

Complementation of Coalgebra Automata

Christian Kissig^{1,*} and Yde Venema^{2,**}

¹ Department of Computer Science, University of Leicester, University Road,
Leicester LE1 7RH, United Kingdom

`christian@mcs.le.ac.uk`

² Institute for Logic, Language, and Computation, Universiteit van Amsterdam,
Plantage Muidergracht 24, 1018 TV Amsterdam, Netherlands

`Y.Venema@uva.nl`

Abstract. Coalgebra automata, introduced by the second author, generalize the well-known automata that operate on infinite words/streams, trees, graphs or transition systems. This coalgebraic perspective on automata lays foundation to a universal theory of automata operating on infinite models of computation.

In this paper we prove a complementation lemma for coalgebra automata. More specifically, we provide a construction that transforms a given coalgebra automaton with parity acceptance condition into a device of similar type, which accepts exactly those pointed coalgebras that are rejected by the original automaton. Our construction works for automata operating on coalgebras for an arbitrary standard set functor which preserves weak pullbacks and restricts to finite sets.

Our proof is coalgebraic in flavour in that we introduce and use a notion of game bisimilarity between certain kinds of parity games.

1 Introduction

Through their close connection with modal fixpoint logics, automata operating on infinite objects such as words/streams, trees, graphs or transition systems, provide an invaluable tool for the specification and verification of the ongoing behavior of infinite systems. Coalgebra automata, introduced by the second author [9, 10], generalize the well-known automata operating on possibly infinite systems of a specific type. The motivation underlying the introduction of coalgebra automata is to gain a deeper understanding of this branch of automata theory by studying properties of automata in a uniform manner, parametric in the type of the recognized models, that is, the coalgebra functor. The aim is thus to contribute to Universal Coalgebra [8] as a mathematical theory of state-based evolving systems.

* This author would like to thank the Institute for Logic, Language, and Computation, and especially Yde Venema and his Algebra|Coalgebra group for their hospitality during his visit, which this work ensued.

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

Operationally, coalgebra automata are devices for scanning pointed coalgebras. Structurally the automata rather closely resemble the coalgebras on which they operate; for instance, in the nondeterministic variant, every state in a coalgebra automaton may be seen as capturing various ‘realizations’ as a T -coalgebra state. The resemblance between coalgebra automata and coalgebras is reflected by the *acceptance game* that is used to determine whether a pointed coalgebra is eventually accepted or rejected by the automaton: This acceptance game, a two-play parity graph game, is a variant of the bisimilarity game [1] that can be played in order to see whether two pointed coalgebras are bisimilar or not.

Earlier work by Kupke and the second author [5] revealed that in fact a large part of the theory of parity automata can be lifted to a coalgebraic level of generality, and thus indeed belongs to the theory of Universal Coalgebra. More specifically, the main result of [5] is a construction transforming a given alternating coalgebra automaton into an equivalent nondeterministic one; this shows that the nondeterministic model is just as powerful as the alternating one. In addition, coalgebra automata satisfy various *closure properties*: under some conditions on the coalgebra type functor, the collection of recognizable languages (that is, classes of pointed coalgebras that are recognized by some coalgebra automaton) are closed under taking unions, intersections, and existential projections. These results have many applications in the theory of coalgebraic fixpoint logics.

The question whether coalgebra automata admit a *Complementation Lemma* was left as an open problem in [5]. Closure under complementation does not obviously hold, since even for alternating automata the role of the two players in the acceptance game is *not* symmetric. Nevertheless, in this paper we provide a positive answer to the complementation problem, under the same conditions¹ on the functor as in [5]. More precisely, we will prove the following theorem.

Theorem 1 (Complementation Lemma). *Let T be a standard set functor that restricts to finite sets and preserves weak pullbacks. Then the class of recognizable T -languages is closed under complementation.*

Our proof of this Complementation Lemma will be based on an explicit construction which transforms a T -automaton \mathbb{A} into a T -automaton \mathbb{A}^c which accepts exactly those pointed T -coalgebras that are rejected by \mathbb{A} . In order to define and use this construction, it will be necessary to move from the nondeterministic format of our automata to a wider setting. First of all, since we want to apply the dualization and complementation method of Muller & Schupp [6], we work with automata that are *alternating*, meaning that we increase the role of Abélard in the acceptance game, and *logical* in the sense that the possible unfoldings of an automaton state are expressed as a logical formula rather than as a set of options. Concretely, we would like to move to a setting where the unfolding of an automaton state is a lattice term over nabla-formulas over automaton states,

¹ The condition that T should restrict to finite sets is not mentioned in [5], but it is needed there as well, in order to guarantee the correctness of the construction of an equivalent nondeterministic automaton for a given alternating one.

with conjunctions denoting choices for Abélard, disjunctions denoting choices for Éloise, and ∇ denoting Moss' coalgebraic modality. However, it will turn out that by the nature of our construction we also need conjunctions and disjunctions *under* (that is, within the scope of) the modality. Because of this enriched transition structure, the resulting automata will be called *transalternating*.

Concerning these transalternating automata we will prove two results. First, we establish a complementation lemma, by providing a construction which transforms a transalternating automaton \mathbb{A} into a transalternating automaton \mathbb{A}^c which behaves as the complement of \mathbb{A} . Note that this construction is linear in size: in fact, \mathbb{A}^c is based on the *same* carrier set as \mathbb{A} . And second, we give a construction transforming an arbitrary transalternating automaton into an equivalent alternating one (of size exponential in the size of the original automaton).

At the heart of our construction of the complement automaton lies a new result in coalgebraic modal logic that we call the *One-step Complementation² Lemma*. This Lemma states that the *Boolean dual* Δ of Moss' modality ∇ can be expressed using disjunctions, Moss' modality itself, and conjunctions. Here 'Boolean dual' refers to the fact that we want for a given formula $\nabla\alpha$ that

$$\Delta\alpha \equiv \neg\nabla(T\neg)(\alpha), \quad (1)$$

For instance, in the case of the power set functor \mathcal{P} , defining, for a nonempty α ,

$$\Delta\alpha := \nabla\emptyset \vee \bigvee\{\nabla\{a\} \mid a \in \alpha\} \vee \nabla\{\bigwedge\alpha, \top\}$$

we indeed obtain that

$$\Delta\{a_1, \dots, a_n\} \equiv \neg\nabla\{\neg a_1, \dots, \neg a_n\}.$$

In the general case, we will see that for each formula $\nabla\alpha$, we can find a set $Q(\alpha)$ (which only uses finite conjunctions over the ingredient formulas from α), such that the following definition

$$\Delta\alpha := \bigvee\{\nabla\beta \mid \beta \in Q(\alpha)\}$$

indeed provides a Boolean dual Δ for ∇ , i.e., for this Δ we may prove (1). Note that in order for this definition to give a proper (finitary!) lattice term, we need the set $Q(\alpha)$ to be finite; it is for this reason that we need the functor T to restrict to finite sets.

Applying the methodology of Muller & Schupp [6] to this coalgebraic dualization, we obtain a very simple definition of the complement automaton. Roughly speaking, we obtain the complement of a transalternating automaton by *dualizing* its transition map, and performing a *role switch* on its priority map.

² Perhaps 'One-step Dualization Lemma' might have been a more adequate name — we chose against this because in the context of coalgebraic logic the word 'dual' has strong connotations towards Stone-type dualities.

Then, in order to *prove* that \mathbb{A}^c is the complement of \mathbb{A} , we compare, for an arbitrary pointed coalgebra (\mathbb{S}, s) , the acceptance games $\mathcal{G}(\mathbb{A}, \mathbb{S})$ and $\mathcal{G}(\mathbb{A}^c, \mathbb{S})$. We will show that

$$(a_I, s) \in \text{Win}_{\exists}(\mathcal{G}(\mathbb{A}^c, \mathbb{S})) \text{ iff } (a_I, s) \in \text{Win}_{\forall}(\mathcal{G}(\mathbb{A}, \mathbb{S})).$$

In order to streamline the proof of this result, we introduce a notion of equivalence between positions in parity games. We base this definition on well-known ideas from game theory (for instance, van Benthem’s notion of power of players to achieve certain outcomes in finite games [2], and Pauly’s bisimulation for coalition logics [7]), adding some features taking care of the acceptance condition. With the resulting notion of ‘basic game bisimilarity’, we may exploit some of the coalgebraic intuitions we have of parity graph games.

2 Preliminaries

For background information on coalgebra automata, we refer to [10, 5]. We will also use notation and terminology from that paper, with the following minor deviations and additions.

Categorical notions. Throughout the paper we assume that T is a set functor, that is, an endofunctor on the category **Set** of sets as objects and functions as arrows. This functor is supposed to be (i) standard (that is to preserve embeddings), (ii) to preserve weak pullbacks, and (iii) to restrict to finite sets. We let T_ω denote the finitary version of T , given on objects by $T_\omega X := \bigcup\{TX_0 \mid X_0 \subseteq_\omega X\}$, while on arrows $T_\omega f := Tf$. Given an object $\alpha \in T_\omega X$, we let $\text{Base}(\alpha)$ denote the smallest subset $X_0 \subseteq_\omega X$ such that $\alpha \in T_\omega(X_0)$.

Preserving weak pullbacks, T extends to a unique relator \overline{T} , which is monotone when T is standard. For a given relation $Z \subseteq X \times Y$, we denote the lifted relation as $\overline{T}Z \subseteq TX \times TY$. Without warning we will assume that \overline{T} is monotone, that it is a relator (that is, it preserves the diagonal relation and distributes over relation composition), and that it commutes with taking relation converse.

Games. Throughout the paper, when discussing parity games, the players will be $0 = \exists$ and $1 = \forall$. We will denote an arbitrary player as Π , and use Σ to denote Π ’s opponent.

Coalgebra automata. As mentioned we will work with coalgebra automata in logical format [10, section 4.5]. It will be convenient to make explicit reference to Moss’ modality.

Definition 1. *The functor $\mathcal{L} : \mathbf{Set} \rightarrow \mathbf{Set}$ takes a set Q to the carrier $\mathcal{L}(Q)$ of the free lattice term algebra over Q . Lattice terms t are respectively of the form*

$$a ::= q \in Q \mid \bigwedge \tau \mid \bigvee \tau,$$

Table 1. Acceptance game for alternating automata

Position	Player	Admissible Moves	$\Omega_{\mathcal{G}}$
$(q, s) \in Q \times S$	-	$\{(\theta(q), s)\}$	$\Omega(q)$
$(\bigwedge \tau, s) \in \mathcal{L}T^{\nabla}Q \times S$	\forall	$\{(a, s) \mid a \in \tau\}$	0
$(\bigvee \tau, s) \in \mathcal{L}T^{\nabla}Q \times S$	\exists	$\{(a, s) \mid a \in \tau\}$	0
$(\nabla \alpha, s) \in T^{\nabla}Q \times S$	\exists	$\{Z \subseteq Q \times S \mid (\alpha, \sigma(s)) \in \overline{TZ}\}$	0
$Z \subseteq Q \times S$	\forall	Z	0

where τ denotes a finite set of lattice terms. Given a set X , let $T_{\omega}^{\nabla}X$ denote the set $\{\nabla \alpha \mid \alpha \in T_{\omega}X\}$, and define $\mathcal{L}_1Q := \mathcal{L}T_{\omega}^{\nabla}\mathcal{L}(Q)$. Elements of $\mathcal{L}Q$ will be called depth-zero formulas and elements of \mathcal{L}_1Q depth-one formulas over Q .

Definition 2 (Alternating Automata). An alternating T -automaton is a structure $\mathbb{A} = \langle Q, \theta, q_I, \Omega \rangle$ consisting of a finite set Q of states, a transition function $\theta : Q \rightarrow \mathcal{L}T_{\omega}^{\nabla}Q$, a state $q_I \in Q$ distinguished as initial, and a priority function $\Omega : Q \rightarrow \mathbb{N}$.

Definition 3 (Acceptance Game). The notion of an alternating T -automaton accepting a pointed T -coalgebra (\mathbb{S}, s) is defined in terms of the parity graph game $\mathcal{G}(\mathbb{A}, \mathbb{S})$ given by Table 1: \mathbb{A} accepts a pointed coalgebra (\mathbb{S}, s) iff \exists has a winning strategy in the acceptance game $\mathcal{G}(\mathbb{A}, \mathbb{S})$ from the initial position (q_I, s) . Positions of the form $(q, s) \in Q \times S$ are called basic.

Remark 1. In a basic position (q, s) , there is exactly one admissible move, namely, to position $(\theta(q), s)$. As a consequence, technically it does not matter to which player the position is assigned. In Table 1 we have not assigned a player to these positions, since this has some conceptual advantages further on.

One-step semantics. It will be convenient to think of elements of \mathcal{L}_1Q as formulas indeed, with the following semantics.

Definition 4. Given sets Q and S , a Q -valuation on S is a map $V : Q \rightarrow \mathcal{P}(S)$. We define relations $\Vdash_0^V \subseteq S \times \mathcal{L}Q$ and $\Vdash_1^V \subseteq TS \times \mathcal{L}_1Q$, as follows. For \Vdash_0^V , we define $s \Vdash_0^V q$ if $s \in V(q)$, $s \Vdash_0^V \bigwedge \tau$ ($\bigvee \tau$, respectively) if $s \Vdash_0^V a$ for all $a \in \tau$ (some $a \in \tau$, respectively); and we define a relation \Vdash_1^V such that $\sigma \Vdash_1^V \nabla \alpha$ if $(\sigma, \alpha) \in \overline{T}(\Vdash_0^V)$, while for \bigwedge and \bigvee the same clauses apply as for \Vdash_0^V .

For clarity of notation we also write $V, s \Vdash_0 a$ for $s \Vdash_0^V a$, and $V, s \Vdash_1 a$ for $s \Vdash_1^V a$.

3 One-Step Complementation

The notion of the (Boolean) *dual* of a connective makes a frequent appearance in various branches of logic. For instance, \wedge & \vee , and \diamond & \square , are well-known pairs of dual operators in propositional and in modal logic, respectively. In this section we introduce the dual, Δ , of Moss' coalgebraic modality, ∇ . Building on this we define the dual of a depth-one formula.

Definition 5. Let T be a standard set functor. Then for each set Q , and each $\alpha \in T_\omega Q$, we define the set $D(\alpha) \subseteq T_\omega \mathcal{P}Q$ as follows:

$$D(\alpha) := \left\{ \Phi \in T_\omega \mathcal{P}_\omega \text{Base}(\alpha) \mid (\alpha, \Phi) \notin (\overline{T\neq}) \right\}$$

In case T restricts to finite sets, define $\Delta\alpha$ as the following formula:

$$\Delta\alpha := \bigvee \left\{ \nabla T\wedge(\Phi) \mid \Phi \in D(\alpha) \right\}.$$

Here we see the connective \wedge as a map $\wedge : \mathcal{P}_\omega Q \rightarrow Q$, so that $T\wedge : T_\omega \mathcal{P}_\omega Q \rightarrow T_\omega Q$.

We claim that Δ is the Boolean dual of ∇ . In case we are dealing with a full Boolean language, we can simply express this fact as in (1). In a setting where negation is not available as a connective, we can formulate the notion of Boolean duality as follows.

Definition 6. Given a valuation $V : Q \rightarrow \mathcal{P}(S)$, we define the complementary valuation V^c by putting $V^c(q) := S \setminus V(q)$. Two formulas a, b in $\mathcal{L}Q$ (in \mathcal{L}_1Q , respectively) are (Boolean) duals if for all sets S and all Q -valuations on S , $S \setminus \widehat{V}^c(a) = \widehat{V}(b)$ (respectively, if $TS \setminus \widehat{V}^c(a) = \widehat{V}(b)$).

Putting it differently, two depth-zero formulas a and b are Boolean duals iff for all sets S , all valuations V , and all $s \in S$: $V^c, s \not\vdash_0 a$ iff $V, s \Vdash_0 b$. Likewise, two depth-one formulas a and b are Boolean duals iff for all sets S , all valuations V , and all $\sigma \in TS$: $V^c, \sigma \not\vdash_1 a$ iff $V, \sigma \Vdash_1 b$.

Theorem 2 (One-Step Complementation Lemma). Let T be a standard set functor preserving weak pullbacks, and let $\alpha \in T_\omega Q$ for some set Q . Then $\nabla\alpha$ and $\Delta\alpha$ are Boolean duals.

Proof. Fix an arbitrary set S , an arbitrary Q -valuation V on S , and an arbitrary element σ of TS .

First assume that $V, \sigma \Vdash_1 \Delta\alpha$, that is, $V, \sigma \Vdash_1 \nabla(T\wedge)(\Phi)$ for some $\Phi \in D(\alpha)$. Then there is some relation $Y \subseteq \mathcal{P}Q \times S$ such that $Y \subseteq (Gr \wedge) \circ (\Vdash_0^V)^\smile$ and $(\Phi, \sigma) \in \overline{TY}$. In order to show that $V^c, \sigma \not\vdash_1 \nabla\alpha$, suppose for contradiction that there is some relation Z such that $(\sigma, \alpha) \in \overline{TZ}$ and $V^c, t \Vdash_0 q$ for all pairs $(t, q) \in Z$. It follows that $(\sigma, \alpha) \in \overline{TZ}$ and $Z \cap \Vdash_0^V = \emptyset$.

Now consider the relation $R := Y \circ Z \subseteq \mathcal{P}Q \times Q$, then clearly $(\Phi, \alpha) \in \overline{TR} = \overline{TY} \circ \overline{TZ}$. On the other hand, it follows from the definition of R that $R \subseteq \neq$, because for any $(\phi, q) \in R$ there is an $s \in S$ such that (i) $(\phi, s) \in Y$ implying $V, s \Vdash_0 p$ for all $p \in \phi$, and (ii) $(s, q) \in Z$ meaning that $V, s \not\vdash_0 q$. But this gives the desired contradiction since $\Phi \in D(\alpha)$.

Conversely, assume that $V^c, \sigma \not\vdash_1 \nabla\alpha$. In order to show that $V, \sigma \Vdash_1 \Delta\alpha$ we need to find some $\Phi \in D(\alpha)$ such that $V, \sigma \Vdash_1 \nabla(T\wedge)(\Phi)$. For this purpose, define a map $\phi : S \rightarrow \mathcal{P}\text{Base}(\alpha)$ by putting, for any $s \in S$, $\phi_s := \{q \in \text{Base}(\alpha) \mid V, s \Vdash_0 q\}$.

We claim that $\Phi := T\phi(\sigma)$ has the required properties. First of all, it follows by construction that $Gr(\bigwedge \circ \phi) \subseteq \Vdash_0^V$, so that $Gr(T\phi) \circ Gr(T\bigwedge) \subseteq \overline{T}\Vdash_0^V$. From this it is immediate that $V, \sigma \Vdash_1 \nabla T\bigwedge(\Phi)$. It remains to show that $\Phi \in D(\alpha)$. For that purpose, consider the relation $Z := Gr(\phi) \circ \not\vdash \subseteq S \times Q$. It is easily verified that $V^c, s \Vdash_0 q$ for all $(s, q) \in Z$. Hence, we may derive from the assumption $V^c, \sigma \not\vdash_1 \nabla \alpha$ that $(\sigma, \alpha) \notin \overline{T}Z = Gr(T\phi) \circ (\overline{T}\not\vdash)$. But then it follows from $(\sigma, \Phi) \in Gr(T\phi)$ that $(\Phi, \alpha) \notin (\overline{T}\not\vdash)$, as required.

It is an almost immediate consequence of this result, that negation can be defined as an abbreviated connective in the finitary version of Moss' language (with finite conjunctions, disjunctions, and ∇). As we will see, the main result of this paper, viz., the complementation lemma for coalgebra automata, is also a direct corollary of the one-step complementation lemma. As a first step towards proving that statement, let us define the dual of a one-step formula.

Definition 7. *Given a set Q we define the base dualization map $\delta_0 : \mathcal{L}Q \rightarrow \mathcal{L}Q$ and the one-step dualization map $\delta_1 : \mathcal{L}_1Q \rightarrow \mathcal{L}_1Q$ as follows:*

$$\begin{array}{ll} \delta_0(q) := q & \delta_1(\nabla \alpha) := \Delta(T\delta_0)\alpha \\ \delta_0(\bigwedge \phi) := \bigvee \delta_0[\phi] & \delta_1(\bigwedge \phi) := \bigvee \delta_1[\phi] \\ \delta_0(\bigvee \phi) := \bigwedge \delta_0[\phi] & \delta_1(\bigvee \phi) := \bigwedge \delta_1[\phi] \end{array}$$

Example 1. With $T = \mathcal{P}$ and $\alpha = \{q_1 \vee q_2, q_3 \wedge q_4\}$, we may calculate that

$$\begin{aligned} \delta_1(\nabla \alpha) &= \Delta\{q_1 \wedge q_2, q_3 \vee q_4\} \\ &= \nabla \emptyset \vee \nabla\{q_1 \wedge q_2\} \vee \nabla\{q_3 \vee q_4\} \vee \nabla\{(q_1 \wedge q_2) \wedge (q_3 \vee q_4), \top\}. \end{aligned}$$

The following corollary of the One-Step Complementation lemma states that these dualization operations indeed send formulas to their one-step Boolean duals. The proof of this result is left for the reader.

Corollary 1. *The formulas a and $\delta_1(a)$ are Boolean duals, for any $a \in \mathcal{L}_1Q$.*

4 Game Bisimulations

In the preliminaries we introduced acceptance games of coalgebra automata as specific parity graph games. In acceptance games we distinguish some positions as basic, which allows us to partition plays into rounds, each delimited by basic positions. From this observation we derive the following definition of basic sets.

Definition 8 (Basic Sets of Positions). *Given a parity graph game $\mathcal{G} = \langle V_0, V_1, E, v_I, \Omega \rangle$, call a set $B \subseteq V_0 \cup V_1$ of positions basic if*

1. *the initial position of \mathcal{G} belongs to B ;*
2. *any full play starting at some $b \in B$ either ends in a terminal position or it passes through another position in B ; and*
3. *$\Omega(v) = 0$ iff $v \notin B$.*

Another way of expressing condition 2 is to say that there are no infinite paths of which the first position is the only basic one. The third condition is there to ensure that who wins an infinite match is determined by the sequence of basic positions induced by the match. It should be clear that the collection of basic positions in acceptance games of coalgebra automata indeed qualifies as a basic set in the sense of Definition 8. By a slight abuse of language, we refer to elements of any basic set as basic positions.

The point behind the introduction of basic positions is that, just as in the special case of acceptance games, we may think of parity games as proceeding in rounds that start and finish at a basic position. Formally we define these rounds, that correspond to subgraphs of the arena, by their unravelling as follows.

Definition 9 (Local Game Trees). *Let B be a basic set related to some parity game $\mathcal{G} = \langle V_0, V_1, E, v_I, \Omega \rangle$. The local game tree associated with a basic position $b \in B$ is defined as the following bipartite tree $T^b = \langle V_0^b, V_1^b, E^b, v_I^b \rangle$. Let V^b be the set of those finite paths β starting at b , of which the only basic positions are $\text{first}(\beta) = b$ and possibly $\text{last}(\beta)$. The bipartition of V^b is given through the bipartition of V , that is $V_\Pi^b := \{\beta \in V^b \mid \text{last}(\beta) \in V_\Pi\}$ for both players $\Pi \in \{0, 1\}$. The root v_I^b is $\langle b \rangle$, the path beginning and ending in b . The edge relation is defined as $E^b := \{(\beta, \beta v) \mid v \in E(\text{last}(\beta))\}$.*

A node $\beta \in V^b$ is a leaf of T^b if $|\beta| > 1$ and $\text{last}(\beta) \in B$; we let $\text{Leaves}(T^b)$ denote the set of leaves of T^b , and put $N(b) := \{\text{last}(\beta) \mid \beta \in \text{Leaves}(T^b)\}$.

Intuitively, T^b can be seen as the tree representation of one round of the game \mathcal{G} , starting at b . $N(b) \subseteq V$ denotes the set of positions one may encounter in a match of \mathcal{G} as the next basic position after b — this intuition will be made more precise in Lemma 1 below. It follows from the definition that all paths in T^b are finite, and hence we may use (bottom-up) induction on the immediate successor relation of the tree, or, as we shall say, on the *height* of nodes.

Once we have established the dissection of a parity graph game through basic positions, we may think of the game being in an iterative strategic normal form, in the sense that in each round players make only one choice, determining their complete strategies for that round right at the beginning. We may formalize this using the game-theoretic notion of *power*, which describes the terminal positions in a finitary graph game which a player can force, for instance by deploying a strategy; we refer to van Benthem [2] for an extensive discussion and pointers to the literature. In the following we define the notion of power of a player Π at a basic position b as a collection of sets of basic nodes. Intuitively, when we say that U is in Π 's power at basic position b , we mean that at position b , Π has a local strategy ensuring that the next basic position belongs to the set U . Hence, the collection of sets that are in Π 's power at a certain basic position is closed under taking supersets. In the definition below we define a collection $P_\Pi(b) \subseteq \mathcal{P}B$. Further on we will see that a set $U \subseteq B$ is in Π 's power at b , in the sense described above, iff $U \supseteq V$ for some $V \in P_\Pi(b)$. It will be convenient to use the following notation.

Definition 10. Given a set $\mathcal{H} \subseteq \mathcal{P}B$, define

$$\uparrow(\mathcal{H}) := \{V \subseteq B \mid V \supseteq U \text{ for some } U \in \mathcal{H}\}.$$

Definition 11 (Powers). Let B be a basic set related to some parity game $\mathcal{G} = \langle V_0, V_1, E, v_I, \Omega \rangle$, and let b be a basic position. By induction on the height of a node $\beta \in V^b$, we define, for each player Π , the power of Π at β as a collection $P_\Pi(\beta)$ of subsets of $N(b)$:

- If $\beta \in \text{Leaves}(T^b)$, we put, for each player,

$$P_\Pi(\beta) := \{\{last(\beta)\}\}.$$

- If $\beta \notin \text{Leaves}(T^b)$, we put

$$P_\Pi(\beta) := \begin{cases} \bigcup \{P_\Pi(\gamma) \mid \gamma \in E^b(\beta)\} & \text{if } \beta \in V_\Pi^b, \\ \left\{ \bigcup_{\gamma \in E^b(\beta)} Y_\gamma \mid Y_\gamma \in P_\Pi(\gamma), \text{ all } \gamma \right\} & \text{if } \beta \in V_\Sigma^b. \end{cases}$$

where Σ denotes the opponent of Π .

Finally, we define the power of Π at b as the set $P_\Pi(b) := P_\Pi(\langle b \rangle)$, where $\langle b \rangle$ is the path beginning and ending at b .

Perhaps some special attention should be devoted to the paths β in T^b such that $E^b(\beta) = \emptyset$. If such a β is a leaf of T^b , then the definition above gives $P_\Pi(\beta) = P_\Sigma(\beta) = \{\{last(\beta)\}\}$. But if β is not a leaf of T^b , then we obtain, by the inductive clause of the definition:

$$P_\Pi(\beta) := \begin{cases} \emptyset & \text{if } \beta \in V_\Pi^b, \\ \{\emptyset\} & \text{if } \beta \in V_\Sigma^b. \end{cases}$$

This indeed confirms our intuition that at such a position, the player who is to move gets stuck and loses the match.

The following definitions make the notion of a local game more precise.

Definition 12 (Local Games). Let $\mathcal{G} = \langle V_0, V_1, E, v_I, \Omega \rangle$ be a parity graph game with basic set $B \subseteq V$. We define a game \mathcal{G}^b local to a basic position $b \in B$ in \mathcal{G} as the (finite length) graph game played by 0 and 1 on the local game tree $T^b = \langle V^b, E^b \rangle$. Matches of this game are won by a player Π if their opponent Σ gets stuck, and end in a tie if the last position of the match is a leaf of T^b .

Definition 13 (Local Strategies). A local strategy of a player Π in a local game $\mathcal{G}^b = \langle V_0^b, V_1^b, E^b, v_I^b \rangle$ is a partial function $f : V_\Pi^b \rightarrow V^b$ defined on such β iff $E^b(\beta)$ is non-empty, then $f(\beta) \in E^b(\beta)$. Such a local strategy for player Π is surviving if it guarantees that Π will not get stuck, and thus does not lose the local game; and a local strategy is winning for Π if it guarantees that her opponent Σ gets stuck.

Consider the match of the local game \mathcal{G}^b in which 0 and 1 play local strategies f_0 and f_1 , respectively. If this match ends in a leaf β of the local game tree, we let $\text{Res}(f_0, f_1)$ denote the basic position $\text{last}(\beta)$; if one of the players gets stuck in this match, we leave $\text{Res}(f_0, f_1)$ undefined. Given a local strategy f_0 for player 0, we define

$$X_{f_0} := \{ \text{Res}(f_0, f_1) \mid f_1 \text{ a local strategy for player 1} \},$$

and similarly we define X_{f_1} for a strategy f_1 for player 1.

Local strategies for Π in \mathcal{G}^b can be linked to (fragments of) strategies for Π in \mathcal{G} . Since these links are generally obvious, we will refrain from introducing notation and terminology here.

The following lemma makes precise the links between players' power and their local strategies.

Proposition 1. *Let \mathcal{G} , B and b as in Definition 11, let $\Pi \in \{0, 1\}$ be a player, and let W be a subset of $N(b)$. Then the following are equivalent:*

1. $W \in P_\Pi(b)$;
2. Π has a surviving strategy f in \mathcal{G}^b such that $W = X_f$.

In the next section we will need the following lemma, which states a determinacy property of local games.

Proposition 2. *Let \mathcal{G} , B and b as in Definition 11, and let $\Pi \in \{0, 1\}$ be a player. For any subset $U \subseteq N(b)$, if $U \not\subseteq \uparrow(P_\Pi(b))$ then there is a $V \in P_\Sigma(b)$ such that $U \cap V = \emptyset$.*

The partitioning of matches of parity games into rounds between basic positions, and the normalization of the players' moves within each round, lay the foundations to the introduction of a structural equivalence between parity games that we refer to as basic game bisimulation. This equivalence combines (i) a structural part that can be seen as an instantiation of Pauly's bisimulation between extensives strategic games [7], or van Benthem's *power bisimulation* [2], with a combinatorial part that takes care of the parity acceptance condition.

Definition 14 (Basic Game Bisimulation). *Let $\mathcal{G} = \langle V_0, V_1, E, \Omega \rangle$ and $\mathcal{G}' = \langle V'_0, V'_1, E', \Omega' \rangle$ be parity graph games with basic sets B and B' , respectively, and let Π and Π' be (not necessarily distinct) players in \mathcal{G} and \mathcal{G}' , respectively.*

A Π, Π' -game bisimulation is a binary relation $Z \subseteq B \times B'$ satisfying for all $v \in V$ and $v' \in V'$ with vZv' the structural conditions

- $(\Pi, \text{forth}) \forall W \in P_\Pi^{\mathcal{G}}(v). \exists W' \in P_{\Pi'}^{\mathcal{G}'}(v'). \forall w' \in W'. \exists w \in W. (w, w') \in Z,$
- $(\Sigma, \text{forth}) \forall W \in P_\Sigma^{\mathcal{G}}(v). \exists W' \in P_{\Sigma'}^{\mathcal{G}'}(v'). \forall w' \in W'. \exists w \in W. (w, w') \in Z,$
- $(\Pi, \text{back}) \forall W' \in P_{\Pi'}^{\mathcal{G}'}(v'). \exists W \in P_\Pi^{\mathcal{G}}(v). \forall w \in W. \exists w' \in W'. (w, w') \in Z,$
- $(\Sigma, \text{back}) \forall W' \in P_{\Sigma'}^{\mathcal{G}'}(v'). \exists W \in P_\Sigma^{\mathcal{G}}(v). \forall w \in W. \exists w' \in W'. (w, w') \in Z,$

and the priority conditions

- (parity) $\Omega(v) \bmod 2 = \Pi$ iff $\Omega'(v') \bmod 2 = \Pi'$,
- (contraction) for all $v, w \in V$ and $v', w' \in V'$ with vZv' and wZw' , $\Omega(v) \leq \Omega(w)$ iff $\Omega'(v') \leq \Omega'(w')$.

In Condition (parity) above we identify players with their characteristic parity. Note that in fact there are only two kinds of game bisimulations: the $(0, 0)$ -bisimulations coincide with the $(1, 1)$ -bisimulations, and the $(0, 1)$ -bisimulations coincide with the $(1, 0)$ -bisimulations.

The following theorem bears witness to the fact that game bisimulation is indeed a good notion of structural equivalence between parity games.

Theorem 3. *Let $\mathcal{G} = \langle V_0, V_1, E, \Omega \rangle$ and $\mathcal{G}' = \langle V'_0, V'_1, E', \Omega' \rangle$ be parity graph games with basic sets B and B' , respectively, and let Π and Π' be (not necessarily distinct) players in \mathcal{G} and \mathcal{G}' , respectively. Whenever v and v' are related through a Π, Π' -bisimulation $Z \subseteq B \times B'$, we have*

$$v \in \text{Win}_\Pi(\mathcal{G}) \text{ iff } v' \in \text{Win}_{\Pi'}(\mathcal{G}'). \quad (2)$$

5 Complementation of Alternating Coalgebra Automata

This section is devoted to our main result:

Theorem 4 (Complementation Lemma for Coalgebra Automata). *The class of alternating coalgebra automata is closed under taking complements.*

As mentioned already, based on the approach by Muller & Schupp [6], we will obtain the complement of an alternating automaton by *dualizing* its transition map, and performing a *role switch* on its priority map. More in detail, consider an alternating automaton $\mathbb{A} = \langle Q, q_I, \theta, \Omega \rangle$, with $\theta : Q \rightarrow \mathcal{L}T^\nabla Q$. Its complement \mathbb{A}^c will be based on the same carrier set Q and will have the same initial state q_I , while the role switch operation on the priority map can be implemented very simply by putting $\Omega^c(q) := 1 + \Omega(q)$. With *dualizing* the transition map we mean that for each state $q \in Q$, we will define $\theta^c(q)$ as the *dual term* of $\theta(q)$ (as given by Definition 7). However, already when dualizing a simple term $\nabla\alpha$ with $\alpha \in TQ$ (rather than $\alpha \in T\mathcal{L}Q$), we see conjunctions popping up *under* the modal operator ∇ . Hence, if we are after a class of automata that is closed under the proposed complementation construction, we need to admit devices with a slightly *richer transition structure*. It follows from the results in section 3 that the set \mathcal{L}_1Q is the smallest set of formulas containing the set $\mathcal{L}T^\nabla Q$ which is closed under the dualization map δ_1 , and hence the ‘richer transition structure’ that we will propose comprises maps of type $\theta : Q \rightarrow \mathcal{L}_1Q$.

We will call the resulting automata *transalternating* to indicate that there is alternation both under and over the modality ∇ . For these *transalternating* automata we devise a simple algorithm to compute complements. Finally we show that the richer transition structure of transalternating automata does not really increase the recognizing power, by providing an explicit construction transforming a transalternating automaton into an equivalent alternating one.

Table 2. Acceptance Games for Transalternating Automata

Position	Player	Admissible Moves	$\Omega_{\mathcal{G}}$
$(q, s) \in Q \times S$	-	$\{(\theta(q), s)\}$	$\Omega(q)$
$(\bigwedge_{i \in I} a_i, s) \in \mathcal{L}_1 Q \times S$	\forall	$\{(a_i, s) \mid i \in I\}$	0
$(\bigvee_{i \in I} a_i, s) \in \mathcal{L}_1 Q \times S$	\exists	$\{(a_i, s) \mid i \in I\}$	0
$(\nabla \alpha, s) \in T^{\forall} \mathcal{L} Q \times S$	\exists	$\{Z \subseteq \mathcal{L} Q \times S \mid (\alpha, \sigma(s)) \in TZ\}$	0
$Z \subseteq \mathcal{L} Q \times S$	\forall	Z	0
$(\bigwedge_{i \in I} a_i, s) \in \mathcal{L} Q \times S$	\forall	$\{(a_i, s) \mid i \in I\}$	0
$(\bigvee_{i \in I} a_i, s) \in \mathcal{L} Q \times S$	\exists	$\{(a_i, s) \mid i \in I\}$	0

5.1 Transalternating Automata

Formally, transalternating automata are defined as follows.

Definition 15 (Transalternating Automata). A transalternating T -automaton $\mathbb{A} = \langle Q, \theta, q_I, \Omega \rangle$ consists of a finite set Q of states, a transition function $\theta : Q \rightarrow \mathcal{L}_1 Q$, a state $q_I \in Q$ distinguished as initial, and a priority function $\Omega : Q \rightarrow \mathbb{N}$.

Just as for alternating automata, acceptance of a pointed coalgebra (\mathbb{S}, s) is defined in terms of a parity game.

Definition 16 (Acceptance Games for Transalternating Automata). Let $\mathbb{A} = \langle Q, q_I, \theta, \Omega \rangle$ be a transalternating automaton and let $\mathbb{S} = \langle S, \sigma \rangle$ be a T -coalgebra. The acceptance game $\mathcal{G}(\mathbb{A}, \mathbb{S})$ is the parity graph game $\mathcal{G}(\mathbb{A}, \mathbb{S}) = \langle V_{\exists}, V_{\forall}, E, v_I, \Omega_{\mathcal{G}} \rangle$ with $v_I = (q_I, s_I)$ and $V_{\exists}, V_{\forall}, E$ and $\Omega_{\mathcal{G}}$ given by Table 2. Positions from $Q \times S$ will be called basic.

In the next subsection we will define the complementation construction for transalternating automata and show the correctness of this transformation. Our proof will be based on a game bisimulation between two acceptance games, and for that purpose we need to understand the powers of the two players in a single round of such an acceptance game. The next Lemma shows that these powers can be conveniently expressed and understood in terms of the one-step semantics defined in section 3. In order to make sense of this remark, recall that the *basic* positions in an acceptance game are of the form $(q, s) \in Q \times S$. Also, note that this terminology is justified because the set $Q \times S$ is indeed basic in the sense of Definition 8. Hence, we may apply the notions defined in section 4, and in particular, given a basic position (q, s) we may consider the question, what it means for a set $Z \subseteq N(q, s)$ to belong to the set $P_{\Pi}(q, s)$ for one of the players $\Pi \in \{\exists, \forall\}$. First note that potential elements of $P_{\Pi}(q, s)$ are sets of basic positions, and as such, binary relations between Q and S . A key observation is that such a binary relation may be identified with a Q -valuation on S .

Definition 17. Given a relation $Z \subseteq Q \times S$, we define the associated Q -valuation on S as the map $V_Z : Q \rightarrow \mathcal{P}(S)$ given by

$$V_Z(q) := \{s \in S \mid (q, s) \in Z\}.$$

Using this perspective on relations between Q and S , we may give the following logical characterization of the players' powers in single rounds of the acceptance game.

Proposition 3. Let \mathbb{A} be a transalternating T -automaton, and \mathbb{S} a T -coalgebra. Given a basic position (q, s) in $\mathcal{G}(\mathbb{A}, \mathbb{S})$, and a relation $Z \subseteq Q \times S$, we have:

$$\begin{aligned} Z \in \uparrow(P_{\exists}(q, s)) &\iff V_Z, \sigma(s) \Vdash_1 \theta(q) \\ Z \in \uparrow(P_{\forall}(q, s)) &\iff V_Z, \sigma(s) \Vdash_1 \delta_1(\theta(q)) \end{aligned}$$

5.2 Complementation of Transalternating Automata

As announced, we define the complement of a transalternating automaton as follows.

Definition 18 (Complements of Transalternating Automata). The complement of a transalternating T -coalgebra automaton $\mathbb{A} = \langle Q, \theta, q_I, \Omega \rangle$ is the transalternating automaton $\mathbb{A}^c = \langle Q, \theta^c, q_I, \Omega^c \rangle$ defined with $\theta^c(q) := \delta_1(\theta(q))$ and $\Omega^c(q) := \Omega(q) + 1$, for all $q \in Q$.

\mathbb{A}^c is indeed the complement of \mathbb{A} .

Proposition 4. For every transalternating T -coalgebra automaton \mathbb{A} , the automaton \mathbb{A}^c accepts precisely those pointed T -coalgebras that are rejected by \mathbb{A} .

Proof. Clearly it suffices to prove, for a given T -coalgebra \mathbb{S} , state q of \mathbb{A} , and point s in \mathbb{S} :

$$(q, s) \in \text{Win}_{\exists}(\mathcal{G}(\mathbb{A}^c, \mathbb{S})) \text{ iff } (q, s) \in \text{Win}_{\forall}(\mathcal{G}(\mathbb{A}, \mathbb{S})).$$

In order to prove this claim we will use a basic game bisimulation.

First we note that $Q \times S$ is a basic set in both acceptance games, $\mathcal{G}(\mathbb{A}, \mathbb{S})$ and $\mathcal{G}(\mathbb{A}^c, \mathbb{S})$. The main observation is that the diagonal relation $\text{Id}_{Q \times S} := \{((q, s), (q, s)) \mid q \in Q, s \in S\}$ is an \forall, \exists -game bisimulation between $\mathcal{G}(\mathbb{A}, \mathbb{S})$ and $\mathcal{G}(\mathbb{A}^c, \mathbb{S})$. Since it is immediate from the definitions that the diagonal relation satisfies the *priority conditions*, it is left to check the structural conditions.

Leaving the other three conditions for the reader, we establish the condition (\exists, forth) as an immediate consequence of the following claim:

$$\text{for all } Z \in P_{\exists}^{\mathcal{G}}(q, s) \text{ there is a } Z_{\forall} \subseteq Z \text{ such that } Z_{\forall} \in P_{\forall}^{\mathcal{G}^c}(q, s), \quad (3)$$

which can be proved via the following chain of implications:

$$\begin{aligned}
& Z \in P_{\exists}^{\mathcal{G}}(q, s) \\
& \Rightarrow V_Z, \sigma(s) \Vdash_1 \theta(q) && \text{(Proposition 3)} \\
& \Rightarrow (V_Z)^c, \sigma(s) \not\Vdash_1 \delta_1(\theta(q)) && \text{(Corollary 1)} \\
& \Rightarrow (V_Z)^c, \sigma(s) \not\Vdash_1 \theta^c(q) && \text{(definition of } \theta^c) \\
& \Rightarrow V_{Z^c}, \sigma(s) \not\Vdash_1 \theta^c(q) && (\dagger) \\
& \Rightarrow Z^c \notin \uparrow(P_{\exists}^{\mathcal{G}^c}(q, s)) && \text{(Proposition 3)} \\
& \Rightarrow \text{there is a } Z_{\forall} \in P_{\forall}^{\mathcal{G}^c}(q, s) \text{ with } Z^c \cap Z_{\forall} = \emptyset && \text{(Proposition 2)} \\
& \Rightarrow \text{there is a } Z_{\forall} \subseteq Z \text{ with } Z_{\forall} \in P_{\forall}^{\mathcal{G}^c}(q, s) && \text{(elementary set theory)}
\end{aligned}$$

Here we let Z^c denote the set $(Q \times S) \setminus Z$, and in the implication marked (\dagger) we use the easily verified fact that $(V_Z)^c = V_{Z^c}$.

5.3 Transalternating and Alternating Automata

In the remainder of this section we prove that the enriched structure of transalternating automata does not add expressivity compared to alternating automata. Since there is an obvious way to see an alternating automaton as a transalternating one, this shows that the transalternating automata have exactly the same recognizing power as the alternating ones. Consequently, the complementation lemma for transalternating automata also yields a complementation lemma for alternating ones.

The transformation of a given transalternating automaton \mathbb{A} into an equivalent alternating automaton will be taken care of in two successive steps. The intuitive idea is to first remove the disjunctions, and then the conjunctions, from under the modal operators. In the intermediate stage we are dealing with automata that only allow (certain) conjunctions under the nabla operator.

Definition 19. *Given a set Q we let $\mathcal{S}Q$ denote the set of conjunctions of the form $\bigwedge Q'$, with $Q' \subseteq Q$. We let $\leq \subseteq \mathcal{S}Q \times Q$ denote the relation given by $\bigwedge Q' \leq q$ iff $q \in Q'$.*

Definition 20. *A semi-transalternating T -automaton is a automaton of the form $\mathbb{A} = \langle Q, \theta, q_I, \Omega \rangle$ with $\theta : Q \rightarrow \mathcal{L}T^{\nabla}\mathcal{S}Q$.*

We omit the (obvious) definition of the acceptance game associated with this type of automaton.

Proposition 5. *There is an effective algorithm transforming a given transalternating automaton into an equivalent semi-transalternating one based on the same carrier.*

The *proof* of this proposition, which we omit for reasons of space limitations, is based on the fact that every formula in \mathcal{L}_1Q can be rewritten into an equivalent formula in $\mathcal{L}T^{\nabla}\mathcal{S}Q$, see [3] for the details.

It is left to prove that every semi-transalternating automaton can be transformed into an equivalent alternating one. Intuitively speaking, one has to carry over the (universal) non-determinism after the ∇ -translation to the branching for the next automaton state.

Definition 21. Let $\mathbb{A} = \langle Q, q_I, \theta, \Omega \rangle$ be a semi-transalternating automaton. We define the alternating automaton $\mathbb{A}^\circ = \langle Q', q'_I, \theta', \Omega' \rangle$ by putting $Q' := Q \times \mathcal{S}Q$, $q'_I := (q_I, q_I)$, $\Omega'(a, b) := \Omega(a)$, while $\theta' : (Q \times \mathcal{S}Q) \rightarrow \mathcal{L}T^\nabla(Q \times \mathcal{S}Q)$ is given by

$$\theta'(q, a) := \bigwedge_{p \geq a} \mathcal{L}T^\nabla \iota_p(\theta(p)).$$

Here $\iota_q : \mathcal{S}Q \rightarrow (Q \times \mathcal{S}Q)$, for $q \in Q$, is the map given by $\iota_q(a) := (q, a)$.

For reasons of space limitations we omit the proof of the equivalence of \mathbb{A}° and \mathbb{A} , confining ourselves to an intuitive explanation of this definition. Note that in principle we would like to base \mathbb{A}° on the carrier $\mathcal{S}Q$, defining a transition map $\theta''(a) := \bigwedge_{p \geq a} \theta(p)$. (Note that this gives a well-defined map $\theta'' : \mathcal{S}Q \rightarrow \mathcal{L}T^\nabla \mathcal{S}Q$.) Unfortunately, this set-up is too simple to give a proper account of the acceptance condition. The problem is that in the acceptance game, when moving from a state $(a, s) \in \mathcal{S}Q \times Q$ to $(\theta''(a), s)$ and then on to $(\theta(p), s)$ for some $p \geq a$, one would ‘bypass’ the position (p, s) , and thus miss the crucial contribution of the priority of the state p to the information determining the winner of the match. A solution to this problem is to tag p to the new states of \mathbb{A}° occurring in $\theta(p)$. This is exactly the role of the function ι_p in the definition of θ' above: note that $\mathcal{L}T^\nabla \iota_p(\theta(p))$ is nothing but the term $\theta(p)$, seen as an element of $\mathcal{L}T^\nabla \mathcal{S}Q$, but with every formula $b \in \mathcal{S}Q$ under the ∇ replaced by $(p, b) \in Q \times \mathcal{S}Q$.

Thus, intuitively, a \mathbb{A}° -state of the form (p, b) represents the conjunction of the states above b (where q is above $b = \bigwedge Q'$ iff it belongs to Q'). However, the *priority* of the state (p, b) is that of p — thus the state (p, b) encodes (‘remembers’) an earlier visit to p .

Proposition 6. \mathbb{A} and \mathbb{A}° are equivalent, for any semi-transalternating automaton \mathbb{A} .

The latter two propositions establish that there is an effective procedure transforming a transalternating automaton into an equivalent alternating one.

5.4 Size Matters

We conclude this section with some remarks about the size of the automaton resulting from the complementation algorithm presented. By the *size* of an automaton we understand the number of its states.

First of all, since every alternating T -automaton can itself be seen as a transalternating T -automaton, there is no size issue here. Also, our complement of a transalternating T -automaton is of the same size as the original. It is

the translation from transalternating to alternating T -automata that introduces exponentially many new states: more specifically, with \mathbb{A} a transalternating automaton of size $|Q| = n$, the equivalent alternating automaton of Definition 21 has $|\mathcal{S}Q \times Q| = 2^n * n$ states. From these observations the following is immediate.

Theorem 5. *For any alternating T -automaton with n states there is a complementing alternating automaton with at most $n * 2^n$ states.*

In case we are dealing with a functor T such that the set $\mathcal{L}T^\nabla Q$ is closed under taking duals, we do not need the concept of transalternation, and we can obtain a much better upper bound. The following result applies for instance to alternating tree automata, see [4] for the details.

Theorem 6. *Let T be such that for any $\alpha \in T_\omega Q$, the formula $\nabla\alpha$ has a dual $\Delta\alpha \in \mathcal{L}T^\nabla Q$. Then for any alternating T -automaton of n states there is a complementing alternating automaton with at most $n + c$ states, for some constant c .*

Acknowledgement. We are grateful to Alessandra Palmigiano and one of the anonymous reviewers for suggesting a simplification of a definition, and to the anonymous reviewers in general for their comments.

References

1. Baltag, A.: A logic for coalgebraic simulation. In: Reichel, H. (ed.) *Coalgebraic Methods in Computer Science (CMCS 2000)*. Electronic Notes in Theoretical Computer Science, vol. 33, pp. 41–60 (2000)
2. van Benthem, J.: Extensive games as process models. *Journal of Logic, Language and Information* 11(3), 289–313 (2002)
3. Kurz, A., Kupke, C., Venema, Y.: Completeness of the finitary Moss logic. In: Areces, C., Goldblatt, R. (eds.) *Advances in Modal Logic*, vol. 7, pp. 193–217. College Publications (2008)
4. Kissig, C.: Satisfiability of S2S. Master’s thesis, Institute of Logic, Language and Computation, Amsterdam (2007)
5. Kupke, C., Venema, Y.: Coalgebraic automata theory: basic results. *Logical Methods in Computer Science* 4, 1–43 (2008)
6. Muller, D., Schupp, P.: Alternating automata on infinite trees. *Theoretical Computer Science* 54, 267–276 (1987)
7. Pauly, M.: Logic for Social Software. PhD thesis, Institute of Logic, Language and Computation, Amsterdam (2001)
8. Rutten, J.: Universal coalgebra: A theory of systems. *Theoretical Computer Science* 249, 3–80 (2000)
9. Venema, Y.: Automata and fixed point logics for coalgebras. *Electronic Notes in Theoretical Computer Science* 106, 355–375 (2004)
10. Venema, Y.: Automata and fixed point logic: a coalgebraic perspective. *Information and Computation* 204, 637–678 (2006)