(s)) \\ &= \sigma_\lambda(s), \end{aligned}$$

and

$$\begin{aligned} (F_\lambda \pi_A)(\zeta_\lambda) &= ((F_\lambda \pi_A) \circ \lambda_Z)(\zeta) \\ &= (\lambda_A \circ F_\lambda \pi_A)(\zeta) \\ &= \lambda_A(\varphi). \end{aligned}$$

From this it is immediate, by definition of \bar{F}_λ , that $(\sigma_\lambda(s), \lambda_A(\varphi)) \in \bar{F}_\lambda Z_\lambda$. \square

The proof of the main result in this section now follows easily.

Proof of Theorem 6.1. For the nontrivial direction of the Theorem, let $\mathbb{A} = \langle A, a_I, \Delta, \Omega \rangle$ be an nondeterministic parity F-automaton with $L(\mathbb{A}) \neq \emptyset$. As an immediate consequence of Proposition 6.12, we see that $L(\mathbb{A}_\lambda) \neq \emptyset$, where F_λ is some standardization of F via a natural isomorphism λ . Since we already showed the theorem to hold for standard functors, this means that there is some F_λ -coalgebra $\langle A, \varphi \rangle$, with $\varphi(a) \in \Delta_\lambda(a)$ for all $a \in A$, such that $(\langle A, \varphi \rangle, a_I)$ is accepted by \mathbb{A}_λ .

Now consider the F-coalgebra $\langle A, \alpha \rangle$ given by defining

$$\alpha(a) := \lambda_A^{-1}(\varphi(a)),$$

where λ_A^{-1} is the inverse of the bijection λ_A . Clearly then, $\varphi = \lambda_A \circ \alpha$, and so it follows from Proposition 6.12, that \mathbb{A} accepts $(\langle A, \alpha \rangle, a_I)$. Finally, by definition of Δ_λ , $\varphi(a) = \lambda_A(\alpha'(a))$ for some $\alpha'(a) \in \Delta(a)$. Since λ_A is a bijection, it follows that $\alpha(a) = \alpha'(a)$, so that $\alpha(a) \in \Delta(a)$, as required. \square

7. CONCLUSIONS & QUESTIONS

There is a long list of issues that need some further discussion. 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, given the connection between formulas of coalgebraic fixed-point logics and coalgebra automata theory, established in [29, 30], it is easy to show that the logics introduced in the mentioned work, have the finite model property. Or, generalizing results in [6], we can show that the coalgebraic fixed-point logics of [29] all have some kind of *uniform interpolation*. We hope to say more 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 *nonsymmetric* 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-point operators) *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, see Corollary 7.1 in [30].

In order to gain a better understanding of coalgebra automata it will also be useful to investigate instances of **F**-coalgebra automata other than word, tree or graph automata. In Definition 2.9 we defined a class of **Set**-functors which all fall into the scope of our work. Future research will show whether e.g. the coalgebra automata for the functors D_ω and M_ω yield reasonable automata for probabilistic transition systems and for directed weighted graphs respectively.

Finally, we are interested to see whether the conditions on the functor are really needed. 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] R. Bird and O. de Moor. *Algebra of Programming*. Prentice Hall, 1997.
- [5] E. Clarke, O. Grumberg, and D. Peled, editors. *Model Checking*. The MIT Press, 2000.
- [6] G. d'Agostino and M. Hollenberg. Logical questions concerning the μ -calculus. *Journal of Symbolic Logic*, 65:310–332, 2000.
- [7] E.P. de Vink and J.J.M.M. Rutten. Bisimulation for probabilistic transition systems: a coalgebraic approach. *Theoretical Computer Science*, 221(1–2):271–293, 1999.
- [8] E.A. Emerson and C.S. Jutla. Tree automata, mu-calculus and determinacy. In *Proceedings of the 32nd IEEE Symposium on Foundations of Computer Science (FoCS'91)*, pages 368–377. IEEE Computer Society Press, 1991.
- [9] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logic, and Infinite Games*, volume 2500 of *LNCS*. Springer, 2002.
- [10] H. P. Gumm and T. Schröder. Monoid-labeled transition systems. In *Coalgebraic Methods in Computer Science (CMCS 2001)*, volume 44 of *Electronic Notes in Theoretical Computer Science*, pages 1–20, 2001.
- [11] B. Jacobs. Distributive laws for the coinductive solution of recursive equations. *Information and Computation*, 204:561–587, 2006.

- [12] 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.
- [13] V. King, O. Kupferman, and M. Vardi. On the complexity of parity word automata. In *Proceedings of the International Conference on Foundations of Software Science and Computation Structures (FOSSACS 2001)*, volume 2030 of *LNCS*, pages 276–286, 2001.
- [14] C. Kissig. Decidability of S2S. Master's thesis, ILLC, University of Amsterdam, 2007.
- [15] C. Kupke and Y. Venema. Closure properties of coalgebra automata. In *Proceedings of the Twentieth Annual IEEE Symposium on Logic in Computer Science (LICS 2005)*, pages 199–208. IEEE Computer Society Press, 2005.
- [16] K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1–28, 1991.
- [17] D. Martin. Borel determinacy. *Annals of Mathematics*, 102:363–371, 1975.
- [18] R. McNaughton. Testing and generating infinite sequences by a finite automaton. *Information and control*, 9:521–530, 1966.
- [19] L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999. (Erratum published *APAL* 99:241–259, 1999).
- [20] A. Mostowski. Games with forbidden positions. Technical Report 78, Instytut Matematyki, Uniwersytet Gdański, Poland, 1991.
- [21] D. Muller and P. Schupp. Simulating alternating tree automata by nondeterministic ones. *Theoretical Computer Science*, 141:69–107, 1995.
- [22] D. Niwiński. Fixed point characterization of infinite behavior of finite-state systems. *Theoretical Computer Science*, 189:1–69, 1997.
- [23] N. Piterman. 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)*, pages 255–264. IEEE Computer Society, 2006.
- [24] M.O. Rabin. Decidability of second-order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141:1–35, 1969.
- [25] J. Rutten. Universal coalgebra: A theory of systems. *Theoretical Computer Science*, 249:3–80, 2000.
- [26] S. Safra. On the complexity of ω -automata. In *Proc. 29th IEEE Symposium on the Foundations of Computer Science*, pages 319–327, 1988.
- [27] V. Trnková. Some properties of set functors. *Commentationes Mathematicae Universitatis Carolinae*, 10(2), 1969.
- [28] V. Trnková. Relational automata in a category and theory of languages. In *Proc. FCT 1977*, pages 340–355. Springer, 1977. LNCS 56.
- [29] Y. Venema. Automata and fixed point logics for coalgebras. *Electronic Notes in Theoretical Computer Science*, 106:355–375, 2004.
- [30] Y. Venema. Automata and fixed point logic: a coalgebraic perspective. *Information and Computation*, 204:637–678, 2006.