

Some model theory for the modal μ -calculus: syntactic characterisations of semantic properties *

Gaëlle Fontaine [†] Yde Venema [‡]

Institute for Logic, Language and Computation
Universiteit van Amsterdam
The Netherlands

Abstract

This paper contributes to the theory of the modal μ -calculus by proving some model-theoretic results. More in particular, we discuss a number of semantic properties pertaining to formulas of the modal μ -calculus. For each of these properties we provide a corresponding syntactic fragment, in the sense that a μ -formula ξ has the given property iff it is equivalent to a formula ξ' in the corresponding fragment. Since this formula ξ' will always be effectively obtainable from ξ , as a corollary, for each of the properties under discussion, we prove that it is decidable in elementary time whether a given μ -calculus formula has the property or not.

The properties that we study have in common that they all concern the dependence of the truth of the formula at stake, on a single proposition letter p . In each case the semantic condition on ξ will be that ξ , if true at a certain state in a certain model, will remain true if we restrict the set of states where p holds, to a subset of the state space which is of a special kind, associated with the semantic condition at stake. Important examples include the properties of full and complete additivity and of continuity, where the special subsets are the singletons and the finite sets, respectively.

Our proofs for these characterization results will be automata-theoretic in nature; we will see that the effectively defined maps on formulas are in fact induced by rather simple transformations on modal automata. Thus our results can also be seen as a contribution to the model theory of modal automata.

Keywords. modal logic, μ -calculus, model theory, characterization results, full additivity, Scott continuity, modal automata.

1 Introduction

This paper is inspired by the model-theoretic tradition in logic of linking semantic properties of formulas to syntactic restrictions on their shape. Such correspondences abound in the model theory of classical (propositional or first-order) logic [14]. Well-known preservation results are the Los-Tarski theorem stating that the models of a formula φ are closed under taking submodels iff φ is equivalent to a universal formula, or Lyndon's theorem stating that a formula φ is monotone with respect to the interpretation of a relation symbol R iff φ is equivalent to a formula in which all occurrences of R are

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

[†]gaille2l@gmail.com

[‡]y.venema@uva.nl

positive. In the last example, the semantic property is monotonicity, and the syntactic restriction is positivity.

Our aim here is to establish such correspondences in the setting of the (modal) μ -calculus μML , the extension of modal logic with least and greatest fixpoint operators [8]. Since its introduction by Kozen in the 1980s [35], the modal μ -calculus has found increasing recognition as an important and natural formalism for specifying properties of processes. The main reason for this is that μML , just like basic modal logic, strikes a very favorable balance between expressiveness and computational feasibility. In particular, it was proved by Janin & Walukiewicz [32] that μML is expressively complete for those monadic second-order properties that are bisimulation invariant, so that most, if not all, interesting properties of processes can be specified in the language. On the other hand, despite this large expressive power, the computational complexity of the satisfiability problem for μML can be solved in exponential time [18], which is basically the same as for any extension of modal logic with fixpoint connectives. Other attractive features of the modal μ -calculus include a semantics that can both be presented in a compositional, algebraic format and in intuitive, game-theoretic terms; a natural axiomatization, formulated by Kozen [35] and proven to be complete for the semantics, partially by Kozen himself and fully by Walukiewicz [46]; a tight link with automata theory, established by Janin & Walukiewicz [31] and Wilke [47] via the introduction of modal automata, of non-deterministic respectively alternating type, as automata-theoretic equivalents to μML -formulas; and a certain internal expressive balance, witnessed by the property of uniform interpolation, proved by D'Agostino & Hollenberg [17].

With all these positive results, the modal μ -calculus has become the canonical modal process logic, and it seems worth while to develop its model theory in full detail. Some results are known: in particular, preservation results, similar to the Los-Tarski and Lyndon theorems, have been shown for the μ -calculus by D'Agostino & Hollenberg [17]. However, in the intended semantics of μML , where models represent computational processes, and accessibility relations, bisimulations, and trees play an important role, there are some specific properties of interest that have not been studied in classical model theory. Important examples that we will study here include the properties of *full* and *complete additivity* and of *continuity* with respect to some fixed propositional variable p .

Full and complete additivity To define the properties of *full* and *complete additivity*, recall that in each Kripke model $\mathbb{S} = (S, R, V)$ we may formalize the dependence of the meaning of a μ -formula ξ in \mathbb{S} on a fixed proposition letter p as a map

$$\xi_p^{\mathbb{S}} : \mathcal{P}S \rightarrow \mathcal{P}S,$$

defined by $\xi_p^{\mathbb{S}}(X) := \{s \in S \mid \mathbb{S}[p \mapsto X] \Vdash \xi\}$, where $\mathbb{S}[p \mapsto X] = (S, R, V[p \mapsto X])$ is the model obtained from \mathbb{S} by modifying V so that $V(p) = X$. Then a formula ξ is fully additive in p if for each model \mathbb{S} , the operation $\xi_p^{\mathbb{S}}$ distributes over arbitrary unions:

$$\xi_p^{\mathbb{S}}\left(\bigcup \mathcal{X}\right) = \bigcup \left\{ \xi_p^{\mathbb{S}}(X) \mid X \in \mathcal{X} \right\},$$

for any collection \mathcal{X} of subsets of S , whereas we say that ξ is completely additive if $\xi_p^{\mathbb{S}}$ distributes over arbitrary *non-empty* unions. Clearly, the difference between the two notions is that for full additivity we require the map $\xi_p^{\mathbb{S}}$ to be *normal*, that is, $\xi_p^{\mathbb{S}}(\bigcup \emptyset) = \bigcup \emptyset$, or equivalently, $\xi_p^{\mathbb{S}}(\emptyset) = \emptyset$.

It is not difficult to see that full additivity is equivalent to requiring that for all $X \subseteq S$,

$$\xi_p^{\mathbb{S}}(X) = \bigcup \left\{ \xi_p^{\mathbb{S}}(\{x\}) \mid x \in X \right\},$$

or, putting it yet differently, to the following constraint, for every pointed Kripke model (\mathbb{S}, s) :

$$\mathbb{S}, s \Vdash \xi \text{ iff } \mathbb{S}[p \mapsto \{x\}], s \Vdash \xi, \text{ for some } x \in V(p).$$

This is a very natural property in the context of modal logic, since for any set S , there is a 1–1 correspondence between the fully additive functions on $\mathcal{P}S$ and the binary (accessibility) relations on S , the relation associated with the map f given as $R_f := \{(s, s') \mid s \in f(\{s'\})\}$. Related to this, a more specific reason for studying full additivity is given by its pivotal role in the characterization of the fragments of first- and monadic second-order logic of formulas that are *safe for bisimulations* (for a detailed discussion of this notion and its relation to full additivity we refer to section 9). Syntactic characterizations of the formulas that are fully additive in a given proposition letter p , were obtained by van Benthem [4] in the setting of basic (i.e., fixpoint-free) modal logic and by Hollenberg [30] in the setting of the modal μ -calculus. As an alternative to Hollenberg’s result, we shall give a different syntactic fragment characterizing full additivity. More precisely, we will prove that a μML -formula is fully additive in p if it is equivalent to a formula in the fragment $\mu\text{ML}_{\{p\}}^A$, where we define, for a set P of propositional variables, the fragment μML_P^A of the modal μ -calculus by the following grammar:

$$\varphi ::= p \mid \varphi \vee \varphi \mid \varphi \wedge \psi \mid \diamond\varphi \mid \mu x.\varphi'.$$

Here p belongs to P , ψ is a P -free formula (i.e., without free occurrences of proposition letters in P), and the formula φ' belongs to the fragment $\mu\text{ML}_{P \cup \{x\}}^A$. For the property of complete additivity we obtain a very similar characterization.

Continuity Another important property featuring in this study is that of *continuity*. We shall call a formula ξ *continuous* in a proposition letter p if for all \mathbb{S} , and all $X \subseteq S$,

$$\xi_p^{\mathbb{S}}(X) = \bigcup \left\{ \xi_p^{\mathbb{S}}(F) \mid F \subseteq_{\omega} X \right\}.$$

That is, where the meaning of $\xi_p(X)$ in the case of ξ being fully additive in p depends on the singleton subsets of X , here the meaning of $\xi_p(X)$ depends on the *finite* subsets of X .

What explains both the name and the importance of this property, is its equivalence to *Scott continuity*: for any model \mathbb{S} , the map $\xi_p^{\mathbb{S}}$ is continuous (in our sense) iff it is continuous with respect to the Scott topology on the powerset algebra. Scott continuity is of key importance in many areas of theoretical computer sciences where ordered structures play a role, such as domain theory (see, e.g., [1]). Another motivation concerns the relation between continuity and another property of computational interest, *constructivity*. A monotone formula ξ is constructive in a proposition letter p if for each model \mathbb{S} , the least fixpoint of the map $\xi_p^{\mathbb{S}}$ is reached in at most ω approximation steps. Locally, this means that a state satisfies a least fixpoint formula if it satisfies one of its finite approximants. While the exact relation between the two properties is not clearly understood, it is well-known that continuity strictly implies constructivity, and we believe that in a sense continuity can be considered as the most natural property to approximate constructivity syntactically. A full discussion of the notion of continuity and its relation with constructivity can be found in section 8

As one of the main results of this paper we will show that a μ -formula ξ is continuous in p iff it is equivalent to a formula φ in the syntactic fragment μML_P^C , where we define, for any set P of propositional variables, the set μML_P^C by a mutual induction based on the following grammar:

$$\varphi ::= p \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond\varphi \mid \mu x.\varphi',$$

where ψ is a P -free μ -formula, and φ' belongs to the fragment $\mu\text{ML}_{P \cup \{x\}}^A$. A first presentation of this result was given by the first author in [24]; here we will give an alternative and more insightful proof of this result.

Aim The general purpose of this paper is to prove, in a uniform framework, syntactic characterization results corresponding to a number of semantic properties, including full and complete additivity, and continuity. What these properties have in common is that they all concern the dependence of the truth of the formula at stake, on a single proposition letter — it will be convenient to fix this letter from now on as ‘ p ’. More precisely, given a Kripke model $\mathbb{S} = (S, R, V)$ (which in some cases we require to be a tree), let $\mathcal{X}_{\mathbb{S}}$ uniformly denote a certain class of subsets of S , such as singletons, finite sets, paths, finitely branching subtrees, etc. Let ξ be a formula of the modal μ -calculus. Assuming that ξ is monotone in p , we say that it has the \mathcal{X} -property with respect to p if for every pointed model (\mathbb{S}, s) we have

$$\mathbb{S}, s \models \xi \text{ iff } \mathbb{S}[p \upharpoonright X], s \models \xi \text{ for some } X \in \mathcal{X}_{\mathbb{S}}, \quad (1)$$

where $\mathbb{S}[p \upharpoonright X] := \mathbb{S}[p \mapsto V(p) \cap X]$ is the model obtained from \mathbb{S} by *restricting* p to the set X . In the examples that we shall consider, $\mathcal{X}_{\mathbb{S}}$ consists of the following sets¹:

- (a) singletons inside $V(p)$, leading to the property of full additivity;
- (b) singletons, leading to the property of *complete additivity*²;
- (c) finite sets, leading to the property of continuity;
- (d) finitely branching subtrees, leading to the *finite width* property;
- (e) noetherian subtrees (i.e., without infinite paths), leading to the *finite depth* property;
- (f) branches, leading to the *single branch* property.

Clearly, there are some interesting relations between some of these properties. For instance, both full and complete additivity imply continuity. Also, continuity can be seen as the combination of a ‘horizontal’ and a ‘vertical’ component: the finite width property and the finite depth property, respectively. The latter equivalence will be put to good use in the paper.

The above abstract presentation allows us to summarize our results in a concise and uniform manner. Basically, for each instance (a–f) of \mathcal{X} we present a *syntactic characterization* of the \mathcal{X} -property, in the form of a syntactically defined fragment $\mu\text{ML}_p^X \subseteq \mu\text{ML}$ such that a μ -formula has the \mathcal{X} -property with respect to p iff it is equivalent to a formula in μML_p^X . Since monotonicity forms part of the definition of each of the properties (a–f), to facilitate the other proofs we first prove a slightly stronger version of D’Agostino & Hollenberg’s Lyndon theorem; in the remainder of this introduction it will be convenient to let \mathcal{X} also cover the property of monotonicity.

Main results & proof method Our proofs, though different in each case, follow a uniform method, which goes back to the proofs of Janin & Walukiewicz [32] and D’Agostino & Hollenberg [17]. For each property \mathcal{X} , we will exhibit an explicit translation which, given a μML -formula ξ , computes a formula $\xi^X \in \mu\text{ML}_p^X$ such that

$$\xi \text{ has the } \mathcal{X}\text{-property iff } \xi \text{ is equivalent to } \xi^X. \quad (2)$$

Since in each case the translation $(\cdot)^X$ is effectively computable, and the equivalence problem for two given μML -formulas is decidable, as a nice corollary of (2) we obtain various new *decidability* results. With \mathcal{X} -PROP being the problem whether a given μML -formula has the \mathcal{X} -property, for each \mathcal{X} as discussed we will show that

$$\text{the problem } \mathcal{X}\text{-PROP is decidable.} \quad (3)$$

Considering that our main interest here is model-theoretic, we have not undertaken an in-depth study of the computational *complexity* of the \mathcal{X} -PROP problems. What we can say is that each of our translations $(\cdot)^X$ will be based on the composition of simple transformations, each of which constructs

¹There are a few subtleties here. For instance the property might only make sense when we investigate (1) on certain (tree) models, we might want the set \mathcal{X} to consist of subsets of $V(p)$ or not, etc.

²In the formulation (1), complete additivity differs from full additivity in that we do not require the singletons in \mathcal{X} to be subsets of $V(p)$.

an output structure of size at most exponential in the size of the input. From this it follows that all of the \mathcal{X} -PROP problems that we study can be solved in *elementary* time.

While in almost all cases pure logic-based proofs for our results are possible (and have been given in the dissertation of the first author [25]), the proofs we provide in this paper are automata-theoretic in nature. Automata for the modal μ -calculus were introduced by Janin & Walukiewicz [31] under the name of *μ -automata*; where these devices are non-deterministic in nature, Wilke [47] came up with an alternating variant. We will use both alternating and non-deterministic devices here, under the names of, respectively, *modal* and *disjunctive modal* automata. The particular shape of our automata is logic-based: the transition map of our structures maps states of the automata to so-called *one-step formulas*, and many of our proofs are based on syntactic manipulations on these very simple modal formulas. In some sense then, our paper is also a contribution to the model theory of modal automata.

Finally, we have formulated our results in the setting of the *mono-modal* μ -calculus; that is, the language that we consider has one diamond only, and correspondingly the Kripke models have only one accessibility relation. This restriction is solely for the purpose of simplifying the presentation of our results and proofs. We want to stress, however, that all of the results in this paper can be generalized to the setting of the polymodal μ -calculus, with no conceptual and little technical complication.

Overview The paper is organized as follows. In order to fix our terminology and notation, we give a review of the syntax and semantics of the modal μ -calculus in the following section, and in section 3 we introduce the modal automata that we will be working with. In section 4 we make some first model-theoretic steps, proving the characterization of monotonicity, and introducing an automata-theoretic construction that will be used in the other parts of the paper. In the subsequent three sections we discuss the finite width property (section 5), the single branch property (section 7) and the finite depth property (section 6). After that we arrive the most important parts of the paper, viz., section 8 on continuity and section 9 on full and complete additivity. In the final section of the paper we draw some conclusions, discuss some related results, and list some open problems for future research.

Acknowledgement We are deeply indebted to the two referees for many helpful comments on an earlier version of this paper.

2 Preliminaries

In this paper we assume familiarity with the syntax and semantics of the modal μ -calculus, as presented in for instance [35, 2, 28, 8, 45], and with the basic notions concerning infinite games [28]. Here we fix some notation and terminology.

Convention 2.1 Throughout the text we fix an infinite set PROP of propositional variables, of which we often single out a finite subset X .

2.1 Parity games

Definition 2.2 A *parity game* is a tuple $\mathbb{G} = (G_{\exists}, G_{\forall}, E, \Omega)$ where G_{\exists} and G_{\forall} are disjoint sets, and, with $G := G_{\exists} \cup G_{\forall}$ denoting the *board* of the game, the binary relation $E \subseteq G^2$ encodes the moves that are *admissible* to the respective players, and the *priority function* $\Omega : G \rightarrow \omega$, which is required to be of finite range, determines the *winning condition* of the game. Elements of G_{\exists} and G_{\forall} are called *positions* for the players \exists and \forall , respectively; given a position p for player $\Pi \in \{\exists, \forall\}$, the set $E[p]$ denotes the set of *moves* that are *legitimate* or *admissible* to Π at p . In case $E[p] = \emptyset$ we say that player Π *gets stuck* at p .

An *initialized board game* is a pair consisting of a board game \mathbb{G} and a *initial* position p , usually denoted as $\mathbb{G}@p$. \triangleleft

Definition 2.3 A *match* of a graph game $\mathbb{G} = (G_{\exists}, G_{\forall}, E, \Omega)$ is a (finite or infinite) path through the graph (G, E) . Such a match Σ is called *partial* if it is finite and $E[\text{last}\Sigma] \neq \emptyset$, and *full* otherwise. We let PM_{Π} denote the collection of partial matches Σ ending in a position $\text{last}(\Sigma) \in G_{\Pi}$, and define $\text{PM}_{\Pi}@p$ as the set of partial matches in PM_{Π} starting at position p .

The *winner* of a full match Σ is determined as follows. If Σ is finite, it means that one of the two players got stuck at the position $\text{last}(\Sigma)$, and so this player loses Σ , while the opponent wins. If $\Sigma = (p_n)_{n \in \omega}$ is infinite, we declare its winner to be \exists if the maximum value occurring infinitely often in the stream $(\Omega p_n)_{n \in \omega}$ is even. \triangleleft

Definition 2.4 A *strategy* for a player $\Pi \in \{\exists, \forall\}$ is a map $f : \text{PM}_{\Pi} \rightarrow G$. A strategy is *positional* if it only depends on the last position of a partial match, i.e., if $f(\Sigma) = f(\Sigma')$ whenever $\text{last}(\Sigma) = \text{last}(\Sigma')$; such a strategy can and will be presented as a map $f : G_{\Pi} \rightarrow G$.

A match $\Sigma = (p_i)_{i < \kappa}$ is *guided* by a Π -strategy f if $f(p_0 p_1 \dots p_{n-1}) = p_n$ for all $n < \kappa$ such that $p_0 \dots p_{n-1} \in \text{PM}_{\Pi}$. A position is *reachable* by a strategy f if there is an f -guided match Σ of which p is the last position. A Π -strategy f is *legitimate* in $\mathbb{G}@p$ if the moves that it prescribes to f -guided partial matches in $\text{PM}_{\Pi}@p$ are always admissible to Π , and *winning for Π* in $\mathbb{G}@p$ if in addition all f -guided full matches starting at p are won by Π .

A position p is a *winning position* for player $\Pi \in \{\exists, \forall\}$ if Π has a winning strategy in the game $\mathbb{G}@p$; the set of these positions is denoted as Win_{Π} . The game $\mathbb{G} = (G_{\exists}, G_{\forall}, E, \Omega)$ is *determined* if every position is winning for either \exists or \forall . \triangleleft

When defining a strategy f for one of the players in a board game, we can and in practice will confine ourselves to defining f for partial matches that are themselves guided by f .

The following fact, independently due to Emerson & Jutla [19] and Mostowski [39], will be quite useful to us.

Fact 2.5 (Positional Determinacy) *Let $\mathbb{G} = (G_{\exists}, G_{\forall}, E, \Omega)$ be a parity game. Then \mathbb{G} is determined, and both players have positional winning strategies.*

In the sequel we will often refer to a ‘positional winning strategy’ for one of the players in a parity game. With this we mean any positional strategy which is winning for that player when starting at any of his/her winning positions.

2.2 Structures

Definition 2.6 Given a set S , an *A-marking* on S is a map $m : S \rightarrow \text{PA}$; an *A-valuation* on S is a map $V : A \rightarrow \text{PS}$. Any valuation $V : A \rightarrow \text{PS}$ gives rise to its *transpose marking* $V^{\dagger} : S \rightarrow \text{PA}$ defined by $V^{\dagger}(s) := \{a \in A \mid s \in V(a)\}$, and dually each marking gives rise to a valuation in the same manner. \triangleleft

Since markings and valuations are interchangeable notions, we will often switch from one perspective to the other, based on what is more convenient in context.

Definition 2.7 A *Kripke structure* over a set X of proposition letters is a triple $\mathbb{S} = (S, R, V)$ such that S is a set of objects called *points*, $R \subseteq S \times S$ is a binary relation called the *accessibility* relation, and V is an X -valuation on S . A *pointed Kripke structure* is a pair (\mathbb{S}, s) where s is a point of \mathbb{S} .

Given a Kripke structure $\mathbb{S} = (S, R, V)$, a propositional variable x and a subset U of S , we define $V[x \mapsto U]$ as the $X \cup \{x\}$ -valuation given by

$$V[x \mapsto U](p) := \begin{cases} V(p) & \text{if } p \neq x \\ U & \text{otherwise,} \end{cases}$$

and we let $\mathbb{S}[x \mapsto U]$ denote the structure $(S, R, V[x \mapsto U])$. The structure $\mathbb{S}[x \mapsto V(x) \cap U]$ is usually denoted as $\mathbb{S}[x|U]$. \triangleleft

It will often be convenient to take a coalgebraic perspective on Kripke structures.

Definition 2.8 Given a Kripke structure $\mathbb{S} = (S, R, V)$ over the set X , we define its (*coalgebraic*) *unfolding map* $\sigma_{\mathbb{S}} : S \rightarrow \text{PX} \times \text{PS}$ given by $\sigma_{\mathbb{S}}(s) = (\sigma_V(s), \sigma_R(s))$, where $\sigma_V(s) := V^\dagger(s)$ and $\sigma_R(s) := R[s]$ are the sets of, respectively, the proposition letters true at s and the successors of s . We will write σ rather than $\sigma_{\mathbb{S}}$ in case no confusion is likely. \triangleleft

Definition 2.9 A *path* through a Kripke structure $\mathbb{S} = (S, R, V)$ is a sequence $(s_i)_{i < \kappa}$ such that $(s_i, s_{i+1}) \in R$ for all i with $i + 1 < \kappa$; here $\kappa \leq \omega$ is the *length* of the path. We let \sqsubseteq denote the prefix (initial segment) relation between paths, and use \sqsubset for the strict (irreflexive) version of \sqsubseteq . \triangleleft

Definition 2.10 Given two models $\mathbb{S} = (S, R, V)$ and $\mathbb{S}' = (S', R', V')$, a relation $Z \subseteq S \times S'$ is a *bisimulation* if it satisfies, for all $(s, s') \in Z$, the conditions

- (prop) $s \in V(p)$ iff $s' \in V'(p)$, for all $p \in X$;
- (forth) for all $t \in \sigma_R(s)$ there is a $t' \in \sigma_{R'}(s')$ with Ztt' ; and
- (back) for all $t' \in \sigma_{R'}(s')$ there is a $t \in \sigma_R(s)$ with Ztt' .

We say that s and s' are *bisimilar*, notation $\mathbb{S}, s \Leftrightarrow \mathbb{S}', s'$ if there is some bisimulation Z with Zss' . A function $f : S \rightarrow S'$ is a *bounded morphism* from \mathbb{S} to \mathbb{S}' , notation $f : \mathbb{S} \rightarrow \mathbb{S}'$, if its graph $\{(s, f(s)) \mid s \in S\}$ is a bisimulation. \triangleleft

Definition 2.11 The reflexive/transitive closure and the transitive closure of R are denoted as R^* and R^+ , respectively; elements of the sets $R^*[s]$ and $R^+[s]$ are called *descendants* and *proper descendants* of s , respectively.

A pointed structure (\mathbb{S}, s) is a *tree* (with *root* s) if $S = R^*[s]$ and every state $t \neq s$ has a unique predecessor. A *branch* of a tree (\mathbb{S}, s) is a maximal path through \mathbb{S} , starting at the root. In a tree model \mathbb{S} , a set $U \subseteq S$ is *downward closed* if for all $s \in U$, the predecessors of s belongs to U . A *sibling* of a node t in a tree is a node $t' \neq t$ with the same predecessor as t .

Given a number $\kappa \leq \omega$, a tree is κ -*expanded* if every node (apart from the root) has $\kappa - 1$ many bisimilar siblings (where $\omega - 1 = \omega$). Given a pointed structure (\mathbb{S}, s) , its κ -*expansion* is the structure $\mathbb{S}_s^\kappa := (S', R', V')$, where S' is the set of all finite sequences $s_0 k_1 s_1 \dots k_n s_n$ ($n \geq 0$) such that $s_0 = s$ and $k_i < \kappa$, $s_i \in S$ and $s_{i-1} R s_i$ for all $i > 0$; where $R' := \{(s_0 k_1 s_1 \dots k_n s_n, s_0 k_1 s_1 \dots k_n s_n k t) \mid k < \kappa \text{ and } s_n R t\}$; and where $V'(p) := \{s_0 k_1 s_1 \dots k_n s_n \mid s_n \in V(p)\}$. The 1-expansion of a pointed model (\mathbb{S}, s) is also called its *unravelling*. \triangleleft

Fact 2.12 Fix an ordinal $\kappa \leq \omega$. Given a pointed model (\mathbb{S}, s) , the structure (\mathbb{S}_s^κ, s) is a κ -*expanded tree* which is bisimilar to (\mathbb{S}, s) via the canonical bounded morphism mapping a path $s_0 k_1 s_1 \dots k_n s_n$ to its last element s_n .

2.3 Syntax

Definition 2.13 The language μML of the modal μ -calculus is given by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \diamond\varphi \mid \mu x.\varphi, \tag{4}$$

where p and x are propositional variables, and the formation of the formula $\mu x.\varphi$ is subject to the constraint that the variable x is *positive* in φ , i.e., all occurrences of x in φ are in the scope of an even number of negations. Elements of μML will be called *modal fixpoint formulas*, μ -*formulas*, or simply *formulas*. \triangleleft

We will often make the assumption that our formulas are in negation normal form.

Definition 2.14 A formula of the modal μ -calculus is in *negation normal form* if it belongs to the language given by the following grammar:

$$\varphi ::= p \mid \neg p \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond \varphi \mid \square \varphi \mid \mu x.\varphi \mid \nu x.\varphi,$$

where p and x are propositional variables, and the formation of the formulas $\mu x.\varphi$ and $\nu x.\varphi$ is subject to the constraint that the variable x is *positive* in φ . We use the symbol η to range over μ and ν . \triangleleft

Convention 2.15 In order to increase readability by reducing the number of brackets, we adopt some standard scope conventions. We let the unary (propositional and modal) connectives, \neg , \diamond and \square , bind stronger than the binary propositional connectives \wedge , \vee and \rightarrow , and use associativity to the left for the connectives \wedge and \vee . Furthermore, we use ‘dot notation’ to indicate that the fixpoint operators preceding the dot have maximal scope. For instance, $\mu x.\neg p \vee \diamond x \vee \nu y.q \wedge \square y$ stands for $\mu x.\left(\left(\neg p\right) \vee \left(\diamond x\right)\right) \vee \nu y.\left(q \wedge \square y\right)$

We gather some definitions pertaining to formulas.

Definition 2.16 We let $Sfor(\xi)$ denote the collection of subformulas of a formula ξ , defined as usual, and we write $\varphi \trianglelefteq \xi$ if φ is a subformula of ξ . The *size* of ξ is defined as its number of subformulas, $|\varphi| := |Sfor(\xi)|$. We write $BV(\xi)$ and $FV(\xi)$ for, respectively, the set of *bound* and *free variables* of a formula ξ . We let $\mu\text{ML}(\mathbf{X})$ denote the set of μ -formulas of which all free variables belong to the set \mathbf{X} .

A μ -formula ξ is *well-named* if $BV(\xi) \cap FV(\xi) = \emptyset$, and with every bound variable x of ξ we may associate a unique subformula of the form $\eta x.\delta$ (with $\eta \in \{\mu, \nu\}$). This unique subformula will be denoted as $\eta_x.\delta_x$, and we call x a μ -variable if $\eta_x = \mu$, and a ν -variable if $\eta_x = \nu$. \triangleleft

As a convention, the free variables of a formula φ are denoted by the symbols p, q, r, \dots , and referred to as *proposition letters*, while we use the symbols x, y, z, \dots for the bound variables of a formula.

Definition 2.17 Let φ and $\{\psi_z \mid z \in Z\}$ be modal fixpoint formulas, where $Z \cap BV(\varphi) = \emptyset$. Then we let

$$\varphi[\psi_z/z \mid z \in Z]$$

denote the formula obtained from φ by simultaneously substituting each formula ψ_z for z in φ (with the usual understanding that no free variable in any of the ψ_z will get bound by doing so). In case Z is a singleton z , we will simply write $\varphi[\psi_z/z]$, or $\varphi[\psi]$ if z is clear from context. \triangleleft

Definition 2.18 Let ξ be a well-named μ -formula. The *dependency order* $<_\xi$ on the bound variables of ξ is defined as the least strict partial order such that $x <_\xi y$ if δ_x is a subformula of δ_y .

We define a map $Act : Sfor(\xi) \rightarrow \mathcal{P}(FV(\xi))$ assigning to each subformula $\varphi \trianglelefteq \xi$ the least set $Act(\varphi)$ such that $Act(p) = \{p\}$ if $p \in FV(\xi)$, $Act(\varphi * \psi) = Act(\varphi) \cup Act(\psi)$ if $*$ $\in \{\wedge, \vee\}$, $Act(\heartsuit \varphi) = Act(\varphi)$ if $\heartsuit \in \{\diamond, \square\}$, $Act(\eta x.\varphi) = Act(\varphi)$ if $\eta \in \{\mu, \nu\}$, and $Act(x) = Act(\delta_x)$ if $x \in BV(\xi)$. If $p \in Act(\varphi)$ we say that p is *active* in φ (relative to ξ). \triangleleft

Position z	Player	Admissible moves $E(z)$
(p, s) with $p \in FV(\xi)$ and $s \in V(p)$	\forall	\emptyset
(p, s) with $p \in FV(\xi)$ and $s \notin V(p)$	\exists	\emptyset
$(\neg p, s)$ with $p \in FV(\xi)$ and $s \in V(p)$	\exists	\emptyset
$(\neg p, s)$ with $p \in FV(\xi)$ and $s \notin V(p)$	\forall	\emptyset
(x, s) with $x \in BV(\xi)$	-	$\{(\delta_x, s)\}$
$(\varphi \vee \psi, s)$	\exists	$\{(\varphi, s), (\psi, s)\}$
$(\varphi \wedge \psi, s)$	\forall	$\{(\varphi, s), (\psi, s)\}$
$(\eta x.\varphi, s)$	-	$\{(\varphi, s)\}$
$(\diamond\varphi, s)$	\exists	$\{(\varphi, t) \mid sRt\}$
$(\square\varphi, s)$	\forall	$\{(\varphi, t) \mid sRt\}$

Table 1: The evaluation game $\mathcal{E}(\xi, \mathbb{S})$

2.4 Semantics

Definition 2.19 By induction on the complexity of modal fixpoint formulas, we define a meaning function $\llbracket \cdot \rrbracket$, which assigns to a formula $\varphi \in \mu\text{ML}(\mathbf{X})$ its *meaning* $\llbracket \varphi \rrbracket^{\mathbb{S}} \subseteq S$ in any Kripke structure $\mathbb{S} = (S, R, V)$ over \mathbf{X} . The clauses of this definition are standard:

$$\begin{aligned}
\llbracket p \rrbracket^{\mathbb{S}} &:= V(p) \\
\llbracket \neg\varphi \rrbracket^{\mathbb{S}} &:= S \setminus \llbracket \varphi \rrbracket^{\mathbb{S}} \\
\llbracket \varphi \vee \psi \rrbracket^{\mathbb{S}} &:= \llbracket \varphi \rrbracket^{\mathbb{S}} \cup \llbracket \psi \rrbracket^{\mathbb{S}} \\
\llbracket \diamond\varphi \rrbracket^{\mathbb{S}} &:= \{s \in S \mid R[s] \cap \llbracket \varphi \rrbracket^{\mathbb{S}} \neq \emptyset\} \\
\llbracket \mu x.\varphi \rrbracket^{\mathbb{S}} &:= \bigcap \{U \in \text{PS} \mid \llbracket \varphi \rrbracket^{\mathbb{S}[x \mapsto U]} \subseteq U\}.
\end{aligned}$$

If a point $s \in S$ belongs to the set $\llbracket \varphi \rrbracket^{\mathbb{S}}$, we write $\mathbb{S}, s \Vdash \varphi$, and say that φ is *true at s* or *holds at s* , or that s *satisfies* φ . Two formulas φ and ψ are *equivalent*, notation: $\varphi \equiv \psi$, if $\llbracket \varphi \rrbracket^{\mathbb{S}} = \llbracket \psi \rrbracket^{\mathbb{S}}$ for any structure \mathbb{S} . \triangleleft

Throughout this paper we will rely on the *bisimulation invariance* of the modal μ -calculus.

Fact 2.20 Let (\mathbb{S}, s) and (\mathbb{S}', s') be pointed Kripke structures such that $(\mathbb{S}, s) \Leftrightarrow (\mathbb{S}', s')$. Then $\mathbb{S}, s \Vdash \varphi$ iff $\mathbb{S}', s' \Vdash \varphi$, for all $\varphi \in \mu\text{ML}$.

We will usually take a *game-theoretic* perspective on the semantics of the modal μ -calculus.

Definition 2.21 Let $\mathbb{S} = (S, R, V)$ be a Kripke model and let ξ be a formula in μML . We define the *evaluation game* $\mathcal{E}(\xi, \mathbb{S})$ as the parity game (G, E, Ω) of which the board consists of the set $Sfor(\xi) \times S$, and the game graph (i.e., the partitioning of $Sfor(\xi) \times S$ into positions for the two players, together with the set $E(z)$ of admissible moves at each position z), is given in Table 1. Note that we do not assign a player to positions that admit a single move only.

To define the priority map Ω of $\mathcal{E}(\xi, \mathbb{S})$, consider an infinite match $\Sigma = (\varphi_n, s_n)_{n \in \omega}$, and let $Inf(\Sigma)$ denote the set of (bound) variables that get unfolded infinitely often during the match. This set contains a highest variable x (with respect to the dependency order $<_{\xi}$), and the winner of Σ is \exists if $\eta_x = \mu$, and \forall if $\eta_x = \nu$. It is not difficult to define a priority map $\Omega : (Sfor \times S) \rightarrow \omega$ that is compatible with this condition, but we do not need the details of the precise definition. \triangleleft

The following fact states the *adequacy* of the game semantics.

Fact 2.22 Let ξ be a well-named formula of the modal μ -calculus, and let (\mathbb{S}, s) be some pointed Kripke structure. Then $\mathbb{S}, s \Vdash \xi$ iff $(\xi, s) \in \text{Win}_{\exists}(\mathcal{E}(\xi, \mathbb{S}))$.

Finally, for our decidability results we will use the following result, due to Emerson & Jutla [18].

Fact 2.23 The question whether a given μML -formula ξ is satisfiable can be decided in time exponential in the size of ξ . As a corollary, the question whether two μML -formula ξ and ξ' are equivalent is decidable in time exponential in the sum of the sizes of ξ and ξ' .

Since this paper is about various syntactic fragments of the modal μ -calculus, we gather here some specific formulas that will be discussed as (non-)instances of these fragments throughout the paper.

Example 2.24 Below we introduce a number of formulas, $\varphi_0, \dots, \varphi_6$, and for each listed formula we explain in words its meaning in an arbitrary tree model \mathbb{S} at the root r .

- $\varphi_0 := p$ simply states that p holds at r ;
- $\varphi_1 := \nu y.(q \wedge \diamond y)$ expresses the existence of an infinite q -path at r ;
- $\varphi_2 := p \wedge \nu y.(q \wedge \diamond y)$ is the conjunction of φ_0 and φ_1 ;
- $\varphi_3 := \mu x.p \vee \diamond \diamond x$ expresses the existence of a p -state at an even distance of r ;
- $\varphi_4 := \mu x.p \vee \square \square x$ states that every path starting at r has a p -state at even distance from r ;
- $\varphi_5 := \nu y.\mu x.((p \wedge \diamond y) \vee \diamond x)$ expresses the existence of a path with infinitely many p -nodes;
- $\varphi_6 := \mu x.(p \vee (\diamond(q \wedge x) \wedge \diamond(\neg q \wedge x)))$ expresses the existence of a finite binary subtree rooted at r , of which all leaves satisfy p , and all inner nodes have both a q - and a $\neg q$ -successor. \triangleleft

3 Modal automata

In this section we give a detailed introduction to the automata that feature in our proofs. Basically, we will be working with the guarded specimens of Wilke's alternating tree automata [47], and with the non-deterministic versions of these which correspond to the μ -automata of Janin & Walukiewicz [32]. As mentioned in the introduction, many of our proofs will be based on manipulating the one-step formulas featuring as the co-domain of the transition map of a modal automaton.

3.1 One-step logic

The transition map of our modal automata will be based on a so-called *modal one-step language*, consisting of modal formulas of rank 1, built up from *proposition letters* (which may have negative occurrences but must appear unguarded) and *variables* (which must appear both guarded and positively).

Definition 3.1 Given a set P , we define the set $\text{Latt}(P)$ of *lattice terms* over P through the following grammar:

$$\pi ::= p \mid \perp \mid \top \mid \pi \wedge \pi \mid \pi \vee \pi,$$

where $p \in P$. Given two sets \mathbf{X} and A , we define the set $\mathbf{1ML}(\mathbf{X}, A)$ of *modal one-step formulas* over A with respect to \mathbf{X} inductively by

$$\alpha ::= p \mid \neg p \mid \diamond \pi \mid \square \pi \mid \perp \mid \top \mid \alpha \wedge \alpha \mid \alpha \vee \alpha,$$

with $p \in \mathbf{X}$ and $\pi \in \text{Latt}(A)$. \triangleleft

These formulas are naturally interpreted in so-called one-step models.

Definition 3.2 Fix sets \mathbf{X} and A . A *one-step frame* is a pair (\mathbf{Y}, S) where S is any set, and $\mathbf{Y} \subseteq \mathbf{X}$. A *one-step model* is a triple (\mathbf{Y}, S, m) such that (\mathbf{Y}, S) is a one-step frame and $m : S \rightarrow \mathcal{P}A$ is an *A-marking* on S . \triangleleft

Observe that with this definition, the coalgebraic representation of a Kripke structure (S, R, V) can now be seen as a function $\sigma_{\mathbb{S}}$ mapping any state $s \in S$ to a one-step frame of which the carrier is a subset of S . This plays a fundamental role in the acceptance game for modal automata, and explains the following *one-step satisfaction relation* \Vdash^1 between one-step models and one-step formulas.

Definition 3.3 Fix a one-step model (\mathbf{Y}, S, m) . First we define the value $\llbracket \pi \rrbracket \subseteq S$ of a lattice formula π over A by induction, setting $\llbracket a \rrbracket := \{s \in S \mid a \in m(s)\}$ for $a \in A$, and treating the boolean connectives in the obvious manner.

The *one-step satisfaction relation* \Vdash^1 is inductively defined as follows. For the literals and modal operators we set:

- $(\mathbf{Y}, S, m) \Vdash^1 p$ iff $p \in \mathbf{Y}$,
- $(\mathbf{Y}, S, m) \Vdash^1 \neg p$ iff $p \notin \mathbf{Y}$,
- $(\mathbf{Y}, S, m) \Vdash^1 \Box \pi$ iff $\llbracket \pi \rrbracket = S$,
- $(\mathbf{Y}, S, m) \Vdash^1 \Diamond \pi$ iff $\llbracket \pi \rrbracket \neq \emptyset$,

while we have the standard clauses for the boolean connectives. \triangleleft

3.2 Modal automata

Definition 3.4 A *modal \mathbf{X} -automaton* \mathbb{A} is a triple (A, Θ, Ω) where A is a non-empty finite set of *states*, $\Omega : A \rightarrow \omega$ is the *priority map*, while the *transition map*

$$\Theta : A \rightarrow \mathbf{1ML}(\mathbf{X}, A)$$

maps states to one-step formulas. An *initialized* modal automaton is pair (\mathbb{A}, a) , usually denoted as $\mathbb{A}\langle a \rangle$, consisting of a modal automaton \mathbb{A} together with a designated *initial* state a . The classes of (initialized) modal automata over the set \mathbf{X} are denoted as $IAut(\mathbf{X})$ and $Aut(\mathbf{X})$, respectively. \triangleleft

The operational semantics of modal automata is defined in terms of *acceptance games*.

Definition 3.5 Let $\mathbb{A} = (A, \Theta, \Omega)$ and $\mathbb{S} = (S, R, V)$ be a modal \mathbf{X} -automaton and a Kripke structure, respectively. The *acceptance game* $\mathcal{A}(\mathbb{A}, \mathbb{S})$ for \mathbb{A} with respect to \mathbb{S} is defined as the parity game given by the following table:

Position	Player	Admissible moves	Priority
$(a, s) \in A \times S$	\exists	$\{m : \sigma_R(s) \rightarrow \mathcal{P}A \mid \sigma(s), m \Vdash^1 \Theta(a)\}$	$\Omega(a)$
m	\forall	$\{(b, t) \mid b \in m(t)\}$	0

We say that $\mathbb{A}\langle a \rangle$ *accepts* the pointed structure (\mathbb{S}, s) if (a, s) is a winning position in the acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{S})$, and we write $\mathbb{S}, s \Vdash \mathbb{A}\langle a \rangle$ to denote this. \triangleleft

Convention 3.6 We will usually identify a match $\Sigma = (a_0, s_0)m_0(a_1, s_1)m_1(a_2, s_2)m_2 \dots$ of the acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{S})$ with the sequence $(a_0, s_0)(a_1, s_1)(a_2, s_2) \dots$ of its *basic positions*.

Some basic concepts concerning modal automata are introduced in the following definition.

Definition 3.7 Fix a modal \mathbf{X} -automaton $\mathbb{A} = (A, \Theta, \Omega)$.

The *size* $|\mathbb{A}|$ of \mathbb{A} is defined as the cardinality $|A|$ of its carrier set, while its *weight* $w(\mathbb{A}) := \max\{|\Theta(a)| \mid a \in A\}$ is given as the maximum size of the one-step formulas in the range of Θ .

Given a state a of \mathbb{A} , we write $\eta_a = \mu$ if $\Omega(a)$ is odd, and $\eta_a = \nu$ if $\Omega(a)$ is even; we call η_a the (*fixpoint*) *type* of a and say that a is an η_a -state.

The *occurrence graph* of \mathbb{A} is the directed graph $(G, E_{\mathbb{A}})$, where $aE_{\mathbb{A}}b$ if a occurs in $\Theta(b)$. We let $\triangleleft_{\mathbb{A}}$ denote the transitive closure of $E_{\mathbb{A}}$ and say that a is *active* in b if $a \triangleleft_{\mathbb{A}} b$. We write $a \triangleright_{\mathbb{A}} b$ if $a \triangleleft_{\mathbb{A}} b$ and $b \triangleleft_{\mathbb{A}} a$. A *cluster* of \mathbb{A} is a cell of the equivalence relation generated by $\triangleright_{\mathbb{A}}$ (i.e., the smallest equivalence relation on A containing $\triangleright_{\mathbb{A}}$); a cluster C is *degenerate* if it is of the form $C = \{a\}$ with $a \not\triangleright_{\mathbb{A}} a$. The unique cluster to which a state $a \in A$ belongs is denoted as C_a . We write $a \sqsubset_{\mathbb{A}} b$ if $\Omega(a) < \Omega(b)$, and $a \sqsubseteq_{\mathbb{A}} b$ if $\Omega(a) \leq \Omega(b)$. \triangleleft

3.3 Disjunctive modal automata

Many of our proofs involve the non-deterministic version of modal automata that we call *disjunctive*. The transition map of these automata makes use of the so-called *cover modality* ∇ , which was independently introduced in coalgebraic logic by Moss [38] and in automata theory by Janin & Walukiewicz [31] (where the authors used a different notation). It is a slightly non-standard connective that takes a finite set of formulas as its argument.

Definition 3.8 Given a finite set Φ , we let $\nabla\Phi$ abbreviate the formula

$$\nabla\Phi := \bigwedge \diamond\Phi \wedge \square \bigvee \Phi,$$

where $\diamond\Phi$ denotes the set $\{\diamond\varphi \mid \varphi \in \Phi\}$. \triangleleft

Remark 3.9 In words, $\nabla\Phi$ holds at s iff Φ and $\sigma_R(s)$ ‘cover’ one another, in the sense that every successor of s satisfies some formula in Φ and every formula in Φ holds in some successor of s . From this observation it is easy to derive that, conversely, the standard modal operators can be expressed in terms of the cover modality:

$$\begin{aligned} \diamond\varphi &\equiv \nabla\{\varphi, \top\} \\ \square\varphi &\equiv \nabla\{\varphi\} \vee \nabla\emptyset, \end{aligned}$$

where we note that $\nabla\emptyset$ holds at a point s iff s is a ‘blind’ world, that is, $R[s] = \emptyset$. \triangleleft

Definition 3.10 Let X be a given set of proposition letters and A any finite set. We define a *literal* over X to be a formula of the form p or $\neg p$ with $p \in X$, and define the language $\text{CL}(X)$ of conjunctions of literals to be generated by π in the grammar:

$$\pi ::= p \mid \neg p \mid \top \mid \pi \wedge \pi$$

where $p \in X$. We now define the set $\mathbf{1DML}(X, A)$ of *disjunctive* formulas in $\mathbf{1ML}(X, A)$ as follows:

$$\alpha ::= \pi \wedge \nabla B \mid \perp \mid \alpha \vee \alpha,$$

where $\pi \in \text{CL}(X)$ and $B \subseteq A$. \triangleleft

Definition 3.11 A modal X -automaton $\mathbb{A} = (A, \Theta, \Omega)$ is *disjunctive* if $\Theta(a) \in \mathbf{1DML}(X, A)$ for all $a \in A$. \triangleleft

The non-deterministic nature of disjunctive automata is exemplified by the following fact, which can easily be verified.

Fact 3.12 Let \mathbb{A} be a disjunctive automaton and \mathbb{S} an ω -unravelling tree. Then without loss of generality we may assume \exists 's positional winning strategy to be such that each marking m picked by the strategy at a position (a, s) satisfies $|m(t)| = 1$, for all $t \in \sigma_R(s)$.

In many branches of automata theory a crucial role is played by a *simulation theorem*, stating that automata of certain given type can be transformed into, or ‘simulated’ by, an equivalent automaton of which the transition structure is of a conceptually simpler kind. In the case of modal automata, such a theorem can be proved along the lines of Janin & Walukiewicz’ characterization of the μ -calculus in terms of the non-deterministic μ -automata [31]. For a detailed proof of the result below we refer to Venema [45].

Fact 3.13 (Simulation Theorem) *There is an effective procedure transforming an initialized modal automaton $\mathbb{A}\langle a \rangle$ into an equivalent initialized disjunctive modal automaton $\mathbb{D}\langle d \rangle$. The size of \mathbb{D} is exponential in the size of \mathbb{A} , and its weight is at most exponential in the product of the size and the weight of \mathbb{A} .*

3.4 Formulas and automata

The automata-theoretic approach towards the modal μ -calculus hinges on the existence of truth-preserving translations between formulas and automata, which testify that automata and formulas have the same expressive power. In the direction from formulas to automata we only need the result as such.

Fact 3.14 (i) *There is an effective procedure providing for any formula $\xi \in \mu\text{ML}(\mathbf{X})$ an equivalent initialized modal automaton $\mathbb{A}_\xi\langle a_\xi \rangle$. The size of \mathbb{A}_ξ is linear and its weight is at most exponential in the size of ξ .*

(ii) *There is an effective procedure providing for any formula $\xi \in \mu\text{ML}(\mathbf{X})$ an equivalent initialized disjunctive modal automaton $\mathbb{D}_\xi\langle d_\xi \rangle$. The size of \mathbb{D}_ξ is exponential and its weight is at most doubly exponential in the size of ξ .*

The translations mentioned in Fact 3.14 originate with Janin & Walukiewicz [32] for part (ii) and with Wilke [47] for part (i). Concerning the size matters, note that we get an exponential (and not polynomial) weight bound in Fact 3.14(1) because our modal automata are guarded (that is, all occurrences of states/variables in one-step formulas must occur in the scope of a modality; see Bruse et alii for details [9]). The size bounds in Fact 3.14(ii) are obtained simply by combining part (i) with Fact 3.13.

In the opposite direction we will need an actual map transforming an initialized modal automaton into an equivalent μ -calculus formula. For our definition of such a map, which is a variation of the one found in [28], we need some preparations. For a proper inductive formulation of this definition it is convenient to extend the class of modal automata, allowing states of the automaton to appear in the scope of a modality in a one-step formula.

Definition 3.15 A *generalized modal automaton* over \mathbf{X} is a triple $\mathbb{A} = (A, \Theta, \Omega)$ such that A is a finite set of states, $\Omega : A \rightarrow \omega$ is a priority map, and $\Theta : A \rightarrow \mathbf{1ML}(\mathbf{X}, A \cup \mathbf{X})$ maps states of A to *generalized one-step formulas*. \triangleleft

To mark the difference with standard modal automata, the formula $\diamond a \wedge \Box(p \vee b)$ is a generalized one-step formula but not a proper one-step formula. Whenever possible, we will apply concepts that have been defined for modal automata to these generalized structures without explicit notification. For the operational semantics of generalized modal automata we may extend the notion of a one-step model in the obvious way. Readers who are interested in the details may consult [22].

Definition 3.16 A (generalized) modal automaton $\mathbb{A} = (A, \Theta, \Omega)$ is called *linear* if the relation $\sqsubset_{\mathbb{A}}$ is a linear order (i.e., the priority map Ω is injective), and satisfies $\Omega(a) > \Omega(b)$ in case b is active in a but not vice versa. A *linearization* of \mathbb{A} is a linear automaton $\mathbb{A}' = (A, \Theta, \Omega')$ such that (1) for all $a \in A$, $\Omega'(a)$ has the same parity as $\Omega(a)$, and (2) for all $a, b \in A$ that belong to the same cluster we have $\Omega'(a) < \Omega'(b)$ iff $\Omega(a) < \Omega(b)$. \triangleleft

The following proposition is easy to verify.

Proposition 3.17 *Let \mathbb{A} be a modal automaton.*

- (1) \mathbb{A} has some linearization.
- (2) If \mathbb{A}' is a linearization of \mathbb{A} , then $\mathbb{A}\langle a \rangle \equiv \mathbb{A}'\langle a \rangle$.

Definition 3.18 We define a map

$$\mathbf{tr}_{\mathbb{A}} : A \rightarrow \mu\text{ML}(\mathbf{X})$$

for any linear generalized \mathbf{X} -automaton $\mathbb{A} = (A, \Theta, \Omega)$. These maps are defined by induction on the size of the automaton \mathbb{A} .

In case $|\mathbb{A}| = 1$, we set

$$\mathbf{tr}_{\mathbb{A}}(a) := \eta_a a. \Theta(a),$$

where a is the unique state of \mathbb{A} .

In case $|\mathbb{A}| > 1$, by linearity there is a unique state m reaching the maximal priority of \mathbb{A} , that is, with $\Omega(m) = \max(\text{Ran}(\Omega))$. Let $\mathbb{A}^- = (A^-, \Theta^-, \Omega^-)$ be the $\mathbf{X} \cup \{m\}$ -automaton given by $A^- := A \setminus \{m\}$, while Θ^- and Ω^- are defined as the restrictions of, respectively, Θ and Ω to A^- . Since $|\mathbb{A}^-| < |\mathbb{A}|$, inductively³ we may assume a map $\mathbf{tr}_{\mathbb{A}^-} : A \rightarrow \mu\text{ML}(\mathbf{X} \cup \{m\})$.

Now we first define

$$\mathbf{tr}_{\mathbb{A}}(m) := \eta_m m. \Theta(m) [\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-],$$

and then set

$$\mathbf{tr}_{\mathbb{A}}(a) := \mathbf{tr}_{\mathbb{A}^-}(a) [\mathbf{tr}_{\mathbb{A}}(m)/m]$$

for the states $a \neq m$. ◁

The following proposition states that this map is truth-preserving, and of exponential size. The proof of the first statement is rather standard; details can be found in [28, 45]; the size bound can be established via a straightforward induction on the size of the automaton.

Fact 3.19 *Let $\mathbb{A}\langle a \rangle$ be an initialized linear modal automaton. Then $\mathbf{tr}_{\mathbb{A}}(a)$ is equivalent to $\mathbb{A}\langle a \rangle$ and its size is at most exponential in the sum of the size and the weight of \mathbb{A} . If \mathbb{A} is positive in $p \in \mathbf{X}$ then so is $\mathbf{tr}_{\mathbb{A}}(a)$.*

Based on Proposition 3.17 and Fact 3.19 we may define a truth-preserving translation $\mathbf{tr}_{\mathbb{A}} : A \rightarrow \mu\text{ML}$ for every modal automaton $\mathbb{A} = (A, \Theta, \Omega)$ by picking a linearization \mathbb{A}^l of \mathbb{A} and defining $\mathbf{tr}_{\mathbb{A}}(a) := \mathbf{tr}_{\mathbb{A}^l}(a)$ for all $a \in A$. Clearly the exact shape of $\mathbf{tr}_{\mathbb{A}}(a)$ will depend on the choice of the linearization.

3.5 Bipartite automata

As mentioned in our introduction, the results in this paper will be based on constructions that transform a given automaton into one of a particular shape. In this final subsection on modal automata we look at the target automata of these constructions in some more detail.

Definition 3.20 A modal automaton \mathbb{A} is called a *bipartite* automaton if its carrier A can be partitioned into an *initial* part A_0 and a *final* part A_1 , in such a way that we have $\triangleleft_{\mathbb{A}} \cap (A_0 \times A_1) = \emptyset$. It will sometimes be convenient to represent a bipartite automaton as a quadruple $\mathbb{A} = (A_0, A_1, \Theta, \Omega)$.

Given such an automaton and a state a of \mathbb{A} , we say that $\mathbb{A}\langle a \rangle$ is an *initialized* bipartite automaton if a belongs to the initial part of \mathbb{A} ; and given a class C of bipartite automata, we let IC denote the class of corresponding initialized bipartite automata. ◁

³Observe that since m is a proposition letter and not a variable in \mathbb{A}^- , the latter structure need not be a modal automaton, even if \mathbb{A} is. It is for this reason that we introduced the notion of a *generalized* modal automaton.

The intuition underlying this definition is that, in any acceptance game related to a bipartite automaton $\mathbb{A} = (A_0, A_1, \Theta, \Omega)$ there will be two kinds of matches: either \mathbb{A} stays in its initial part throughout the match, or at some stage it moves to its final part, where it remains throughout the remainder of the match. The bipartite automata that we will meet in this paper will be such that its initial and final part behave differently with respect to the designated propositional variable p .

Recall that each of our results concerns a certain syntactic fragment $F(P)$ of the modal μ -calculus, consisting of formulas where we have imposed restrictions on the occurrence of the proposition letters belonging to a certain set $P \subseteq \text{PROP}$. (Note that while these restrictions generally concern the free variables of the formula, not the bound ones, there is no relation between the fragment $F(P)$ and the set $\mu\text{ML}(\mathbf{X})$ of formulas of which the free variables belong to \mathbf{X}). In each case considered, part of our proof will consist in showing that given a bipartite automaton \mathbb{A} of a certain kind, one can find a translation $\text{tr} : A \rightarrow \mu\text{ML}$, such that for all states a in the initial part of \mathbb{A} , the formula $\text{tr}(a)$ belongs to the fragment $F(P)$. Since these proofs are all quite similar, we have extracted their pattern in the form of the following definition and proposition, so that in each specific case we may confine our attention to a verification that the fragment and the automata at stake meet the required conditions.

Definition 3.21 Suppose that we have defined, for each finite set P of proposition letters, a fragment $F(P) \subseteq \mu\text{ML}$ of μ -calculus formulas. In the sequel we shall consider the following properties that are applicable to such a family $F = \{F(P) \mid P \subseteq_{\omega} \text{PROP}\}$:

- (EP) *extension property*: $\varphi \in F(P \cup \{q\})$ whenever $\varphi \in F(P)$ and $q \notin FV(\varphi)$;
- (SP1) *first substitution property*: $\varphi[\psi/x] \in F(P)$ whenever $\varphi \in F(P \cup \{x\})$ and $\psi \in F(P)$;
- (SP2) *second substitution property*: $\varphi[\psi/x] \in F(P)$ whenever $\varphi \in F(P)$, $x \notin P$ and $FV(\psi) \cap P = \emptyset$;
- (CA $_{\eta}$) *η -fixpoint closure property*: $\eta x.\varphi \in F(P)$ whenever $\varphi \in F(P \cup \{x\})$. \triangleleft

Proposition 3.22 Let $F = \{F(Q) \mid Q \subseteq \text{PROP}\}$ be a family of fragments of μML , and let P, \mathbf{X} be two finite sets of proposition letters. Let $\mathbb{A} = (A, B, \Theta, \Omega)$ be a bipartite \mathbf{X} -automaton, and let \mathbb{B} denote the automaton $(B, \Theta_{\uparrow B}, \Omega_{\uparrow B})$. Assume that

- (i) F has the properties (EP), (SP1), (SP2) and, for all $a \in A$, (CA $_{\eta_a}$);
- (ii) $\Theta(a) \in F(P \cup A)$ for all $a \in A$;
- (iii) there is a translation $\text{tr} : B \rightarrow \mu\text{ML}(\mathbf{X})$ such that $\mathbb{B}(b) \equiv \text{tr}(b) \in F(P)$, for all $b \in B$;

Then there is a translation $\text{tr} : A \rightarrow \mu\text{ML}(\mathbf{X})$ such that $\text{tr}_{\mathbb{A}}(a) \equiv \mathbb{A}(a)$ and $\text{tr}_{\mathbb{A}}(a) \in F(P)$ for all $a \in A$.

Proof. By Proposition 3.17 we may assume without loss of generality that \mathbb{A} is linear itself and that $\Omega(a) > \Omega(b)$ for all $a \in A$ and $b \in B$. We will also assume that the translation tr of (iii) is identical to the map $\text{tr}_{\mathbb{B}}$ obtained by applying Definition 3.18 to the automaton \mathbb{B} . (The more general case can be proved by a straightforward modification of the proof given below.)

Enumerate $A = \{a_1, \dots, a_n\}$ in such a way that $\Omega(a_i) < \Omega(a_j)$ iff $i < j$. For $0 \leq k \leq n$ we define $A_k := \{a_i \mid 0 < i \leq k\}$, and we let \mathbb{A}_k be the linear generalized automaton $(B \cup A_k, \Theta_{\uparrow B \cup A_k}, \Omega_{\uparrow B \cup A_k})$. It is straightforward to verify that $\mathbb{A}_0 = \mathbb{B}$, $\mathbb{A}_n = \mathbb{A}$, and that $\mathbb{A}_i = (\mathbb{A}_{i+1})^-$ for all $i < n$. We abbreviate $\text{tr}_k := \text{tr}_{\mathbb{A}_k}$.

Our task is then to show that $\text{tr}_n(a) \in F(P)$ for all $a \in A$, and our approach will be to prove, by induction on k , that

$$\text{tr}_k(c) \in F(P \cup \{a_i \mid k < i \leq n\}), \text{ for all } c \in A \cup B \text{ and } k \leq n. \quad (5)$$

In the base step of the induction, where $k = 0$, we are dealing with the automaton $\mathbb{A}_0 = \mathbb{B}$, corresponding to the final part of \mathbb{A} , and we only have to worry about states $b \in \mathbb{B}$. It is easy to see that

$$\text{tr}_0(b) = \text{tr}_{\mathbb{B}}(b) \text{ for all } b \in B,$$

and so by assumption (iii) and the assumption (i) that F has the extension property, we obtain that $\mathbf{tr}_0(b) \in F(P) \subseteq F(P \cup A)$. This suffices to prove (5) in the base case.

In the inductive case, where $k > 0$, we consider the automaton \mathbb{A}_k , with maximum-priority state $a_k \in A$. Observe that by definition we have

$$\mathbf{tr}_k(a_k) := \eta_{a_k} a_k \cdot \Theta(a_k) [\mathbf{tr}_{k-1}(a_i)/a_i \mid i < k, \mathbf{tr}_{k-1}(b)/b \mid b \in B] \quad (6)$$

and, for $i < k$,

$$\mathbf{tr}_k(a_i) := \mathbf{tr}_{k-1}(a_i) [\mathbf{tr}_k(a_k)/a_k]. \quad (7)$$

Observe that we may conclude from (6) that

$$\mathbf{tr}_k(a_k) := \eta_{a_k} a_k \cdot \Theta(a_k) [\mathbf{tr}_{k-1}(a_i)/a_i \mid i < k] [\mathbf{tr}_{k-1}(b)/b \mid b \in B], \quad (8)$$

since no $b \in B$ has a free occurrence in any formula $\mathbf{tr}_{k-1}(a_i)$.

We now turn to the proof of (5). Starting with the states in B , it is easy to see that

$$\mathbf{tr}_k(b) = \mathbf{tr}_{\mathbb{B}}(b) \text{ for all } b \in B \text{ and } k \leq n,$$

and so we find $\mathbf{tr}_k(b) \in F(P \cup \{a_i \mid k < i \leq n\})$ by assumption (iii) and the assumption (i) that F has the extension property.

We now consider the formula $\mathbf{tr}_k(a_k)$. By the inductive hypothesis we have $\mathbf{tr}_{k-1}(a_i) \in F(P \cup \{a_i \mid k \leq i \leq n\})$, while by assumption (iii) we have $\Theta(a_k) \in F(P \cup A) = F((P \cup \{a_i \mid k \leq i \leq n\}) \cup \{a_i \mid i < k\})$. But then by the first substitution property of F we find that

$$\Theta(a_k) [\mathbf{tr}_{k-1}(a_i)/a_i \mid i < k] \in F(P \cup \{a_i \mid k \leq i \leq n\}),$$

and by the fact that F is closed under the application of η_{a_k} -fixpoint operators, we obtain

$$\eta_{a_k} a_k \cdot \Theta(a_k) [\mathbf{tr}_{k-1}(a_i)/a_i \mid i < k] \in F(P \cup \{a_i \mid k < i \leq n\}). \quad (9)$$

Now observe that for all $b \in B$ we have $b \notin P$ and $\mathbf{tr}_{k-1}(b) \in F(P \cup \{a_i \mid k < i \leq n\})$, so that by successive applications of the second substitution property we obtain from (9) and (8) that

$$\mathbf{tr}_k(a_k) \in F(P \cup \{a_i \mid k < i \leq n\}), \quad (10)$$

as required.

Finally we consider the formula $\mathbf{tr}_k(a_j)$ for some fixed but arbitrary $j < k$. From the inductive observation that $\mathbf{tr}_{k-1}(a_j)$ belongs to $F(P \cup \{a_i \mid k \leq i \leq n\}) = F((P \cup \{a_i \mid k < i \leq n\}) \cup \{a_k\})$, we may then conclude by (10), (7) and the first substitution property that

$$\mathbf{tr}_k(a_j) \in F(P \cup \{a_i \mid k < i \leq n\}),$$

and so we are done. QED

4 First steps

In this section we provide both an automata-theoretic construction and a first preservation result that that will be of use throughout the paper. We start with the first, introducing a simple operation on automata that will feature throughout the remainder of the paper.

Convention 4.1 As mentioned in the introduction, in this paper we shall focus on the contribution of one specific proposition letter in the semantics of formulas and automata. It will be convenient to fix this letter from now on, and reserve the name ‘ p ’ for it.

Definition 4.2 Given a modal automaton $\mathbb{A} = (A, \Theta, \Omega)$, we define $\mathbb{A}^\perp = (A^\perp, \Theta^\perp, \Omega^\perp)$ to be the automaton where $A^\perp := \{a^\perp \mid a \in A\}$, Θ^\perp is given by putting

$$\Theta^\perp(a) := \Theta(a)[\perp/p],$$

and Ω^\perp is simply defined by $\Omega^\perp(a^\perp) := \Omega(a)$. \triangleleft

Proposition 4.3 *Let \mathbb{A} be a modal automaton which is positive in p . Then for any state a in \mathbb{A} and any pointed Kripke model (\mathbb{S}, s) we have*

$$\mathbb{S}, s \Vdash \mathbb{A}^\perp \langle a \rangle \text{ iff } \mathbb{S}[p \mapsto \emptyset], s \Vdash \mathbb{A} \langle a \rangle.$$

Proof. We will show that the two respective acceptance games, $\mathcal{A}^\perp := \mathcal{A}(\mathbb{A}^\perp, \mathbb{S})$ and $\mathcal{A}_\emptyset := \mathcal{A}(\mathbb{A}, \mathbb{S}[p \mapsto \emptyset])$ are in fact, *identical*. The key observation here is that for any position $(a, t) \in A \times S$, the set of moves available to \exists in both games is the same. To see this, it suffices to prove that, for any point $t \in S$ and any one-step formula $\alpha \in \mathbf{1ML}$ which is positive in p , we have

$$\sigma_V(t), \sigma_R(t), m \Vdash^1 \alpha[\perp/p] \text{ iff } \sigma_V(t) \setminus \{p\}, \sigma_R(t), m \Vdash^1 \alpha,$$

and the proof of this statement proceeds by a straightforward induction on α . \square

We now turn to our first characterization result, which concerns the notion of *monotonicity*. D'Agostino and Hollenberg [17] already proved a Lyndon theorem for the modal μ -calculus, stating that monotonicity with respect to a propositional variable p is captured by the formulas that are syntactically *positive* in p . Here we strengthen their result, providing both an explicit translation and a decidability result. Our proof proceeds via some auxiliary lemmas on automata that we shall need further on.

Definition 4.4 A formula $\xi \in \mu\mathbf{ML}(X)$ is *monotone* in $p \in X$ if for all pointed Kripke structures (\mathbb{S}, s) and for all pairs of sets $U, U' \subseteq S$ such that $U \subseteq U'$, it holds that $\mathbb{S}[p \mapsto U], s \Vdash \xi$ implies $\mathbb{S}[p \mapsto U'], s \Vdash \xi$.

Syntactically, the set $\mu\mathbf{ML}_p^M$ of formulas that are *positive* in p , is defined by the following grammar:

$$\varphi ::= p \mid q \mid \neg q \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond \varphi \mid \square \varphi \mid \mu x. \varphi \mid \nu x. \varphi,$$

where $q \neq p$ and x are propositional variables, and the formation of the formulas $\mu x. \varphi$ and $\nu x. \varphi$ is subject to the constraint that the formula φ is *positive* in x . \triangleleft

Definition 4.4 gives an explicit grammar for generating the formulas that are positive in p . It is easy to see that this definition coincides with the one we gave earlier, viz., that a formula is positive in p iff all occurrences of p are in the scope of an even number of negations.

The theorem below can be seen as a strong version of D'Agostino & Hollenberg's characterization result.

Theorem 4.5 *There is an effective translation which maps a given $\mu\mathbf{ML}$ -formula ξ to a formula $\xi^M \in \mu\mathbf{ML}_p^M$ such that*

$$\xi \text{ is monotone in } p \text{ iff } \xi \equiv \xi^M, \tag{11}$$

and it is decidable in elementary time whether a given formula ξ is monotone in p .

It is routine to prove that all formulas in $\mu\mathbf{ML}_p^M$ are monotone in p , so we focus on the hard part of Theorem 4.5, for which we shall involve automata.

Definition 4.6 A modal automaton $\mathbb{A} = (A, \Theta, \Omega)$ is *positive in p* if for all $a \in A$ the one-step formula $\Theta(a)$ is positive in p . \triangleleft

It is easy to see that for any linear modal automaton that is positive in p , the translation map given in Definition 3.18 produces formulas that are also positive in p . From this the following is immediate.

Proposition 4.7 *Let $\mathbb{A} = (A, \Theta, \Omega)$ be a modal automaton which is positive in p . Then there is a truth-preserving translation $\mathbf{tr}_{\mathbb{A}} : A \rightarrow \mu\text{ML}_p^M$.*

We now turn to the key lemma underlying the proof of Theorem 4.5, for which we need the following definition.

Definition 4.8 Let $(\cdot)^M : \text{CL}(\mathbf{X}) \rightarrow \text{CL}(\mathbf{X})$ be the translation which replaces every occurrence of the literal $\neg p$ with \top , and let $(\cdot)^M : \mathbf{1DML}(\mathbf{X}, A) \rightarrow \mathbf{1ML}(\mathbf{X}, A)$ be the one-step translation given by the following inductive definition:

$$\begin{aligned} (\pi \wedge \nabla B)^M &:= \pi^M \wedge \nabla B \\ \perp^M &:= \perp \\ (\alpha \vee \beta)^M &:= \alpha^M \vee \beta^M. \end{aligned}$$

Let $\mathbb{A} = (A, \Theta, \Omega)$ be a disjunctive modal automaton. We define the automaton \mathbb{A}^M as the structure (A, Θ^M, Ω) where the map Θ^M is given by putting

$$\Theta^M(a) := \Theta(a)^M$$

for every state $a \in A$. \triangleleft

The following proposition is fairly obvious; we list it for future reference.

Proposition 4.9 *Let \mathbb{A} be a disjunctive automaton. Then its transformation \mathbb{A}^M is disjunctive as well, and positive in p .*

Proposition 4.10 *Let $\mathbb{A}\langle a_I \rangle$ be an initialized disjunctive modal automaton. If $\mathbb{A}\langle a_I \rangle$ is monotone in p , then $\mathbb{A}\langle a_I \rangle \equiv \mathbb{A}^M\langle a_I \rangle$.*

Proof. Given the nature of the translation $(\cdot)^M$ at the level of conjunctions of literals, it is easy to see that $\mathbb{A}\langle a_I \rangle$ implies $\mathbb{A}^M\langle a_I \rangle$, and so we focus on the opposite implication. It suffices to take an arbitrary ω -unravalled tree model (\mathbb{S}, r) and show that

$$\mathbb{S}, r \Vdash \mathbb{A}^M\langle a_I \rangle \text{ only if } \mathbb{S}, r \Vdash \mathbb{A}\langle a_I \rangle. \quad (12)$$

Assume that $\mathbb{S}, r \Vdash \mathbb{A}^M\langle a_I \rangle$. Our aim is to find a subset $U \subseteq S$ such that

$$\mathbb{S}[p \upharpoonright U], r \Vdash \mathbb{A}\langle a_I \rangle, \quad (13)$$

from which it will immediately follow by monotonicity that $\mathbb{S}, r \Vdash \mathbb{A}\langle a_I \rangle$.

By Fact 3.12 \exists has a positional winning strategy f in $\mathcal{A}(\mathbb{A}^M, \mathbb{S})$ such that each marking m picked by the strategy at a position (a, s) satisfies $|m(t)| = 1$, for all $t \in \sigma_R(s)$. From this it easily follows that for all $t \in S$ there is exactly one state $a_t \in A$ such that the position (a_t, t) may occur in an f -guided match of $\mathcal{A}(\mathbb{A}^M, \mathbb{S})$ starting at (a_I, r) . It easily follows that the pair (a_t, t) is a winning position for \exists in $\mathcal{A}(\mathbb{A}^M, \mathbb{S})$; in particular, with $m_t : \sigma_R(t) \rightarrow \text{PA}$ being the marking picked by f at this position, by the legitimacy of this move we have that

$$\sigma_V(t), \sigma_R(t), m_t \Vdash^1 \Theta^M(a_t).$$

In order to determine whether this t should be a member of U or not, we check whether the one-step formula $\Theta(a_t)$ holds at the one-step modal $(\sigma_V(t), \sigma_R(t), m_t)$. If so, we are happy with $\sigma_V(t)$ as it is and put $t \in U$, but if not, we will want to make p false at t , claiming that

$$\sigma_V(t) \setminus \{p\}, \sigma_R(t), m_t \Vdash^1 \Theta(a_t). \quad (14)$$

To see why this is the case, observe that the only reason why we can have $\sigma_V(t), \sigma_R(t), m_t \Vdash^1 \Theta^M(a_t)$ but $\sigma_V(t), \sigma_R(t), m_t \not\Vdash^1 \Theta(a_t)$ is that $\Theta(a_t)$ has a disjunct $\pi \wedge \nabla B$ such that $\neg p$ is a conjunct of π and $p \in \sigma_V(t)$. But then it is immediate that $\sigma_V(t) \setminus \{p\}, \sigma_R(t), m_t \Vdash^1 \pi \wedge \nabla B$, which implies (14).

These observations reveal that if we define

$$U := \{t \in S \mid (a_t, t) \in \text{Win}_{\exists}(\mathcal{A}(\mathbb{A}^M, \mathbb{S})) \text{ and } \sigma_V(t), \sigma_R(t), m_t \Vdash^1 \Theta(a_t)\},$$

we obtain for the valuation $V[p|U]$ that

$$\sigma_{V[p|U]}(t), \sigma_R(t), m_t \Vdash^1 \Theta(a_t),$$

whenever (a_t, t) is a winning position for \exists in $\mathcal{A}(\mathbb{A}, \mathbb{S})$. From this it is easy to derive that f itself is a winning strategy for \exists in the acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{S}[p|U])@(\alpha_I, r)$, as required to prove (13). QED

We now have all material needed to prove Theorem 4.5.

Proof of Theorem 4.5. Let ξ be an arbitrary modal μ -formula; and let $\mathbb{D}_{\xi}\langle d_{\xi} \rangle$ be an initialized disjunctive automaton that is equivalent to ξ . Such a structure exists by Fact 3.14, and clearly ξ is monotone in p iff $\mathbb{D}_{\xi}\langle d_{\xi} \rangle$ is so. It then follows by Proposition 4.10 that $\mathbb{D}_{\xi}\langle d_{\xi} \rangle$ (and hence ξ) is monotone in p iff it is equivalent to the initialized disjunctive automaton $\mathbb{D}_{\xi}^M\langle d_{\xi} \rangle$. Now define

$$\xi^M := \text{tr}_{\mathbb{D}_{\xi}^M}(d_{\xi}).$$

It is easy to verify that $\xi^M \in \mu\text{ML}_p^M$, while we have $\mathbb{D}_{\xi}^M\langle d_{\xi} \rangle \equiv \xi^M$ by Fact 3.14. Putting these observations together we obtain that, indeed, ξ is monotone in p iff $\xi \equiv \xi^M$, where ξ^M is effectively obtained from ξ . This establishes the first part of the theorem.

For the statement concerning decidability, it suffices to observe that all constructions that are involved in the definition of the map $(\cdot)^M : \mu\text{ML}(\mathbf{X}) \rightarrow \mu\text{ML}_p^M$ have uniformly elementary size bounds, and that the problem, whether two μ -calculus formulas are equivalent or not, can be decided in exponential time. From this the decidability claim is immediate. QED

5 Finite width property

The first new property that we consider is that of the *finite width property*.

Definition 5.1 A formula $\xi \in \mu\text{ML}(\mathbf{X})$ has the *finite width property* for $p \in \mathbf{X}$ if ξ is monotone in p , and, for every tree model (\mathbb{S}, s) ,

$$\mathbb{S}, s \Vdash \xi \text{ iff } \mathbb{S}[p|U], s \Vdash \xi, \text{ for some finitely branching subtree } U \subseteq S,$$

where a subset $U \subseteq S$ is a *finitely branching subtree* if U is downward closed and the set $R(u) \cap U$ is finite for every $u \in U$. \triangleleft

We will associate the following syntactic fragment of the modal μ -calculus with this property.

Definition 5.2 Given a set $P \subseteq \mathbf{X}$, we define the fragment μML_P^W by the following grammar:

$$\varphi ::= p \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond \varphi \mid \mu x. \varphi' \mid \nu x. \varphi',$$

where $p \in P$, $\psi \in \mu\text{ML}(\mathbf{X} \setminus P)$ is a P -free formula and $\varphi' \in \mu\text{ML}_{P \cup \{x\}}^W$. In case P is a singleton, say, $P = \{p\}$, we will write μML_p^W rather than $\mu\text{ML}_{\{p\}}^W$. \triangleleft

In words, the fragment μML_p^W consists of those formulas that are positive in p and do not admit any occurrence of a p -active subformula in the scope of a box modality. All formulas from Example 2.24 belong to μML_p^W , except φ_4 .

The following theorem states that modulo equivalence, μML_p^W is the syntactic fragment of the modal μ -calculus that captures the finite width property, and that it is decidable whether a given μ -formula has this property.

Theorem 5.3 *There is an effective translation which maps a given μML -formula ξ to a formula $\xi^W \in \mu\text{ML}_p^W$ such that*

$$\xi \text{ has the finite width property for } p \text{ iff } \xi \equiv \xi^W, \quad (15)$$

and it is decidable in elementary time whether a given formula ξ has the finite width property for p .

First we prove the easy part of the theorem, stating that formulas in the fragment μML_p^W indeed have the required semantic property.

Proposition 5.4 *Every formula $\xi \in \mu\text{ML}_p^W$ has the finite width property with respect to p .*

Proof. Let ξ be a formula in μML_p^W , then ξ is obviously positive, and hence, monotone in p . Fix a tree model \mathbb{S} with root $r \in S$. We have to prove

$$\mathbb{S}, r \Vdash \xi \text{ iff } \mathbb{S}[p|U], r \Vdash \xi, \text{ for some finitely branching subtree } U \subseteq S. \quad (16)$$

The direction from right to left follows from the fact that ξ is monotone in p . For the opposite direction, suppose that $\mathbb{S}, r \Vdash \xi$. We need to find a finitely branching subtree U of S that is downward closed and such that $\mathbb{S}[p|U], r \Vdash \xi$. Let f be a positional winning strategy of \exists in the game $\mathcal{E}_0 := \mathcal{E}(\xi, \mathbb{S})@(\xi, r)$. We define $U \subseteq S$ such that

$$u \in U \text{ iff there is a } \varphi \text{ such that } (\varphi, u) \text{ is } f\text{-reachable in } \mathcal{E}_0 \text{ and } p \text{ is active in } \varphi.$$

It is easy to see that the set U is downward closed. Indeed, if a position (φ, t) is reached during an \mathcal{E}_0 -match Σ and p is not active in φ , then all positions occurring after (φ, t) will be of the form (ψ, u) , where p is not active in ψ .

Hence it suffices to show that U is finitely branching. Fix $u \in U$ and let us show that $\sigma_R(u) \cap U$ is finite. Let $t \in U$ be a successor of u . Since u is the only predecessor of t , by definition of U , there must be an f -guided match during which a move occurs from $(\Delta\varphi_t, u)$ to (φ_t, t) , where $\Delta \in \{\square, \diamond\}$ and φ_t is a p -active subformula of ξ . Because of the syntactic constraints on μML_p^W , this can only happen if $\Delta = \diamond$. But then $(\diamond\varphi_t, u)$ is a position which belongs to \exists and so t is her choice as dictated by f . From this, it follows that for all t and t' in $R(u) \cap U$, we have $\varphi_t \neq \varphi_{t'}$ if $t \neq t'$. Putting this together with the fact that $Sfor(\xi)$ is finite, we obtain that $R(u) \cap U$ is finite. This finishes the proof that U is downward closed and finitely branching.

It remains to show that $\mathbb{S}[p|U], r \Vdash \xi$. Let \mathcal{E} be the game $\mathcal{E}(\xi, \mathbb{S}[p|U])@(\xi, r)$. We show that f itself is a winning strategy for \exists in the game \mathcal{E} . The winning conditions for \mathcal{E}_0 and \mathcal{E} are the same. Moreover, the rules of the two games are the same, except when we reach a position of the form (p, t) . So to prove that f is a winning strategy for \exists in \mathcal{E} , it suffices to show that if an f -guided \mathcal{E} -match

Σ arrives at a position (p, t) , then $\mathbb{S}[p|U], t \Vdash p$, that is, $t \in V(p) \cap U$. Suppose that we are in this situation. Since Σ is also an f -guided \mathcal{E}_0 -match and since f is a winning strategy for \exists in \mathcal{E}_0 , t belongs to $V(p)$. It remains to show that t belongs to U . That is, we have to find a p -active formula φ such that (φ, t) is f -reachable in \mathcal{E}_0 . Clearly, the formula p itself satisfies these two conditions. This proves that f is a winning strategy for \exists in the game \mathcal{E} and hence shows that $\mathbb{S}[p|U], r \Vdash \xi$. QED

For the hard part of the theorem, we involve automata. The particular modal automata that we associate with the finite width property are given below; recall that we introduced *bipartite* automata in Definition 3.20.

Definition 5.5 A bipartite modal automaton $\mathbb{A} = (A, B, \Theta, \Omega)$ belongs to the class Aut_p^W of *finite-width automata* if the one-step language associated with B is the language $\mathbf{1ML}(X \setminus \{p\}, B)$, and the one-step language associated with A is given by the following grammar:

$$\alpha ::= p \mid \diamond \pi_0 \mid \beta \mid \alpha \wedge \alpha \mid \top \mid \alpha \vee \alpha \mid \perp \quad (17)$$

where $\pi_0 \in \mathbf{Latt}(A)$ and $\beta \in \mathbf{1ML}(X \setminus \{p\}, B)$. \triangleleft

In words, an initialized modal automaton $\mathbb{A}\langle a_I \rangle$, with $\mathbb{A} = (A, \Theta, \Omega)$, belongs to the class Aut_p^W if A can be partitioned as $A = A_0 \uplus A_1$ such that (0) a_I belongs to A_0 , (1) p occurs only positively in $\Theta(a)$, for $a \in A_0$, (2) p does not occur in any $\Theta(a)$, $a \in A_1$, (3) if $a, b \in A_0$ then a may only occur in $\Theta(b)$ in the scope of a diamond (not of a box) modality, and (4) if $a \in A_0$ and $b \in A_1$ then a may not occur in $\Theta(b)$.

Proposition 5.6 Let $\mathbb{A} = (A, B, \Theta, \Omega)$ be a bipartite modal automaton in Aut_p^W . Then there is a translation $\mathbf{tr}_{\mathbb{A}} : A \rightarrow \mu\mathbf{ML}$ such that $\mathbf{tr}_{\mathbb{A}}(a) \in \mu\mathbf{ML}_p^W$ for every state $a \in A$.

Proof. It suffices to check that the fragment $\mu\mathbf{ML}^W := \{\mu\mathbf{ML}_P^W \mid P \subseteq_{\omega} \mathbf{PROP}\}$ and the automaton \mathbb{A} satisfy the conditions (i) – (iii) of Proposition 3.22, with $P = \{p\}$.

Starting with (i), we need to show that $\mu\mathbf{ML}^W$ satisfies the properties (EP), (SP1), (SP2) and (AC $_{\eta}$) for $\eta \in \{\mu, \nu\}$. All these results can be established by routine proofs; we consider the property (SP1) as an example (the other properties are easier to show). By a straightforward formula induction on φ one may show that if $\varphi \in \mu\mathbf{ML}_{P \cup \{x\}}^W$, then for all $\psi \in \mu\mathbf{ML}_P^W$ the formula $\varphi[\psi/x]$ belongs to the fragment $\mu\mathbf{ML}_P^W$.

Confining ourselves to the inductive case where φ is of the form $\mu y. \varphi'$, we reason as follows. Without loss of generality we may assume that $y \notin FV(\psi)$, so that $\psi \in \mu\mathbf{ML}_{P \cup \{y\}}^W$ by the extension property. Also, if $\mu y. \varphi' \in \mu\mathbf{ML}_{P \cup \{x\}}^W$, then by the formulation rules of $\mu\mathbf{ML}_P^W$ it must be the case that $\varphi' \in \mu\mathbf{ML}_{P \cup \{x, y\}}^W$. It then follows by the inductive hypothesis that $\varphi'[\psi/x] \in \mu\mathbf{ML}_{P \cup \{y\}}^W$, so that $\varphi[\psi/x] \in \mu\mathbf{ML}_P^W$, again by definition of the fragment.

To check that (ii) the formula $\Theta(a)$ belongs to $\mu\mathbf{ML}_{\{p\} \cup A}^W$ for all $a \in A$, one may proceed via a straightforward induction on the complexity of the one-step formulas generated by the grammar (17).

Finally, it is easy to verify that there is a translation $\mathbf{tr} : B \rightarrow \mu\mathbf{ML}(X \setminus \{p\})$; from this it is immediate that (iii) $\mathbf{tr}(b) \in \mu\mathbf{ML}_{\{p\}}^W$, for all $b \in B$. QED

Note that it follows from Proposition 5.6 and Proposition 5.4 that initialized automata in $IAut_p^W$ have the finite width property.

Preparing for the main technical result of this section, we define the following transformation of automata. Recall from Definition 3.10 that $\mathbf{1DML}(X, A)$ denotes the set of *disjunctive* one-step formulas over X and A , and that \mathbb{A}^{\perp} is the automaton given in Definition 4.2.

Definition 5.7 Let $(\cdot)^W : 1\text{DML}(\mathbb{X}, A) \rightarrow 1\text{ML}(\mathbb{X}, A \uplus A^\perp)$ be the one-step translation given by the following inductive definition:

$$\begin{aligned} (\pi \wedge \nabla B)^W &:= \bigvee \left\{ \pi \wedge \bigwedge \diamond B_1 \wedge \nabla B_2^\perp \mid B_1 \cup B_2 = B \right\} \\ \perp^W &:= \perp \\ (\alpha \vee \beta)^W &:= \alpha^W \vee \beta^W, \end{aligned}$$

where B_2^\perp denotes the set $B_2^\perp := \{b^\perp \mid b \in B_2\}$.

Let $\mathbb{A} = (A, \Theta, \Omega)$ be a disjunctive modal automaton which is positive in p . We define \mathbb{A}^W as the bipartite automaton $(A^W, \Theta^W, \Omega^W)$ where $A^W := A \uplus A^\perp$, and the maps Θ^W and Ω^W are given by putting

$$\begin{aligned} \Theta^W(a) &:= \Theta(a)^W & \text{and} & & \Omega^W(a) &:= \Omega(a) \\ \Theta^W(a^\perp) &:= \Theta^\perp(a) & & & \Omega^W(a^\perp) &:= \Omega(a), \end{aligned}$$

for an arbitrary state $a \in A$. \triangleleft

In words, \mathbb{A}^W is a bipartite automaton that we obtain from \mathbb{A} by putting a copy of \mathbb{A} ‘in front of’ a copy of \mathbb{A}^\perp , changing the transition map Θ of the initial part A of \mathbb{A}^W via the one-step translation $(\cdot)^W$. The final part of the structure \mathbb{A}^W is isomorphic to the automaton \mathbb{A}^\perp , so that we have

$$\mathbb{A}^W \langle a^\perp \rangle \equiv \mathbb{A}^\perp \langle a^\perp \rangle \quad (18)$$

for every state $a \in A$. Finally, observe that, while we define the transformation $(\cdot)^W$ for disjunctive automata only, the resulting structures are generally *not* disjunctive.

The following proposition is easy to verify, we leave the details for the reader.

Proposition 5.8 *Let $\mathbb{A} \langle a \rangle$ be an initialized disjunctive modal automaton which is positive in p . Then its transformation $\mathbb{A}^W \langle a \rangle$ belongs to the class IAut_p^W .*

We are now ready for the main technical lemma of this section.

Proposition 5.9 *Let $\mathbb{A} \langle a \rangle$ be an initialized disjunctive modal automaton which is positive in p . If $\mathbb{A} \langle a_I \rangle$ has the finite width property, then $\mathbb{A} \langle a_I \rangle \equiv \mathbb{A}^W \langle a_I \rangle$.*

Proof. Let $\mathbb{A} = (A, \Theta, \Omega)$ be a disjunctive modal automaton, and assume that, for some state $a_I \in A$, the initialized automaton $\mathbb{A} \langle a_I \rangle$ has the finite width property. In order to prove the equivalence of $\mathbb{A} \langle a_I \rangle$ and $\mathbb{A}^W \langle a_I \rangle$, it suffices to take an arbitrary ω -unravalled Kripke tree (\mathbb{S}, r) and prove that

$$\mathbb{S}, r \Vdash \mathbb{A} \langle a_I \rangle \text{ iff } \mathbb{S}, r \Vdash \mathbb{A}^W \langle a_I \rangle. \quad (19)$$

We first consider the direction from left to right of (19). Assume that $\mathbb{S}, r \Vdash \mathbb{A} \langle a_I \rangle$, then it follows from the finite width property of $\mathbb{A} \langle a \rangle$ that there is a finitely branching subtree $U \subseteq S$ such that $\mathbb{S}[p|U], r \Vdash \mathbb{A} \langle a_I \rangle$. Without loss of generality we may assume that $U \neq \emptyset$, or equivalently, that $r \in U$. By monotonicity of \mathbb{A}^W it suffices to show that $\mathbb{S}[p|U], r \Vdash \mathbb{A}^W \langle a_I \rangle$; that is, we need to supply \exists with a winning strategy h in the game $\mathcal{A}^W := \mathcal{A}(\mathbb{A}^W, \mathbb{S}[p|U]) @ (a_I, r)$. In order to define this strategy, we will make use of two auxiliary strategies: let f and g be positional winning strategies for \exists in the acceptance games $\mathcal{A}(\mathbb{A}, \mathbb{S}[p|U])$ and $\mathcal{A}(\mathbb{A}^W, \mathbb{S}[p|U])$ itself, respectively.

It will be \exists 's goal to maintain the following condition throughout any match of the acceptance game \mathcal{A}^W :

$$\begin{aligned} &\text{With } \Sigma = (a_n, s_n)_{n \leq k} \text{ a partial match of } \mathcal{A}^W, \text{ one of the following holds:} \\ (\dagger_W) \quad & \begin{cases} (\dagger_W^1) & (a_k, s_k) \in A \times U \text{ and } \Sigma \text{ corresponds to an } f\text{-guided match of } \mathcal{A}(\mathbb{A}, \mathbb{S}[p|U]), \\ (\dagger_W^2) & (a_l, s_l) \in A^\perp \times S \text{ for some } l \leq k \text{ such that } \mathbb{S}[p|U], s_l \Vdash \mathbb{A}^W \langle a_l \rangle, \\ & \text{and } (a_i, s_i)_{l \leq i \leq k} \text{ is a } g\text{-guided } \mathcal{A}(\mathbb{A}^W, \mathbb{S}[p|U])\text{-match.} \end{cases} \end{aligned}$$

In words, \exists will make sure that the match either stays in $A \times U$ and corresponds to an f -guided match of $\mathcal{A}(\mathbb{A}, \mathbb{S}[p|U])$, or it moves to $A^\perp \times S$ at a moment when it is safe for her to follow the strategy g . Let us first see that \exists can keep this condition during one single round of the game.

CLAIM 1 Let Σ be a partial match of \mathcal{A}^W satisfying (\dagger_W) . Then \exists has a legitimate move guaranteeing that, after any response move by \forall , (\dagger_W^2) holds again.

PROOF OF CLAIM Let $\Sigma = (a_n, s_n)_{n \leq k}$ be as in the claim, and distinguish cases.

The easiest case is where Σ satisfies (\dagger_W^2) : here \exists can simply continue to use her winning strategy g .

If Σ satisfies (\dagger_W^1) , then obviously its final position (a_k, s_k) is a winning position for \exists in $\mathcal{A}(\mathbb{A}, \mathbb{S}[p|U])$. Note that since $s_k \in U$, for its coalgebraic unfolding we have $\sigma_{\mathbb{S}[p|U]}(s_k) = \sigma_{\mathbb{S}}(s_k)$, so we may simply denote this object as $\sigma(s_k)$ without causing confusion. Let $m : \sigma_R(s_k) \rightarrow PA$ be the marking given by her positional winning strategy f . By the legitimacy of this move we have that $\sigma(s_k), m \Vdash^1 \pi \wedge \nabla B$ for some disjoint $\pi \wedge \nabla B$ of $\Theta(a_k)$.

If $B = \emptyset$ it immediately follows that $\sigma_R(s_k) = \emptyset$. It is easy to see that in this case, taking \exists 's move to be the *empty* map m^W and with $B_1 = B_2 = \emptyset$, we find $\sigma(s_k), m^W \Vdash^1 \pi \wedge \bigwedge \diamond B_1 \wedge \nabla B_2^\perp$. This means that m^W is a legitimate move for her, while \forall has no legitimate response to it. Thus \exists immediately wins (and the condition of the claim is satisfied).

Assuming in the sequel that $B \neq \emptyset$, we arrive at the heart of the proof. Define

$$\begin{aligned} B_1 &:= \{a \in A \mid a \in m(u) \text{ for some } u \in \sigma_R(s_k) \cap U\}, \\ B_2 &:= \{a \in A \mid a \in m(t) \text{ for some } t \in \sigma_R(s_k) \setminus U\}. \end{aligned}$$

Then clearly it follows from $\sigma(s_k), m \Vdash^1 \nabla B$ that $B_1 \cup B_2 = B$. A crucial observation is that since \mathbb{S} is ω -unravelled, and U is finitely branching, any $u \in \sigma_R(s_k) \cap U$ has a sibling $\underline{u} \in \sigma_R(s_k) \setminus U$ such that $\mathbb{S}, u \Leftrightarrow \mathbb{S}, \underline{u}$. Note that this bisimilarity does not hold for the structure $\mathbb{S}[p|U]$, but for the structure $\mathbb{S}[p \mapsto \emptyset]$ it does follow that

$$\mathbb{S}[p \mapsto \emptyset], u \Leftrightarrow \mathbb{S}[p \mapsto \emptyset], \underline{u}. \quad (20)$$

We can now define the desired move for \exists as the A^W -marking $m^W : \sigma_R(s_k) \rightarrow PA^W$ given by

$$m^W(t) := \begin{cases} m(t) \cup m(\underline{t})^\perp & \text{if } t \in U \\ m(t)^\perp & \text{if } t \notin U, \end{cases}$$

where $m(t)^\perp$ denotes the set $\{a^\perp \mid a \in m(t)\}$. We claim that m^W is a legitimate move for \exists at position (a_k, s_k) in $\mathcal{A}(\mathbb{A}, \mathbb{S}[p|U])$, and to prove this we need to show that

$$\sigma(s_k), m^W \Vdash^1 \Theta^W(a_k). \quad (21)$$

To do so, it clearly suffices to prove that

$$\sigma(s_k), m^W \Vdash^1 \bigwedge \diamond B_1 \wedge \nabla B_2^\perp. \quad (22)$$

For this purpose, first consider an arbitrary state $b \in B_1$. It is immediate by the definitions of B_1 and m^W that $b \in m^W(u)$ for some $u \in \sigma_R(s_k) \cap U$; hence we find that $\sigma(s_k), m^W \Vdash^1 \bigwedge \diamond B_1$, as required.

Now we consider the other conjunct, viz., the formula ∇B_2^\perp . First take an arbitrary element $b \in B_2$; it is immediate from the definitions of B_2 and m^W that $b \in m^W(u)$ for some $u \in \sigma_R(s_k) \setminus U$, so that $\sigma(s_k), m^W \Vdash^1 \diamond b$, as required. Conversely, take an arbitrary successor t of s_k . If $t \notin U$, it follows from $\sigma(s_k), m \Vdash^1 \nabla B$, the non-emptiness of B and the definition of B_2 , that $m(t) \cap B_2^\top \neq \emptyset$. If, on the other hand, t belongs to U , then by the same reasoning, but now applied to its *sibling* $\underline{t} \notin U$, we find that $m(\underline{t}) \cap B_2^\top \neq \emptyset$. In both cases it is immediate by the definition of m^W that $m^W(t) \cap B_2^\top \neq \emptyset$, so that

$\sigma(s_k), m^W \Vdash^1 \square \bigvee B_2^\perp$. Thus we obtain $\sigma(s_k), m^W \Vdash^1 \nabla B_2^\perp$, which means that we have established both (22) and (21), showing that m^W is a legitimate move indeed.

It remains to show that, playing m^W , \exists ensures that (\dagger_W^1) continues to hold after any response by \forall . So suppose that \forall picks a basic position (b, t) such that $b \in m^W(t)$. There are three cases to distinguish. First, if $b \in m(t)$ and $t \in U$, then by our assumption on m the continuation $\Sigma \cdot (b, t)$ of Σ clearly satisfies condition (\dagger_W^1) .

Second, suppose that $t \in \sigma_R(s_k) \setminus U$. Then by definition of m^W , b belongs to the A^\perp -part of \mathbb{A}^W , say $b = a^\perp$ for some $a \in A$. From the definition of m^W it follows that $a \in m(t)$, and since m is part of \exists 's winning strategy f , this means that $\mathbb{S}[p|U], t \Vdash \mathbb{A}\langle a \rangle$. Now observe that since $t \notin U$ and U is downward closed, the entire subtree generated by t is disjoint from U , so that we find $\mathbb{S}[p \mapsto \emptyset], t \Vdash \mathbb{A}\langle a \rangle$. We may now use Proposition 4.3 and obtain $\mathbb{S}[p|U], t \Vdash \mathbb{A}^\perp\langle a^\perp \rangle$. By (18) it is immediate from this that $\mathbb{S}[p|U], t \Vdash \mathbb{A}^W\langle a^\perp \rangle$. In other words, in this case the continuation match $\Sigma \cdot (a^\perp, t)$ satisfies condition (\dagger_W^2) .

Third, we consider the case where $t \in \sigma_R(s_k) \cap U$, and b , belonging to the A^\perp -part of \mathbb{A}^W , is of the form $b = a^\perp$ for some $a \in A$ such that $a \in m(t)$. Reasoning as in the previous case, but now for the sibling \underline{t} of t , we find that $\mathbb{S}[p \mapsto \emptyset], \underline{t} \Vdash \mathbb{A}\langle a \rangle$. But then it follows from (20) that $\mathbb{S}[p \mapsto \emptyset], t \Vdash \mathbb{A}\langle a \rangle$ as well. And, again reasoning as before, we find that in this case the match $\Sigma \cdot (a^\perp, t)$ satisfies condition (\dagger_W^2) as well. \blacktriangleleft

Based on Claim 1 we may provide \exists with the following strategy h . Given a partial match Σ , h picks any move for \exists as given by Claim 1 in case Σ satisfies (\dagger_W^1) , while h picks a random move if Σ does not meet mentioned condition. We will now prove that this is in fact a winning strategy.

CLAIM 2 Any h -guided full match Σ of \mathcal{A}^W is won by \exists .

PROOF OF CLAIM Let Σ be an h -guided full match of \mathcal{A}^W . Note that (a_I, r) is the first position of Σ , and that by our assumption that $r \in U$, the partial match $\Sigma_0 := (a_I, r)$ satisfies condition (\dagger_W^1) . It then easily follows from Claim 1 that, playing h , \exists will never get stuck.

To show that h is a winning strategy, we may thus focus on the case where $\Sigma = (a_n, s_n)_{n < \omega}$ is infinite; again by Claim 1 it follows that every initial part $\Sigma_l := (a_n, s_n)_{n \leq l}$ of Σ satisfies (\dagger_W^1) . From this it is obvious that we may distinguish the following two cases.

If every Σ_l satisfies (\dagger_W^1) , then Σ itself corresponds to an f -guided full match of $\mathcal{A}(\mathbb{A}^W, \mathbb{S}[p|U])$. Since f was assumed to be winning for \exists in $\mathcal{A}(\mathbb{A}, \mathbb{S}[p|U])$ from position (a_0, s_0) , this means that the A -stream $(a_n)_{n \in \omega}$ satisfies the acceptance condition of \mathbb{A} . But then this stream satisfies the acceptance condition of \mathbb{A}^W as well, which means that Σ is won by \exists indeed.

Alternatively, there is a first $l \in \omega$ such that Σ_l satisfies (\dagger_W^2) . It easily follows from Claim 1 that in this case, the match $(a_n, s_n)_{l \leq n < \omega}$ is a g -guided full match of $\mathcal{A}(\mathbb{A}^W, \mathbb{S}[p|U])$, which starts at a position, viz., (a_l, s_l) , which is winning for \exists . Clearly then the A -stream $(a_n)_{l \leq n \in \omega}$ satisfies the acceptance condition of \mathbb{A}^W ; from this it is immediate that the full A -stream $(a_n)_{n \in \omega}$ induced by Σ does so as well. \blacktriangleleft

Clearly, it follows from Claim 2 that h is a winning strategy for \exists in the game $\mathcal{A}^W = \mathcal{A}(\mathbb{A}^W, \mathbb{S}[p|U]) @ (a_I, r)$. This proves the direction from left to right of (19).

In order to prove the opposite, right-to-left, direction of (19), our line of reasoning is similar (but simpler). Assume that $\mathbb{S}, r \Vdash \mathbb{A}^W\langle a_I \rangle$, with $a_I \in A$. We will supply \exists with a winning strategy h in the game $\mathcal{A}(\mathbb{A}, \mathbb{S}) @ (a_I, r)$. For this purpose we will make use of arbitrary but fixed positional winning strategies f and g for \exists in the acceptance games $\mathcal{A}(\mathbb{A}^W, \mathbb{S})$ and $\mathcal{A}(\mathbb{A}, \mathbb{S})$, respectively.

\exists 's strategy in $\mathcal{A}(\mathbb{A}, \mathbb{S}) @ (a_I, r)$ will be based on maintaining the following condition:

- With $\Sigma = (a_n, s_n)_{n \leq k}$ a partial match of $\mathcal{A}(\mathbb{A}, \mathbb{S}) @ (a_I, r)$, one of the following holds:
- (\dagger_W^1) $(\dagger_W^1) \Sigma$ corresponds to an f -guided match of $\mathcal{A}(\mathbb{A}^W, \mathbb{S})$,
 - (\dagger_W^2) $\mathbb{S}, s_l \Vdash \mathbb{A}\langle a_l \rangle$ for some $l \leq k$, and $(a_i, s_i)_{l \leq i \leq k}$ is a g -guided $\mathcal{A}(\mathbb{A}, \mathbb{S})$ -match.

As before, our main claim is that \exists can keep condition (\dagger_W) during one single round of the game.

CLAIM 3 Let Σ be a partial match of $\mathcal{A}(\mathbb{A}, \mathbb{S})@(a_I, r)$ satisfying (\dagger_W) . Then \exists has a legitimate move guaranteeing that, after any response move by \forall , (\dagger_W) holds again.

PROOF OF CLAIM Let $\Sigma = (a_n, s_n)_{n \leq k}$ be as in the claim, and distinguish cases. If Σ satisfies (\dagger_W^2) , \exists can simply continue to use the strategy g .

If Σ satisfies (\dagger_W^1) , then obviously its final position (a_k, s_k) is a winning position for \exists in $\mathcal{A}(\mathbb{A}^W, \mathbb{S})$. Let $m : \sigma_R(s_k) \rightarrow PA^W$ be the marking given by her positional winning strategy f . By the legitimacy of this move we have that $\sigma(s_k), m \Vdash^1 \Theta^W(a_k)$, which means that

$$\sigma(s_k), m \Vdash^1 \pi \wedge \bigwedge \diamond B_1 \wedge \nabla B_2^\perp, \quad (23)$$

where $\pi \wedge \nabla(B_1 \cup B_2)$ is some disjunct of $\Theta(a_k)$.

We claim that the A -marking $m_W : \sigma_R(s_k) \rightarrow PA$, defined by

$$m_W(t) := \{a \in A \mid a \in m(t) \text{ or } a^\perp \in m(t)\}$$

is the right move for \exists in the partial match Σ . In order to prove the legitimacy of m_W , we show that:

$$\sigma(s_k), m_W \Vdash^1 \pi \wedge \nabla(B_1 \cup B_2). \quad (24)$$

which clearly implies that $\sigma(s_k), m_W \Vdash^1 \Theta(a_k)$. It easily follows from (23) that $\sigma(s_k), m_W \Vdash^1 \pi$, and so we may focus on the formula $\nabla(B_1 \cup B_2)$. First take an arbitrary successor t of s_k ; it follows from (23) that $m(t)$ contains some variable $b^\perp \in B_2^\perp$. But this immediately gives that $b \in m_W(t)$, showing that $\sigma(s_k), m_W \Vdash^1 \square \nabla(B_1 \cup B_2)$. Conversely, take an arbitrary state $b \in B_1 \cup B_2$; we need to show that $\sigma(s_k), m_W \Vdash^1 \diamond b$. Given the definition of m_W , this easily follows from $\sigma(s_k), m \Vdash^1 \bigwedge \diamond B_1$ if $b \in B_1$, and from $\sigma(s_k), m \Vdash^1 \nabla B_2^\perp$ if $b \in B_2$. Thus we have proved (24).

It is left to show that each response (b, t) of \forall to \exists 's move m_W constitutes a continuation $\Sigma \cdot (b, t)$ of Σ that satisfies (\dagger_W) . Let (b, t) be an arbitrary such move, that is, an arbitrary pair such that $b \in m_W(t)$. Again we distinguish cases: if $b \in m(t)$ then by our assumption of m being provided by \exists 's strategy f , the partial match $\Sigma \cdot (b, t)$ satisfies condition (\dagger_W^1) . Alternatively, if $b^\perp \in m(t)$, then by the assumption that f is a winning strategy for \exists in $\mathcal{A}(\mathbb{A}^W, \mathbb{S})@(a_I, r)$, we find that $\mathbb{S}, t \Vdash \mathbb{A}^W \langle b^\perp \rangle$. By (18) this obviously implies $\mathbb{S}, t \Vdash \mathbb{A}^\perp \langle b^\perp \rangle$, whence by Proposition 4.3 and monotonicity of \mathbb{A} we obtain, subsequently, $\mathbb{S}[p \mapsto \emptyset], t \Vdash \mathbb{A} \langle b \rangle$ and $\mathbb{S}, t \Vdash \mathbb{A} \langle b \rangle$. But then the continuation match $\Sigma \cdot (b, t)$ satisfies condition (\dagger_W^2) . \blacktriangleleft

We now define a strategy for \exists in the game $\mathcal{A}(\mathbb{A}, \mathbb{S})@(a_I, r)$: Given a partial match Σ , h picks any move for \exists as given by Claim 3 in case Σ satisfies (\dagger_W) , while h picks a random move if Σ does not meet mentioned condition. We claim that h is in fact a winning strategy at position (a_I, r) , and since one may prove this claim in the same manner as above, we leave the details for the reader. **QED**

Proof of Theorem 5.3. To define the required map $(\cdot)^W : \mu\text{ML}(\mathbf{X}) \rightarrow \mu\text{ML}_p^W(\mathbf{X})$, fix a μML -formula ξ .

By Fact 3.14(2) there is an initialized disjunctive modal automaton $\mathbb{D}_\xi \langle d_\xi \rangle$ that is equivalent to ξ ; using the constructions of Definition 4.8 and Definition 5.7, we transform \mathbb{D}_ξ into a finite-width automaton \mathbb{D}_ξ^{MW} . Finally, we apply the translation map of Definition 3.18, instantiated at the automaton \mathbb{D}_ξ^{MW} , to its state d_ξ , and we let ξ^W be the resulting formula. Summarizing, we define

$$\xi^W := \text{tr}_{\mathbb{D}_\xi^{MW}}(d_\xi).$$

It then follows by the Propositions 4.9, 5.8 and 5.6 that the formula ξ^W belongs to the fragment μML_p^W , and by the Propositions 4.10 and 5.9, together with the equivalences $\xi \equiv \mathbb{D}_\xi \langle d_\xi \rangle$ and $\mathbb{D}_\xi^{MW} \langle d_\xi \rangle \equiv \text{tr}_{\mathbb{D}_\xi^{MW}}(d_\xi)$, that ξ has the finite width property for p iff $\xi \equiv \xi^W$.

Finally, all constructions that are involved in the definition of the map $(\cdot)^W : \mu\text{ML}(\mathbf{X}) \rightarrow \mu\text{ML}_p^W(\mathbf{X})$ are effective, and can be performed in elementary time. In addition, the problem whether two formulas of the modal μ -calculus are equivalent or not can be solved in exponential time. From this it is immediate that the problem whether a given μ -formula has the finite width property, is decidable, and its time complexity is bounded by an elementary function on the size of the formula. QED

6 Finite depth property

The property that we consider in this section can be informally classified as ‘vertical’ in the sense that its definition involves subtrees containing no infinite paths.

Definition 6.1 A formula $\xi \in \mu\text{ML}(\mathbf{X})$ has the *finite depth property* for $p \in \mathbf{X}$ if ξ is monotone in p , and, for every tree model (\mathbb{S}, s) ,

$$\mathbb{S}, s \Vdash \xi \text{ iff } \mathbb{S}[p \upharpoonright U], s \Vdash \xi, \text{ for some noetherian subtree } U \subseteq S,$$

where we call U a *noetherian subtree* of \mathbb{S} if it is downward closed and contains no infinite paths. ◁

We will associate the following syntactic fragment of the modal μ -calculus with this property.

Definition 6.2 Given a set $P \subseteq \mathbf{X}$, we define the fragment μML_P^D by the following grammar:

$$\varphi ::= p \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond \varphi \mid \square \varphi \mid \mu x. \varphi',$$

where ψ is a p -free formula and $\varphi' \in \mu\text{ML}_{P \cup \{x\}}^D$. In case P is a singleton, say, $P = \{p\}$, we will write μML_p^D rather than $\mu\text{ML}_{\{p\}}^D$. ◁

In words, the fragment μML_p^D consists of those formulas of which no p -active subformula is in the scope of a greatest fixpoint operator ν . All formulas from Example 2.24 belong to μML_p^D , except φ_5 .

The following theorem states that modulo equivalence, μML_p^D is the syntactic fragment of the modal μ -calculus that captures the finite depth property, and that it is decidable whether a given μ -formula has this property.

Theorem 6.3 *There is an effective translation which maps a given μML -formula ξ to a formula $\xi^D \in \mu\text{ML}_p^D$ such that*

$$\xi \text{ has the finite depth property for } p \text{ iff } \xi \equiv \xi^D, \tag{25}$$

and it is decidable in elementary time whether a given formula ξ has the finite depth property for p .

First we prove the easy part of the theorem, stating that formulas in the fragment μML_p^D indeed have the required semantic property.

Proposition 6.4 *Every formula $\xi \in \mu\text{ML}_p^D$ has the finite depth property with respect to p .*

Proof. Let ξ be a formula in $\mu\text{ML}_W(p)$, then ξ is obviously positive, and hence, monotone in p . Fix a tree model \mathbb{S} with root $r \in S$. We have to prove

$$\mathbb{S}, r \Vdash \xi \text{ iff } \mathbb{S}[p \upharpoonright U], r \Vdash \xi, \text{ for some noetherian subtree } U \subseteq S. \tag{26}$$

The direction from right to left follows from the fact that ξ is monotone in p . For the opposite direction, suppose that $\mathbb{S}, r \Vdash \xi$. We need to find a noetherian subtree U of S such that $\mathbb{S}[p|U], r \Vdash \xi$.

Let f be a positional winning strategy of \exists in the game $\mathcal{E}_0 := \mathcal{E}(\xi, \mathbb{S})@(\xi, r)$. We define $U \subseteq S$ such that

$$u \in U \text{ iff there is } \varphi \text{ such that } (\varphi, u) \text{ is } f\text{-reachable in } \mathcal{E}_0 \text{ and } p \text{ is active in } \varphi.$$

As before it is easy to see that the set U is downward closed, so we omit the details. Suppose for contradiction that U contains an infinite path P . We let \mathcal{A} be the set of all finite f -guided \mathcal{E}_0 -matches Σ such that for all positions (φ, u) occurring in Σ , u belongs to P and p is active in φ . Recall that \sqsubseteq denotes the initial-segment relation on paths (including matches).

Clearly, the structure $(\mathcal{A}, \sqsubseteq)$ is a tree. Moreover, it is finitely branching as P is a single path, and all the formulas occurring in matches in \mathcal{A} belong to the finite set $Sfor(\xi)$. Next we show that the set \mathcal{A} is infinite. It suffices to define an injective map h from P to \mathcal{A} . Fix t in P . In particular, t belongs to U and by definition of U , there is a formula φ such that (φ, t) is f -reachable in \mathcal{E}_0 and p is active in φ . We let $h(t)$ be a finite f -guided \mathcal{E}_0 -match with last position (φ, t) . It is easy to check that any such map h is an injection from P to \mathcal{A} .

By König's lemma, since $(\mathcal{A}, \sqsubseteq)$ is infinite and finitely branching, it must contain an infinite path. This infinite path corresponds to an infinite f -guided \mathcal{E}_0 -match Σ such that for all positions (φ, t) occurring in Σ , t belongs to P and p is active in φ . Since ξ belongs to the fragment μML_p^D , this can only happen if all the variables unfolded in Σ are μ -variables. This implies that Σ is lost by \exists and thus contradicts the fact that f is a winning strategy for \exists in \mathcal{E}_0 . As a consequence, U contains no infinite path.

It remains to show that $\mathbb{S}[p|U], r \Vdash \xi$. Here we omit the details since the proof is similar to that of the corresponding statement in the proof of Proposition 5.4. QED

For the hard part of the theorem, we involve so-called finite-depth automata.

Definition 6.5 A bipartite modal automaton $\mathbb{A} = (A, B, \Theta, \Omega)$ belongs to the class Aut_p^D of *finite-depth automata* if the one-step language associated with B is the language $\mathbf{1ML}(X \setminus \{p\}, B)$, and the one-step language associated with A is given by the following grammar:

$$\alpha ::= p \mid \alpha_0 \mid \beta \mid \alpha \wedge \alpha \mid \top \mid \alpha \vee \alpha \mid \perp \tag{27}$$

where $\alpha_0 \in \mathbf{1ML}(X \setminus \{p\}, A)$ and $\beta \in \mathbf{1ML}(X \setminus \{p\}, B)$. Most importantly, we require that $\Omega(a)$ is *odd*, for every $a \in A$. \triangleleft

In words, an initialized bipartite modal automaton $\mathbb{A}\langle a_I \rangle$, with $\mathbb{A} = (A, B, \Theta, \Omega)$, belongs to the class $IAut_p^D$ if (1) p occurs only positively in $\Theta(a)$, for $a \in A$, (2) p does not occur in any $\Theta(b)$, $b \in B$, (3) $\Theta(a)$ is a (negation-free) propositional formulas obtained p and p -free one-step formulas over A and B , respectively, (4) all states in A have an odd priority. Note that it follows from (4) that in order to win a match of the satisfiability game, unless \forall gets stuck, \exists has to make sure that the automaton leaves its initial part at some moment.

Proposition 6.6 Let $\mathbb{A} = (A_0, A_1, \Theta, \Omega)$ be a bipartite modal automaton in Aut_p^D . Then there is a translation $\text{tr}_{\mathbb{A}} : A \rightarrow \mu\text{ML}$ such that $\text{tr}_{\mathbb{A}}(a) \in \mu\text{ML}_p^D$ for every state $a \in A_0$.

Proof. It suffices to check that the fragment $\mu\text{ML}^D := \{\mu\text{ML}_p^D \mid P \subseteq_{\omega} \text{PROP}\}$ and the automaton \mathbb{A} satisfy the conditions (i) – (iii) of Proposition 3.22, with $P = \{p\}$. In all cases, the proof is routine. Note that we only need to show that μML^D satisfies the condition (AC_{μ}) , not (AC_{ν}) . QED

It follows from Proposition 6.6 and Proposition 6.4 that any initialized automaton in $IAut_p^D$ has the finite depth property for p .

The main technical result of this section will be based on the following transformation of automata. Recall that \mathbb{A}^\perp is the automaton given in Definition 4.2.

Definition 6.7 Let $\mathbb{A} = (A, \Theta, \Omega)$ be a modal automaton which is positive in p . We define the automaton \mathbb{A}^D as the structure $(A^D, \Theta^D, \Omega^D)$ where $A^D := A \uplus A^\perp$, and the maps Θ^D and Ω^D are given by putting

$$\begin{aligned} \Theta^D(a) &:= \Theta(a) \vee \Theta(a)[b^\perp/b \mid b \in A] & \text{and} & & \Omega^D(a) &:= 1 \\ \Theta^D(a^\perp) &:= \Theta^\perp(a) & & & \Omega^D(a^\perp) &:= \Omega(a), \end{aligned}$$

for an arbitrary state $a \in A$. ◁

To obtain the automaton \mathbb{A}^D from \mathbb{A} we put a modified copy of \mathbb{A} ‘in front of’ a copy of \mathbb{A}^\perp . The modifications consist of changing the priority map Ω on A by assigning each state $a \in A$ priority 1, and of allowing the same transitions from A to the final part A^\perp of \mathbb{A}^D as to the initial part A itself. As before, this makes the final part of the structure \mathbb{A}^D isomorphic to the automaton \mathbb{A}^\perp , so that again we have

$$\mathbb{A}^D\langle a^\perp \rangle \equiv \mathbb{A}^\perp\langle a^\perp \rangle \quad (28)$$

for every state $a \in A$. Finally, observe that we define the transformation $(\cdot)^D$ for *arbitrary* (that is, not necessarily disjunctive) modal automata.

The following proposition is easy to verify.

Proposition 6.8 *Let $\mathbb{A}\langle a \rangle$ be an initialized modal automaton which is positive in p . Then its transformation $\mathbb{A}^D\langle a \rangle$ belongs to the class $IAut_p^D$.*

We are now ready for the main technical lemma of this section.

Proposition 6.9 *Let $\mathbb{A}\langle a \rangle$ be an initialized modal automaton which is positive in p . If $\mathbb{A}\langle a_I \rangle$ has the finite depth property with respect to p , then $\mathbb{A}\langle a_I \rangle \equiv \mathbb{A}^D\langle a_I \rangle$.*

Proof. Let $\mathbb{A} = (A, \Theta, \Omega)$ be a modal automaton, and assume that $\mathbb{A}\langle a_I \rangle$ has the finite depth property for some state $a_I \in A$. In order to prove the equivalence of $\mathbb{A}\langle a_I \rangle$ and $\mathbb{A}^D\langle a_I \rangle$, it suffices to take an arbitrary Kripke tree (\mathbb{S}, r) and prove that

$$\mathbb{S}, r \Vdash \mathbb{A}\langle a_I \rangle \text{ iff } \mathbb{S}, r \Vdash \mathbb{A}^D\langle a_I \rangle. \quad (29)$$

We first consider the direction from left to right of (29). Assume that $\mathbb{S}, r \Vdash \mathbb{A}\langle a_I \rangle$, then it follows from the finite depth property of $\mathbb{A}\langle a \rangle$ that there is a noetherian subtree $U \subseteq S$ such that $\mathbb{S}[p|U], r \Vdash \mathbb{A}\langle a_I \rangle$. We may assume that $r \in U$, and that U is *balanced* in the sense that for any $u \in U$, either $\sigma_R(u) \cap U = \emptyset$ or $\sigma_R(u) \subseteq U$. (This is without loss of generality: should U itself not meet these conditions, then we may proceed with the set $U' := U \cup \{r\} \cup \bigcup \{\sigma_R(u) \mid u \in U, \sigma_R(u) \cap U \neq \emptyset\}$.) By monotonicity of \mathbb{A}^D it suffices to show that $\mathbb{S}[p|U], r \Vdash \mathbb{A}^D\langle a_I \rangle$; that is, we need to supply \exists with a winning strategy h in the game $\mathcal{A}^D := \mathcal{A}(\mathbb{A}^D, \mathbb{S}[p|U]) @ (a_I, r)$. In order to define this strategy, we will make use of two auxiliary strategies: let f and g be positional winning strategies for \exists in the acceptance games $\mathcal{A}(\mathbb{A}, \mathbb{S}[p|U])$ and $\mathcal{A}(\mathbb{A}^D, \mathbb{S}[p|U])$, respectively.

\exists 's strategy in \mathcal{A}^D will be based on maintaining the following condition:

With $\Sigma = (a_n, s_n)_{n \leq k}$ a partial match of \mathcal{A}^D , one of the following holds:

$$(\dagger_D) \quad \begin{aligned} (\dagger_D^1) & (a_k, s_k) \in A \times U \text{ and } \Sigma \text{ corresponds to an } f\text{-guided match of } \mathcal{A}(\mathbb{A}, \mathbb{S}[p|U]), \\ (\dagger_D^2) & (a_l, s_l) \in A^\perp \times S \text{ for some } l \leq k \text{ such that } \mathbb{S}[p|U], s_l \Vdash \mathbb{A}^D\langle a_l \rangle, \\ & \text{and } (a_i, s_i)_{l \leq i \leq k} \text{ is a } g\text{-guided } \mathcal{A}(\mathbb{A}^D, \mathbb{S}[p|U])\text{-match.} \end{aligned}$$

In words, \exists will make sure that the match either stays in $A \times U$ while she can play f , or it moves to $A^\perp \times S$ at a moment when it is safe for her to follow the strategy g .

Let us first see that \exists can maintain this condition during one single round of the game.

CLAIM 1 Let Σ be a partial match of \mathcal{A}^D satisfying (\dagger_D) . Then \exists has a legitimate move guaranteeing that, after any response move by \forall , (\dagger_D) holds again.

PROOF OF CLAIM Let $\Sigma = (a_n, s_n)_{n \leq k}$ be as in the claim, and distinguish cases. Leaving the easy case, where Σ satisfies (\dagger_D^2) , for the reader, we assume that Σ satisfies (\dagger_D^1) . Then obviously its final position (a_k, s_k) is a winning position for \exists in $\mathcal{A}(\mathbb{A}, \mathbb{S}[p \upharpoonright U])$. Since $s_k \in U$ we have $\sigma_{\mathbb{S}[p \upharpoonright U]}(s_k) = \sigma_{\mathbb{S}}(s_k)$, so we may denote this object as $\sigma(s_k)$ without causing confusion. Let $m : \sigma_R(s_k) \rightarrow \text{PA}$ be the marking given by her positional winning strategy f at position (a_k, s_k) .

By the legitimacy of this move we have that $\sigma(s_k), m \Vdash^1 \Theta(a_k)$. In order to define \exists 's move m^D in $\mathcal{A}(\mathbb{A}^D, \mathbb{S}[p \upharpoonright U])$, we distinguish cases.

If $\sigma_R(s_k) \subseteq U$ we simply put $m^D := m$, and we leave it for the reader to verify that this move of \exists satisfies the requirements formulated in the claim.

Alternatively, by our assumption on U we have that $\sigma_R(s_k) \cap U = \emptyset$; in this case we put

$$m^D(t) := m(t)^\perp,$$

for each $t \in \sigma_R(s_k)$. It then easily follows from $\sigma(s_k), m \Vdash^1 \Theta(a)$ that $\sigma(s_k), m^D \Vdash^1 \Theta^\perp(a)$, and so $\sigma(s_k), m^D \Vdash^1 \Theta^D(a)$ by definition of Θ^D . In other words, m^D is a legitimate move.

To see that by picking this move \exists maintains the condition (\dagger_D) , consider an arbitrary response (b^\perp, t) of \forall such that $b^\perp \in m^D(t)$. By definition of m^D it follows that $b \in m(t)$, and so by the assumption that m is part of \exists 's winning strategy f , we obtain $\mathbb{S}[p \upharpoonright U], t \Vdash \mathbb{A}\langle b \rangle$. But now we may reason as before (cf. Claim 1 in the proof of Proposition 5.9): Since $t \notin U$ and U is downward closed, the entire subtree generated by t is disjoint from U , so that we find $\mathbb{S}[p \mapsto \emptyset], t \Vdash \mathbb{A}\langle b \rangle$. We may now use Proposition 4.3 and (28) to obtain $\mathbb{S}[p \upharpoonright U], t \Vdash \mathbb{A}^D\langle b^\perp \rangle$. In other words, in this case the continuation match $\Sigma \cdot (b^\perp, t)$ satisfies condition (\dagger_D^2) . \blacktriangleleft

Based on Claim 1 we may provide \exists with a winning strategy h in \mathcal{A}^D , exactly as in the proof of Proposition 5.9. This proves the direction from left to right of (29).

We now turn to the right-to-left direction of (29), for which we assume that $\mathbb{S}, r \Vdash \mathbb{A}^D\langle a_I \rangle$, with $a_I \in A$. In order to supply \exists with a winning strategy h in the game $\mathcal{A}(\mathbb{A}, \mathbb{S})@_I(a_I, r)$, we will make use of arbitrary but fixed positional winning strategies f and g for \exists in the acceptance games $\mathcal{A}(\mathbb{A}^D, \mathbb{S})$ and $\mathcal{A}(\mathbb{A}, \mathbb{S})$, respectively.

\exists 's strategy in $\mathcal{A}(\mathbb{A}, \mathbb{S})@_I(a_I, r)$ will be based on maintaining the following condition:

- With $\Sigma = (a_n, s_n)_{n \leq k}$ a partial match of $\mathcal{A}(\mathbb{A}, \mathbb{S})@_I(a_I, r)$, one of the following holds:
- (\dagger_D^1) $(\dagger_D^1) \Sigma$ corresponds to an f -guided match of $\mathcal{A}(\mathbb{A}^D, \mathbb{S})$,
 - (\dagger_D^2) $\mathbb{S}, s_l \Vdash \mathbb{A}\langle a_l \rangle$ for some $l \leq k$, and $(a_i, s_i)_{i \leq k}$ is a g -guided $\mathcal{A}(\mathbb{A}, \mathbb{S})$ -match.

Once more our main claim is that \exists can maintain condition (\dagger_D) during one single round of the game.

CLAIM 2 Let Σ be a partial match of $\mathcal{A}(\mathbb{A}, \mathbb{S})@_I(a_I, r)$ satisfying (\dagger_D) . Then \exists has a legitimate move guaranteeing that, after any response move by \forall , (\dagger_D) holds again.

PROOF OF CLAIM Let $\Sigma = (a_n, s_n)_{n \leq k}$ be as in the claim, and distinguish cases. If Σ satisfies (\dagger_D^2) , \exists can simply continue to use the strategy g .

If Σ satisfies (\dagger_D^1) , then obviously its final position $(a_k, s_k) \in A \times S$ is a winning position for \exists in $\mathcal{A}(\mathbb{A}^D, \mathbb{S})$. Let $m : \sigma_R(s_k) \rightarrow \text{PA}^D$ be the marking given by her positional winning strategy f ,

then we have $\sigma(s_k), m \Vdash^1 \Theta^D(a_k)$ by legitimacy of m . Now distinguish cases, as to which disjunct of $\Theta^D(a_k) = \Theta(a_k) \vee \Theta(a_k)[b^\perp/b \mid b \in A]$ holds at the one-step model $(\sigma(s_k), m)$.

In case $\sigma(s_k), m \Vdash^1 \Theta(a_k)$, we may assume without loss of generality that $m(t) \in P(A)$, for all $t \in \sigma_R(a_k)$. We now simply define \exists 's move m_D in $\mathcal{A}(\mathbb{A}, \mathbb{S})@_{(a_I, r)}$ by setting $m_D := m$. In this case it is straightforward to verify that m_D is legitimate and that for any response move (b, t) of \forall , the resulting partial match $\Sigma \cdot (b, t)$ satisfies condition (\ddagger_D^1) .

In case $\sigma(s_k), m \Vdash^1 \Theta(a_k)[b^\perp/b \mid b \in A]$, we may assume without loss of generality that $m(t) \in P(A^\perp)$, for all $t \in \sigma_R(a_k)$. We define \exists 's move m_D by putting

$$m_D(t) := \{b \in A \mid b^\perp \in m(t)\},$$

for each $t \in \sigma_R(s_k)$. It is straightforward to verify that $\sigma(s_k), m_D \Vdash^1 \Theta(a_k)$, which shows that m_D is a legitimate move for \exists in $\mathcal{A}(\mathbb{A}, \mathbb{S})$ at position (a_k, s_k) .

Now consider an arbitrary (legitimate) response (b, t) of \forall ; we claim that the resulting continuation $\Sigma \cdot (b, t)$ of Σ satisfies (\ddagger_D^2) . To see this, observe that by definition of m_D it follows from $b \in m_D(t)$ that $b^\perp \in m(t)$, which means that (b^\perp, t) is a legitimate move for \forall in $\mathcal{A}(\mathbb{A}^D, \mathbb{S})$ at position m . But since m is part of \exists 's winning strategy f , this means that $\mathbb{S}, t \Vdash \mathbb{A}^D \langle b^\perp \rangle$. It then follows from (28) that $\mathbb{S}, t \Vdash \mathbb{A}^\perp \langle b^\perp \rangle$. Hence by Proposition 4.3 we obtain that $\mathbb{S}[p \mapsto \emptyset], t \Vdash \mathbb{A} \langle b \rangle$, and then by monotonicity of \mathbb{A} that $\mathbb{S}, t \Vdash \mathbb{A} \langle b \rangle$. This means that $\Sigma \cdot (b, t)$ satisfies (\ddagger_D^2) indeed. \blacktriangleleft

Based on this claim, we define the usual strategy h for \exists in $\mathcal{A}(\mathbb{S}, \mathbb{S})$. To see why this strategy is winning for her, we consider an arbitrary h -guided match Σ of $\mathcal{A}(\mathbb{S}, \mathbb{S})@_{(a_I, r)}$, and focus on the case where $\Sigma = (a_n, s_n)_{n \in \omega}$ is infinite. Observe that because Ω^D assigns an *odd* priority to states in the initial part A of \mathbb{A}^D , it cannot be the case that Σ corresponds to a *full* f -guided $\mathcal{A}^D(\mathbb{A}^D, \mathbb{S})$ -match. Hence, there must be a stage $k \in \omega$ when the match moves out of the initial part of \mathbb{A}^D , which means that $\mathbb{S}, s_k \Vdash \mathbb{A} \langle a_k \rangle$ and $(a_n, s_n)_{k \leq n < \omega}$ is a g -guided match of $\mathcal{A}(\mathbb{A}, \mathbb{S})$. From this it is straightforward to derive that Σ is won by \exists . QED

Proof of Theorem 6.3. To define the required map $(\cdot)^D : \mu\text{ML}(\mathbf{X}) \rightarrow \mu\text{ML}_p^D(\mathbf{X})$, fix a μML -formula ξ . We define

$$\xi^D := \text{tr}_{\mathbb{A}_\xi^{MD}}(a_\xi),$$

where $\mathbb{A}_\xi \langle a_\xi \rangle$ is an initialized modal automaton that is equivalent to ξ , and $\text{tr}_{\mathbb{A}_\xi^{MD}}$ is the translation associated with the automaton \mathbb{A}_ξ^{MD} . We leave it for the reader to verify that this map satisfies the requirements stated by the theorem — the proof follows the same lines as that of Theorem 5.3. QED

7 The single branch property

As a variation of the finite-width property, we consider the single-branch property.

Definition 7.1 A formula $\xi \in \mu\text{ML}(\mathbf{X})$ has the *single branch property* for $p \in \mathbf{X}$ if ξ is monotone in p , and, for every tree model (\mathbb{S}, r) ,

$$\mathbb{S}, r \Vdash \xi \text{ iff } \mathbb{S}[p \upharpoonright U], r \Vdash \xi, \text{ for some branch } U \subseteq S,$$

where a subset $U \subseteq S$ is a *branch* if $U = \{s_n \mid n < \kappa\}$ for some (finite or infinite) path $(s_n)_{n < \kappa}$ starting at $s_0 = r$. \blacktriangleleft

The syntactic fragment corresponding to this property is the following.

Definition 7.2 Given a set $Q \subseteq X$ such that $p \in Q$, we define the fragment $\mu\text{ML}_{p,Q}^B$ by the following grammar:

$$\varphi ::= p \mid q \mid \psi \mid \varphi \vee \varphi \mid p \wedge \varphi \mid \psi \wedge \varphi \mid \diamond \varphi \mid \mu x. \varphi' \mid \nu x. \varphi',$$

where $q \in Q$, $\psi \in \mu\text{ML}(X \setminus Q)$ is a Q -free formula and $\varphi' \in \mu\text{ML}_{p,Q \cup \{x\}}^B$. In case Q is the singleton $\{p\}$, we will write μML_p^B rather than $\mu\text{ML}_{\{p\},\{p\}}^B$. \triangleleft

Formulas ξ in this fragment share with those in μML_p^W the property that no p -active subformula of ξ may be in the scope of a box modality. The difference with μML_p^W lies in the role of conjunctions: if $\varphi_0 \wedge \varphi_1$ is a subformula of $\xi \in \mu\text{ML}_p^B$, then p cannot be active in both conjuncts φ_i , unless one of these conjuncts is actually identical to p . All formulas from Example 2.24, except φ_4 and φ_6 , belong to μML_p^S .

A further difference with the earlier (and later) languages is that μML_p^B is defined using fragments of the form $\mu\text{ML}_{p,Q}^B$, where the propositional variables in Q have a slightly different role than p . This difference can be best explained by means of an example: where the formula $\nu x. \diamond(x \wedge \diamond p)$ does not belong to the fragment μML_p^W (and does not have the single-branch property), the formula $\nu x. \diamond(p \wedge \diamond x) \in \mu\text{ML}_p^W$ does have the property.

Our characterization theorem for the single-branch property reads as follows.

Theorem 7.3 *There is an effective translation which maps a given μML -formula ξ to a formula $\xi^B \in \mu\text{ML}_p^B$ such that*

$$\xi \text{ has the single branch property for } p \text{ iff } \xi \equiv \xi^B, \quad (30)$$

and it is decidable in elementary time whether a given formula ξ has the single branch property for p .

As before we start with the easy part of the proof.

Proposition 7.4 *Every formula $\xi \in \mu\text{ML}_p^B$ has the single branch property with respect to p .*

Proof. Let ξ be a formula in μML_p^B , then clearly ξ is monotone in p . Fix a tree model \mathbb{S} with root r . We have to show

$$\mathbb{S}, r \Vdash \xi \text{ iff } \mathbb{S}[p|U], r \Vdash \xi, \text{ for some branch } U \subseteq S. \quad (31)$$

The direction from right to left follows from the monotonicity of ξ in p . For the direction from left to right, suppose that $\mathbb{S}, r \Vdash \xi$. We need to find a branch U such that $\mathbb{S}[p|U], r \Vdash \xi$.

Since ξ is true at r in \mathbb{S} , \exists has a positional winning strategy f in the game $\mathcal{E}_0 = \mathcal{E}(\xi, \mathbb{S})@(\xi, s)$.

CLAIM 1 For every $k < \omega$ there is at most one f -guided match $\Sigma = (\varphi_n, s_n)_{n \leq k}$ of length k , such that p is active in every φ_n but $\varphi_k \neq p$.

PROOF OF CLAIM We prove the claim by induction on k . Call a match $\Sigma = (\varphi_n, s_n)_{n \leq k}$ p -active if p is active in every φ_n , $n \leq k$.

For the base case there is nothing to prove since there is only one match of length 0 altogether, viz., the match consisting of the starting position (ξ, r) .

For the inductive case, it suffices to show that if $\Sigma = (\varphi_n, s_n)_{n \leq k}$ is a p -active match, then it has at most one p -active continuation $\Sigma \cdot (\varphi_{k+1}, s_{k+1})$. To show this, we distinguish cases as to the shape of φ_k .

If φ_k is a proposition letter, then Σ is a full match, and thus has no continuation at all. (Observe that by assumption we must have $\varphi_k = p$, but this is not relevant for the argument.) If $\varphi_k = x$ for some bound variable x , then $(\varphi_{k+1}, s_{k+1}) = (\delta_x, s_k)$, and p is active in δ_k . If φ_k is either a disjunction or a formula of the form $\diamond \varphi'$, then, since we fixed \exists 's strategy, there is exactly one continuation $\Sigma \cdot (\varphi_{k+1}, s_{k+1})$, and therefore at most one p -active such continuation. Finally, if $\varphi = \varphi' \wedge \varphi''$ is a

conjunction, then by the constraints of well-formed formulas of μML_p^B , at most one of the conjuncts can be active in p , unless φ' or φ'' is identical to p . From this it is immediate that Σ has at most one p -active continuation $\Sigma \cdot (\varphi_{k+1}, s_{k+1})$ where $\varphi_{k+1} \neq p$.

This finishes the proof of the claim. \blacktriangleleft

We define $U \subseteq S$ such that

$$u \in U \text{ iff there is a } \varphi \text{ such that } (\varphi, u) \text{ is } f\text{-reachable in } \mathcal{E}_0 \text{ and } p \text{ is active in } \varphi.$$

It easily follows from the claim that U is a branch of \mathbb{S} . It is now a routine exercise to verify that (ξ, r) is a winning position for \exists in $\mathcal{E}(\xi, \mathbb{S}[p|U])$ — she may in fact use the very same strategy f . QED

For the hard part of the proof we involve the following class of automata.

Definition 7.5 A bipartite modal automaton $\mathbb{A} = (A, B, \Theta, \Omega)$ belongs to the class Aut_p^B of *single-branch automata* if the one-step language associated with B is the language $1\text{ML}(\mathbf{X} \setminus \{p\}, B)$, and the one-step language associated with A is given by the following grammar:

$$\alpha ::= p \mid \diamond a \mid \beta \mid \alpha \wedge \beta \mid \top \mid \alpha \vee \alpha \mid \perp \quad (32)$$

where $a \in A$ and $\beta \in 1\text{ML}(\mathbf{X} \setminus \{p\}, B)$. \triangleleft

In words, the two key conditions on a bipartite automaton $\mathbb{A} = (A, B, \Theta, \Omega)$ are that (i) the initial part A of \mathbb{A} can only be accessed from itself through a formula of the form $\diamond a$ (with $a \in A$), and that (ii), as before, the propositional variable p may only occur positively in the initial part, and not at all in the final part of \mathbb{A} .

Proposition 7.6 Let $\mathbb{A} = (A_0, A_1, \Theta, \Omega)$ be a bipartite modal automaton in Aut_p^B . Then there is a translation $\text{tr}_{\mathbb{A}} : A \rightarrow \mu\text{ML}$ such that $\text{tr}_{\mathbb{A}}(a) \in \mu\text{ML}_p^B$ for every state $a \in A_0$.

Proof. We prove this proposition using a (minor) variation of Proposition 3.22, where we define the fragment $F(Q) := \mu\text{ML}_{p,Q}^B$ only for subsets $Q \subseteq \text{PROP}$ such that $p \in Q$. (It is easy to check that this modification does not have an effect on the proof of the proposition.) The verification that (i) $\mu\text{ML}^B := \{\mu\text{ML}_{p,Q}^W \mid p \in Q \subseteq \text{PROP}\}$ satisfies the properties (EP), (SP1), (SP2) and (AC_η) for $\eta \in \{\mu, \nu\}$, can be established by routine proofs. We may prove that (ii) for all $a \in A$, the formula $\Theta(a)$ belongs to the set $\mu\text{ML}_{\{p\},A}^B$ by a straightforward formula induction on the formulas generated by the grammar (32). And, finally, as in the earlier cases it is easy to show the existence of a translation $\text{tr} : B \rightarrow \mu\text{ML}(\mathbf{X} \setminus \{p\})$; from this it is immediate that (iii) $\text{tr}(b) \in \mu\text{ML}_{\{p\}}^W$, for all $b \in B$. QED

Note that it follows from Proposition 7.6 and Proposition 7.4 that initialized automata in IAut_p^B have the single branch property.

We now turn to the key definition of this section, viz., the transformation $(\cdot)^B$ of a disjunctive automaton into an Aut_p^B -structure. This construction is very similar to the one we used in the section on the finite-width property; the difference lies in the one-step translation, where now, in order to translate a disjunct $\pi \wedge \nabla B$, we do not consider all ways to write B as a union $B = B_1 \cup B_2$, but only the ones where B_1 is either empty or a singleton.

Definition 7.7 Let $(\cdot)^B : 1\text{DML}(\mathbf{X}, A) \rightarrow 1\text{ML}(\mathbf{X}, A \uplus A^\perp)$ be the one-step translation given by the following inductive definition:

$$\begin{aligned} (\pi \wedge \nabla B)^B &:= (\pi \wedge \nabla B^\perp) \vee \bigvee \{ \pi \wedge \diamond b \wedge \nabla B_2^\perp \mid \{b\} \cup B_2 = B \} \\ \perp^B &:= \perp \\ (\alpha \vee \beta)^B &:= \alpha^B \vee \beta^B. \end{aligned}$$

Let $\mathbb{A} = (A, \Theta, \Omega)$ be a disjunctive modal automaton which is positive in p . We define the automaton \mathbb{A}^B as the bipartite structure $(A^B, \Theta^B, \Omega^B)$ where $A^B := A \uplus A^\perp$, and the maps Θ^B and Ω^B are given by putting

$$\begin{array}{ll} \Theta^B(a) & := \Theta(a)^B \\ \Theta^B(a^\perp) & := \Theta^\perp(a) \end{array} \quad \text{and} \quad \begin{array}{ll} \Omega^B(a) & := \Omega(a) \\ \Omega^B(a^\perp) & := \Omega(a), \end{array}$$

for an arbitrary state $a \in A$. ◁

We leave it for the reader to verify the following proposition.

Proposition 7.8 *Let $\mathbb{A}\langle a \rangle$ be an initialized disjunctive modal automaton which is positive in p . Then its transformation $\mathbb{A}^B\langle a \rangle$ belongs to the class $IAut_p^B$.*

We are now ready for the main technical lemma of this section.

Proposition 7.9 *Let $\mathbb{A}\langle a \rangle$ be an initialized disjunctive modal automaton which is positive in p . If $\mathbb{A}\langle a_I \rangle$ has the single-branch property, then $\mathbb{A}\langle a_I \rangle \equiv \mathbb{A}^B\langle a_I \rangle$.*

Proof. Let $\mathbb{A} = (A, \Theta, \Omega)$ be a disjunctive modal automaton, and assume that $\mathbb{A}\langle a_I \rangle$ has the single branch property, where $a_I \in A$. In order to prove the equivalence of $\mathbb{A}\langle a_I \rangle$ and $\mathbb{A}^B\langle a_I \rangle$, it suffices to take an arbitrary ω -unravalled Kripke tree (\mathbb{S}, r) and prove that

$$\mathbb{S}, r \Vdash \mathbb{A}\langle a_I \rangle \text{ iff } \mathbb{S}, r \Vdash \mathbb{A}^B\langle a_I \rangle. \quad (33)$$

Our proof of (33) is very similar to that of the corresponding statement (19) in the proof of Proposition 5.9. For this reason we will focus on the details where the two proofs differ.

We first consider the direction from left to right of (33). Assume that $\mathbb{S}, r \Vdash \mathbb{A}\langle a_I \rangle$, then it follows from the single branch property of $\mathbb{A}\langle a \rangle$ that there is a path $(u_n)_{n < \kappa}$ starting at r and such that, with $U := \{u_n \mid n < \kappa\}$, we have $\mathbb{S}[p \upharpoonright U], r \Vdash \mathbb{A}\langle a_I \rangle$. Without loss of generality we may assume that $(u_n)_{n < \kappa}$ is of *maximal* length, that is, either $\kappa = \omega$, or $\kappa < \omega$ and $\sigma_R(u_{\kappa-1}) = \emptyset$; it follows that in particular, $r \in U$. By monotonicity of \mathbb{A}^B it suffices to show that $\mathbb{S}[p \upharpoonright U], r \Vdash \mathbb{A}^B\langle a_I \rangle$; that is, we need to supply \exists with a winning strategy h in the game $\mathcal{A}^B := \mathcal{A}(\mathbb{A}^B, \mathbb{S}[p \upharpoonright U]) @ (a_I, r)$. In order to define this strategy, we will make use of two auxiliary strategies: let f and g be positional winning strategies for \exists in the acceptance games $\mathcal{A}(\mathbb{A}, \mathbb{S}[p \upharpoonright U])$ and $\mathcal{A}(\mathbb{A}^B, \mathbb{S}[p \upharpoonright U])$ itself, respectively. By Fact 3.12 and the disjunctivity of \mathbb{A} we may assume without loss of generality that at any position (a, s) that is winning for \exists in $\mathcal{A}(\mathbb{A}, \mathbb{S}[p \upharpoonright U])$, the marking picked by f assigns a *singleton* to each $t \in \sigma_R(s)$.

The condition that \exists will maintain when playing \mathcal{A}^B is the following.

With $\Sigma = (a_n, s_n)_{n \leq k}$ a partial match of \mathcal{A}^B , one of the following holds:

$$(\dagger_B) \quad \begin{array}{l} (\dagger_B^1) \ (a_k, s_k) \in A \times U \text{ and } \Sigma \text{ is an } f\text{-guided match of } \mathcal{A}(\mathbb{A}, \mathbb{S}[p \upharpoonright U]), \\ (\dagger_B^2) \ (a_l, s_l) \in A^\perp \times S \text{ for some } l \leq k \text{ such that } \mathbb{S}[p \upharpoonright U], s_l \Vdash \mathbb{A}^B\langle a_I \rangle, \\ \text{and } (a_i, s_i)_{l \leq i \leq k} \text{ is a } g\text{-guided } \mathcal{A}(\mathbb{A}^B, \mathbb{S}[p \upharpoonright U])\text{-match.} \end{array}$$

Observe that this corresponds seamlessly to the condition (\dagger_B) featuring in the finite-width case.

CLAIM 1 Let Σ be a partial match of \mathcal{A}^B satisfying (\dagger_B) . Then \exists has a legitimate move guaranteeing that, after any response move by \forall , (\dagger_B) holds again.

PROOF OF CLAIM Let $\Sigma = (a_n, s_n)_{n \leq k}$ be as in the claim. Leaving the easy case, where Σ satisfies (\dagger_B^2) , as a exercise for the reader, we focus on the case where (\dagger_B^1) holds. Here the final position (a_k, s_k) of Σ is a winning position for \exists in $\mathcal{A}(\mathbb{A}, \mathbb{S}[p \upharpoonright U])$, and with $m : \sigma_R(s_k) \rightarrow PA$ being the marking given by \exists 's positional winning strategy f , we have $\sigma(s_k), m \Vdash^1 \pi \wedge \nabla B$ for some disjunct $\pi \wedge \nabla B$ of $\Theta(a_k)$.

If s_k is the last point of the branch U , we define $m^B : \sigma_R(s_k) \rightarrow \mathcal{P}A^B$ by putting $m^B(t) := m(t)^\perp$, and we leave it for the reader to verify that m^B , as a move for \exists in $\mathcal{A}(\mathbb{A}^B, \mathbb{S}[p|U])$, satisfies the conditions of the claim.

Alternatively, let $u \in \sigma_R(s_k) \cap U$ be the (unique) successor of s_k on the branch U . By our assumption on m , there is a unique state $b \in A$ such that $m(u) = \{b\}$. As in the finite-width case, u has a sibling $\underline{u} \in \sigma_R(s_k) \setminus U$ such that $\mathbb{S}, u \leftrightarrow \mathbb{S}, \underline{u}$, and hence, $\mathbb{S}[p \mapsto \emptyset], u \leftrightarrow \mathbb{S}[p \mapsto \emptyset], \underline{u}$. Now define $m^B : \sigma_R(s_k) \rightarrow \mathcal{P}A^B$ by putting

$$m^B(t) := \begin{cases} m(t) \cup m(t)^\perp & \text{if } t = u \\ m(t)^\perp & \text{if } t \neq u. \end{cases}$$

In order to show that m^B is a legitimate move for \exists in $\mathcal{A}(\mathbb{A}^B, \mathbb{S}[p|U])$ at position (a_k, s_k) , we define $B_2 := \bigcup_{t \neq u} m(t)$. Then clearly $B = \{b\} \cup B_2$ and so it suffices to prove that

$$\sigma(s_k), m^B \Vdash^1 \diamond b \wedge \nabla B_2^\perp. \quad (34)$$

The proof of (34) proceeds exactly as that of (22) in the finite-width case. We omit the details, as we do for the proof that, for any response move (t, c) by \forall to \exists 's move m^B , the resulting continuation $\Sigma \cdot (t, c)$ satisfies condition (\dagger_B) . \blacktriangleleft

Based on this claim it is straightforward to define a winning strategy h for \exists in \mathcal{A}^B . This proves the direction from left to right of (33).

In order to prove the opposite, right-to-left, direction of (33), our line of reasoning is the same as in the finite-width case. Again we leave the details for the reader. QED

Proof of Theorem 7.3. Given a formula $\xi \in \mu\text{ML}(\mathbf{X})$, we define

$$\xi^B := \text{tr}_{\mathbb{D}_{\xi}^{MB}}(d_\xi).$$

It is not difficult to verify that this map $(\cdot)^B : \mu\text{ML}(\mathbf{X}) \rightarrow \mu\text{ML}_p^B(\mathbf{X})$ satisfies the requirements stated by the theorem — the proof follows exactly the same lines as that of Theorem 5.3. QED

8 Continuity

In this section we prove one of our main results, namely, we give a syntactic characterization of the continuous fragment of the modal μ -calculus. We recall the definition of continuity.

Definition 8.1 A μ -formula $\xi \in \mu\text{ML}(\mathbf{X})$ is *continuous* in $p \in \mathbf{X}$ if

$$\mathbb{S}, s \Vdash \xi \text{ iff } \mathbb{S}[p|U], s \Vdash \xi, \text{ for some finite subset } U \subseteq S$$

for every pointed model (\mathbb{S}, s) . \blacktriangleleft

We leave it for the reader to verify that continuity implies monotonicity: Any formula that is continuous in p is also monotone in p .

The property of continuity is of interest for at least two reasons: its link to the well-known topological notion of *Scott continuity* [27] (which also explains the name ‘continuity’), and its connection with the notion of *constructivity*.

Scott topology Recall that a *complete lattice* \mathbb{P} is a partially ordered set (P, \leq) in which each subset has a greatest lower bound (called the *meet*) and a least upper bound (called the *join*). If U is a subset of the lattice, we denote by $\bigwedge U$ the meet of U and by $\bigvee U$ the join of U . For example, for all sets S , the power set $\mathsf{P}(S)$, ordered by set inclusion, is a complete lattice.

Given a complete lattice (P, \leq) , a subset $D \subseteq P$ is *directed* if for every pair $d_1, d_2 \in D$ there is a $d \in D$ such that $d_1 \leq d$ and $d_2 \leq d$. A subset $U \subseteq P$ is called *Scott open* if it is upward closed (that is, if $u \in U$ and $u \leq v$ then $v \in U$), and satisfies, for any directed $D \subseteq P$, the property that $U \cap D \neq \emptyset$ whenever $\bigvee D \in U$. It is not hard to prove that the Scott open sets indeed provide a topology, the so-called *Scott topology*. For the associated topological notion of continuity, let (P, \leq) and (P', \leq') be two complete lattices. A map $f : P \rightarrow P'$ is *Scott continuous* if for all Scott opens $U' \subseteq P'$, the set $f^{-1}[U']$ is Scott open. It is a standard result that a map $f : P \rightarrow P'$ is Scott continuous iff f preserves *directed joins* (that is, if $D \subseteq P$ is directed, then $f(\bigvee D) = \bigvee' f[D]$).

To connect this to our notion of continuity, recall from the introduction that, given a formula ξ , a proposition letter p and a model $\mathbb{S} = (S, R, V)$, the map $\xi_p^{\mathbb{S}} : \mathsf{P}(S) \rightarrow \mathsf{P}(S)$ is defined by

$$\xi_p^{\mathbb{S}}(U) = \{s \in S \mid \mathbb{S}[p \mapsto U], s \Vdash \xi\}. \quad (35)$$

Now the link is given by the following Proposition. The (routine) proof of this Proposition is left as an exercise for the reader.

Proposition 8.2 *A μ -formula ξ is continuous in p iff for all models $\mathbb{S} = (S, R, V)$, the map $\xi_p^{\mathbb{S}} : \mathsf{P}(S) \rightarrow \mathsf{P}(S)$ is Scott continuous.*

Constructivity Basically, a formula is constructive if the ordinal approximation of its least fixpoint is always reached in at most ω steps.

Given a formula ξ and a model $\mathbb{S} = (S, R, V)$, we define, by induction on $i < \omega$, a map $\xi_p^i : \mathsf{P}(S) \rightarrow \mathsf{P}(S)$. We let ξ_p^0 be the identity map, and for $i < \omega$ define $\xi_p^{i+1} := \xi_p^{\mathbb{S}} \circ \xi_p^i$, where $\xi_p^{\mathbb{S}}$ is as in (35). A monotone formula ξ is *bounded* in p if for some natural number n , the least fixpoint of ξ is always reached in n steps (that is, $\xi_p^n(\emptyset) = \xi_p^{n+1}(\emptyset)$ for all models \mathbb{S}), and *constructive* in p if the least fixpoint is always reached in ω many steps (that is, for all models $\mathbb{S} = (S, R, V)$, the least fixpoint of the map ξ_p is equal to $\bigcup\{\xi_p^i(\emptyset) \mid i < \omega\}$).

In [40], Otto proved that it is decidable in exponential time whether a basic modal formula is bounded (and whether a given formula of the modal μ -calculus is equivalent to a basic modal formula). But to the best of our knowledge, decidability of constructivity (that is, the problem whether a given μ -formula is constructive in p) is an open problem. In passing we mention that Czarnecki [16] found, for each ordinal $\beta < \omega^2$, a formula ξ_β for which β is the least ordinal such that the least fixpoint of ξ_β is always reached in β steps.

The connection between the notions of continuity and constructivity is an intriguing one. It is a routine exercise to prove that continuity implies constructivity: if ξ is continuous in p , then it is also constructive in p . The opposite inclusion does not hold, as the examples $\xi_1(p) = \Box p \wedge \Box \Box \perp$ and $\xi_2(p) = \nu x.p \wedge \Diamond x$ testify. However, in the previous examples, we have $\mu p.\xi_1 \equiv \mu p.\Box \Box \perp$ and $\mu p.\xi_2 \equiv \mu p.\perp$. That is, in each case there is a continuous formula ψ_i that is equivalent to ξ_i ‘modulo an application of the least fixpoint operation’. This suggests the following question concerning the link between continuity and constructivity. Can we find, for any formula $\xi \in \mu\text{ML}$ which is constructive in p , a continuous formula ψ such that $\mu p.\xi \equiv \mu p.\psi$? We leave this as an open problem, as we do with the broader question whether there is a ‘nice’ syntactic fragment of the modal μ -calculus that captures constructivity in the sense that a formula ξ is constructive in p iff it is equivalent to a formula ξ' which belongs to the given fragment (and which would preferably be effectively obtainable from ξ).

We now turn to the main result of this section, namely, our characterization result for continuity. Our approach here is based on the observation that continuity can be seen as the combination of monotonicity, the finite depth property and the finite width property.

Proposition 8.3 *A μ -formula ξ is continuous in p iff it is monotone in p and has both the finite depth and the finite width property with respect to p .*

Proof. It is fairly easy to see that if a μ -calculus formula ξ is continuous in p , then it has both the finite depth and the finite width property for p . For some detail, let (\mathbb{S}, r) be a tree model such that $\mathbb{S}, r \Vdash \xi$. Assuming that ξ is continuous in p , we can find a finite set $U \subseteq S$ such that $\mathbb{S}[p \upharpoonright U] \Vdash \xi$. Now consider the set $\uparrow U := \{s \in S \mid U \subseteq R^*[s]\}$. Since U is finite, the set $\uparrow U$ is a subtree of \mathbb{S} that is both noetherian and finitely branching; and since $U \subseteq \uparrow U$, it follows by monotonicity that $\mathbb{S}[p \upharpoonright (\uparrow U)] \Vdash \xi$. This suffices to show that φ has both the finite depth property and the finite width property for p .

For the opposite implication, assume that $\xi \in \mu\text{ML}$ has all three properties mentioned. Fix a Kripke structure $\mathbb{S} = (S, R, V)$ and a point $s \in S$. We have to show

$$\mathbb{S}, s \Vdash \xi \text{ iff } \mathbb{S}[p \upharpoonright U], s \Vdash \xi, \text{ for some finite subset } U \subseteq S.$$

The direction from right to left follows from the fact that ξ is monotone in p . For the opposite direction, suppose that $\mathbb{S}, s \Vdash \xi$. Let \mathbb{T} , with root r , be a tree unravelling of (\mathbb{S}, s) , with $f : \mathbb{T} \rightarrow \mathbb{S}$ denoting the canonical bounded morphism.

Since ξ has the finite width property with respect to p , there is a downward closed subset $U_1 \subseteq T$ which is finitely branching and such that $\mathbb{T}[p \upharpoonright U_1], r \Vdash \xi$. But ξ also has the finite depth property with respect to p . Hence there is a subset U_2 of T such that U_2 is downward closed, does not contain any infinite path and satisfies $\mathbb{T}[p \upharpoonright U_1 \cap U_2], r \Vdash \xi$. By König's Lemma, the set $U := U_1 \cap U_2$ is finite.

We claim that

$$\mathbb{S}[p \upharpoonright f[U]] \Vdash \xi.$$

To see this, define $U' := f^{-1}[f[U]]$, then clearly $U \subseteq U'$. By monotonicity it follows from $\mathbb{T}[p \upharpoonright U], r \Vdash \xi$ that $\mathbb{T}[p \upharpoonright U'], r \Vdash \xi$. It is straightforward to verify that $\mathbb{T}[p \upharpoonright U'], r \Leftrightarrow \mathbb{S}[p \upharpoonright f[U]], s$, and so we find that $\mathbb{S}[p \upharpoonright f[U]], s \Vdash \xi$ by invariance under bisimilarity (Fact 2.20). This suffices, since clearly $f[U]$ is finite. QED

The syntactic fragment corresponding to continuity can be defined as follows.

Definition 8.4 Given a set $P \subseteq \mathbf{X}$, we define the fragment μML_P^C by the following grammar:

$$\varphi ::= p \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond \varphi \mid \mu x. \varphi',$$

where $\psi \in \mu\text{ML}(\mathbf{X} \setminus P)$ is a P -free formula and $\varphi' \in \mu\text{ML}_{P \cup \{x\}}^C$. In case P is a singleton, say, $P = \{p\}$, we will write μML_p^C rather than $\mu\text{ML}_{\{p\}}^C$. \triangleleft

Observe that this fragment is the intersection of the fragments μML_p^W and μML_p^D (defined in, respectively, Definition 5.2 and 6.2). That is, μML_p^C consists of those formulas $\xi \in \mu\text{ML}_p^M$ such that no p -active subformula ψ of ξ occurs in the scope of either a box modality or a greatest fixpoint operator. For instance, all formulas from Example 2.24 belong to μML_p^C , except φ_4 and φ_5 .

From these observations, the earlier results on μML_p^M , μML_p^W and μML_p^D , and Proposition 8.3 the following is immediate.

Proposition 8.5 *Every formula $\xi \in \mu\text{ML}_p^C$ is continuous in p .*

Our characterization then is as follows. As mentioned in the introduction, an earlier presentation of this result was given by the first author in [24]. Note that van Benthem [3] gave a characterization of the *first-order* formulas that are continuous (or ‘finitely distributive’, in his terminology) in a given predicate letter P .

Theorem 8.6 *There is an effective translation which, given a μML -formula ξ , computes a formula $\xi^C \in \mu\text{ML}_C(p)$ such that*

$$\xi \text{ is continuous in } p \text{ iff } \xi \equiv \xi^C, \quad (36)$$

and it is decidable in elementary time whether a given formula ξ is continuous in p .

The automata characterizing continuity are defined as follows.

Definition 8.7 A bipartite modal automaton $\mathbb{A} = (A_0, A_1, \Theta, \Omega)$ belongs to the class Aut_p^C of *p-continuous automata* if, relative to this partition, \mathbb{A} is both a finite-width and a finite-depth automaton, with respect to p . \triangleleft

Proposition 8.8 *Let $\mathbb{A} = (A_0, A_1, \Theta, \Omega)$ be a bipartite modal automaton in Aut_p^C . Then there is a translation $\text{tr}_{\mathbb{A}} : A \rightarrow \mu\text{ML}$ such that $\text{tr}_{\mathbb{A}}(a) \in \mu\text{ML}_p^C$ for every state $a \in A_0$.*

The proof of Proposition 8.8 is straightforward, on the basis of the corresponding proofs for the finite-depth and the finite-width automata.

We are now ready for the main technical lemma of this section.

Proposition 8.9 *Let $\mathbb{A}\langle a_I \rangle$ be an initialized disjunctive modal automaton which is positive in p . If $\mathbb{A}\langle a_I \rangle$ is continuous in p , then $\mathbb{A}\langle a_I \rangle \equiv \mathbb{A}^{WD}\langle a_I \rangle$.*

Proof. Let $\mathbb{A}\langle a_I \rangle$ be as in the formulation of the proposition, then $\mathbb{A}\langle a_I \rangle$ obviously is monotone in p . The proposition then follows from the following chain of equivalences and implications:

$$\begin{aligned} & \mathbb{A}\langle a_I \rangle \text{ is continuous in } p \\ \iff & \mathbb{A}\langle a_I \rangle \text{ has both the finite width and the finite depth property for } p && \text{(Proposition 8.3)} \\ \implies & \mathbb{A}\langle a_I \rangle \equiv \mathbb{A}^W\langle a_I \rangle \text{ and } \mathbb{A}\langle a_I \rangle \text{ has the finite depth property for } p && \text{(Proposition 5.9)} \\ \iff & \mathbb{A}\langle a_I \rangle \equiv \mathbb{A}^W\langle a_I \rangle \text{ and } \mathbb{A}^W\langle a_I \rangle \text{ has the finite depth property for } p && \text{(obvious)} \\ \implies & \mathbb{A}\langle a_I \rangle \equiv \mathbb{A}^W\langle a_I \rangle \text{ and } \mathbb{A}^W\langle a_I \rangle \equiv \mathbb{A}^{WD}\langle a_I \rangle && \text{(Proposition 6.9)} \\ \implies & \mathbb{A}\langle a_I \rangle \equiv \mathbb{A}^{WD}\langle a_I \rangle && \text{(obvious)} \end{aligned}$$

QED

Proof of Theorem 8.6. Given a formula $\xi \in \mu\text{ML}(X)$, define

$$\xi^C := \text{tr}_{\mathbb{D}_\xi^{MWD}}(d_\xi).$$

It is then straightforward to show that $\xi \equiv \xi^C$ iff ξ is continuous in p , and this characterization can also be used to prove the decidability of continuity. (Alternatively and more efficiently, we can use Proposition 8.3, together with the decidability of each of the three properties mentioned there that are, jointly taken, equivalent to continuity.) QED

9 Full and complete additivity

The last two properties of formulas that we look at both concern the way in which the semantics of the formula depends on the proposition letter p being true at some single point.

Definition 9.1 A formula $\xi \in \mu\text{ML}$ is *fully additive* in p if we have

$$\mathbb{S}, s \Vdash \xi \text{ iff } \mathbb{S}[p \upharpoonright u], s \Vdash \xi, \text{ for some } u \in V(p),$$

completely additive in p if

$$\mathbb{S}, s \Vdash \xi \text{ iff } \mathbb{S}[p \upharpoonright u], s \Vdash \xi, \text{ for some } u \in S,$$

and *normal* if $\mathbb{S}[p \mapsto \emptyset], s \not\Vdash \xi$, for every pointed model (\mathbb{S}, s) . Here and in the sequel we write $\mathbb{S}[p \upharpoonright u]$ instead of $\mathbb{S}[p \upharpoonright \{u\}]$. \triangleleft

The properties defined in Definition 9.1 go back to Jónsson & Tarski [33, 34], as do the terms ‘completely additive’ and ‘normal’; we have introduced the term ‘fully additive’ here. It is easy to see (especially using the definitions in terms of the semantic map $\varphi_p^{\mathbb{S}}$ that we gave in the introduction), that full additivity is equivalent to the combination of normality and complete additivity.

In the context of modal logic, the property of full additivity is of interest for at least two reasons: its role in the *duality* theory of modal logic [44], and its link with the notion of *safety* for bisimulations.

Discrete duality In the algebraic approach to modal logic, two dualities feature prominently: a topological duality linking modal algebras to certain topological Kripke frames consisting of a relational structure which is expanded with a nicely fitting Stone topology, and a *discrete* duality connecting ordinary Kripke frames to so-called *perfect* modal algebras. The latter structures consist of a complete and atomic Boolean algebra, which is expanded with an additional operation that is fully additive (that is, preserves all joins of the algebra). This discrete duality, formulated by Thomasson [43], is based on a 1–1 correspondence, going back to Jónsson & Tarski [33, 34], between the fully additive maps on the power set of S and the binary relations on S . Here, the relation associated with a fully additive map f on PS is given as

$$Q_f := \{(s, s') \mid s \in f(\{s'\})\}, \quad (37)$$

while conversely, every binary relation R on S gives rise to the fully additive map $\langle R \rangle$ defined by

$$\langle R \rangle(U) := \{s \in S \mid R[s] \cap U \neq \emptyset\}. \quad (38)$$

In other words, the discrete duality concerns the semantics of the modal *diamond*.

Safety for bisimulation A second and more specific reason for studying full additivity concerns its key role in the characterization of formulas that are safe for bisimulation. To define this notion, consider a formula $\alpha(x, y)$ in some appropriate language for describing Kripke models. This formula induces, on every Kripke model \mathbb{S} , a binary relation $R_\alpha^{\mathbb{S}} := \{(s, t) \mid \mathbb{S} \models \alpha(s, t)\}$. Given two models \mathbb{S} and \mathbb{S}' , we call a relation $Z \subseteq S \times S'$ an α -*bisimulation* if it is a bisimulation for the relations $R_\alpha^{\mathbb{S}}$ and $R_\alpha^{\mathbb{S}'}$ (in the sense of Definition 2.10, with $R_\alpha^{\mathbb{S}}$ and $R_\alpha^{\mathbb{S}'}$ replacing the relations R and R' , respectively), and we say that α is *safe for bisimulation* if every ordinary bisimulation is also an α -bisimulation. This notion was introduced by van Benthem [4], who also gave a characterization of the safe fragment of first-order logic, that is, the set of first-order formulas $\alpha(x, y)$ that are safe for bisimulation.

The link with the notion of full additivity is provided by the discrete duality just described: the idea is that we can encode the transformations (37) and (38) in the syntax of the ambient logic, which in our context is monadic second-order logic (MSO). More specifically, we already saw how a formula $\alpha(x, y)$ induces a binary relation $R_\alpha^{\mathbb{S}}$ on every Kripke model. Now consider a formula $\beta(x; p)$, where x is an individual variable and p is a monadic predicate, which, in the style of modal correspondence theory, we will also think of as a proposition letter. Such an MSO-formula β induces a map $f_\beta^{\mathbb{S}}$ on any Kripke model \mathbb{S} , given by $f_\beta^{\mathbb{S}}(U) := \{s \in S \mid \mathbb{S}[p \mapsto U] \models \beta(s; p)\}$.

Now, given an MSO-formula $\alpha(x, y)$, define the formula

$$\alpha^*(x; p) := \exists y(\alpha(x, y) \wedge p(y)), \quad (39)$$

where p is a fresh monadic predicate, and conversely, given an MSO-formula $\beta(x; p)$ which is fully additive in p , define the formula

$$\beta_*(x, y) := \beta(x)[\lambda z.z=y/p(z)], \quad (40)$$

where y is a fresh individual variable, and $[\lambda z.z=y/p(z)]$ is the substitution replacing all atomic formulas of the form $p(z)$ by $z = y$. It is not hard to show that $\alpha^*(x; p)$ is always fully additive in p , and we leave it for the reader to verify that (39) and (40) encode, respectively, (38) and (37), in the sense that

$$\langle R_\alpha^S \rangle = f_{\alpha^*}^S \text{ and } Q_{f_\beta^S} = R_{\beta_*}^S.$$

As a manifestation of the discrete duality, we find that $\alpha \equiv (\alpha^*(x; p))_*(y)$ and $\beta \equiv (\beta_*(x, y))^*(x; p)$.

The key observation, which can be proved by a routine argument, is now that

$$\alpha(x, y) \text{ is safe for bisimulation iff } \alpha^*(x; p) \text{ is bisimulation invariant.} \quad (41)$$

At this point we invoke the Janin-Walukiewicz Theorem [32], which states that an MSO-formula $\gamma(x)$ is bisimulation invariant iff is equivalent to (the standard translation of) a μML -formula γ^\diamond , which may be effectively obtained from γ . Combining this with (41), we obtain that

$$\alpha(x, y) \text{ is safe for bisimulation iff } \alpha^*(x; p) \equiv (\alpha^*(x; p))^\diamond. \quad (42)$$

In particular, an MSO-formula $\alpha(x, y)$ is safe for bisimulation iff the MSO-formula $\alpha^*(x; p)$ is equivalent to a μML -formula that is fully additive in p .

Thus, a syntactic characterization of the fully additive modal μ -formulas also yields a syntactic characterization of the safe fragment of monadic second-order logic: Suppose that $F_p \subseteq \mu\text{ML}$ characterizes (modulo equivalence) the fragment of the modal μ -calculus that characterizes full additivity in p , then the set $\{(ST_x(\varphi))_*(x, y) \mid \varphi \in F_p\}$ characterizes (modulo equivalence) the bisimulation-safe fragment of MSO; here $ST_x : \mu\text{ML} \rightarrow \text{MSO}$ denotes some standard truth-preserving translation mapping μML -formulas to MSO-formulas with one free individual variable x . Such a characterization was first obtained by Hollenberg [30]; we will come back to his results in Remark 9.11.

We now turn to our syntactic characterizations of the fully and completely additive modal fixpoint formulas.

Definition 9.2 Given a set $P \subseteq \mathbf{X}$, we define the fragment μML_P^F by the following grammar:

$$\varphi ::= p \mid \perp \mid \varphi \vee \varphi \mid \varphi \wedge \psi \mid \diamond\varphi \mid \mu x.\varphi',$$

where $p \in P$, $\psi \in \mu\text{ML}(\mathbf{X} \setminus P)$ is a P -free formula and $\varphi' \in \mu\text{ML}_{P \cup \{x\}}^F$. Similarly, we define the fragment μML_p^A by induction in the following way:

$$\varphi ::= p \mid \perp \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \psi \mid \diamond\varphi \mid \mu x.\varphi,$$

where $p \in P$, $\psi \in \mu\text{ML}(\mathbf{X} \setminus P)$ is a P -free formula and $\varphi' \in \mu\text{ML}_{P \cup \{x\}}^A$. In case P is a singleton, say, $P = \{p\}$, we will write μML_p^F and μML_p^A rather than, respectively, $\mu\text{ML}_{\{p\}}^F$ and $\mu\text{ML}_{\{p\}}^A$. \triangleleft

We let the grammars of Definition 9.2 speak for itself. Of the formulas from Example 2.24, φ_0 , φ_2 and φ_3 belong to μML_p^F ; these formulas also belong to μML_p^A , as does φ_1 . The difference between the fragments μML_p^A and μML_p^F is that p -free formulas belong to μML_p^A but not to μML_p^F (except the formula \perp).

Theorem 9.3 (i) *There is an effective translation which, given a μML -formula ξ , computes a formula $\xi^F \in \mu\text{ML}_p^F$ such that*

$$\xi \text{ is fully additive in } p \text{ iff } \xi \equiv \xi^F, \quad (43)$$

and it is decidable in elementary time whether a given formula ξ is fully additive in p .

(ii) *Similarly, there is an effective translation which, given a μML -formula ξ , computes a formula $\xi^A \in \mu\text{ML}_p^A$ such that*

$$\xi \text{ is completely additive in } p \text{ iff } \xi \equiv \xi^A, \quad (44)$$

and it is decidable in elementary time whether a given formula ξ is completely additive in p .

In the sequel we will only prove the first part of Theorem 9.3, the proof for complete additivity is a variant of this. We first consider the easy direction of Theorem 9.3(i).

Proposition 9.4 *Every formula $\xi \in \mu\text{ML}_p^F$ is fully additive in p .*

Proof. Let ξ be a formula in μML_p^F , then clearly ξ is monotone in p . Fix a tree model \mathbb{S} with root r . We have to show

$$\mathbb{S}, r \Vdash \xi \text{ iff } \mathbb{S}[p \upharpoonright u], r \Vdash \xi, \text{ for some point } u \in V(p). \quad (45)$$

The direction from right to left follows from the monotonicity of ξ in p . For the direction from left to right, suppose that $\mathbb{S}, r \Vdash \xi$. We need to find a point $u \in V(p)$ such that $\mathbb{S}[p \upharpoonright u], r \Vdash \xi$.

Since ξ is true at r in \mathbb{S} , \exists has a positional winning strategy f in the game $\mathcal{E}_0 = \mathcal{E}(\xi, \mathbb{S}) @ (\xi, s)$. Similar to the proof of Proposition 7.4, we can prove the following claim.

CLAIM 1 For every $k < \omega$ there is at most one f -guided match $\Sigma = (\varphi_n, s_n)_{n \leq k}$ of length k , such that p is active in every φ_n .

In proving this claim, the difference with the single-branch case is that now, the only p -active conjunctions are of the form $p \wedge \psi$, where ψ is not p -active. Hence, a partial match ending in a position with such a conjunction, will have exactly one p -active continuation.

It follows from Claim 1 that there is a unique *maximal* p -active match $\Sigma = (\varphi_n, s_n)_{n < \kappa}$.

CLAIM 2 Σ is finite and its last position is the unique position in Σ of the form (p, u) .

PROOF OF CLAIM It is easy to see that Σ must be finite, since otherwise, being f -guided, it should be won by \exists , while the only p -active bound variables of ξ are *least* fixpoint variables (cf. the proof of Proposition 6.4 in the finite-depth case). It should also be clear that Σ can have at most one position of the form (p, s) (since at such a position the match will be over).

We may thus consider the final position (φ, r) of Σ . It follows by maximality of Σ that φ cannot be of the form x (with x a bound variable of ξ), $\diamond \varphi'$, $\varphi_0 \wedge \varphi_1$, or $\varphi_0 \vee \varphi_1$ — in the latter case, both disjuncts φ_i would be p -active. Hence the only possibility left is that $\varphi = p$ indeed. \blacktriangleleft

To finish the proof, let $u \in S$ be the state such that the pair (p, u) is the final position of Σ . It is easy to check that (ξ, r) is a winning position for \exists in $\mathcal{E}(\xi, \mathbb{S}[p \upharpoonright u])$ — she may use the very same strategy f as in $\mathcal{E}(\xi, \mathbb{S})$. QED

For the proof of the hard direction of Theorem 9.3, we introduce the following class of automata.

Definition 9.5 A bipartite modal automaton $\mathbb{A} = (A, B, \Theta, \Omega)$ belongs to the class Aut_p^F of *finite-width automata* if the one-step language associated with B is the language $\mathbf{1ML}(\mathbf{X} \setminus \{p\}, B)$, and the one-step language associated with A is given by the following grammar:

$$\alpha ::= p \mid \diamond a \mid \perp \mid \beta \wedge \alpha \mid \alpha \vee \alpha \quad (46)$$

where $a \in A$ and $\beta \in \mathbf{1ML}(\mathbf{X} \setminus \{p\}, B)$. \triangleleft

Proposition 9.6 Let $\mathbb{A} = (A, B, \Theta, \Omega)$ be a bipartite modal automaton in Aut_p^F . Then there is a translation $\text{tr}_{\mathbb{A}} : A \rightarrow \mu\text{ML}$ such that $\text{tr}_{\mathbb{A}}(a) \in \mu\text{ML}_p^F$ for every state $a \in A$.

Proof. Once more we will apply Proposition 3.22. Since it is fairly obvious that \mathbb{A} satisfies the conditions (ii) and (iii) of mentioned proposition, and that $\mu\text{ML}^F := \{\mu\text{ML}_P^F \mid P \subseteq_{\omega} \text{PROP}\}$ satisfies the properties (EP), (SP2) and (CA_{μ}) , we only show that μML^F satisfies the first substitution property.

For this purpose we will prove, by induction on φ , that $\varphi[\psi/x]$ belongs to the fragment μML_P^F whenever $\varphi \in \mu\text{ML}_{P \cup \{x\}}^F$ and $\psi \in \mu\text{ML}_P^F$. We leave the easy cases as exercises to the reader, and note that the case where $\varphi = \mu y.\varphi'$ is dealt with exactly as in the proof of Proposition 5.6. We focus on the case where φ is of the form $\chi \wedge \varphi'$ because $\varphi' \in \mu\text{ML}_{P \cup \{x\}}^F$ and χ is $P \cup \{x\}$ -free. It then follows that $\chi[\psi/x] = \chi$ is P -free, and by induction that $\varphi'[\psi/x] \in \mu\text{ML}_P^F$. But from this it is immediate that $\varphi[\psi/x] = \chi[\psi/x] \wedge \varphi'[\psi/x]$ belongs to the fragment μML_P^F indeed. \square

As before, our main result is based on a transformation of an arbitrary disjunctive automaton into an automaton in the class Aut_p^F .

Definition 9.7 Let $(\cdot)^F : 1\text{DML}(X, A) \rightarrow 1\text{ML}(X, A \uplus A^{\perp})$ be the one-step translation given by the following inductive definition:

$$\begin{aligned} (\pi \wedge \nabla B)^F &:= \begin{cases} \pi \wedge \nabla B^{\perp} & \text{if } p \in \pi \\ \perp & \text{if } p \notin \pi \text{ and } B = \emptyset \\ \bigvee \{\pi \wedge \diamond b \wedge \nabla B_2^{\perp} \mid \{b\} \cup B_2 = B\} & \text{if } p \notin \pi \text{ and } B \neq \emptyset \end{cases} \\ \perp^F &:= \perp \\ (\alpha \vee \beta)^F &:= \alpha^F \vee \beta^F. \end{aligned}$$

Let $\mathbb{A} = (A, \Theta, \Omega)$ be a disjunctive modal automaton which is positive in p . Without loss of generality we may assume that for every $a \in A$ and every disjunct $\pi \wedge \nabla B$ of $\Theta(a)$, either p or $\neg p$ (but not both) is a conjunct of π . (If not, we may replace $\pi \wedge \nabla B$ with the formula $(\pi \wedge p) \wedge \nabla B \vee (\pi \wedge \neg p) \wedge \nabla B$.) We define the automaton \mathbb{A}^F as the structure $(A^F, \Theta^F, \Omega^F)$, where $A^F := A \uplus A^{\perp}$, and the maps Θ^F and Ω^F are given by putting

$$\begin{aligned} \Theta^F(a) &:= \Theta(a)^F & \text{and} & & \Omega^F(a) &:= 1 \\ \Theta^F(a^{\perp}) &:= \Theta^{\perp}(a) & & & \Omega^F(a^{\perp}) &:= \Omega(a), \end{aligned}$$

for an arbitrary state $a \in A$. \triangleleft

Proposition 9.8 Let $\mathbb{A}\langle a \rangle$ be an initialized disjunctive modal automaton which is positive in p . Then its transformation $\mathbb{A}^F\langle a_I \rangle$ belongs to the class IAut_p^F .

Proposition 9.9 Let $\mathbb{A}\langle a \rangle$ be an initialized disjunctive modal automaton which is positive in p . If $\mathbb{A}\langle a_I \rangle$ is fully additive in p , then $\mathbb{A}\langle a_I \rangle \equiv \mathbb{A}^F\langle a_I \rangle$.

Proof. Let $\mathbb{A} = (A, \Theta, \Omega)$ be a disjunctive modal automaton, and assume that, for some state $a_I \in A$, $\mathbb{A}\langle a_I \rangle$ is fully additive in p . In order to prove the equivalence of $\mathbb{A}\langle a_I \rangle$ and $\mathbb{A}^F\langle a_I \rangle$, it suffices to take an arbitrary ω -unravalled Kripke tree (\mathbb{S}, r) and prove that

$$\mathbb{S}, r \Vdash \mathbb{A}\langle a_I \rangle \text{ iff } \mathbb{S}, r \Vdash \mathbb{A}^F\langle a_I \rangle. \quad (47)$$

Our proof of (47) is very similar to that of the corresponding statements in the proofs of Proposition 7.9 and Proposition 6.9; for this reason we will be brief, often referring for details to these earlier proofs.

We first consider the direction from left to right of (47). Assume that $\mathbb{S}, r \Vdash \mathbb{A}\langle a_I \rangle$, then it follows by the full additivity of $\mathbb{A}\langle a \rangle$ in p that there is a point $u \in V(p)$ such that $\mathbb{S}[p \upharpoonright u], r \Vdash \mathbb{A}\langle a_I \rangle$. Let $\rho = (u_n)_{n \leq m}$ be the unique path from $r = u_0$ to $u = u_m$, and define $U := \{t_n \mid n \leq m\}$ to be the set of points on this path. By monotonicity of \mathbb{A}^F it suffices to show that $\mathbb{S}[p \upharpoonright u], r \Vdash \mathbb{A}^F\langle a_I \rangle$; that is, we need to supply \exists with a winning strategy h in the game $\mathcal{A}^F := \mathcal{A}(\mathbb{A}^F, \mathbb{S}[p \upharpoonright u]) @ (a_I, r)$. We will write $\sigma'_V := \sigma_{\mathbb{S}[p \upharpoonright u]}$ and $\sigma' := (\sigma'_V, \sigma_R)$ for the coalgebraic unfolding map of the Kripke structure $\mathbb{S}[p \upharpoonright u]$.

In order to define the strategy h , we will make use of positional winning strategies f and g for \exists in the acceptance games $\mathcal{A}(\mathbb{A}, \mathbb{S}[p \upharpoonright u])$ and $\mathcal{A}(\mathbb{A}^F, \mathbb{S}[p \upharpoonright u])$ itself, respectively. By Fact 3.12 we may assume without loss generality that at any position (a, s) that is winning for \exists in $\mathcal{A}(\mathbb{A}, \mathbb{S}[p \upharpoonright u])$, the marking picked by f assigns a *singleton* to each $t \in \sigma_R(s)$.

The condition that \exists will maintain when playing \mathcal{A}^F is the following.

- With $\Sigma = (a_n, s_n)_{n \leq k}$ a partial match of \mathcal{A}^F , one of the following holds:
- (\dagger_F^1) $(a_k, s_k) \in A \times U$ and Σ is an f -guided match of $\mathcal{A}(\mathbb{A}, \mathbb{S}[p \upharpoonright u])$,
 - (\dagger_F^2) $(a_l, s_l) \in A^\perp \times S$ for some $l \leq k$ such that $\mathbb{S}[p \upharpoonright u], s_l \Vdash \mathbb{A}^F\langle a_l \rangle$,
and $(a_i, s_i)_{l \leq i \leq k}$ is a g -guided $\mathcal{A}(\mathbb{A}^F, \mathbb{S}[p \upharpoonright u])$

CLAIM 1 Let Σ be a partial match of \mathcal{A}^F satisfying (\dagger_F). Then \exists has a legitimate move guaranteeing that, after any response move by \forall , (\dagger_F) holds again.

PROOF OF CLAIM The proof of this claim is a subtle variation on that of the corresponding claim in the single-branch case. Let $\Sigma = (a_n, s_n)_{n \leq k}$ be as in the claim. Leaving the easy case, where Σ satisfies (\dagger_F^2), as an exercise for the reader, we focus on the case where (\dagger_F^1) holds.

Here the path $(s_n)_{n \leq k} = (u_n)_{n \leq k}$ is an initial segment of the branch U , and the final position (a_k, s_k) of Σ is a winning position for \exists in $\mathcal{A}(\mathbb{A}, \mathbb{S}[p \upharpoonright u])$. With $m : \sigma_R(s_k) \rightarrow PA$ being the marking given by \exists 's positional winning strategy f , we have $\sigma'(s_k), m \Vdash^1 \Theta(a_k)$. We now distinguish cases.

If $s_k = u$ is the last point of the branch U , then $\mathbb{S}[p \upharpoonright u], s_k \Vdash p$, and so we have $\sigma'(s_k), m \Vdash^1 \pi \wedge \nabla B$ for some disjunct $\pi \wedge \nabla B$ of $\Theta(a_k)$ such that p is a conjunct of π . (Here we use the fact that for every disjunct $\pi \wedge \nabla B$ of $\Theta(a)$, either p or $\neg p$ is a conjunct of π). We define the \mathbb{A}^F -marking m^F on $\sigma_R(s_k)$ by putting $m^F(t) := m(t)^\perp$, and we leave it for the reader to verify that m^F , as a move for \exists in $\mathcal{A}(\mathbb{A}^F, \mathbb{S}[p \upharpoonright u])$, satisfies the conditions of the claim.

If $s_k \neq u$, then p is false at s_k . Let $u_{k+1} \in \sigma_R(s_k)$ be the (unique) successor of $s_k = u_k$ on the branch U . By our assumption on m , there is a unique state $b \in A$ such that $m(u_{k+1}) = \{b\}$. Note that in this situation, we have that $\sigma(s_k), m \Vdash^1 \pi \wedge \nabla B$ for some disjunct $\pi \wedge \nabla B$ of $\Theta(a_k)$ such that $B \neq \emptyset$ and $\neg p$ is a conjunct of π . From this we may infer that the translation $(\pi \wedge \nabla B)^F$ is not equal to \perp , but of the form $\bigvee \{ \pi \wedge \diamond b \wedge \nabla B_2^\perp \mid \{b\} \cup B_2 = B \}$. In order to define a suitable marking m^F , we continue as in the single-branch case. Let u be the unique successor of s_k in U , let b be the unique state in A such that $m(u) = \{b\}$, and let $\underline{u} \in \sigma_R \setminus U$ be some sibling of u such that $\mathbb{S}, u \leftrightarrow \mathbb{S}, \underline{u}$. Define

$$m^F(t) := \begin{cases} m(t) \cup m(t)^\perp & \text{if } t = u \\ m(t)^\perp & \text{if } t \neq u. \end{cases}$$

The verification, that with this definition the marking m^F meets all the specifications of the claim, is exactly as in the proof of Proposition 7.9, and so we omit the details. \blacktriangleleft

On the basis of Claim 1 we may define, in the by now familiar way, a strategy h for \exists which is winning for her in $\mathcal{A}(\mathbb{A}^F, \mathbb{S}[p \upharpoonright u])$. We omit the details.

The opposite direction of (47) can be proved by a similar argument as in the proof of Proposition 6.9, so again we omit the details. QED

Proof of Theorem 9.3. As mentioned earlier on, we only cover part (i) of the theorem explicitly, part (ii) can be proved by a fairly obvious variation of this.

To define the required map $(\cdot)^F : \mu\text{ML}(\mathbf{X}) \rightarrow \mu\text{ML}_P^F(\mathbf{X})$, fix a μML -formula ξ . We define

$$\xi^F := \text{tr}_{\mathbb{D}_\xi^{MF}}(d_\xi),$$

where $\mathbb{D}_\xi\langle d_\xi \rangle$ is an initialized disjunctive modal automaton that is equivalent to ξ , and $\text{tr}_{\mathbb{D}_\xi^{MF}}$ is the translation associated with the automaton \mathbb{D}_ξ^{MF} . We leave it for the reader to verify that this map satisfies the requirements stated by the theorem — the proof follows the same lines as that of Theorem 5.3. QED

We finish this section with a series of remarks that provide some context to our results.

Remark 9.10 There are interesting connections between the fragments μML_P^F and μML_P^A , and the language PDL of *propositional dynamic logic* [29]. Since PDL is by nature a *poly-modal* language, to make our point we momentarily switch to the poly-modal μ -calculus. Carreiro & Venema [13] showed that PDL has the same expressive power as the fragment of μML in which the formula construction $\mu x.\varphi$ is allowed only if φ is completely additive with respect to x . More precisely, define the set μ_{AML} of formulas by the following grammar:

$$\varphi ::= q \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle a \rangle \varphi \mid \mu q.\varphi',$$

where q is an arbitrary proposition letter, and φ' belongs to the fragment $\mu\text{ML}_q^A \cap \mu_{\text{AML}}$. Then there are inductive, truth-preserving translations from PDL to μ_{AML} and vice versa [13]. ◁

Remark 9.11 Hollenberg’s characterization [30] of the fully additive fragment⁴ of the modal μ -calculus has a strong connection with propositional dynamic logic as well. He defines the sets of extended μ -formulas φ and so-called μ -programs π by the following simultaneous induction (again we take a poly-modal perspective):

$$\begin{aligned} \varphi &::= q \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle \pi \rangle \varphi \mid \mu q.\varphi \\ \pi &::= a \mid \varphi? \mid \pi + \pi \mid \pi; \pi \mid \pi^*, \end{aligned}$$

where q is an arbitrary propositional variable, and a is an atomic program; in $\mu q.\varphi$, q may only occur positively in φ . Hollenberg proves that a formula ξ is fully additive in p iff ξ is equivalent to a formula of the form $\langle \pi \rangle p$, where π is a p -free μ -program.

Comparing Hollenberg’s result to ours, while his characterization is clearly well-suited to find the safe fragment of monadic second-order logic, our result has the advantage of directly providing a characterizing fragment inside the modal μ -calculus. But in any case, there are direct translations between our fragment and Hollenberg’s.

From Hollenberg’s fragment to ours, by a simultaneous induction on formulas and programs one may define a translation $(\cdot)^\tau$ mapping a formula φ in Hollenberg’s language to a formula $\varphi^\tau \in \mu\text{ML}$, and, for each μ -program π , a function $f_\pi : \mu\text{ML} \rightarrow \mu\text{ML}$, in such a way that f_π restricts to the fragment μML_P^F if π is P -free. Some key clauses in this definition are $(\langle \pi \rangle \varphi)^\tau := f_\pi(\varphi^\tau)$, $f_{\psi?}(\varphi) := \psi \wedge \varphi$, and $f_{\pi^*}(\varphi) := \mu x.\varphi \vee f_\pi(x)$, where x is a fresh variable.

Conversely, by a straightforward formula induction one may provide, for each formula $\xi \in \mu\text{ML}_P^F$, a collection $\{\pi_p \mid p \in P\}$ of P -free μ -programs such that

$$\xi \equiv \bigvee \{ \langle \pi_p \rangle p \mid p \in P \}.$$

⁴Note that Hollenberg’s terminology clashes with ours: what he calls ‘completely additive’ is what we call ‘fully additive’.

The key induction step here is for a formula of the form $\mu x.\xi$, where we may infer from the above equivalence that $\mu x.\xi \equiv \bigvee \{ \langle \pi_x^*; \pi_p \rangle p \mid p \in P \setminus \{x\} \}$.

We refrain from giving more details here, referring the interested reader to section 3 of Carreiro & Venema [13], where very similar translations between PDL and $\mu_F\text{ML}$ are defined (cf. Remark 9.10), or to Carreiro [10]. \triangleleft

Remark 9.12 As another variation of the properties of *full* and *complete* additivity, a formula $\xi \in \mu\text{ML}$ is *finitely additive* in $p \in X$ if, for every Kripke model \mathbb{S} ,

$$\xi_p^{\mathbb{S}}\left(\bigcup \mathcal{X}\right) = \bigcup \left\{ \xi_p^{\mathbb{S}}(X) \mid X \in \mathcal{X} \right\},$$

for any *finite* collection \mathcal{X} of subsets of S . This condition can be equivalently expressed by requiring that the formula ξ is both *additive* ($\xi(p \vee p') \equiv \xi(p) \vee \xi(p')$) and *normal* ($\xi(\perp) \equiv \perp$) in p . In the case of basic modal logic, the two properties can be shown to be equivalent (through a straightforward argument based on finite trees of depth not exceeding the modal depth of ξ), but this is not so in the case of the modal μ -calculus. For instance, consider the formula $\nu y.\mu x.(p \wedge \diamond y) \vee \diamond x$, expressing the existence of an infinite path, starting at the current state, where p holds infinitely often. It is easy to see that this formula is finitely but neither fully nor completely additive in p .

We leave it as an open problem to characterize the finitely p -additive fragment of μML . \triangleleft

10 Conclusions

We finish the paper with drawing some conclusions, listing some issues for discussion, and suggesting some questions for further research.

This paper contributes to the theory of the modal μ -calculus by proving some model-theoretic results. For a number of semantic properties pertaining to formulas of the modal μ -calculus, we provided a corresponding syntactic fragment, showing that a μ -formula ξ has the given property iff it is equivalent to a formula ξ' in the corresponding fragment. Since this formula ξ' will always be effectively obtainable from ξ , as a corollary, for each of the properties under discussion, we prove that it is decidable in elementary time whether a given μ -calculus formula has the property or not.

The properties that we study have in common that they all concern the dependence of the truth of the formula at stake, on a single proposition letter p . In each case the semantic condition on ξ will be that ξ , if true at a certain state in a certain model, will remain true if we restrict the set of states where p holds, to a special subset of the state space. Important examples include the properties of full additivity and continuity, where the special subsets are the singletons and the finite sets, respectively.

Our proofs for these characterization results will be fairly uniform in nature, employing the well-known correspondence between formulas of the modal μ -calculus, and *modal automata*. In fact, the effectively defined maps on formulas are induced by rather simple transformations on modal automata, based on composing a bipartite automaton \mathbb{A}' from an arbitrary (disjunctive) automaton \mathbb{A} , where the final part of \mathbb{A} consists of the automaton \mathbb{A}^\perp and its initial part of another modification of \mathbb{A} . This modification is always obtained by applying a straightforward one-step translation to the transition map of \mathbb{A} , by redefining its priority map, or by a combination of these operations.

Discussion

1. As mentioned in the introduction, pure logic-based proofs for our results are possible in almost all cases — the exception being the single-branch property where we only have automata-theoretic proofs. In fact, logic-based proofs were given in the dissertation of the first author [25]. The main advantage of the automata-theoretic approach is that it allows for transparent and uniform proofs based on simple transformations of automata.

In any case, the difference between the two approaches should not be exaggerated. Recall that the particular shape of our automata is logic-based: the transition map of our structures uses so-called *one-step formulas*, and many of our proofs are based on semantic properties of and syntactic manipulations on these very simple modal formulas. In some sense then, our paper is also a contribution to the model theory of modal automata.

2. As mentioned in the introduction we have not undertaken an in-depth study of the computational *complexity* of the various problems of which we established the decidability. It should be clear that the algorithms that we have presented here are not optimal. In particular, in order to find out whether a formula ξ has, say, the finite width property, it is not needed to compute its translation ξ^W : it suffices to check whether the initialized automata $\mathbb{D}(d_\xi)$ and $\mathbb{D}_\xi^W(d_\xi)$ are equivalent. To obtain a good upper bound here, one should know the exact size and weight of the automaton \mathbb{D}_ξ in terms of the size $|\xi|$ of ξ . Fact 3.14(ii) gives a doubly exponential weight for \mathbb{D}_ξ , but we conjecture that a tighter bound is possible. We leave this, and other complexity-theoretic matters, as questions for further research.
3. There are some variations of our results that are not hard to prove. To start with, all characterization results (and their proofs) can be easily restricted to the setting of basic (i.e., fixpoint-free) modal logic. For instance, if we define the fragment ML_p^C by the following grammar:

$$\varphi ::= p \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond \varphi$$

where $\psi \in \mu\text{ML}(X \setminus \{p\})$ is a p -free formula, then we can show that our map $(\cdot)^C$ maps formulas in ML to ML_p^C . As a result we find that a basic modal logic formula ξ is continuous in p iff $\xi \equiv \xi^C$, so that ML_p^C characterizes continuity-in- p for basic modal logic.

4. Recall that in the presentation of the language μML (as in Definition 2.13), the standard restriction on the occurrence of the least fixpoint operator μx is that it can be applied only to formulas that are positive in x , i.e., belong to the language μML_x^M . We get interesting logics by restricting the application of μ -operators even further. This applies in particular to the fragments μML^D , μML^C and μML^A discussed in this paper. For $Q \in \{D, C, A\}$, let $\mu_Q\text{ML}$ be the version of the modal μ -calculus of which the formulas are given by the following grammar:

$$\varphi ::= p \mid \neg \varphi \mid \varphi \vee \varphi \mid \diamond \varphi \mid \mu x. \varphi', \quad (48)$$

where p is a propositional variable, and the formation of the formula $\mu x. \varphi$ is subject to the constraint that the formula φ' belongs to the fragment μML_x^Q . We already saw in Remark 9.10 that the language $\mu_A\text{ML}$ is effectively equivalent to PDL. It is not hard to prove that the logic $\mu_D\text{ML}$ is effectively equivalent to the *alternation-free* fragment of the modal μ -calculus, whereas it seems that the logic $\mu_C\text{ML}$ has not been used or studied much (although it was mentioned under the name ‘ ω - μ -calculus’ by van Benthem [5], and it is related, and perhaps equivalent in expressive power, to the logic CPDL of *concurrent propositional dynamic logic*, cf. Carreiro [10, section 3.2] for more information).

These logics become particularly interesting in the light of the Janin-Walukiewicz Theorem [32]. Recall that this result states that the modal μ -calculus is the bisimulation-invariant fragment of monadic second-order logic (MSO), in brief: $\mu\text{ML} \equiv \text{MSO}/\simeq$. For each of the logics $\mu_Q\text{ML}$, with $Q \in \{D, C, A\}$ we can prove the following version of this result:

$$\mu_Q\text{ML} \equiv \text{MSO}_Q/\simeq,$$

where MSO_Q is a variant of MSO where we quantify over a restricted collection $P_Q(S)$ of subsets of the model \mathbb{S} . More specifically, $P_D(S)$ consists of the so-called *noetherian* sets of a Kripke model [23], $P_C(S)$ is the collection of *finite* subsets of S [12] (so that MSO_D is *weak* monadic second-order logic), and $P_A(S)$ is the set of so-called *generalized finite chains* in \mathbb{S} [11].

5. In section 2 we proved a strengthened version of the Lyndon Theorem for the modal μ -calculus proved by D'Agostino and Hollenberg [17]. In the same vein as the other results in this paper, we can also strengthen their *Los-Tarski Theorem*.

We say that a formula $\xi \in \mu\text{ML}$ is *preserved under substructures* if $\mathbb{S}, s \Vdash \xi$ implies $\mathbb{S}', s \Vdash \xi$, whenever \mathbb{S}' is a substructure of \mathbb{S} (in the standard model-theoretic sense). D'Agostino and Hollenberg proved that a μ -formula ξ is preserved under substructures iff it is equivalent to a *universal* formula, that is, a formula in the \diamond -free fragment μML^U given by the following grammar:

$$\varphi ::= q \mid \neg q \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond \varphi \mid \mu x. \varphi \mid \nu x. \varphi.$$

We can reprove this result by our means; given a disjunctive automaton $\mathbb{D} = (D, \Theta, \Omega)$, define the automaton $\mathbb{D}^U = (D, \Theta^U, \Omega)$ where Θ^U is given by the one-step translation based on the clause $(\pi \bullet \nabla B)^U := \pi \wedge \square \bigvee B$. One may then show that an initialized disjunctive automaton $\mathbb{D}\langle d \rangle$ is preserved under taking substructures iff $\mathbb{D}\langle d \rangle \equiv \mathbb{D}^U\langle d \rangle$. From this the result of D'Agostino and Hollenberg easily follows, and as a bonus we find that it is decidable in elementary time whether a given formula $\xi \in \mu\text{ML}$ has this property.

Questions

Finally, we mention some open problems for further research.

1. It would be interesting to find out the *exact* complexity of the problems discussed in this paper. This would include establishing suitable *lower bounds*.
2. As mentioned in the section on continuity, it would be good to know whether the μ -calculus formulas that are *constructive* in p admit a good syntactic characterization. In particular, we would like to clarify the connection between the notions of continuity and constructivity. Can we find, for any formula $\xi \in \mu\text{ML}$ which is constructive in p , a continuous formula ψ such that $\mu p. \xi \equiv \mu p. \psi$?
3. Similarly, we would be curious to see a syntactic characterization of the finitely p -additive fragment of the modal μ -calculus (cf. Remark 9.12).
4. While, as already mentioned, some variations of our results are easy to obtain, there are some *interesting variations* of the problems considered here as well. For instance, it is not so clear how to adapt our characterisation results to other fixpoint logics like PDL or CTL. A second direction to take here would be to look for *coalgebraic generalisations* of our results. In recent years it has been shown that many results on the modal μ -calculus, including the link with automata theory, can be generalized to the far wider setting of coalgebraic modal logic [36, 15, 26, 21].
5. Not directly related to the results in this paper, but in our opinion one of the most interesting open model-theoretic problems concerning the modal μ -calculus is whether μML admits a natural abstract characterization in the form of a *Lindström theorem*. Going back to de Rijke [42], there are various Lindström-type characterizations of basic modal logic (see for instance [6, 41, 37, 20]) but to the best of our knowledge no abstract characterizations of fixpoint logics have been established yet.

References

- [1] S. Abramsky and A. Jung. Domain theory. In S. Abramsky, D.M. Gabbay, and T.S.M Maibaum, editors, *Handbook for Logic in Computer Science*. Clarendon Press, Oxford, 1994.

- [2] A. Arnold and D. Niwiński. *Rudiments of μ -calculus*, volume 146 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Publishing Co., Amsterdam, 2001.
- [3] J. van Benthem. Logic as programming. *Fundamenta Informaticae*, 17(4):285–317, 1992.
- [4] J. van Benthem. *Exploring logical dynamics*. CSLI publications, Stanford, 1996.
- [5] J. van Benthem. Modal frame correspondences and fixed-points. *Studia Logica*, 83:133–155, 2006.
- [6] J. van Benthem. A new modal Lindström theorem. *Logica Universalis*, 1:125–138, 2007.
- [7] J. van Benthem, P. Blackburn, and F. Wolter, editors. *Handbook of Philosophical Logic*. Elsevier, 2006.
- [8] J. Bradfield and C. Stirling. Modal μ -calculi. In Benthem et al. [7], pages 721–756.
- [9] F. Bruse, O. Friedmann, and M. Lange. On guarded transformation in the modal μ -calculus. *Logic Journal of the IGPL*, 23:194–216, 2015.
- [10] F. Carreiro. *Fragments of fixpoint logics*. PhD thesis, Institute for Logic, Language and Computation, Universiteit van Amsterdam, 2015.
- [11] F. Carreiro. PDL is the bisimulation-invariant fragment of weak chain logic. In *Proceedings of the 30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2015)*, pages 341–352, 2015.
- [12] F. Carreiro, A. Facchini, Y. Venema, and F. Zanasi. Weak MSO: automata and expressiveness modulo bisimilarity. In T.A. Henzinger and D. Miller, editors, *Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science, CSL-LICS 2014*, pages 27:1–27:27. ACM, 2014.
- [13] F. Carreiro and Y. Venema. PDL inside the μ -calculus: A syntactic and an automata-theoretic characterization. In R. Goré, B.P. Kooi, and A. Kurucz, editors, *Advances in Modal Logic (AiML 10)*, pages 74–93. College Publications, 2014.
- [14] C.C. Chang and H.J. Keisler. *Model Theory*. North Holland, 1973.
- [15] C. Cirstea, C. Kupke, and D. Pattinson. Exptime tableaux for the coalgebraic μ -calculus. In *Proceedings of the 18th Annual Conference of the EACSL on Computer Science Logic (CSL 2009)*, pages 179–193. Springer, 2009.
- [16] M. Czarnecki. How fast can the fixpoints in modal μ -calculus be reached? Manuscript accepted at FICS10, 2010.
- [17] G. D’Agostino and M. Hollenberg. Logical questions concerning the μ -calculus. *Journal of Symbolic Logic*, 65:310–332, 2000.
- [18] E.A. Emerson and C.S. Jutla. The complexity of tree automata and logics of programs (extended abstract). In *Proceedings of the 29th Symposium on the Foundations of Computer Science*, pages 328–337. IEEE Computer Society Press, 1988.
- [19] E.A. Emerson and C.S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *Proceedings of the 32nd Symposium on the Foundations of Computer Science*, pages 368–377. IEEE Computer Society Press, 1991.
- [20] S. Enqvist. A general lindström theorem for some normal modal logics. *Logica Universalis*, 7(2):233–264, 2013.
- [21] S. Enqvist, F. Seifan, and Y. Venema. Completeness for coalgebraic fixpoint logic. In *Proceedings of the 25th EACSL Annual Conference on Computer Science Logic (CSL 2016)*, volume 62 of *LIPICs*, pages 7:1–7:19, 2016.
- [22] S. Enqvist, F. Seifan, and Y. Venema. Completeness for the modal μ -calculus: separating the combinatorics from the dynamics. Technical Report PP-2016-33, Institute for Logic, Language and Computation, Universiteit van Amsterdam, 2016.
- [23] A. Facchini, Y. Venema, and F. Zanasi. A characterization theorem for the alternation-free fragment of the modal μ -calculus. In *Proceedings of the 28th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2013)*, pages 478–487. IEEE Computer Society, 2013.

- [24] G. Fontaine. Continuous fragment of the mu-calculus. In M. Kaminski and S. Martini, editors, *Computer Science Logic 2008*, volume 5213 of *LNCS*, pages 139–153. Springer, 2008.
- [25] G. Fontaine. *Modal fixpoint logic: some model-theoretic questions*. PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, 2010.
- [26] G. Fontaine, R. Leal, and Y. Venema. Automata for coalgebras: An approach using predicate liftings. In *Automata, Languages and Programming: 37th International Colloquium ICALP'10*, pages 381–392, 2010.
- [27] G. Gierz, K.H. Hofman, K. Keimel, J.D. Lawson, M. Mislove, and D.S Scott. *A Compendium of Continuous Lattices*. Springer-Verlag, 1980.
- [28] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logic, and Infinite Games*, volume 2500 of *LNCS*. Springer, 2002.
- [29] D. Harel, D. Kozen, and J. Tiuryn. *Dynamic Logic*. The MIT Press, 2000.
- [30] M. Hollenberg. *Logic and Bisimulation*. PhD thesis, Zeno Institute of Philosophy, Utrecht University, 1998.
- [31] D. Janin and I. Walukiewicz. Automata for the modal μ -calculus and related results. In *Proceedings of the Twentieth International Symposium on Mathematical Foundations of Computer Science, MFCS'95*, volume 969 of *LNCS*, pages 552–562. Springer, 1995.
- [32] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ -calculus w.r.t. monadic second-order logic. In *Proceedings of the Seventh International Conference on Concurrency Theory, CONCUR '96*, volume 1119 of *LNCS*, pages 263–277, 1996.
- [33] B. Jónsson and A. Tarski. Boolean algebras with operators, Part I. *American Journal of Mathematics*, 73:891–939, 1952.
- [34] B. Jónsson and A. Tarski. Boolean algebras with operators, Part II. *American Journal of Mathematics*, 74:127–162, 1952.
- [35] D. Kozen. Results on the propositional μ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
- [36] C. Kupke and Y. Venema. Coalgebraic automata theory: basic results. *Logical Methods in Computer Science*, 4, 2008. (An earlier version of this paper appeared in LICS 2005.).
- [37] A. Kurz and Y. Venema. Coalgebraic lindström theorems. In Lev D. Beklemishev, Valentin Goranko, and Valentin B. Shehtman, editors, *Advances in Modal Logic (AiML 8)*, pages 292–309. College Publications, 2010.
- [38] L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999. (Erratum published *Ann.P.Appl.Log.* 99:241–259, 1999).
- [39] A.M. Mostowski. Games with forbidden positions. Technical Report 78, Instytut Matematyki, Uniwersytet Gdański, Poland, 1991.
- [40] M. Otto. Eliminating recursion in the mu-calculus. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS '99)*, volume 1563 of *LNCS*, pages 531–540. Springer, 1999.
- [41] M. Otto and R. Piro. A Lindström characterization of the guarded fragment and of modal logic with a global modality. In C. Areces and R. Goldblatt, editors, *Advances in Modal Logic (AiML 7)*, pages 273–287. College Publications, 2008.
- [42] M. de Rijke. A Lindström theorem for modal logic. In A. Ponse, M. de Rijke, and Y. Venema, editors, *Modal Logic and Process Algebra: A Bisimulation Perspective*, volume 53 of *Lecture Notes*, pages 217–230. CSLI Publications, 1995.
- [43] S. Thomason. Categories of frames for modal logics. *Journal of Symbolic Logic*, 40:439–442, 1975.
- [44] Y. Venema. Algebras and coalgebras. In Benthem et al. [7], pages 331–426.
- [45] Y. Venema. Lectures on the modal μ -calculus. Lecture Notes, ILLC, University of Amsterdam, 2012.
- [46] I. Walukiewicz. Completeness of Kozen’s axiomatisation of the propositional μ -calculus. *Information and Computation*, 157:142–182, 2000.
- [47] T. Wilke. Alternating tree automata, parity games, and modal μ -calculus. *Bulletin of the Belgian Mathematical Society*, 8:359–391, 2001.