Weak MSO: Automata and Expressiveness Modulo Bisimilarity

Facundo Carreiro

U. Amsterdam

fcarreiro@dc.uba.ar

Alessandro Facchini

U. Warsaw

facchini@mimuw.edu.pl

Yde Venema

U. Amsterdam

y.venema@uva.nl

Fabio Zanasi

ENS Lyon, U. Lyon, CNRS, INRIA, UCBL

fabio.zanasi@ens-lyon.fr

Abstract

We prove that the bisimulation-invariant fragment of weak monadic second-order logic (WMSO) is equivalent to the fragment of the modal μ-calculus where the application of the least fixpoint operator $\mu p.\varphi$ is restricted to formulas φ that are continuous in p. Our proof is automata-theoretic in nature; in particular, we introduce a class of automata characterizing the expressive power of WMSO over tree models of arbitrary branching degree. The transition map of these automata is defined in terms of a logic FOE_1^∞ that is the extension of first-order logic with a generalized quantifier $\exists^\infty$, where $\exists^\infty x.\varphi$ means that there are infinitely many objects satisfying φ. An important part of our work consists of a model-theoretic analysis of FOE_1^∞.

Categories and Subject Descriptors   F.4.1 [*Mathematical Logic and Formal Languages*]: Mathematical Logic; F.1.1 [*Computation by Abstract Devices*]: Models of Computation—automata

General Terms   Theory

Keywords   weak monadic second-order, characterisation theorem, automata theory, modal mu-calculus, MSO-automata, WMSO, bisimulation invariance, Janin-Walukiewicz theorem

1. Introduction

1.1 Expressiveness modulo bisimilarity

This paper concerns the relative expressive power of some languages used for describing properties of pointed labelled transitions systems (LTSs), or Kripke models. The interest in such expressiveness questions stems from applications where these structures model computational processes, and bisimilar pointed structures represent the *same* process. Seen from this perspective, properties of transition structures are relevant only if they are invariant under bisimilarity. This explains the importance of bisimulation invariance results of the form

$$M \equiv L/\!\!\leftrightarrow \quad \text{(over } K\text{)}$$

stating that, if one restricts attention to a certain class K of transition structures, one language M is expressively complete with respect to the relevant (i.e., bisimulation-invariant) properties that can be formulated in another language L. In this setting, generally

L is some rich yardstick formalism such as first-order or monadic second-order logic, and M is some modal-style fragment of L, usually displaying much better computational behavior than the full language L.

A seminal result in the theory of modal logic is van Benthem's Characterization Theorem [15], stating that every bisimulation-invariant first-order formula $\alpha(x)$ is actually equivalent to (the standard translation of) a modal formula, that is:

$$\mathrm{ML} \equiv \mathrm{FO}/\!\!\leftrightarrow \quad \text{(over the class of all LTSs)}.$$

Over the years, a wealth of variants of the Characterization Theorem have been obtained. For instance, Rosen proved that van Benthem's theorem is one of the few preservation results that transfers to the setting of finite models [12]. A recent, rich source of van Benthem-style characterization results is given by Dawar & Otto [3]. Our main point of reference for this article is the work of Janin & Walukiewicz [8], who extended van Benthem's result to the setting of fixpoint logics, by proving that the modal μ-calculus ($\mu\mathrm{ML}$) is the bisimulation-invariant fragment of monadic second-order logic (MSO):

$$\mu\mathrm{ML} \equiv \mathrm{MSO}/\!\!\leftrightarrow \quad \text{(over the class of all LTSs)}.$$

1.2 Bisimulation invariance for WMSO

The yardstick logic that we consider in this paper is *weak* monadic second-order logic (WMSO), a variant of monadic second-order logic where the second-order quantifiers range over *finite* subsets of the transition structure rather than over arbitrary ones. Our target will be to identify the bisimulation-invariant fragment of WMSO.

In the case of finitely branching models, it is not very hard to show that WMSO is a (proper) fragment of MSO, and it seems to be folklore that WMSO/$\leftrightarrow$ corresponds to AFMC, the alternation-free fragment of the modal μ-calculus. For binary trees, this result was proved by Arnold & Niwiński [1]. Over structures of arbitrary branching degree, however, WMSO and MSO have *incomparable* expressive power [13, 18]. For this reason, the relative expressive power of WMSO/$\leftrightarrow$ and $\mu\mathrm{ML} \equiv \mathrm{MSO}/\!\!\leftrightarrow$ is not a priori clear. In any case, we have that WMSO/$\leftrightarrow \not\equiv$ AFMC: the class of well-founded trees, definable by the simple AFMC-formula $\mu p.\Box p$, is not definable in WMSO, see e.g. [13]. Incidentally, three of the present authors proved that:

$$\mathrm{AFMC} \equiv \mathrm{WFMSO}/\!\!\leftrightarrow \quad \text{(over the class of all LTSs)},$$

for yet another variant WFMSO of MSO [5].

The main result that we shall prove in this paper states that

$$\mu_c\mathrm{ML} \equiv \mathrm{WMSO}/\!\!\leftrightarrow \quad \text{(over the class of all LTSs)}, \qquad (1)$$

where $\mu_c\mathrm{ML}$ is a certain fragment of AFMC, characterized by a restriction on the application of fixpoint operators which involves the notion of *(Scott) continuity*.

Continuity, an interesting property that features naturally in the semantics of many (fixpoint) logics, plays a key role in our work. For its definition, we consider how the meaning $[\![\varphi]\!]^{\mathbb{T}} \subseteq T$ of a

formula φ in a structure $\mathbb{T}$ (with domain T) depends on the meaning of a fixed proposition letter or monadic predicate symbol p. This dependence can be formalized as a map $\varphi_p^{\mathbb{T}} : \wp(T) \to \wp(T)$, and if $\varphi_p^{\mathbb{T}}$ satisfies the condition

$$\varphi_p^{\mathbb{T}}(X) = \bigcup \left\{ \varphi_p^{\mathbb{T}}(X') \mid X' \text{ is a finite subset of } X \right\}, \quad (2)$$

we say that φ is *continuous in* p. The topological terminology stems from the observation that (2) expresses the continuity of $\varphi_p^{\mathbb{T}}$ with respect to the Scott topology on $\wp(T)$. It is easy to see that continuous maps are constructive, meaning that the least fixpoint of a continuous map F can be obtained as the union of the finite approximations $\emptyset, F\emptyset, F^2\emptyset, \ldots$

As we will see, the link with WMSO lies in the fact that if such a map is given by a modal-like formula, then any point in a model belongs to its least fixpoint iff it belongs to a *finite* prefixpoint.

If we look at concrete cases, this definition can be given a different reading: if φ is a formula of the modal μ-calculus, (2) means that φ holds at some state s of $\mathbb{T}$ iff we can shrink the interpretation of the proposition letter p to some finite subset of the original interpretation, in such a way that φ holds at s in the modified version of $\mathbb{T}$.

A syntactic *characterization* of continuity for μML was obtained by Fontaine [6], who proved that a μML-formula φ is continuous in a proposition letter p iff φ is equivalent to a formula in the fragment $\mu\text{MLC}_{\{p\}}$ of μML given as follows.

Definition 1. *Let* $\mathsf{Q} \subseteq \mathsf{P}$ *be sets of proposition letters. The fragment* $\mu\text{MLC}_{\mathsf{Q}}$ *of formulas continuous in* Q *is given by:*

$$\varphi ::= p \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Diamond\varphi \mid \mu q.\alpha$$

where $p \in \mathsf{Q}$, $q \in \mathsf{P}$, $\psi \in \mu\text{ML}$ *is* Q-*free,* $\alpha \in \mu\text{MLC}_{\mathsf{Q}\cup\{q\}}$.

The definition of our fragment $\mu_c\text{ML}$ uses this characterization as follows: whereas in the full language of μML the only syntactic condition on the formation of a formula $\mu p.\varphi$ is that φ is *positive* in p, for the fragment $\mu_c\text{ML}$ this condition is strengthened to the requirement that φ is (syntactically) *continuous* in p. More precisely, $\mu_c\text{ML}$ is defined as follows:

Definition 2. *The formulas of the language* $\mu_c\text{ML}$ *are given by the following induction:*

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \Diamond\varphi \mid \mu q.\alpha$$

where $p, q \in \mathsf{P}$, *and* $\alpha \in \mu\text{MLC}_{\{q\}} \cap \mu_c\text{ML}$.

In fact we will prove, analogous to the result by Janin & Walukiewicz, the following strong version of the characterization result (1).

Theorem 1. *There are effective translations* $(\cdot)^{\bullet} : \text{WMSO} \to \mu_c\text{ML}$ *and* $(\cdot)_{\bullet} : \mu_c\text{ML} \to \text{WMSO}$ *such that*

(i) $\varphi \in \text{WMSO}$ *is bisimulation invariant iff* $\varphi \equiv \varphi^{\bullet}$, *and*
(ii) $\psi \equiv \psi_{\bullet}$ *for every formula* $\psi \in \mu_c\text{ML}$.

To see how this theorem implies (1) observe that part (i) shows that $\text{WMSO}/{\leftrightarrow} \leq \mu_c\text{ML}$. Part (ii) states that $\mu_c\text{ML} \leq \text{WMSO}$. Combined with the fact that every formula in $\mu_c\text{ML} \subseteq \mu\text{ML}$ is bisimulation invariant, this gives $\mu_c\text{ML} \leq \text{WMSO}/{\leftrightarrow}$.

1.3 Automata

As usual in this research area, our proof will be automata-theoretic in nature. In particular, as the second main contribution of this paper, we introduce a new class of parity automata that exactly captures the expressive power of WMSO over the class of tree models of arbitrary branching degree. We introduce a similar class of automata for the language $\mu_c\text{ML}$, and as in [8], the key argument

in our proof of Theorem 1 revolves around a comparison of these two kinds of automata.

Before we turn to a description of these automata, we have a look at the ones introduced by Walukiewicz [18], corresponding to MSO (over tree models). Fix a set P of proposition letters and think of $\wp\mathsf{P}$ as a set of colors or labels. An MSO-automaton is a tuple $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$, where A is a finite set of states, a_I an initial state, and $\Omega : A \to \mathbb{N}$ a parity function. The transition function $\Delta : A \times \wp\mathsf{P} \to \text{FOE}_1^+(A)$ takes as its codomain the set of positive sentences in the first-order language (with equality) over a signature of which the states in A provide the (monadic) predicates. This language plays a key role in the acceptance game that we associate with an MSO-automaton $\mathbb{A}$ and an LTS $\mathbb{T}$: essentially, at each round of the game, the two players focus on one specific sentence in $\text{FOE}_1^+(A)$. For this reason, we shall refer to FOE_1 as the *one-step language* of MSO-automata, and denote this class by $Aut(\text{FOE}_1)$. Walukiewicz' result states that

$$\text{MSO} \equiv Aut(\text{FOE}_1) \quad \text{(over the class of all trees).} \quad (3)$$

Similarly, Janin & Walukiewicz show that over the class of all models, μML can be captured by the class $Aut(\text{FO}_1)$ of automata that take first-order logic without equality as their one-step language. In order to adapt this approach to the setting of WMSO, observe that by König's lemma a subset of a tree $\mathbb{T}$ is finite iff it is both a subset of a finitely branching subtree of $\mathbb{T}$ and *noetherian*, that is, a subset of a subtree of $\mathbb{T}$ that has no infinite branches. This suggests two kinds of modifications to MSO-automata, roughly speaking corresponding to a horizontal and a vertical 'dimension' of trees.

For the 'vertical modification' we turn to weak automata [11]. The acceptance condition Ω of a parity automaton $\mathbb{A}$ is *weak* if $\Omega(a) = \Omega(a')$ whenever the states a and a' belong to the same SCC (strongly connected component, Definition 4) of the automaton. Let $Aut_w(\text{FOE}_1)$ denote the set of MSO-automata with a weak parity condition. It was proved in [5, 19] that

$$\text{WFMSO} \equiv Aut_w(\text{FOE}_1) \quad \text{(over the class of all trees),}$$

with WFMSO denoting the earlier mentioned variant of MSO where second-order quantification is restricted to noetherian subsets of trees. From this it easily follows that

$$\text{WMSO} \equiv Aut_w(\text{FOE}_1) \quad \text{(over finitely branching trees).}$$

Over the class of all trees, however, this equivalence does not hold, as is witnessed by the earlier mentioned class of well-founded trees, which can be defined in $\text{AFMC} \leq \text{WFMSO}$, but not in WMSO.

The hurdle to take, in order to shape automata for WMSO on trees of *arbitrary* branching degree, concerns the horizontal dimension; the main problem lies in finding the right one-step language for these automata. An obvious candidate would be WMSO itself, or more precisely, its variant WMSO_1 over the signature of monadic predicates (corresponding to the automata states). A very helpful observation by Väänänen [14] implies that WMSO_1 is equivalent to the logic FOE_1^∞ we obtain by extending FOE_1 with the generalized quantifier $\exists^\infty$, the intended meaning of $\exists^\infty x.\varphi$ being that there are *infinitely* many objects satisfying φ.

As it turns out, taking the *full* language FOE_1^∞ as our one-step language would give too much expressive power, it is here that we will crucially involve the notion of *continuity*.

The automata corresponding to WMSO will be of the form $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$, where $\Delta : A \times \wp\mathsf{P} \to \text{FOE}_1^{\infty+}(A)$ is subject to the following two constraints, for all $a, a' \in A$ belonging to the same SCC:

(weakness) $\Omega(a) = \Omega(a')$, and

(continuity) if $\Omega(a)$ is odd/even, then for each label $c \in \wp\mathsf{P}$, $\Delta(a, c)$ is continuous/co-continuous in a'.

Here co-continuity is a dual notion to continuity.

A proper definition of these WMSO-automata requires a *syntactic* characterization of the $\mathrm{FOE}_1^\infty(A)$-sentences that are (co-)continuous in one (or more) monadic predicates of A. For this purpose, we will conduct a fairly detailed model-theoretic study of the logic FOE_1^∞, which we consider to be the third main contribution of our work. Similar to the results for FOE_1, we provide normal forms for the sentences of $\mathrm{FOE}_1^\infty(A)$, and syntactic characterizations of the fragments whose sentences are monotone (respectively continuous) in some monadic predicate $a \in A$.

Based on this analysis, we give a precise definition of the class $Aut_{cw}(\mathrm{FOE}_1^\infty)$ of WMSO-automata in Definition 12. Our second main theorem states the following.

Theorem 2. *There are effective transformations from* WMSO-*formulas to* WMSO-*automata and vice-versa, witnessing*

$$\mathrm{WMSO} \equiv Aut_{cw}(\mathrm{FOE}_1^\infty) \quad \textit{(on tree models)}. \qquad (4)$$

Turning to the proof of Theorem 1, we provide an analogous characterization result for $\mu_c\mathrm{ML}$, based on the class $Aut_{cw}(\mathrm{FO}_1)$ (see Definition 24) of those automata in $Aut(\mathrm{FO}_1)$ that satisfy similar weakness and continuity conditions as $Aut_{cw}(\mathrm{FOE}_1^\infty)$:

$$\mu_c\mathrm{ML} \equiv Aut_{cw}(\mathrm{FO}_1) \quad \text{(on all LTSs)}. \qquad (5)$$

Following a similar approach as Janin & Walukiewicz [8], our proof will then revolve around a map from $Aut_{cw}(\mathrm{FOE}_1^\infty)$ to $Aut_{cw}(\mathrm{FO}_1)$, guided by the additional insight from [16] that such a link can already be determined by a translation on the level of the one-step languages FOE_1^∞ and FO_1.

1.4 Overview of the paper

In Section 2 we give a precise definition of the preliminaries required to read this article. In Section 3 we analyze several one-step logics and give normal forms and syntactic characterizations of their monotone and continuous fragments. In Section 4 we formally define WMSO-automata and we prove Theorem 2. In Section 5 we finally prove Theorem 1.

1.5 Extended version

Due to space limitations, in this paper we confine ourselves to brief discussions and proof sketches. Readers that are interested in the details may consult the extended version [2].

2. Preliminaries

2.1 Transition Systems and Trees

Throughout this article we fix a set P of *proposition letters* (or monadic predicate symbols) and call $C := \wp(\mathsf{P})$ its set of *labels*. Given $R \subseteq X \times Y$ and $x \in X$, we denote the successors of x as $R[x] := \{y \mid (x, y) \in R\}$ and $\mathsf{Ran}(R) := \bigcup_{x \in X} R[x]$. We write $\subseteq_\omega$ for the finite subset relation and $\vec{s}$ for a sequence $(s_1, \ldots, s_n)$.

A *C-labeled transition system* (LTS, for short) is a tuple $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$ where T is the domain of $\mathbb{T}$, $\sigma : T \to \wp(\mathsf{P})$ is a marking, $R \subseteq T^2$ is the accessibility relation and $s_I \in T$ is a distinguished node. A *C-tree* is an LTS in which every node can be reached from s_I (called the *root*) and every node except s_I has a unique predecessor. We use the term *tree language* as a synonym for a class of C-trees.

The ω-*unravelling* $\mathbb{T}^\omega$ of an LTS $\mathbb{T}$ is defined as the LTS $\mathbb{T}^\omega := \langle \widehat{T}, \widehat{R}, \widehat{\sigma}, s_I \rangle$ where $\widehat{T}$ is the set of finite sequences $s_0(k_1, s_1)(k_2, s_2) \cdots (k_n, s_n)$ such that $s_0 = s_I$, and $k_i \in \omega$, $s_i \in T$ and $R(s_{i-1}, s_i)$ for each $i \geq 1$; $\widehat{R}(t, t')$ iff t' extends t with a single pair (k_{n+1}, s_{n+1}); and $\widehat{\sigma}$ labels a sequence $t \in \widehat{T}$ with the color of its last node in T.

Given a proposition letter $p \notin P$, a *p-variant* of a $\wp(\mathsf{P})$-transition system $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$ is a $\wp(\mathsf{P} \cup \{p\})$-transition system $\langle T, R, \sigma', s_I \rangle$ such that $\sigma'(s) \setminus \{p\} = \sigma(s)$ for all $s \in T$. Given a set $S \subseteq T$, we let $\mathbb{T}[p \mapsto S]$ denote the p-variant where $p \in \sigma'(s)$ iff $s \in S$. A p-variant $\mathbb{T}[p \mapsto S]$ called is *finitary* when S is a finite set.

Given a formula φ of some logic $\mathcal{L}$, $[\![\varphi]\!]$ denotes the class of LTSs that make φ true. A class K of LTSs is $\mathcal{L}$-*definable* if $[\![\varphi]\!] = K$ for some $\varphi \in \mathcal{L}$. The notation $\varphi \equiv \psi$ means that $[\![\varphi]\!] = [\![\psi]\!]$ and given two logics $\mathcal{L}, \mathcal{L}'$ we use $\mathcal{L} \equiv \mathcal{L}'$ when the $\mathcal{L}$-definable and $\mathcal{L}'$-definable classes of models coincide.

Convention. Throughout this paper, we will only consider LTSs $\mathbb{T}$ in which $R[s]$ is non-empty for every node $s \in T$. All our results, however, can be lifted to the general case.

2.2 Parity Games

A match of a parity game $\mathcal{G}$ consists of two players, $\exists$ and $\forall$, moving a token from one position to another over a partitioned board $G = G_\exists \cup G_\forall$. For every position, players have a set of available moves. If during a match a player reaches a position with no admissible moves, he loses the match. If the match goes on forever then a parity map $\Omega : G \to \mathbb{N}$ is used to choose a winner. The winner is $\exists$ if the *minimum* parity which occurs infinitely often in the match is even, otherwise $\forall$ wins.

A *strategy* for player $\Pi \in \{\exists, \forall\}$ is, intuitively, a specification of choices to be made in the positions belonging to Π. Strategies for parity games can be taken to be *positional* or memory-free (see e.g. [4]) and therefore can be represented as a function $f_\Pi : G_\Pi \to G$. A match is f_Π-*guided* if for each position $u \in G_\Pi$ player Π chooses $f_\Pi(u)$ as next position.

We say that f is a *winning strategy* for Π if (i) for each f-guided match, the moves suggested by f are always available to Π and (ii) Π wins each f-guided match of the game. A *winning position* is one from which Π has a winning strategy.

2.3 Monadic Second-Order Logics

We present *monadic second-order logic* (MSO) and its variant *weak monadic second-order logic* (WMSO) in a form which is best suited to work with in the context of automata. Those two logics share the same syntax, defined on P by:

$$\varphi \quad ::= \quad p \sqsubseteq q \mid \Downarrow p \mid R(p, q) \mid \neg\varphi \mid \varphi \vee \varphi \mid \exists p.\varphi,$$

where p and q are letters from P. We adopt the standard convention that no letter is both free and bound in φ.

MSO and WMSO are distinguished by their semantics. Given an LTS $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$, the interpretation of the atomic formulas and the boolean connectives is standard:

$$
\begin{aligned}
\mathbb{T} &\models p \sqsubseteq q && \text{iff} && \forall s \in T.p \in \sigma(s) \Rightarrow q \in \sigma(s) \\
\mathbb{T} &\models R(p, q) && \text{iff} && \forall s \in T.p \in \sigma(s) \Rightarrow \exists t \in R[s].q \in \sigma(t) \\
\mathbb{T} &\models \Downarrow p && \text{iff} && \forall s \in T.p \in \sigma(s) \Rightarrow s = s_I.
\end{aligned}
$$

The formula $\mathbb{T} \models \exists p.\varphi$ is true if and only if

- **MSO**: $\mathbb{T}[p \mapsto S] \models \varphi$ for some $S \subseteq T$
- **WMSO**: $\mathbb{T}[p \mapsto S] \models \varphi$ for some finite $S \subseteq_\omega T$.

Equivalently, MSO and WMSO can be given in a more standard two-sorted syntax generated by

$$\varphi ::= p(x) \mid R(x, y) \mid x \approx y \mid \neg\varphi \mid \varphi \vee \varphi \mid \exists x.\varphi \mid \exists p.\varphi$$

where $p \in \mathsf{P}$ and $x, y \in \mathsf{iVar}$ (individual variables).

2.4 Parity Automata and their One-Step Languages

We recall the definition of a parity automaton, adapted to our setting. Since we will be comparing parity automata defined in

terms of various one-step languages, it makes sense to make the following abstraction.

Definition 3. *Given a set A of elements called* monadic predicates, *an A-structure is a pair (D, V) with a domain D and a valuation $V : A \to \wp D$. A one-step language is a map $\mathcal{L}_1$ assigning to A a set $\mathcal{L}_1(A)$ of* one-step formulas *over A. One-step languages are interpreted over A-structures, a formula $\varphi \in \mathcal{L}_1$ being either* true *or* false *in (D, V).*

We call $\varphi \in \mathcal{L}_1(A)$ monotone in $a \in A$ if $(D, V) \models \varphi$ and $V(a) \subseteq E$ implies $(D, V[a \mapsto E]) \models \varphi$. We assume that every $\mathcal{L}_1$ has a positive fragment $\mathcal{L}_1^+$ *characterizing monotonicity in the sense that a formula $\varphi \in \mathcal{L}_1(A)$ is monotone in all $a \in A$ iff it is equivalent to some $\varphi' \in \mathcal{L}_1^+(A)$.*

Definition 4. *Let $\mathcal{L}_1$ be some one-step language. A* parity automaton *based on $\mathcal{L}_1$ and alphabet C is a tuple $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ where A is a finite set of states (called* carrier*), $a_I \in A$ is the initial state, $\Delta : A \times C \to \mathcal{L}_1^+(A)$ is the transition map, and $\Omega : A \to \mathbb{N}$ is the parity map. The collection of such automata will be denoted by $Aut(\mathcal{L}_1)$.*

Given states $a, b \in A$, let $a \rightsquigarrow b$ mean that b occurs in $\Delta(a, c) \in \mathcal{L}_1^+(A)$ for some $c \in C$ and $a \preceq b$ mean that (a, b) is in the reflexive-transitive closure of the relation $\rightsquigarrow$. A strongly connected component *(SCC) is a set $B \subseteq A$ of states where $a \preceq b$ and $b \preceq a$ for all $a, b \in B$.*

Acceptance and rejection of an LTS by an automaton is defined in terms of the following parity game.

Definition 5. *Given $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ in $Aut(\mathcal{L}_1)$ and an LTS $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$, the* acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{T})$ *of $\mathbb{A}$ on $\mathbb{T}$ is the parity game defined according to the rules of Table 1. An LTS $\mathbb{T}$ is* accepted *by $\mathbb{A}$ iff (a_I, s_I) is a winning position for $\exists$ in $\mathcal{A}(\mathbb{A}, \mathbb{T})$.*

We use $\mathcal{T}(\mathbb{A})$ to denote the class of *trees* accepted by $\mathbb{A}$. When considering automata for WMSO we will be primarily interested in the following game-theoretical properties.

Definition 6. *Given $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ in $Aut(\mathcal{L}_1)$ and an LTS $\mathbb{T}$, a strategy f for $\exists$ in $\mathcal{A}(\mathbb{A}, \mathbb{T})$ is* functional *in $B \subseteq A$ if for each node s in $\mathbb{T}$ there is at most one $b \in B$ such that (b, s) is a reachable position in an f-guided match. Also f is* finitary *in B if there are only finitely many nodes s in $\mathbb{T}$ for which a position (b, s) with $b \in B$ is reachable in an f-guided match.*

Many properties of parity automata are determined at the one-step level. An important example concerns the notion of complementation.

Definition 7. *Two one-step formulas φ and ψ are each other's* Boolean dual *if for every structure (D, V) we have*

$$(D, V) \models \varphi \text{ iff } (D, V^c) \not\models \psi,$$

where $V^c(a) := D \setminus V(a)$, for all a. A one-step language $\mathcal{L}_1$ is closed under Boolean duals *if for every set A, each formula $\varphi \in \mathcal{L}_1(A)$ has a Boolean dual $\varphi^\delta \in \mathcal{L}_1(A)$.*

Following ideas from [9, 10], we can use Boolean duals, together with a *role switch* between $\forall$ and $\exists$, in order to define a negation or complementation operation on automata.

Definition 8. *Let $\mathcal{L}_1$ be closed under Boolean duals. Given an automaton $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ in $Aut(\mathcal{L}_1)$, we define its* complement *as $\overline{\mathbb{A}} := \langle A, \Delta^\delta, \Omega^\delta, a_I \rangle$, where $\Delta^\delta(a, c) := (\Delta(a, c))^\delta$, and $\Omega^\delta(a) := 1 + \Omega(a)$, for all $a \in A$ and $c \in C$.*

Proposition 1. *For each automaton $\mathbb{A} \in Aut(\mathcal{L}_1)$ and transition system $\mathbb{T}$ we have that $\mathbb{A}$ accepts $\mathbb{T}$ iff $\overline{\mathbb{A}}$ rejects $\mathbb{T}$.*

The proof of Proposition 1 is based on the fact that the power of $\exists$ in $\mathcal{A}(\overline{\mathbb{A}}, \mathbb{T})$ is the same as that of $\forall$ in $\mathcal{A}(\mathbb{A}, \mathbb{T})$.

2.5 First-Order Logic with Infinity Quantifiers

As observed in the introduction, when defining the one-step language for WMSO-automata it will be convenient to work with an extension of first-order logic with the generalized quantifier $\exists^\infty$. Formally, its semantics is defined as follows:

$$\mathbb{T} \models \exists^\infty x. \varphi(x) \text{ iff } \{s \mid \mathbb{T} \models \varphi(s)\} \text{ is infinite.}$$

The dual of $\exists^\infty$ is $\forall^\infty$ and the intended meaning of $\forall^\infty x. \varphi$ is that there are *at most finitely many* elements *falsifying* φ.

Definition 9. *The set $\mathrm{FOE}_1(A)$ of one-step first-order sentences (with equality) is given by the sentences formed by:*

$$\varphi ::= a(x) \mid x \approx y \mid \neg\varphi \mid \varphi \vee \varphi \mid \exists x. \varphi$$

where $x, y \in \mathsf{iVar}$, $a \in A$. The one-step logic $\mathrm{FO}_1(A)$ is as $\mathrm{FOE}_1(A)$ but without equality. The set $\mathrm{FOE}_1^\infty(A)$ of one-step first-order sentences with $\exists^\infty$ and equality is defined analogously by just adding the clause $\exists^\infty x. \varphi$.

2.6 The Modal μ-Calculus.

The language of the modal μ-calculus (μML) on the set of propositions P is given by the following grammar:

$$\varphi ::= q \mid \varphi \vee \varphi \mid \neg\varphi \mid \Diamond\varphi \mid \mu p. \varphi$$

where $p, q \in \mathsf{P}$, and φ is positive in p. We refer to the introduction for the definition of the fragments $\mu\mathrm{MLC}_\mathsf{Q}$ (depending on a set $\mathsf{Q} \subseteq \mathsf{P}$) and $\mu_c\mathrm{ML}$ of $\mu\mathrm{ML}$.

Let $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$ be an LTS and $\varphi \in \mu\mathrm{ML}$. We inductively define the *meaning* $[\![\varphi]\!]^{\mathbb{T}}$ which includes the following clauses for the least (μ) fixpoint operator:

$$[\![\mu p. \psi]\!]^{\mathbb{T}} := \bigcap \{S \subseteq T \mid S \supseteq [\![\psi]\!]^{\mathbb{T}[p \mapsto S]}\}.$$

We say that φ is *true* in $\mathbb{T}$ (notation $\mathbb{T} \Vdash \varphi$) iff $s_I \in [\![\varphi]\!]^{\mathbb{T}}$.

Continuity for $\mu\mathrm{ML}$-formulas boils down to the following.

Proposition 2. *A formula $\varphi \in \mu\mathrm{ML}$ is continuous in $p \in \mathsf{P}$ if and only if for every LTS $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$ there exists some finite $S \subseteq_\omega \{s \in T \mid p \in \sigma(s)\}$ such that $\mathbb{T} \Vdash \varphi$ iff $\mathbb{T}[p \mapsto S] \Vdash \varphi$.*

Continuity can be syntactically characterized as follows.

Fact 1 ([6])**.** *A $\mu\mathrm{ML}$-formula is continuous in p iff it is equivalent to a formula in the fragment $\mu\mathrm{MLC}_{\{p\}}$.*

The next property is easily verified.

Proposition 3. *For each $\mu p. \varphi \in \mu_c\mathrm{ML}$, φ is continuous in p.*

Finally, we remark that $\mu_c\mathrm{ML}$ is strictly included in the alternation free-fragment AFMC of $\mu\mathrm{ML}$.

2.7 Bisimulation

Bisimulation is a notion of behavioral equivalence between processes. For the case of LTSs it is defined as follows.

Definition 10. *Given two LTSs $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$ and $\mathbb{T}' = \langle T', R', \sigma', s_I' \rangle$ a* bisimulation *is a relation $Z \subseteq T \times T'$ such that for all $(t, t') \in Z$ the following holds:*

- $\sigma(t) = \sigma'(t')$;
- *for all $s \in R[t]$ there is $s' \in R'[t']$ such that $(s, s') \in Z$;*
- *for all $s' \in R'[t']$ there is $s \in R[t]$ such that $(s, s') \in Z$.*

Two pointed LTSs $\mathbb{T}$ and $\mathbb{T}'$ are bisimilar *(denoted $\mathbb{T} \leftrightarrow \mathbb{T}'$) if there is a bisimulation $Z \subseteq T \times T'$ containing (s_I, s_I').*

The following fact will be used in the proof of Theorem 1.

Position	Player	Admissible moves	Parity
$(a, s) \in A \times T$	$\exists$	$\{V : A \to \wp(R[s]) \mid (R[s], V) \models \Delta(a, \sigma(s))\}$	$\Omega(a)$
$V : A \to \wp(T)$	$\forall$	$\{(b, t) \mid t \in V(b)\}$	$\max(\Omega[A])$

Table 1. Acceptance game for parity automata.

Fact 2. $\mathbb{T} \leftrightarrow \mathbb{T}^\omega$, *for any LTS* $\mathbb{T}$.

A formula $\varphi \in \mathcal{L}$ is *bisimulation invariant* if $\mathbb{T} \leftrightarrow \mathbb{T}'$ implies that $\mathbb{T} \Vdash \varphi$ iff $\mathbb{T}' \Vdash \varphi$, for all $\mathbb{T}$ and $\mathbb{T}'$. We use $\mathcal{L}/{\leftrightarrow}$ for the class of bisimulation-invariant $\mathcal{L}$-formulas.

Fact 3. *All formulas of* $\mu\mathrm{ML}$ *are bisimulation invariant.*

3. Normal Forms and Continuity

In this section we study the properties of the one-step languages on which the automata introduced in Section 4 and 5 will be based. For each language $\mathcal{L}_1(A)$ that we consider, we focus on $\mathcal{L}_1^+(A)$, as it is the relevant fragment for defining the transition of parity automata (*cf.* Definition 4). However, all our results can be extended to the full language $\mathcal{L}_1(A)$.

Definition 11. *The positive fragment* $\mathrm{FOE}_1^{\infty+}(A)$ *of* $\mathrm{FOE}_1^\infty(A)$ *is given by the sentences generated by:*

$$\varphi ::= a(x) \mid x \approx y \mid x \not\approx y \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \mathcal{Q}x.\varphi(x)$$

where $x, y \in \mathsf{iVar}$, $a \in A$ *and* $\mathcal{Q} \in \{\exists, \forall, \exists^\infty, \forall^\infty\}$.

As required by Definition 3, the above fragment can be shown to characterize monotonicity for FOE_1^∞. Moreover, the characterization naturally restricts to FOE_1 and FO_1.

3.1 Normal Forms

Given a set of names A and $S \subseteq A$, we call $\tau_S^+(x) := \bigwedge_{a \in S} a(x)$ a *positive A-type*. We usually blur the distinction between $\tau_S^+(x)$ and S and call S a positive A-type as well.

As FOE_1^∞ is the one-step language of WMSO-automata, the following normal form theorem will be pivotal.

Theorem 3. *Every sentence* $\varphi \in \mathrm{FOE}_1^{\infty+}(A)$ *is equivalent to a sentence in the basic form* $\bigvee \nabla_{\mathrm{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$ *where*

$$\nabla_{\mathrm{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma) := \nabla_{\mathrm{FOE}}^+(\vec{T}, \Pi \cup \Sigma) \wedge \nabla_\infty^+(\Sigma)$$

$$\nabla_{\mathrm{FOE}}^+(\vec{T}, \Lambda) := \exists \vec{x}.\big(\mathrm{diff}(\vec{x}) \wedge \bigwedge_i \tau_{T_i}^+(x_i) \wedge$$

$$\forall z.(\mathrm{diff}(\vec{x}, z) \to \bigvee_{S \in \Lambda} \tau_S^+(z))\big)$$

$$\nabla_\infty^+(\Sigma) := \bigwedge_{S \in \Sigma} \exists^\infty y.\tau_S^+(y) \wedge \forall^\infty y. \bigvee_{S \in \Sigma} \tau_S^+(y)$$

for some sets of types $\Pi, \Sigma \subseteq \wp A$, *each* T_i *a subset of* A *and* $\mathrm{diff}(y_1, \ldots, y_n) := \bigwedge_{1 \leq m < m' \leq n}(y_m \not\approx y_{m'})$.

A simple argument reveals that, intuitively, every disjunct of the basic form above expresses that any one-step model satisfying it admits a partition of its domain in three parts:

(i) distinct elements $t_1, \ldots, t_n$ with type $T_1, \ldots, T_n$,

(ii) finitely many elements whose types belong to Π, and

(iii) for each $S \in \Sigma$, infinitely many elements with type S.

Proof. The proof of this theorem can be seen as a non-trivial extension of [7, Lemma 16.23]. Given $S \subseteq A$, and a one-step model $\mathbf{D}$, let $|S|_\mathbf{D}$ be the number of elements in $\mathbf{D}$ of type S. We say that two such structures $\mathbf{D}$ and $\mathbf{D}'$ are k-equivalent (notation $\mathbf{D} \sim_k^\infty \mathbf{D}'$) iff for all $S \subseteq A$, either $|S|_\mathbf{D} = |S|_{\mathbf{D}'} \leq k$ or both $|S|_\mathbf{D}, |S|_{\mathbf{D}'} > k$

while at the same time $|S|_\mathbf{D} < \omega$ iff $|S|_{\mathbf{D}'} < \omega$. The equivalence relation $\sim_k^\infty$ clearly has finite index, and each equivalence class of $\sim_k^\infty$ is described by a sentence of FOE_1^∞. Using an extension of Ehrenfeucht-Fraïssé games that takes into account the generalized quantifiers it is not difficult to show that $\mathbf{D} \sim_k^\infty \mathbf{D}'$ implies that they satisfy the same FOE_1^∞-sentences of quantifier rank at most k. It follows that the class of models of such a sentence φ is the union of a (finite) number of $\sim_k^\infty$-cells, and that φ is thus equivalent to the disjunction of the sentences associated with these cells. □

3.2 Continuity and Co-continuity

Next we give a syntactic characterization of (co-)continuity for both FO_1^+ and $\mathrm{FOE}_1^{\infty+}$. This is instrumental in a proper implementation of the **(continuity)** condition on the automata that we seek to define, as explained in the introduction.

We say that $\varphi \in \mathcal{L}_1(A)$ is *continuous in* $a \in A$ if φ is monotone in a and additionally, for every (D, V) we have that $(D, V) \models \varphi$ implies the existence of some $U \subseteq_\omega V(a)$ such that $(D, V[a \mapsto U]) \models \varphi$. Dually, φ is *co-continuous in* $a \in A$ if φ is monotone in a and, for all (D, V) we have that if $(D, V) \not\models \varphi$ then for some $U \subseteq_\omega V(a)$ it holds that $(D, V[a \mapsto D \setminus U]) \not\models \varphi$.

Proposition 4. *A sentence of* $\mathrm{FO}_1^+(A)$ *is continuous in* $a \in A$ *iff it is equivalent to a sentence given by*

$$\varphi ::= \psi \mid a(x) \mid \exists x.\varphi(x) \mid \varphi \wedge \varphi \mid \varphi \vee \varphi$$

where $\psi \in \mathrm{FO}_1^+(A \setminus \{a\})$. *We name this fragment* $\mathrm{FO}_1^+\mathrm{C}_a(A)$.

Theorem 4. *A sentence of* $\mathrm{FOE}_1^{\infty+}(A)$ *is continuous in* $a \in A$ *iff it is equivalent to a sentence of the fragment given by*

$$\varphi ::= \psi \mid a(x) \mid \exists x.\varphi(x) \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \mathbf{W}x.(\varphi, \psi)$$

where $\mathbf{W}x.(\varphi, \psi) := \forall x.(\varphi(x) \vee \psi(x)) \wedge \forall^\infty x.\psi(x)$ *and* $\psi \in \mathrm{FOE}_1^{\infty+}(A \setminus \{a\})$. *We name this fragment* $\mathrm{FOE}_1^{\infty+}\mathrm{C}_a(A)$.

Universal quantification is usually problematic for preserving continuity because of its potentially infinite nature. However, in the case of $\mathbf{W}x.(\varphi, \psi)$, the combination of both quantifiers ensures that all the elements are covered by $\varphi \vee \psi$ but only *finitely many* are required to make φ (in which $a \in A$ may occur) true. This gives no trouble for continuity in a.

Proof. For the challenging direction (from left to right) we prove a stronger claim. We give a translation $(\cdot)^\Delta : \mathrm{FOE}_1^{\infty+}(A) \to \mathrm{FOE}_1^{\infty+}\mathrm{C}_a(A)$ by setting

$$\big(\nabla_{\mathrm{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)\big)^\Delta := \begin{cases} \bot & \text{if } a \in \bigcup \Sigma \\ \nabla_{\mathrm{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma) & \text{otherwise} \end{cases}$$

and $(\bigvee \nabla_{\mathrm{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma))^\Delta := \bigvee (\nabla_{\mathrm{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma))^\Delta$. Then we prove: $\varphi \in \mathrm{FOE}_1^{\infty+}(A)$ is continuous in a iff $\varphi \equiv \varphi^\Delta$. The key observation is that φ^Δ can be rewritten using $\mathbf{W}$. □

As the translation $(\cdot)^\Delta$ preserves the normal form, we get a normal form result for the continuous fragment of FOE_1^∞.

Corollary 1. *A sentence* $\varphi \in \mathrm{FOE}_1^{\infty+}(A)$ *is continuous in* $a \in A$ *if and only if it is equivalent to a sentence in the basic form* $\bigvee \nabla_{\mathrm{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$, *for some sets of types* $\Pi, \Sigma \subseteq \wp A$ *and each* T_i *a subset of* A, *such that* $a \notin \bigcup \Sigma$.

The logics $\mathrm{FOE}_1^{\infty+}(A)$ and $\mathrm{FO}_1^+(A)$ are easily seen to be closed under Boolean duals – e.g., $(\exists^\infty x.\varphi)^\delta = \forall^\infty x.\varphi^\delta$. Combining the previous results with the properties of duals we get a syntactic characterization of co-continuity.

Corollary 2. *The fragments of* $\mathrm{FOE}_1^{\infty+}(A)$ *and* $\mathrm{FO}_1^+(A)$ *which are* co-continuous *in* $a \in A$ *are characterized by*

$$\mathrm{FOE}_1^{\infty+}\overline{\mathrm{C}}_a(A) := \{\varphi \mid \varphi^\delta \in \mathrm{FOE}_1^{\infty+}\mathrm{C}_a(A)\}$$
$$\mathrm{FO}_1^+\overline{\mathrm{C}}_a(A) := \{\varphi \mid \varphi^\delta \in \mathrm{FO}_1^+\mathrm{C}_a(A)\}.$$

4. Automata for WMSO

Using the syntactic characterization of (co-)continuity given in the previous section for $\mathrm{FOE}_1^{\infty+}(A)$, we are now able to state a proper definition of the automata for WMSO outlined in the introduction.

Definition 12. *A* WMSO-*automaton* $\langle A, \Delta, \Omega, a_I \rangle$ *is an automaton in* $Aut(\mathrm{FOE}_1^\infty)$ *such that, for all states* $a, b \in A$ *with* $a \preceq b$ *and* $b \preceq a$, *and for all* $c \in C$, *the following holds:*

(weakness) $\Omega(a) = \Omega(b)$;
(continuity) *if* $\Omega(a)$ *is odd then* $\Delta(a, c)$ *is in* $\mathrm{FOE}_1^{\infty+}\mathrm{C}_b(A)$; *if* $\Omega(a)$ *is even then* $\Delta(a, c)$ *is in* $\mathrm{FOE}_1^{\infty+}\overline{\mathrm{C}}_b(A)$.

We use $Aut_{cw}(\mathrm{FOE}_1^\infty)$ *to denote the class of such automata.*

The rest of this section will be devoted to prove that WMSO-automata characterize WMSO on tree models, as expressed in Theorem 2. First, we focus on showing the direction from formulas to automata. In subsections 4.1 and 4.2 we provide the automata constructions handling the challenging case, that is the translation of an existential formula $\exists p.\psi$ of WMSO into an equivalent WMSO-automaton. To this aim, we define a closure operation on tree languages corresponding to the semantics of WMSO quantification.

Definition 13. *Let* P *be a set of proposition letters,* $p \notin P$ *be a proposition letter, and* $\mathcal{T}$ *be a tree language of* $\wp(\mathsf{P} \cup \{p\})$-*labeled trees. The* finitary projection *of* $\mathcal{T}$ *over* p *is the language* $\exists_F p.\mathcal{T}$ *of* $\wp(\mathsf{P})$-*labeled trees given as follows:*

$$\exists_F p.\mathcal{T} = \{\mathbb{T} \mid \exists \text{ a finite } p\text{-variant } \mathbb{T}' \text{ of } \mathbb{T} \text{ with } \mathbb{T}' \in \mathcal{T}\}.$$

A class K *of tree languages is* closed under finitary projection over p *if, for any language* $\mathcal{T}$ *in* K, *also* $\exists_F p.\mathcal{T}$ *is in* K.

4.1 A Simulation Theorem for WMSO-Automata

Our next goal is a *projection construction* that, given a WMSO-automaton $\mathbb{A}$, provides one recognizing $\exists_F p.\mathcal{T}(\mathbb{A})$. For MSO-automata, an analogous construction crucially uses the following *simulation theorem*: every MSO-automaton $\mathbb{A}$ is equivalent to a *non-deterministic* one $\mathbb{A}'$ [18]. Semantically, non-determinism yields the appealing property that any strategy f for player $\exists$ in the acceptance game $\mathcal{A}(\mathbb{A}', \mathbb{T})$ can be assumed to be functional in A' (*cf.* Definition 6). This is particularly helpful because, to define a p-variant of $\mathbb{T}$ that is accepted by the projection construct on $\mathbb{A}'$, we can infer whether a node s should be labeled with p by the value $f(a, s)$, where a is the unique state of $\mathbb{A}'$ (by functionality) that f associates with s. Now, in the case of WMSO-automata we are interested in guessing *finitary* p-variants, which requires f to be functional only on a *finite* set of nodes. Thus the idea of our simulation theorem is to turn a WMSO-automaton $\mathbb{A}$ into an equivalent one $\mathbb{A}^F$ that behaves non-deterministically on a *finite* portion of any accepted tree.

For MSO-automata, the simulation theorem is based on a powerset construction: if the starting automaton has carrier A, the resulting non-deterministic automaton is based on "macro-states"

from the set $A^\sharp := \wp(A \times A)$.[1] Analogously, for WMSO-automata we will associate the non-deterministic behavior with macro-states. However, as explained above, the automaton $\mathbb{A}^F$ that we construct has to be non-deterministic just on finitely many nodes of the input and may behave as $\mathbb{A}$ (i.e. in "alternating mode") on the others. To this aim, $\mathbb{A}^F$ will be "two-sorted", roughly consisting of one copy of $\mathbb{A}$ (with carrier A) and a variant of its powerset construction, based both on A and $A^\sharp$. For any accepted $\mathbb{T}$, the idea is to make any match π of $\mathcal{A}(\mathbb{A}^F, \mathbb{T})$ consist of two parts:

(Non-deterministic mode) for finitely many rounds of π, each visited basic position has shape $(q, s) \in A^\sharp \times T$. The valuation $V : A \cup A^\sharp \to \wp(R[s])$ picked by player $\exists$ assigns macro-states (from $A^\sharp$) only to a *finite* subset of $R[s]$ and states (from A) to the rest of $R[s]$. Also, she assigns *at most one macro-state* to each node.

(Alternating mode) At a certain round, π visits the last position with a macro-state and turns into a match of the game $\mathcal{A}(\mathbb{A}, \mathbb{T})$, i.e. all next positions are from $A \times T$.

Therefore successful runs of $\mathbb{A}^F$ will have the property of processing only a *finite* amount of the input with $\mathbb{A}^F$ being in a macro-state and all the rest with $\mathbb{A}^F$ behaving exactly as $\mathbb{A}$.

We now proceed in steps towards the construction of $\mathbb{A}^F$. The following is a notion of lifting for types on states that is instrumental in defining a translation to types on macro-states.

Definition 14. *Given a set* A *and* $\Sigma \subseteq \wp A$, *we define the* lifting $\Sigma^\wp \subseteq \wp\wp A$ *as* $\{\{S\} \mid S \in \Sigma \wedge S \neq \emptyset\} \cup \{\emptyset \mid \emptyset \in \Sigma\}$.

The next definition is standard (see e.g. [17, 18]) as an intermediate step to define the transition function of the powerset construct for parity automata. It simply *tags* the (potential next) states occurring in $\Delta(a, c)$ with the information of the current state.

Definition 15. *Let* $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ *be a* WMSO-*automaton. Fix* $a \in A$, $c \in C$. *We define*

$$\Delta^\star(a, c) \quad := \quad \Delta(a, c)[b \mapsto (a, b) \mid b \in A]$$

where $\Delta(a, c)[b \mapsto (a, b) \mid b \in A] \in \mathrm{FOE}_1^{\infty+}(A \times A)$ *is the sentence obtained by replacing each monadic predicate* $b \in A$ *in* $\Delta(a, c)$ *with the monadic predicate* $(a, b) \in A \times A$.

We now define a translation for the one-step language of WMSO-automata, which can be thought as based on carrier $A \times A$ by effect of the transformation of Definition 15.

Definition 16. *Let* $\varphi \in \mathrm{FOE}_1^{\infty+}(A \times A)$ *be of the shape*

$$\varphi = \nabla_{\mathrm{FOE}}^+(\vec{T}, \Pi \cup \Sigma) \wedge \nabla_\infty^+(\Sigma)$$

where $\Pi, \Sigma \subseteq A^\sharp$ *and each* $T_i \subseteq A \times A$ *(provided by Theorem 3). Let* $\widetilde{\Sigma} \subseteq \wp A$ *be* $\widetilde{\Sigma} := \{\mathsf{Ran}(S) \mid S \in \Sigma\}$. *Define the translation* $\varphi^F \in \mathrm{FOE}_1^{\infty+}(A \cup A^\sharp)$ *as:*

$$\varphi^F := \nabla_{\mathrm{FOE}}^+(\vec{T}^\wp, \Pi^\wp \cup \Sigma^\wp \cup \widetilde{\Sigma}) \wedge \nabla_\infty^+(\widetilde{\Sigma}).$$

The idea of translation $(\cdot)^F$ is to encode at the one-step level the non-deterministic mode of $\mathbb{A}^F$. As no macro-state occurs in $\nabla_\infty^+(\widetilde{\Sigma})$, by Corollary 1 the sentence φ^F is continuous in each $R \in A^\sharp$, i.e. it can be made true in a domain D by assigning macro-states to *finitely many* elements of D. Moreover, macro-states occur in φ^F only inside lifted types in $\vec{T}^\wp$, $\Pi^\wp$ or $\Sigma^\wp$: then, by definition of $(\cdot)^\wp$, φ^F can be made true in D by assigning *at most one* macro-state to any element.

---

[1] The use of carrier $\wp(A \times A)$ instead of the more obvious $\wp A$ is needed to correctly associate with a run on macro-states the corresponding bundle of runs of the original automaton $\mathbb{A}$ (*cf.* [18]).

Next we combine the previous definitions to characterize the transition function associated with the macro-states.

Definition 17. *Let* $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ *be a* WMSO-*automaton. Fix any* $c \in C$ *and* $Q \in A^\sharp$. *By Theorem 3 there is a sentence* $\Psi_{Q,c} \in \mathrm{FOE}_1^{\infty+}(A \times A)$ *in the basic form* $\bigvee \nabla_{\mathrm{FOE}\infty}^+ (\vec{T}, \Pi, \Sigma)$, *for some* $\Pi, \Sigma \subseteq A^\sharp$ *and* $T_i \subseteq A \times A$, *such that*

$$\bigwedge_{a \in \mathrm{Ran}(Q)} \Delta^\star(a, c) \equiv \Psi_{Q,c}.$$

By definition, $\Psi_{Q,c} = \bigvee_n \varphi_n$, *with each* φ_k *of shape*

$$\nabla_{\mathrm{FOE}\infty}^+ (\vec{T}, \Pi, \Sigma) = \nabla_{\mathrm{FOE}}^+ (\vec{T}, \Pi \cup \Sigma) \wedge \nabla_\infty^+(\Sigma).$$

We put $\Delta^\sharp(Q, c) := \bigvee_n \varphi_n^F$, *where the translation* $(\cdot)^F$ *is as in Definition 16. Observe that* $\Delta^\sharp(Q, c) \in \mathrm{FOE}_1^{\infty+}(A \cup A^\sharp)$.

We have now all the ingredients for our two-sorted automaton.

Definition 18. *Let* $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ *be a* WMSO-*automaton. We define the* finitary construct *over* $\mathbb{A}$ *as the automaton* $\mathbb{A}^F = \langle A^F, \Delta^F, \Omega^F, a_I^F \rangle$ *given by*

$$
\begin{aligned}
A^F &:= A \cup A^\sharp \\
a_I^F &:= \{(a_I, a_I)\} \\
\Delta^F(q, c) &:= \begin{cases} \Delta(q, c) & q \in A \\ \Delta^\sharp(q, c) \vee \bigwedge_{a \in \mathrm{Ran}(q)} \Delta(a, c) & q \in A^\sharp \end{cases} \\
\Omega^F(q) &:= \begin{cases} \Omega(q) & q \in A \\ 1 & q \in A^\sharp. \end{cases}
\end{aligned}
$$

The definition of $\mathbb{A}^F$ enforces its behavior to be split according to the non-deterministic and alternating mode. Indeed, for any accepted $\mathbb{T}$, a match π of $\mathcal{A}(\mathbb{A}^F, \mathbb{T})$ will visit positions involving macro-states only for finitely many initial rounds, because $\Omega^F[A^\sharp] = \{1\}$. The alternating mode will be entered when, at a certain position $(R, s) \in A^\sharp \times T$, the winning strategy for $\exists$ makes the disjunct $\bigwedge_{a \in \mathrm{Ran}(R)} \Delta(a, c)$ of $\Delta^F(R, c)$ true and then all successive positions only involve states from A. The next proposition fixes our desiderata on $\mathbb{A}^F$.

Proposition 5 (**Simulation Theorem for** WMSO-**automata**)**.** *Let* $\mathbb{A}$ *be a* WMSO-*automaton and* $\mathbb{A}^F$ *its finitary construct.*

(i) $\mathbb{A}^F$ *is a* WMSO-*automaton.*
(ii) For any $\mathbb{T}$, *if* $\exists$ *has a winning strategy in* $\mathcal{A}(\mathbb{A}^F, \mathbb{T})$ *from position* (a_I^F, s_I) *then she has one that is functional in* $A^\sharp$ *and finitary in* $A^\sharp$ (cf. *Definition 6*).
(iii) $\mathbb{A} \equiv \mathbb{A}^F$.

Proof. For (i), observe that any SCC of $\mathbb{A}^F$ involves states of exactly one sort, either A or $A^\sharp$. For SCCs on sort A, **(weakness)** and **(continuity)** of $\mathbb{A}^F$ follow by the ones of $\mathbb{A}$. For SCCs on sort $A^\sharp$, **(weakness)** follows by observing that all macro-states in $\mathbb{A}^F$ have the same parity value. Concerning **(continuity)**, by definition of Δ^F any macro-state can only appear inside a formula of the form φ^F, which as observed above is continuous in each $Q \in A^\sharp$. Statement (ii) again follows by properties of the translation $(\cdot)^F$ and the observation that a winning strategy will eventually let $\mathbb{A}^F$ enter the alternating mode. The argument for (iii) goes as follows. For the direction from left to right, if player $\exists$ has a winning strategy f for $\mathcal{A}(\mathbb{A}, \mathbb{T})$ from position (a_I, s_I), then f will also be winning in $\mathcal{A}(\mathbb{A}^F, \mathbb{T})$ from position (a_I^F, s_I) — intuitively, playing f amounts to immediately let $\mathbb{A}^F$ enter the alternating mode by making the disjunct $\bigwedge_{a \in \mathrm{Ran}(a_I^F)} \Delta(a, \sigma(s_I)) = \Delta(a_I, \sigma(s_I))$ of $\Delta^F(a_I^F, \sigma(s_I))$ true. Conversely, suppose that $\mathbb{A}^F$ accepts $\mathbb{T}$ and let g be the corresponding winning strategy for $\exists$. We can make

$\exists$ win any match π of $\mathcal{A}(\mathbb{A}, \mathbb{T})$ played from position (a_I, s_I) as follows. While playing π, we maintain a "shadow match" π' of $\mathcal{A}(\mathbb{A}^F, \mathbb{T})$ from position (a_I^F, s_I) where $\exists$ plays according to g. By suitably defining the strategy of $\exists$ in π in terms of g, we can enforce the following relation between π and π' at each round: either (I) the same position of the form $(a, s) \in A \times T$ occurs in both matches or (II) a position $(a, s) \in A \times T$ in π corresponds to a position $(Q, s) \in A^\sharp \times T$ in π', with $a \in \mathrm{Ran}(Q)$. Since g is winning, eventually $\mathbb{A}^F$ will enter the alternating mode and thus situation (I) will be the case for all but finitely many initial rounds, implying that $\exists$ wins the match π. $\square$

Remark 1. Albeit similar to the finitary construction, the *two-sorted construction* (*cf.* [19, Def. 3.7], [5]) used for weak MSO-automata would have not been suitable for our purposes, as it fails to preserve the **(continuity)** condition when applied to WMSO-automata. Similarly, the powerset construction used in the simulation theorem for MSO-automata preserves neither the **(weakness)** nor the **(continuity)** condition.

4.2 Closure Properties of WMSO-Automata

We are now ready to introduce our projection construction for WMSO-automata and show that the class of tree languages that they recognize is closed under finitary projection.

Definition 19. *Let* $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ *be a* WMSO-*automaton on alphabet* $\wp(\mathsf{P} \cup \{p\})$, *with* $p \notin \mathsf{P}$. *Let* $\mathbb{A}^F$ *denote its finitary construct. We define the* WMSO-*automaton* $\exists_F p.\mathbb{A} := \langle A^F, a_I^F, \widetilde{\Delta}, \Omega^F \rangle$ *on alphabet* $\wp(\mathsf{P})$ *by putting*

$$\widetilde{\Delta}(q, c) := \begin{cases} \Delta^F(q, c) & q \in A \\ \Delta^F(q, c) \vee \Delta^F(q, c \cup \{p\}) & q \in A^\sharp. \end{cases}$$

Proposition 6. *For every* WMSO-*automaton* $\mathbb{A}$ *on alphabet* $\wp(\mathsf{P} \cup \{p\})$, *with* $p \notin \mathsf{P}$, *we have that* $\mathcal{T}(\exists_F p.\mathbb{A}) = \exists_F p.\mathcal{T}(\mathbb{A})$.

Proof. For the inclusion from left to right, first observe that $\exists_F p.\mathbb{A}$ is defined in terms of $\mathbb{A}^F$ and thus the properties stated in Proposition 5 hold for $\exists_F p.\mathbb{A}$ as well. In particular, given a $\wp(\mathsf{P})$-tree $\mathbb{T}$, any winning strategy f for $\exists$ in $\mathcal{A}(\exists_F p.\mathbb{A}, \mathbb{T})$ from position (a_I^F, s_I) can be assumed to be functional and finitary in $A^\sharp$. We can use such a strategy to guess a finitary p-variant of $\mathbb{T}$ as follows. First, by functionality for each node s there is at most one position (Q_s, s), with $Q_s \in A^\sharp$, that is reachable in any f-guided match. From each such position, let $V_{Q_s, s} : A^F \to \wp(R[s])$ be the valuation suggested by f. We let X_p be the set of nodes s for which $V_{Q_s, s}$ makes the disjunct $\Delta^F(Q_s, \sigma(s) \cup \{p\})$ of $\widetilde{\Delta}(Q_s, \sigma(s))$ true: intuitively, these are the nodes on which $\exists_F p.\mathbb{A}$ behaves "as if they were labeled with p". Since f is finitary in $A^\sharp$, the p-variant $\mathbb{T}'$ of $\mathbb{T}$ given by labeling the nodes in X_p with p is finitary. One can readily verify that $\mathbb{A}^F$ (and thus $\mathbb{A}$ by Proposition 5) accepts $\mathbb{T}'$ by letting $\exists$ playing the strategy f in $\mathcal{A}(\mathbb{A}^F, \mathbb{T}')$.

For the inclusion from right to left, let $\mathbb{T}'$ be a finitary p-variant of some $\wp(\mathsf{P})$-tree $\mathbb{T}$ and suppose that $\exists$ has a winning strategy f for $\mathcal{A}(\mathbb{A}, \mathbb{T}')$ from position (a_I, s_I). We now sketch how $\exists$ is able to win any match π of $\mathcal{A}(\exists_F p.\mathbb{A}, \mathbb{T})$ from position (a_I^F, s_I). The idea is to enforce that, at each round of π, $\exists$ assigns macro-states only to the nodes rooting a subtree of $\mathbb{T}'$ where the labeling p appears, and A-states to the others. Using the information given by f, $\exists$ can make this assignment so that any visited position of shape $(Q, s) \in A^\sharp \times T$ is such that (a, s) is winning for $\exists$ in $\mathcal{A}(\mathbb{A}, \mathbb{T}')$, for each $a \in \mathrm{Ran}(Q)$. In particular, the assignment of $\exists$ will make true the disjunct $\Delta^F(Q, \sigma(s) \cup \{p\})$ of $\widetilde{\Delta}(Q, \sigma(s))$ if s is labeled with p in $\mathbb{T}'$, and the disjunct $\Delta^F(Q, \sigma(s))$ otherwise. Since $\mathbb{T}'$ is a *finitary p-variant*, player $\exists$ will be required to assign macro-states to only finitely many nodes encountered along the play, and

π will eventually arrive to a position from which no node labeled with p is reachable. At that point, $\exists$ allows $\exists_F p.\mathbb{A}$ to switch from the non-deterministic to the alternating mode. By construction, the match π now moves to a position $(a, s) \in A \times T$ that is winning for $\exists$ in $\mathcal{A}(\mathbb{A}, \mathbb{T}')$. It is also winning in $\mathcal{A}(\exists_F p.\mathbb{A}, \mathbb{T})$, because $\mathbb{T}'$ agrees with $\mathbb{T}$ on nodes without labeling p and by definition $\widetilde{\Delta}(a, \sigma(s)) = \Delta(a, \sigma(s))$. $\square$

We have now in position to show our characterization result.

Proof of Theorem 2, direction $(\Rightarrow)$. By induction we prove that for every $\varphi \in$ WMSO there is a WMSO-automaton $\mathbb{A}_\varphi$ such that $[\![\varphi]\!] = \mathcal{T}(\mathbb{A}_\varphi)$. We focus on two inductive cases.

If $\varphi = \neg\psi$, let $\mathbb{A}_\psi$ be the WMSO-automaton for ψ given by inductive hypothesis. As FOE_1^∞ is closed under Boolean duals, we can define the complementation $\overline{\mathbb{A}_\psi}$ of $\mathbb{A}_\psi$ as in Definition 8. Notice that $\overline{\mathbb{A}_\psi}$ is indeed a WMSO-automaton, satisfying the **(weakness)** and **(continuity)** conditions in virtue of their self-dual nature. Proposition 1 yields the complementation lemma allowing to conclude that on trees $[\![\neg\psi]\!] = \mathcal{T}(\overline{\mathbb{A}_\psi})$.

If $\varphi = \exists p.\psi$, let $\mathbb{A}_\psi$ be the automaton given by inductive hypothesis. By semantics of WMSO, on trees $[\![\exists p.\psi]\!] = \exists_F p.[\![\psi]\!]$ and thus $[\![\exists p.\psi]\!] = \mathcal{T}(\exists_F p.\mathbb{A}_\psi)$ by Proposition 6. $\square$

4.3 From WMSO-Automata to WMSO-Formulas

In this section we show the other direction of Theorem 2, completing the automata characterization of WMSO on tree models. The argument is reminiscent of the one showing that MSO-automata can be translated into equivalent formulas of MSO [18]. We start by introducing a fragment of a fixpoint extension of FOE^∞ and show how it embeds into WMSO.

Definition 20. *The fixed point logic μFOE^∞ on* P *is given by:*
$$\varphi ::= p(x) \mid x = y \mid R(x, y) \mid \neg\varphi \mid \varphi \wedge \varphi \mid$$
$$\exists x.\varphi \mid \exists^\infty x.\varphi \mid \mu p.\varphi(p, x)$$

where $p \in$ P*, $x, y \in$* iVar*; moreover p occurs only positively in $\varphi(p, x)$ and x is the only free variable in $\varphi(p, x)$.*

The semantics of $\mu p.\varphi(p, x)$ is the expected one: given an LTS $\mathbb{T}$ and $s \in T$, $\mathbb{T} \models \mu p.\varphi(p, s)$ iff s is in the least fixpoint of the operator $\varphi_p^{\mathbb{T}}(S) := \{t \in T \mid \mathbb{T}[p \mapsto S] \models \varphi(p, t)\}$.

Definition 21. *Given $p \in$* P*, we say that $\varphi \in \mu\text{FOE}^\infty$ is*

- *monotone in p iff for every LTS $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$ and assignment g, if $\mathbb{T}, g \models \varphi$ and $\{s \in T \mid p \in \sigma(s)\} \subseteq E$, then $\mathbb{T}[p \mapsto E], g \models \varphi$,*
- *continuous in p iff for every LTS $\mathbb{T} = \langle T, R, \sigma, s_I \rangle$ and assignment g there is some finite $S \subseteq_\omega \{s \in T \mid p \in \sigma(s)\}$ such that $\mathbb{T}, g \models \varphi$ iff $\mathbb{T}[p \mapsto S], g \models \varphi$.*

We provide a definition of a fragment of μFOE^∞ reminiscent of the one in Theorem 4.

Definition 22. *Given a set $Q \subseteq$* P*, the fragment $\mu\text{FOE}^\infty C_Q$ is defined by the following rules:*
$$\varphi ::= \psi \mid q(x) \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid$$
$$\exists x.\varphi \mid \mathbf{W}x.(\varphi, \psi) \mid \mu p.\varphi'(p, x)$$

where $q \in Q$, $\psi \in \mu\text{FOE}^\infty$ is Q-free, $\mathbf{W}x.(\varphi, \psi) := \forall x.(\varphi(x) \vee \psi(x)) \wedge \forall^\infty x.\psi(x)$ and $\varphi'(p, x)$ is a formula, with only x free and $p \in$ P*, which belongs to $\mu\text{FOE}^\infty C_{Q \cup \{p\}}$.*

By combining the argument for the proof of Proposition 4 and the one used in proving the analogous Lemma 1 in [6], we can thence obtain the following:

Proposition 7. *If $\varphi \in \mu\text{FOE}^\infty C_Q$ then φ is continuous in (each element of) Q.*

Definition 23. *The fragment $\mu_c\text{FOE}^\infty$ of μFOE^∞ is given by restricting the fixpoint operator to the continuous fragment:*
$$\varphi ::= p(x) \mid x = y \mid R(x, y) \mid \neg\varphi \mid \varphi \wedge \varphi \mid$$
$$\exists x.\varphi \mid \exists^\infty x.\varphi \mid \mu p.\varphi'(p, x)$$

where $p \in$ P*, $x, y \in$* iVar*, $\varphi'(p, x) \in \mu\text{FOE}^\infty C_{\{p\}} \cap \mu_c\text{FOE}^\infty$ is positive in p and x is its only free variable.*

Given an LTS $\mathbb{T}$ and $p \in$ P, for every ordinal α we define:

- $\varphi_p^0(\emptyset) := \emptyset$,
- $\varphi_p^{\alpha+1}(\emptyset) := \{s \in T \mid \mathbb{T}[p \mapsto \varphi_p^\alpha(\emptyset)] \models \varphi(p, s)\}$,
- $\varphi_p^\lambda(\emptyset) := \bigcup_{\alpha < \lambda} \varphi_p^\alpha(\emptyset)$, with λ limit.

If φ is monotone in p, then $\varphi_p^{\beta+1}(\emptyset) = \varphi_p^\beta(\emptyset)$, for some ordinal β. Also, the set $\varphi_p^\beta(\emptyset)$ is the least fixpoint of $\varphi_p^{\mathbb{T}}$ (see e.g. [1]).

A formula $\varphi(p, x)$ is *constructive* in p if for every model $\mathbb{T}$, every node $s \in T$, if $\mathbb{T} \models \mu p.\varphi(p, s)$, then $s \in \varphi_p^{i+1}(\emptyset)$, for some $i < \omega$. The next proposition is easily verified:

Proposition 8. *Let $\varphi(p, x)$ be a μFOE^∞-formula with only x free. If $\varphi(p, x)$ is continuous in p, then for every LTS $\mathbb{T}$, and every node $s \in T$, there is $i < \omega$ such that*
$$\mathbb{T} \models \mu p.\varphi(p, s) \text{ iff } s \in \varphi_p^{i+1}(\emptyset).$$

By Proposition 8 and the fact that sets $\varphi_p^{i+1}(\emptyset)$ are essentially defined as finite unfoldings, we obtain the following.

Proposition 9. *Let $\varphi(p, x)$ be a μFOE^∞-formula continuous in p with only x free. Let $\mathbb{T}$ be an LTS, and $s \in T$. Then $\mathbb{T} \models \mu p.\varphi(p, s)$ iff there is a finite set $p^{\mathbb{T}} \subseteq T$ such that $s \in p^{\mathbb{T}}$ and $\mathbb{T}[p \mapsto p^{\mathbb{T}}] \models \varphi(p, t)$ for every $t \in p^{\mathbb{T}}$.*

Proposition 9 naturally suggests the following translation $(\cdot)_\circ$: $\mu\text{FOE}^\infty \to$ WMSO. It is given homomorphically in predicates, Booleans and first-order quantifiers and:

- $(\exists^\infty x.\varphi)_\circ := \forall p.\exists x.(\neg p(x) \wedge (\varphi)_\circ)$,
- $(\mu p.\varphi(p, x))_\circ := \exists p(p(x) \wedge \forall y(p(y) \to (\varphi(p, y))_\circ))$.

Proposition 10. *Let φ be a $\mu_c\text{FOE}^\infty$-formula, $\mathbb{T}$ an LTS and g an assignment. Then $\mathbb{T}, g \models \varphi$ iff $\mathbb{T}, g \models (\varphi)_\circ$.*

Proof. The proof is by induction on φ. The least fixpoint case is handled by applying Proposition 9. $\square$

By virtue of Proposition 10, we are able to conclude the proof of Theorem 2 by showing the following statement.

Proposition 11. *Every* WMSO*-automaton can be effectively translated into an equivalent $\mu_c\text{FOE}^\infty$-formula.*

Proof. The argument is essentially a refinement of the standard proof showing that any automaton in $Aut(\text{FO}_1)$ can be translated into an equivalent μ-formula $\xi_{\mathbb{A}}$ (*cf.* e.g. [17]). The idea is the following. We see a WMSO-automaton as a system of equations expressed in terms of FOE^∞-formulas: each state corresponds to a monadic predicate variable and the odd/even parity of a state corresponds to the least/greatest fixpoint that we seek for the associated variable, etc. One then solves this system of equations via the same inductive procedure used to obtain the formula of the modal μ-calculus from the system associated with an automaton in $Aut(\text{FO}_1)$ (see e.g. [1] for a description of the solution procedure). Because of the **(weakness)** and **(continuity)** conditions on the starting WMSO-automaton $\mathbb{A}$, it is thence possible to verify that the resulting fixpoint formula $\xi_{\mathbb{A}}$ belongs to $\mu_c\text{FOE}^\infty$. $\square$

5. Modal Characterization of WMSO

In this section we prove the main result of this paper, Theorem 1. As mentioned in the introduction, our proof is based on a comparison of $Aut_{cw}(\mathrm{FOE}_1^\infty)$ to the following class of automata corresponding to the fragment $\mu_c\mathrm{ML}$.

Definition 24. *A $\mu_c\mathrm{ML}$-automaton $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ is an automaton in $Aut(\mathrm{FO}_1)$ such that for all states $a, b \in A$ with $a \preceq b$ and $b \preceq a$ the following conditions hold:*

(weakness) $\Omega(a) = \Omega(b)$,
(continuity) *if $\Omega(a)$ is odd (resp. even) then, for each $c \in C$ $\Delta(a, c) \in \mathrm{FO}_1^+\mathrm{C}_b(A)$ (resp. $\Delta(a, c) \in \mathrm{FO}_1^+\overline{\mathrm{C}}_b(A)$).*

We use $Aut_{cw}(\mathrm{FO}_1)$ to denote the class of such automata.

Theorem 5. *There are effective transformations from $\mu_c\mathrm{ML}$ to $Aut_{cw}(\mathrm{FO}_1)$ and vice-versa, witnessing* (5).

Proof. The direction from left to right can be proved by a fairly routine argument, based on the observation that the standard construction of an automaton from a $\mu\mathrm{ML}$-formula transforms $\mu_c\mathrm{ML}$-formulas into automata in $Aut_{cw}(\mathrm{FO}_1)$. In the other direction, the argument is essentially a special case of the one showing Proposition 11. Here the system of equations associated with the automaton $\mathbb{A}$ is expressed in terms of $\mathrm{FO}_1(A)$ formulas and thus can be turned into a system of modal equations. One then *solves* this system of equations in the modal μ-calculus via the standard inductive procedure. The **(weakness)** and **(continuity)** conditions on the strongly connected components of $\mathbb{A}$ ensure that when we execute a step in solving the equations we may work within $\mu_c\mathrm{ML}$. $\square$

We now turn to the proof of Theorem 1. As a key technical result of our paper, in subsection 5.1 we will provide a construction $(\cdot)^\bullet : Aut_{cw}(\mathrm{FOE}_1^\infty) \to Aut_{cw}(\mathrm{FO}_1)$, such that for all $\mathbb{A}$ and $\mathbb{T}$:

$$\mathbb{A}^\bullet \text{ accepts } \mathbb{T} \text{ iff } \mathbb{A} \text{ accepts } \mathbb{T}^\omega, \tag{6}$$

where $\mathbb{T}^\omega$ is the ω-unravelling of $\mathbb{T}$.

Proof of Theorem 1: **(i)** Given a WMSO-formula φ, let $\mathbb{A}_\varphi$ in $Aut_{cw}(\mathrm{FOE}_1^\infty)$ be equivalent to φ, and let $\varphi^\bullet \in \mu_c\mathrm{ML}$ be equivalent to $\mathbb{A}_\varphi$. We verify that φ is bisimulation invariant iff φ and $\varphi^\bullet$ are equivalent. The direction from right to left is immediate by Fact 3 and the observation that $\varphi^\bullet$ is a formula in $\mu\mathrm{ML}$. The opposite direction follows from the following chain of equivalences:

$$\mathbb{T} \models \varphi \text{ iff } \mathbb{T}^\omega \models \varphi \qquad (\varphi \text{ bisimulation invariant})$$
$$\text{iff } \mathbb{A}_\varphi \text{ accepts } \mathbb{T}^\omega \qquad (\varphi \equiv \mathbb{A}_\varphi \text{ on trees})$$
$$\text{iff } \mathbb{A}_\varphi^\bullet \text{ accepts } \mathbb{T} \tag{6}$$
$$\text{iff } \mathbb{T} \Vdash \varphi^\bullet \qquad (\varphi^\bullet \equiv \mathbb{A}_\varphi^\bullet, \text{ Theorem 5})$$

(ii) The second part of Theorem 1 is immediate by Theorem 5, the observation that $Aut_{cw}(\mathrm{FO}_1) \subseteq Aut_{cw}(\mathrm{FOE}_1^\infty)$, and Theorem 2. Alternatively, we can give a direct, truth-preserving translation ST_x from $\mu_c\mathrm{ML}$ to $\mu\mathrm{FOE}^\infty$, of which the key inductive clauses are

- $ST_x(\Diamond \varphi) := \exists y (R(x, y) \wedge ST_y(\varphi))$,
- $ST_x(\mu p. \varphi) := \mu p. ST_x(\varphi)$.

The result then follows by composing ST_x with the translation $(\cdot)_\circ$ of Proposition 10. $\square$

5.1 One-step Translations

In this subsection we will define a construction that transforms an automaton $\mathbb{A}$ in $Aut_{cw}(\mathrm{FOE}_1^\infty)$ into an automaton $\mathbb{A}^\bullet$ in $Aut_{cw}(\mathrm{FO}_1)$, such that $\mathbb{A}$ and $\mathbb{A}^\bullet$ are related as in (6). This construction is completely determined by the following translation at the one-step level.

Definition 25. *Using the fact that by Theorem 3, any sentence in $\mathrm{FOE}_1^{\infty+}(A)$ is equivalent to a disjunction of sentences of the form $\nabla_{\mathrm{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$, we define the following translation $(\cdot)^\bullet : \mathrm{FOE}_1^{\infty+}(A) \to \mathrm{FO}_1^+(A)$. We set*

$$\left(\nabla_{\mathrm{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma) \right)^\bullet := \bigwedge_i \exists x_i. \tau_{T_i}^+(x_i) \wedge \forall x. \bigvee_{S \in \Sigma} \tau_S^+(x)$$

and for $\alpha = \bigvee_i \alpha_i$ we define $\alpha^\bullet := \bigvee \alpha_i^\bullet$.

The key property of this translation is the following.

Proposition 12. *For every one-step model (D, V) and every sentence $\alpha \in \mathrm{FOE}_1^{\infty+}(A)$ we have*

$$(D, V) \models \alpha^\bullet \text{ iff } (D \times \omega, V_\pi) \models \alpha, \tag{7}$$

where V_π is given by $V_\pi(a) := \{(d, k) \mid d \in V(a), k \in \omega\}$.

Proof. Clearly it suffices to prove (7) for sentences of the form $\alpha = \nabla_{\mathrm{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$.

$\boxed{\Rightarrow}$ Assume $(D, V) \models \alpha^\bullet$, we will show that $(D \times \omega, V_\pi) \models \nabla_{\mathrm{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$. Let $d_i \in D$ be such that $\tau_{T_i}^+(d_i)$. It is clear that the (d_i, i) provide *distinct* elements satisfying the first-order existential part of α. The argument for the generalized quantifier part of α is similar. For the universal parts of α it is enough to observe that *every* $d \in D$ realizes a positive type in Σ. The same applies to $(D \times \omega, V_\pi)$, therefore this takes care of both universal quantifiers.

$\boxed{\Leftarrow}$ Assuming that $(D \times \omega, V_\pi) \models \nabla_{\mathrm{FOE}^\infty}^+(\vec{T}, \Pi, \Sigma)$, we show $(D, V) \models \alpha^\bullet$. The existential part of $\alpha^\bullet$ is trivial. For the universal part suppose towards a contradiction that some $d \in D$ is such that $\neg \tau_S^+(d)$ for all $S \in \Sigma$. Then we have $(D \times \omega, V_\pi) \not\models \tau_S^+((d, k))$ for all k. Hence we have $(D \times \omega, V) \not\models \forall^\infty y. \bigvee_{S \in \Sigma} \tau_S^+(y)$. Absurd. $\square$

As a consequence of Proposition 12 we obtain the following.

Definition 26. *Given $\mathbb{A} = \langle A, \Delta, \Omega, a_I \rangle$ in $Aut(\mathrm{FOE}_1^\infty)$, define the automaton $\mathbb{A}^\bullet := \langle A, \Delta^\bullet, \Omega, a_I \rangle$ in $Aut(\mathrm{FO}_1)$ by putting $\Delta^\bullet(a, c) := (\Delta(a, c))^\bullet$ for each $(a, c) \in A \times C$.*

Proposition 13. *For any automaton $\mathbb{A} \in Aut(\mathrm{FOE}_1^\infty)$, and any model $\mathbb{T}$ we have that $\mathbb{A}^\bullet$ accepts $\mathbb{T}$ iff $\mathbb{A}$ accepts $\mathbb{T}^\omega$.*

Proof. The proof is based on a fairly routine comparison of the acceptance games $\mathcal{A}(\mathbb{A}^\bullet, \mathbb{T})$ and $\mathcal{A}(\mathbb{A}, \mathbb{T}^\omega)$. In a slightly more general setting, the details can be found in [16]. $\square$

It only remains to check that the construction $(\cdot)^\bullet$ transforms WMSO-automata into automata of the right class.

Proposition 14. *Let $\mathbb{A}$ be an automaton in $Aut(\mathrm{FOE}_1^\infty)$. If in particular $\mathbb{A} \in Aut_{cw}(\mathrm{FOE}_1^\infty)$, then $\mathbb{A}^\bullet \in Aut_{cw}(\mathrm{FO}_1)$.*

Proof. This proposition can be verified by a straightforward inspection, at the one-step level, that if a sentence $\alpha \in \mathrm{FOE}_1^{\infty+}(A)$ belongs to the fragment $\mathrm{FOE}_1^{\infty+}\mathrm{C}_a(A)$, then its translation $\alpha^\bullet$ lands in the fragment $\mathrm{FO}_1^+\mathrm{C}_a(A)$. $\square$

Remark 2. From Proposition 13 and 14 we can prove:

- $Aut(\mathrm{FO}_1) \equiv Aut(\mathrm{FOE}_1^\infty)/\underline{\leftrightarrow}$, and
- $Aut_{cw}(\mathrm{FO}_1) \equiv Aut_{cw}(\mathrm{FOE}_1^\infty)/\underline{\leftrightarrow}$.

More generally, those equivalences may be obtained for one-step languages $\mathcal{L}_1$ and $\mathcal{L}_1'$ connected by a translation $(\cdot)^\bullet : \mathcal{L}_1' \to \mathcal{L}_1$ satisfying a condition similar to (7) (see [16]).

6. Conclusion

6.1 Overview

In this work we have presented three main contributions. First, we proved that the bisimulation-invariant fragment of WMSO is the fragment μ_cML of μML where the application of the fixpoint operator μp is restricted to formulas that are continuous in p. Our result sheds light on the relationship between MSO, WMSO and μML. In particular, it provides a positive answer to the question whether $\text{WMSO}/\!\leftrightarrow\ \leq$ AFMC on trees of arbitrary branching degree, left open in [5]. This may also be read as the statement that the formulas that separate WMSO from MSO are not bisimulation invariant (and hence, irrelevant in the sense mentioned in the introduction).

To achieve this result, we shaped WMSO-automata, a special kind of parity automata satisfying additional *continuity* and *weakness* conditions, with transition map given by the monadic logic FOE_1^∞. Our second main contribution was to show that they characterize WMSO on tree models.

As our third main contribution we gave a detailed model-theoretic analysis of the monotone and continuous fragments of FOE_1^∞. We provide strong normal forms and syntactic characterizations that, besides being of independent interest, are critical for the development of the aforementioned results.

6.2 Future Work

A first line of research is directly inspired by the methods employed in this work. WMSO-automata and μ_cML-automata are essentially obtained by imposing conditions on the appropriate one-step logic $\mathcal{L}_1$ and transition structure of automata belonging to $Aut(\mathcal{L}_1)$. Following this approach, one could take aim at the automata-theoretic counterpart of other fragments of the modal μ-calculus, like PDL, CTL or CTL*.

Another direction of investigation is based on the observation that, from a topological point of view, all WMSO-definable properties are Borel. Since we do not have examples of Borel MSO-definable properties that are not WMSO-definable, it is tempting to conjecture that WMSO is the Borel fragment of MSO and analogously for μ_cML and μML.

6.3 Acknowledgments

References

[1] A. Arnold and D. Niwiński. *Rudiments of μ-calculus*, volume 146 of *Studies in Logic and the Found. of Mathematics*. North-Holland, 2001.

[2] F. Carreiro, A. Facchini, Y. Venema, and F. Zanasi. Weak MSO: Automata and expressiveness modulo bisimilarity. Technical report, 2014. URL http://arxiv.org/abs/1401.4374.

[3] A. Dawar and M. Otto. Modal characterisation theorems over special classes of frames. *An. Pure App. Logic*, 161(1):1–42, 2009.

[4] E. A. Emerson and C. S. Jutla. Tree automata, μ-calculus and determinacy (extended abstract). In *FOCS*, pages 368–377, 1991.

[5] A. Facchini, Y. Venema, and F. Zanasi. A characterization theorem for the alternation-free fragment of the modal μ-calculus. In *LICS*, pages 478–487, 2013.

[6] G. Fontaine. Continuous fragment of the μ-calculus. In *CSL*, pages 139–153, 2008.

[7] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *LNCS*, 2002. Springer. ISBN 3-540-00388-6.

[8] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ-calculus with respect to monadic second order logic. In *CONCUR*, pages 263–277, 1996.

[9] C. Kissig and Y. Venema. Complementation of coalgebra automata. In *CALCO*, pages 81–96, 2009.

[10] D. Muller and P. Schupp. Alternating automata on infinite trees. *Theoretical Computer Science*, 54:267–276, 1987.

[11] D. E. Muller, A. Saoudi, and P. E. Schupp. Alternating automata, the weak monadic theory of trees and its complexity. *Theor. Comput. Sci.*, 97(2):233–244, 1992.

[12] E. Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6:427–439, 1997.

[13] B. ten Cate and A. Facchini. Characterizing EF over infinite trees and modal logic on transitive graphs. In *MFCS*, pages 290–302, 2011.

[14] J. Väänänen. Remarks on generalized quantifiers and second-order logics. In *Set theory and hierarchy theory*, pages 117–123. Prace Naukowe Instytutu Matematyki Politechniki Wroclawskiej, Wroclaw, 1977.

[15] J. van Benthem. *Modal Correspondence Theory*. PhD thesis, Universiteit van Amsterdam, 1977.

[16] Y. Venema. Expressiveness modulo bisimilarity: a coalgebraic perspective. To appear in: A. Baltag and S. Smets (eds.), *Johan van Benthem, Outstanding Contributions Series*, Springer, 2014.

[17] Y. Venema. Lectures on the modal μ-calculus. Lecture Notes, ILLC, University of Amsterdam, October 2012.

[18] I. Walukiewicz. Monadic second order logic on tree-like structures. In *STACS*, pages 401–413, 1996.

[19] F. Zanasi. Expressiveness of monadic second order logics on infinite trees of arbitrary branching degree. Master's thesis, ILLC, Universiteit van Amsterdam, the Netherlands, 2012.