

AN EXPRESSIVE COMPLETENESS THEOREM FOR COALGEBRAIC MODAL μ -CALCULI

SEBASTIAN ENQVIST^a, FATEMEH SEIFAN^b, AND YDE VENEMA^c

^a Institute for Logic, Language and Computation, University of Amsterdam, PO Box 94242, 1090 GE Amsterdam, The Netherlands, and, Department of Philosophy, Stockholm University, SE - 10691 Stockholm, Sweden
e-mail address: thesebastianenqvist@gmail.com

^{b,c} Institute for Logic, Language and Computation, University of Amsterdam, PO Box 94242, 1090 GE Amsterdam, The Netherlands
e-mail address: fateme.sayfan@gmail.com, Y.Venema@uva.nl

ABSTRACT. Generalizing standard monadic second-order logic for Kripke models, we introduce monadic second-order logic interpreted over coalgebras for an arbitrary set functor. We then consider invariance under behavioral equivalence of MSO-formulas. More specifically, we investigate whether the coalgebraic μ -calculus is the bisimulation-invariant fragment of the monadic second-order language for a given functor. Using automata-theoretic techniques and building on recent results by the third author, we show that in order to provide such a characterization result it suffices to find what we call an adequate uniform construction for the coalgebraic type functor. As direct applications of this result we obtain a partly new proof of the Janin-Walukiewicz Theorem for the modal μ -calculus, avoiding the use of syntactic normal forms, and bisimulation invariance results for the bag functor (graded modal logic) and all exponential polynomial functors (including the “game functor”). As a more involved application, involving additional non-trivial ideas, we also derive a characterization theorem for the monotone modal μ -calculus, with respect to a natural monadic second-order language for monotone neighborhood models.

1. INTRODUCTION

1.1. Logic, automata and coalgebra. The aim of this paper is to strengthen the link between the areas of logic, automata and coalgebra. More precisely, we provide a coalgebraic generalization of the automata-theoretic approach towards monadic second-order logic (MSO), and we address the question of whether the Janin-Walukiewicz Theorem can be generalized from Kripke structures to the setting of arbitrary coalgebras¹.

Fatemeh Seifan was supported by *Vrije Competitie* grant 612.001.115 of the Netherlands Organisation for Scientific Research (NWO).

¹The paper is a substantially extended version of a previous manuscript that was presented at the 2015 conference for Logic in Computer Science (LICS 2015) in Kyoto, Japan [7].

The connection between *monadic second-order logic* and *automata* is classic, going back to the seminal work of Büchi, Rabin, and others. Rabin’s decidability result for the monadic second-order theory of binary trees, or *S2S*, makes use of a translation of monadic second-order logic into a class of automata, thus reducing the satisfiability problem for *S2S* to the non-emptiness problem for the corresponding automata [28]. The link between *MSO* and automata over trees with arbitrary branching was further explored by Walukiewicz [36]. Janin and Walukiewicz considered monadic second-order logic interpreted over Kripke structures, and used automata-theoretic techniques to obtain a van Benthem-like characterization theorem for monadic second-order logic, identifying the modal μ -calculus as the bisimulation invariant fragment of *MSO* [13]. Given the fact that in many applications bisimilar models are considered to represent the *same* process, one has little interest in properties of models that are *not* bisimulation invariant. Thus the Janin-Walukiewicz Theorem can be seen as an expressive completeness result, stating that all *relevant* properties in monadic second-order logic can be expressed in the modal μ -calculus.

Coalgebra enters naturally into this picture. Recall that Universal Coalgebra [29] provides the notion of a *coalgebra* as the natural mathematical generalization of state-based evolving systems such as streams, (infinite) trees, Kripke models, (probabilistic) transition systems, and many others. This approach combines simplicity with generality and wide applicability: many features, including input, output, nondeterminism, probability, and interaction, can easily be encoded in the coalgebra type \mathbb{T} (formally an endofunctor on the category **Set** of sets as objects with functions as arrows). Starting with Moss’ seminal paper [23], coalgebraic logics have been developed for the purpose of specifying and reasoning about *behavior*, one of the most fundamental concepts that allows for a natural coalgebraic formalization. And with Kripke structures constituting key examples of coalgebras, it should come as no surprise that most coalgebraic logics are some kind of modification or generalization of *modal logic* [5].

The coalgebraic modal logics that we consider here originate with Pattinson [27]; they are characterized by a completely standard syntax, in which the semantics of each modality is determined by a so-called *predicate lifting* (see Definition 3.8 below). Many well-known variations of modal logic in fact arise as the coalgebraic logic associated with a set Λ of such predicate liftings; examples include both standard and (monotone) neighborhood modal logic, graded and probabilistic modal logic, coalition logic, and conditional logic. Extensions of coalgebraic modal logics with fixpoint operators, needed for describing *ongoing* behavior, were developed in [33, 4].

The link between coalgebra and automata theory is by now well-established. For instance, finite state automata operating on finite words have been recognized as key examples of coalgebra from the outset [29]. More relevant for the purpose of this paper is the link with precisely the kind of automata mentioned earlier, since the (potentially infinite) objects on which these devices operate, such as streams, trees and Kripke frames, usually are coalgebras. Thus, the automata-theoretic perspective on modal fixpoint logic may be lifted to the abstraction level of coalgebra [33, 8]. In fact, many key results in the theory of automata operating on infinite objects, such as Muller & Schupp’s Simulation Theorem [25] can in fact be seen as instances of more general theorems in Universal Coalgebra [15].

1.2. Coalgebraic monadic second-order logic. Missing from this picture is, to start with, a coalgebraic version of (*monadic*) *second-order logic*. Filling this gap is the first aim of the current paper, which introduces a notion of *monadic second-order logic* MSO_{\top} for

coalgebras of type T . Our formalism combines two ideas from the literature. First of all, we looked for inspiration to the coalgebraic versions of *first-order logic* of Litak & alii [17]. These authors introduced Coalgebraic Predicate Logic as a common generalisation of first-order logic and coalgebraic modal logic, combining first-order quantification with coalgebraic syntax based on predicate liftings. Our formalism $\mathsf{MSO}_{\mathsf{T}}$ will combine a similar syntactic feature with second-order quantification. In fact it can be shown to contain the Coalgebraic Predicate Logic of Litak & alii as a fragment, though we will omit the verification of this fact.

Second, following the tradition in automata-theoretic approaches towards monadic second-order logic, our formalism will be *one-sorted*. That is, we *only* allow second-order quantification in our language, relying on the fact that individual quantification, when called for, can be encoded as second-order quantification relativized to singleton sets. Since predicate liftings are defined as families of maps on powerset algebras, these two ideas fit together very well, to the effect that our second-order logic is in some sense simpler than the first-order formalism of [17].

In section 4 we will define, for any set Λ of monotone² predicate liftings, a formalism MSO_{Λ} , and we let $\mathsf{MSO}_{\mathsf{T}}$ denote the logic obtained by taking for Λ the set of *all* monotone predicate liftings. Clearly we will make sure that this definition generalizes the standard case, in the sense that the standard version of MSO for Kripke structures instantiates the logic $\mathsf{MSO}_{\{\diamond\}}$ and is equivalent to the coalgebraic logic $\mathsf{MSO}_{\mathcal{P}}$ (where \mathcal{P} denotes the power set functor).

The introduction of a monadic second-order logic $\mathsf{MSO}_{\mathsf{T}}$ for T -coalgebras naturally raises the question, for which T does the coalgebraic modal μ -calculus for T , denoted $\mu\mathsf{ML}_{\mathsf{T}}$, correspond to the bisimulation-invariant fragment of $\mathsf{MSO}_{\mathsf{T}}$.

Question 1.1. Which functors T satisfy $\mu\mathsf{ML}_{\mathsf{T}} \equiv \mathsf{MSO}_{\mathsf{T}}/\sim$?

1.3. Automata for coalgebraic monadic second-order logic. In order to address Question 1.1, we take an automata-theoretic perspective on the logics $\mathsf{MSO}_{\mathsf{T}}$ and $\mu\mathsf{ML}_{\mathsf{T}}$, and as the second contribution of this paper we introduce a class of parity automata for $\mathsf{MSO}_{\mathsf{T}}$.

As usual, the operational semantics of our automata is given in terms of a two-player acceptance game, which proceeds in *rounds* moving from one basic position to another, where a basic position is a pair consisting of a state of the automaton and a point in the coalgebra structure under consideration. In each round, the two players, \exists and \forall , focus on a certain local ‘window’ on the coalgebra structure. This ‘window’ takes the shape of a *one-step T -model*, that is, a triple (X, α, V) consisting of a set X , a *chosen object* $\alpha \in \mathsf{T}X$ (representing partial information about the transition map of the coalgebra as seen from the perspective of some given state), and a valuation V interpreting the states of the automaton as subsets of X . More specifically, during each round of the game it is the task of \exists to come up with a valuation V that creates a one-step model in which a certain *one-step formula* δ (determined by the current basic position in the game) is true.

Generally, our automata will have the shape $\mathbb{A} = (A, \Delta, \Omega, a_I)$ where A is a finite carrier set with initial state $a_I \in A$, and Ω and Δ are the parity and transition map of \mathbb{A} , respectively. The flavour of such an automaton is largely determined by the co-domain of its

²In the most general case, restricting to monotone predicate liftings is not needed, one could define $\mathsf{MSO}_{\mathsf{T}}$ as the logic obtained by taking for Λ the set of *all* predicate liftings. However, in the context of this paper, where we take an automata-theoretic perspective on MSO , this restriction makes sense.

transition map Δ , the so-called *one-step language* which consists of the one-step formulas that feature in the acceptance game as described.

Each one-step language L induces its own class of automata $Aut(L)$. For instance, the class of automata corresponding to the coalgebraic fixpoint logic μML_Λ can be given as $Aut(ML_\Lambda)$, where ML_Λ is the set of positive modal formulas of depth one that use modalities from Λ [8]. Basically then, the problem of finding the right class of automata for the coalgebraic monadic second-order logic MSO_Λ consists in the identification of an appropriate one-step language. Our proposal comprises a one-step *second-order logic* SO_Λ which uses predicate liftings to describe the chosen object of the one-step model.

Finally, note that similar to the case of standard MSO , the equivalence between formulas in MSO_\top and automata in $Aut(SO_\Lambda)$ is only guaranteed to hold for coalgebras that are ‘tree-like’ in some sense (to be defined further on).

Theorem 1.2 (Automata for coalgebraic MSO). *For any set Λ of monotone predicate liftings for \top there is an effective construction mapping any formula $\varphi \in MSO_\Lambda$ into an automaton $\mathbb{A}_\varphi \in Aut(SO_\Lambda)$, which is equivalent to φ over \top -tree models.*

The proof of Theorem 1.2 proceeds by induction on the complexity of MSO_\top -formulas, and thus involves various *closure properties* of automata, such as closure under complementation, union and projection. In order to establish these results, it will be convenient to take an *abstract* perspective revealing how closure properties of a class of automata are completely determined at the level of the one-step language, similar to the development in [36].

1.4. Bisimulation Invariance. With automata-theoretic characterizations in place for both coalgebraic MSO and the coalgebraic μ -calculus μML , we can address Question 1.1 by considering the following question:

Question 1.3. Which functors \top satisfy $Aut(ML_\top) \equiv Aut(SO_\top)/\sim$?

Continuing the program of the third author [35], we will approach this question *at the level of the one-step languages*, SO and ML . To start with, observe that any translation (from one-step formulas in) SO to (one-step formulas in) ML naturally induces a translation from SO -automata to ML -automata. A new observation we make here is that any so-called *uniform construction* on the class of one-step models for the functor \top that satisfies certain *adequacy* conditions, provides (1) a translation $(\cdot)^* : SO \rightarrow ML$, together with (2) a construction $(\cdot)_*$ transforming a pointed \top -model (\mathbb{S}, s) into a tree model (\mathbb{S}_*, s_*) which is a coalgebraic pre-image of (\mathbb{S}, s) satisfying

$$\mathbb{A} \text{ accepts } (\mathbb{S}_*, s_*) \text{ iff } \mathbb{A}^* \text{ accepts } (\mathbb{S}, s).$$

where \mathbb{A}^* is obtained by applying to \mathbb{A} the lifting of the one-step translation $(\cdot)^* : SO \rightarrow ML$ to the level of automata. From this it easily follows that an SO -automaton \mathbb{A} is bisimulation invariant iff it is equivalent to the ML -automaton \mathbb{A}^* .

On the basis of these observations we can prove the following generalisation of the Janin-Walukiewicz Theorem.

Theorem 1.4 (Coalgebraic Bisimulation Invariance). *Let \top be any set functor. If \top admits an adequate uniform construction for every finite set of second-order one-step formulas Γ , then*

$$MSO_\top/\sim \equiv \mu ML_\top.$$

In our view, the proof of this theorem separates the ‘clean’, abstract part of bisimulation-invariance results from the more functor-specific parts. As a consequence, Theorem 1.4 can be used to obtain immediate results in particular cases. Examples include the power set functor (standard Kripke structures), where the adequate uniform construction roughly consists of taking ω -fold products (see Example 5.15), the bag functor (Example 5.16), and all exponential polynomial functors (Corollary 5.18).

However, it turns out that there are cases of functors that provably do *not* admit an adequate uniform construction. A concrete example is the *monotone neighborhood functor* \mathcal{M} (see the next section for its definition), which provides the semantics for *monotone modal logic* [10] and its fixpoint extension, the *monotone μ -calculus*. The monotone μ -calculus is interesting for at least two reasons: first, it has a purely theoretical interest as a natural candidate for a *minimal* modal μ -calculus, since the only constraint on the box operator is monotonicity rather than the stronger constraint of complete multiplicativity of the box operator in normal modal logic. In other words, the base modal logic has precisely the logical strength it needs so that Knaster-Tarski-style extremal fixpoints may be smoothly introduced, and no more. Second, a more application driven reason to study the monotone μ -calculus is that it is the natural fixpoint logic in which Parikh’s dynamic logic of games [26] sits as a fragment in the same manner that PDL is a fragment of the modal μ -calculus. (More precisely, game logic is a fragment of a multi-modal variant of the monotone modal μ -calculus, but we will stick to the monomodal case for the sake of notational convenience.) As the final contribution of this paper we show how, with some extra work, a characterization result for the monotone μ -calculus can be derived using our main result 1.4, with respect to a natural monadic second-order language for monotone neighborhood structures that we call “monotone MSO”.

Theorem 1.5. *The monotone μ -calculus is the fragment of monotone MSO that is invariant for neighborhood bisimulations.*

2. SOME TECHNICAL BACKGROUND

In this section we provide some of the necessary technical background. We shall assume familiarity with coalgebras, and at least some basic acquaintance with modal logic and the modal μ -calculus. We will provide insofar as possible a self-contained presentation of the required material on automata and infinite games.

2.1. Automata. Throughout most of its history, the mathematical theory surrounding the modal μ -calculus has relied heavily on the theory of *automata* operating on infinite objects and, what is closely related, *infinite games*. It turns out that these frameworks fit nicely into the coalgebraic world view, which is to be expected as infinite games and automata operating on infinite objects both have a coinductive flavour. This connection is in fact a big part of what makes the coalgebraic modal μ -calculus tick, and will be central to this paper. In this section we give a brief overview of this field. The material that we cover is entirely standard, so readers who are already familiar with these topics will want to merely skim the section to fix notation and terminology.

We start with something that probably most readers have already come across at some point: the theory of word automata recognizing regular languages over some alphabet. (For an introduction to the basic concepts of automata theory, see [11].)

Definition 2.1. A (non-deterministic) *word automaton* is a tuple:

$$(A, \Delta, a_I, F)$$

where A is a finite set, $\Delta : A \times \Sigma \rightarrow \mathcal{P}(A)$, $a_I \in A$ and $F \subseteq A$.

The general mold into which many types of automata fit is of this shape: they consist of a set of *states* (A), a *transition structure* (Δ) telling us how to move between states of the automaton, an *initial state* (a_I) and, finally, some data structure over A that lets us know whether or not a given computation carried out by an automaton accepts or rejects its input. In the case of word automata, the input \vec{w} of an automaton is a finite word $c_0 \dots c_n$ over the alphabet Σ (i.e. an element of Σ^n for some n).

Definition 2.2. A *run* of the automaton \mathbb{A} on a word \vec{w} is a tuple $a_0 \dots a_{n+1}$ of states in A , such that $a_0 = a_I$ and $a_{i+1} \in \Delta(a_i, c_i)$ for $i \leq n$. The run is said to be *accepting* if $a_{n+1} \in F$, and we say that \mathbb{A} *accepts* the word \vec{w} if there is some accepting run of \mathbb{A} on the input \vec{w} . The *language* $L(\mathbb{A})$ of an automaton \mathbb{A} is simply the collection of all the words that it accepts.

A central technique in automata theory is to show that one can simplify the transition structure of some class of automata, without losing any expressive power. In the case of word automata, it is well known that the *non-deterministic* automata that we have defined here have the same expressive power as *deterministic* word automata, which are defined as ones that satisfy the condition that $\Delta(a, c)$ is a *singleton* for each $a \in A$. This means that there is a *unique* run of the automaton on every input.

Fact 2.3 (Simulation of word automata). For every word automaton \mathbb{A} , there exists a deterministic word automaton \mathbb{A}' with $L(\mathbb{A}) = L(\mathbb{A}')$.

The proof consists of an elegant construction known as the *powerset construction*: given an automaton $\mathbb{A} = (A, \Delta, a_I, F)$, we construct a deterministic automaton $\mathbb{A}' = (A', \Delta', a'_I, F')$ recognizing the same language by setting $A' = \mathcal{P}(A)$, $a'_I = \{a_I\}$, the accepting states are defined by

$$F' = \{B \subseteq A \mid B \cap F \neq \emptyset\}$$

and the transition structure is defined

$$\Delta'(B, c) = \bigcup \{\Delta(b, c) \mid b \in B\}.$$

Intuitively, the states of \mathbb{A}' consist of “macro-states” over \mathbb{A} , consisting of several states of \mathbb{A} being visited at once, and a run of \mathbb{A}' on a word can be viewed as a “bundle” of runs of \mathbb{A} on the same word. These sorts of simulation theorems turn up in various places, for various types of automata. In particular, one of our main technical results of this paper is a simulation theorem, allowing us to simplify the transition structure of the automata that we introduce for monadic second-order logic interpreted on coalgebras.

In this paper we are mainly interested in automata that, unlike word automata, classify *infinite* objects. The simplest non-trivial example of this is the theory of automata operating on infinite words, or *streams* as we shall call them. For a given alphabet Σ , a Σ -*stream* is an ω -sequence of elements of Σ . Automata that are used to classify streams are called *stream automata*. The simplest kind of stream automata, known as *Büchi automata*, are simply the same as word automata: they consist of a set of states A , a transition relation $\Delta : A \times \Sigma \rightarrow \mathcal{P}(A)$, a start state a_I and a set F of accepting states. The difference is in the definition of acceptance.

Definition 2.4. Given a stream $c_0c_1c_2c_3\dots$ over Σ , a *run* of the automaton \mathbb{A} is a stream $a_0a_1a_2a_3\dots$ over A , such that $a_0 = a_I$ and $a_{i+1} \in \Delta(a_i, c_i)$ for all $i < \omega$. The run is said to be *accepting* if the following acceptance condition is met: for infinitely many $i \in I$, we have $a_i \in F$.

In other words, a run is accepting if from each given point in the computation, the automaton will eventually end up in an accepting state. The language $L(\mathbb{A})$ consists of the set of Σ -streams accepted by \mathbb{A} .

Again, one of the central results about stream automata is a simulation theorem. However, here the Büchi condition is not flexible enough - it is a standard fact that deterministic Büchi automata have strictly less expressive power than non-deterministic ones. To obtain a simulation theorem, we need to switch to a more sophisticated acceptance condition, and several are available (see [34] for an overview). For the purposes of the present paper, the most natural condition is the *parity* condition.

Given a finite set A , a *priority map* is simply a map $\Omega : A \rightarrow \omega$, assigning a natural number to each element of A . Such a map encodes, in a compact way, two pieces of information that we want to keep track of: first, it partitions the states in A into two sets, the states with even priority and the ones with odd priority. Intuitively, think of the states with even priority as the accepting states. Second, the priority map encodes an *ordering* of the states, where states of higher priority are regarded as more important when determining whether or not a run is accepting. A *parity stream automaton* is then a tuple (A, Δ, a_I, Ω) , where Ω is a priority map. A run $a_0a_1a_2\dots$ on some stream is now said to be *accepting* if the highest number k such that $\Omega(a_i) = k$ for infinitely many i is an even number. In other words: the run visits some accepting state a infinitely many times, and any non-accepting state that is also visited infinitely many times has lower priority than a .

Switching from Büchi automata to parity automata does not give any more expressive power: the two automaton models are expressively equivalent. The point of the parity condition is that we now have a simulation theorem available:

Fact 2.5 (Simulation of stream automata). For every parity stream automaton \mathbb{A} , there is a deterministic parity stream automaton \mathbb{A}' (i.e. such that $\Delta'(a', c)$ is a singleton for all $a' \in A'$) with $L(\mathbb{A}) = L(\mathbb{A}')$.

A standard proof of this simulation theorem for stream automata uses a construction due to Safra, which is quite involved compared to the powerset construction for word automata. We refer to [34] for the details.

Simulation theorems are particularly useful for proving *closure* properties of automata; for example, given a deterministic parity stream automaton \mathbb{A} , it is trivial to define its *complement*, i.e. an automaton that accepts all and only those streams that are *not* accepted by \mathbb{A} . The construction just adjusts the priority map: if Ω was the priority map of \mathbb{A} , then define the complement of \mathbb{A} to be just like \mathbb{A} but with the priority map Ω' defined by $\Omega'(a) = \Omega(a) + 1$.

2.2. Infinite games. The automata we considered in the previous section are all of a fairly simple character: when processing a word or a stream, these automata simply change their state depending on the input that they read, in a deterministic or non-deterministic manner, and the sequence of states visited in this manner is what we have called a run. For more complex automaton models, this notion of a “run” of the automaton on some input will

no longer be adequate - more elaborate models of computation will be required. Here, the theory of infinite games has turned out to be a powerful mathematical framework, providing simple and perspicuous models of how automata process their input. The idea is to think of the computation carried out by an automaton as a game between two players with opposite goals: one player tries to defend the view that the input should be accepted, the other tries to show that it should be rejected.

Definition 2.6. A *board game* \mathbb{G} is a tuple $(B_{\exists}, B_{\forall}, E, W)$ where B_{\exists} and B_{\forall} are disjoint sets, $E \subseteq B \times B$ where $B = B_{\exists} \cup B_{\forall}$ (this union is called the *board*) and W is a map from the set B^{ω} of B -streams into the set $\{\exists, \forall\}$. Elements of B are called *positions*.

Intuitively, \exists (or Eloise) and \forall (or Abelard) are the *players*, B_{\exists} consists of the positions on the board at which it is \exists 's turn to move, and similarly for B_{\forall} . Finally, the relation E tells us what the legal moves are at a given position of the game. It can happen that there are no legal moves at all at a position of the game.

There is nothing to prevent that a match of a board game could go on indefinitely. In fact this is crucial to the connection between infinite games and automata operating on non-wellfounded structures, in which the computation carried out by the automaton processing some input can be infinitely long. In these cases, the map W will be used to decide a winner. Formally:

Definition 2.7. Let $\mathbb{G} = (B_{\exists}, B_{\forall}, E, W)$ be any board game. A *match* in \mathbb{G} is either a finite tuple of positions $\pi_0 \dots \pi_n$ or an infinite sequence $\pi_0 \pi_1 \pi_2 \dots$ of positions in B , such that $a_i E a_{i+1}$ for all $i \in \omega$ (or for all $i < n$, if the match is finite.) A match ρ is called *complete* if it is infinite, or if $E[\pi] = \{\pi' \mid \pi E \pi'\} = \emptyset$ where π is the last position on ρ . If ρ is not complete it is called a *partial* match. The set of partial matches ρ of \mathbb{G} for which the last position π of ρ belongs to Π , and such that $E[\pi] \neq \emptyset$, is denoted by $\text{PM}(\mathbb{G}, \Pi)$. The *winner* of an infinite match ρ is defined to be $W(\rho)$, and the winner of a complete but finite match is \exists if the last position is in B_{\forall} , and \forall if the last position of the match is in B_{\exists} . In these cases, we say that the losing player got stuck.

Definition 2.8. A *strategy* for a player Π in a board game \mathbb{G} is a map $\chi : \text{PM}(\mathbb{G}, \Pi) \rightarrow B$. A match ρ (complete or partial) is said to be χ -*guided* if, for every proper initial segment ρ' of ρ that belongs to $\text{PM}(\mathbb{G}, \Pi)$, the tuple $\rho' \chi(\rho')$ is also an initial segment of ρ .

Finally, a strategy χ for Π is said to be a *winning strategy* at a position π in \mathbb{G} if every χ -guided match beginning with the position π is won by Π . If there is a winning strategy for Π at the position π then we say that π is a *winning position* for Π .

Remark 2.9. The reader may be more familiar with a different terminology here, where the word “play” is used for what we call a “match”, and a play is said to “respect the strategy χ ” if it is χ -guided in our sense.

In practice, when we define a winning strategy χ for either player in a game \mathbb{G} for some given position π , we will usually only take care to define the strategy on those partial matches that are actually reachable from the position π by a χ -guided match.

A special role in connection with automata is played by *parity games*, where the winning condition W is defined via a map $\Omega : B \rightarrow \omega$, such that the range $\Omega[B]$ of this map is finite. We then set $W(\rho) = \exists$ if the highest natural number k such that $\Omega(\pi_i) = k$ for infinitely many $i \in \omega$ is an even number, and we set $W(\rho) = \forall$ otherwise. It follows straightforwardly from a classic result in the theory of infinite games, the *Borel determinacy theorem* [18], that

all parity games are *determined*, meaning that the winning positions for \exists and \forall partition the board: every position is winning for one of the two players. However, a much stronger and quite useful property holds for parity games:

Definition 2.10. A strategy χ for Π is said to be *positional* if, for any two partial matches $\rho_1, \rho_2 \in \text{PM}(\mathbb{G}, \Pi)$ such that the last positions of ρ_1 and ρ_2 are identical, we have $\chi(\rho_1) = \chi(\rho_2)$.

A positional strategy for Π can thus be represented equivalently as a map $\chi : B_{\Pi}^+ \rightarrow B$, where

$$B_{\Pi}^+ = \{\pi \in B_{\Pi} \mid E[\pi] \neq \emptyset\}.$$

Fact 2.11 (Positional Determinacy). Let \mathbb{G} be a parity game. Then for every position $\pi \in B$, there exists a unique player $\Pi \in \{\exists, \forall\}$ and a positional winning strategy for Π at the position π .

Several proofs of the positional determinacy theorem are known, see [6, 24, 37]. We will briefly consider games with non-parity winning condition later; however, all the games we consider will be determined (although not necessarily by positional strategies), as an easy consequence of the Borel determinacy theorem.

2.3. Logics for Kripke models and their automata. With the required machinery in place, we are ready to explain the connection between infinite games, automata, monadic second-order logic and the modal μ -calculus.

Let us start by briefly reviewing monadic second-order logic and the modal μ -calculus, viewed as languages for reasoning about Kripke models. For simplicity we restrict our attention to Kripke models with just one accessibility relation, i.e. models for modal logic in the basic similarity type with just one diamond and one box. Let Var be a fixed infinite supply of variables.

Definition 2.12. A *Kripke model* is a structure $\mathbb{S} = (S, R, V)$ where S is a set, $R \subseteq S \times S$ and $V : Var \rightarrow \mathcal{P}(S)$ is a Var -valuation. A *pointed Kripke model* is a structure (\mathbb{S}, s) where \mathbb{S} is a Kripke model and s is a point in S .

Associated with a valuation V , we define the *conjugate coloring* $V^\dagger : S \rightarrow \mathcal{P}(Var)$ by $V^\dagger(s) := \{p \in Var \mid s \in V(p)\}$. Given a subset $Z \subseteq S$, the valuation $V[p \mapsto Z]$ is as V except that it maps the variable p to Z .

Turning to syntax, we define the formulas of monadic second-order logic **MSO** through the following grammar:

$$\varphi ::= \text{sr}(p) \mid p \subseteq q \mid R(p, q) \mid \neg\varphi \mid \varphi \vee \varphi \mid \exists p.\varphi$$

with $p, q \in Var$. (Here, sr stands for ‘‘source’’.) We define $\top := \forall p.p \subseteq p$ and $\perp := \neg\top$. Formulas are evaluated over pointed Kripke models by the following induction:

- $(S, R, V, s) \Vdash \text{sr}(p)$ iff $V(p) = \{s\}$
- $(S, R, V, s) \Vdash p \subseteq q$ iff $V(p) \subseteq V(q)$
- $(S, R, V, s) \Vdash R(p, q)$ iff for all $v \in V(p)$ there is $w \in V(q)$ with vRw
- standard clauses for the boolean connectives
- $(S, R, V, s) \Vdash \exists p.\varphi$ iff $(S, R, V[p \mapsto Z], s) \Vdash \varphi$ for some $Z \subseteq S$.

We present the language of the *modal μ -calculus* μML in negation normal form, by the following grammar:

$$\varphi ::= p \mid \neg p \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Box\varphi \mid \Diamond\varphi \mid \eta p.\varphi$$

where $p \in \text{Var}$, $\eta \in \{\mu, \nu\}$, and in the formula $\eta p.\varphi$ no free occurrence of the variable p in φ may be in the scope of a negation. We define $\top := \nu p.p$ and $\perp := \mu p.p$.

The satisfaction relation between pointed Kripke models and formulas in μML is defined by the usual induction, with, e.g.

- $(S, R, V, s) \Vdash \mu p.\varphi$ iff $s \in \bigcap \{Z \subseteq S \mid \varphi_p(Z) \subseteq Z\}$ where $\varphi_p(Z)$ denotes the truth set of the formula φ in the model $(S, R, V[p \mapsto Z])$.

We assume familiarity with the notion of bisimilarity between two (pointed) Kripke models, and say that a formula of MSO is *bisimulation invariant* if it has the same truth value in any pair of bisimilar pointed Kripke models.

Fact 2.13 (Janin-Walukiewicz Theorem [13]). A formula φ of MSO is equivalent to a formula of μML iff φ is invariant for bisimulations.

Generally, given two languages L and L' we shall write $L \equiv L'$ to say that every formula of L is semantically equivalent to a formula of L' , and vice versa. We will denote the bisimulation invariant fragment of L by L/\sim . With this notation we can state the previous fact more succinctly as:

$$\mu\text{ML} \equiv \text{MSO}/\sim$$

The proof of the Janin-Walukiewicz theorem proceeds by comparing suitable automaton models for MSO and the modal μ -calculus. To get an understanding of how MSO -automata on Kripke models work, the main conceptual step one needs to take from the perspective of word and stream automata is to take a *logical* view on the transition structure of an automaton, i.e. the transition map of an automaton will be constructed from formulas in a logical system. Let A be a set (of variables). The language of *first-order logic with equality* over A is given by the grammar:

$$a(x) \mid x = y \mid \varphi \wedge \varphi \mid \neg\varphi \mid \exists x.\varphi$$

where $a \in A$ and x, y belong to some fixed, countably infinite set of individual variables. The set of *sentences* in this language is denoted $\text{FOE}(A)$, and the set of sentences in which each $a \in A$ appears positively is denoted by $\text{FOE}^+(A)$. The fragment of $\text{FOE}(A)$ in which the equality symbol does not occur is denoted by $\text{FO}(A)$, and we write $\text{FO}^+(A)$ for $\text{FOE}^+(A) \cap \text{FO}(A)$.

A model for this language, in the usual sense, is simply a set X together with an assignment of a subset of X to each $a \in A$, viewed as a predicate - in other words, a valuation $V : A \rightarrow \mathcal{P}(X)$.

Definition 2.14. Let P be a set of propositional variables. A *P -chromatic MSO -automaton* \mathbb{A} is a tuple (A, Δ, a_I, Ω) , where A is a finite set, $a_I \in A$, $\Omega : A \rightarrow \omega$ is a priority map and

$$\Delta : A \times \mathcal{P}(P) \rightarrow \text{FOE}^+(A)$$

is the transition map, assigning a sentence in $\text{FOE}^+(A)$ to each pair (a, Q) with $a \in A$ and $Q \subseteq P$.

Note that, in this automaton model, the powerset of P takes on the role of an alphabet of the automaton. In order to decide whether or not an automaton \mathbb{A} accepts or rejects a

given pointed Kripke model \mathbb{S}, s , we define a parity game called the *acceptance game* for \mathbb{A} with respect to the model \mathbb{S} . The role of \exists is to act as a “defense attorney”, in favor of acceptance of the model, and the goal of \forall is to act as a “prosecutor” trying to show that the model should be rejected. A “round” of this game will have the following shape: first, the automaton is in a state a and is visiting some point s in the model \mathbb{S} . At such a *basic position* of the acceptance game, it will be \exists ’s turn to move, and her moves are constrained by the transition map. Formally, she needs to come up with a valuation $U : A \rightarrow \mathcal{P}(R[s])$ over the immediate successors of s with respect to the accessibility relation, such that the one-step Kripke model $(R[s], U)$ satisfies the $\text{FOE}^+(A)$ -sentence $\Delta(a, V^\dagger(s) \cap P)$. Here, we think of $V^\dagger(s) \cap P$ - the collection of propositional variables in P true at the point s - as the input from the alphabet currently being read by the automaton. In practice, we will assume that we are dealing with a model in which the valuation V assigns an empty value to each variable not in P , so that we may just write $V^\dagger(s)$ rather than $V^\dagger(s) \cap P$.

For a more precise description of the game, we provide the positions, player assignments, admissible moves and priorities of the acceptance game in a table:

Position	Player	Admissible moves	Priority
$(a, s) \in A \times S$	\exists	$\{U : A \rightarrow \mathcal{P}(R[s]) \mid (R[s], U) \models \Delta(a, V^\dagger(s))\}$	$\Omega(a)$
$U : A \rightarrow \mathcal{P}S$	\forall	$\{(b, t) \mid t \in U(b)\}$	0

Table 1: Acceptance game for MSO-automata

Given $s \in S$, we say that \mathbb{A} accepts the pointed model (\mathbb{S}, s) if (a, s) is a winning position for \exists in the acceptance game. In practice, we will often identify a partial match of the form:

$$(a_0, s_0)U_0(a_1, s_1)U_1 \dots U_{n-1}(a_n, s_n)$$

of the acceptance game belonging with the sequence of basic positions

$$(a_0, s_0) \dots (a_n, s_n)$$

appearing in the match.

It may be helpful at this point to compare the acceptance game for these automata to the notion of acceptance we presented for stream automata in Section 2.1. In particular, the reader can may check that an accepting run of a non-deterministic parity stream automaton can be viewed as a winning strategy for \exists in a parity game, but one in which the role of the opposing player \forall becomes trivialized: *all* choices in this game will belong to \exists .

MSO-formulas are not equivalent to MSO-automata in general. To obtain the desired connection between formulas of MSO and MSO-automata, it is necessary to restrict attention to pointed Kripke models (S, R, V, u) that are *tree-like*, i.e. such that the underlying frame (S, R) is a tree rooted at u . The following result is proved in [36]:

Fact 2.15. For every formula φ of MSO with free second-order variables P , there exists a P -chromatic MSO-automaton accepting all and only those tree-like pointed Kripke models that satisfy φ .

A converse of this result also holds, so that any MSO-automaton can be translated to a sentence of MSO that is equivalent to it over tree-like models. Furthermore, a key insight behind the Janin-Walukiewicz theorem is that MSO automata are very closely related to an automaton model that is suitable to model μ -calculus formulas, on arbitrary Kripke models.

Fact 2.16. A class of pointed Kripke models is definable by a formula in μML with free variables P if, and only if, it is recognized by some P -chromatic MSO-automaton such that, for all $Q \subseteq P$ and $a \in A$, we have $\Delta(a, Q) \in \text{FO}^+(A)$.

The proof of Fact 2.15 involves establishing a number of closure properties of MSO-automata, corresponding to the logical connectives. For example, negation corresponds to *complementation* on the automata side, where the complement of an automaton \mathbb{A} accepts precisely those inputs that are rejected by \mathbb{A} . Similarly, the *union* of \mathbb{A} and \mathbb{B} accepts precisely those inputs accepted by either \mathbb{A} or \mathbb{B} , and corresponds to disjunction on the logical side. The most difficult property to establish is projection, corresponding to second-order quantification. The solution lies in a simulation theorem for MSO-automata, simplifying the formulas appearing in the range of the transition map.

Definition 2.17. A formula in $\text{FOE}^+(A)$ is said to be a *disjunctive formula* if it is a disjunction of formulas of the form:

$\exists x_1 \dots \exists x_n (\text{diff}(x_1, \dots, x_n) \wedge a_1(x_1) \wedge \dots \wedge a_n(x_n) \wedge \forall y (\text{diff}(x_1, \dots, x_n, y) \rightarrow b_1(y) \vee \dots \vee b_k(y)))$
 where $a_1, \dots, a_n, b_1, \dots, b_k \in A$ and $\text{diff}(z_1, \dots, z_m)$ is an abbreviation for the formula:

$$\bigwedge_{1 \leq i < j \leq m} \neg(z_i = z_j)$$

A P -chromatic MSO-automaton \mathbb{A} is said to be *non-deterministic* if, for all $a \in A$ and all $Q \subseteq P$, the formula $\Delta(a, Q)$ is disjunctive.

Fact 2.18 (Simulation of MSO-automata). For every MSO-automaton \mathbb{A} , there exists a non-deterministic MSO-automaton \mathbb{A}' that accepts the same tree-like pointed Kripke models as \mathbb{A} .

One of our contributions in this paper is a purely *semantic* account of what makes disjunctive formulas useful, that is independent of their rather specific syntax and thus suitable for our more abstract setting.

3. COALGEBRAS AND GENERIC MODAL FIXPOINT LOGICS

3.1. Coalgebras and models. In the general case our basic semantic structures will consist of coalgebras together with valuations. We only consider coalgebras over the base category \mathbf{Set} with sets as objects and functions as arrows. The co- and contravariant power set functors will be denoted by $\mathcal{P} : \mathbf{Set} \rightarrow \mathbf{Set}$ and $\mathcal{Q} : \mathbf{Set} \rightarrow \mathbf{Set}^{op}$, respectively. We recall the actions of these functors on morphisms: given $f : X \rightarrow Y$, we have $\mathcal{P}f(Z) = f[Z]$ for $Z \subseteq X$, and $\mathcal{Q}f(Z) = f^{-1}(Z)$ for $Z \subseteq Y$. Covariant endofunctors on \mathbf{Set} will be called *set functors*.

Convention 3.1. We will make an additional assumption on all set functors here: we require that they preserve all monics in \mathbf{Set} , i.e. that $\mathbb{T}f$ is an injective map whenever f is. This is a very mild constraint, since it very nearly holds for all set functors already, the only possible exception being maps with empty domain [32]. So throughout this paper, “set functor” will be taken to mean: “set functor that preserves all monics”. All the functors we consider here will satisfy the constraint.

Definition 3.2. Let \mathbb{T} be a set functor. A \mathbb{T} -coalgebra is a pair (S, σ) consisting of a set S , together with a map $\sigma : S \rightarrow \mathbb{T}S$. A \mathbb{T} -model is a structure $\mathbb{S} = (S, \sigma, V)$ where (S, σ) is a \mathbb{T} -coalgebra and $V : \text{Var} \rightarrow \mathcal{P}S$. A *pointed* \mathbb{T} -model is a structure (\mathbb{S}, s) where \mathbb{S} is a \mathbb{T} -model and $s \in S$.

The usual notion of a p -morphism between Kripke models can be generalized as follows:

Definition 3.3. Let $\mathbb{S}_1 = (S_1, \sigma_1, V_1)$ and $\mathbb{S}_2 = (S_2, \sigma_2, V_2)$ be two \mathbb{T} -models and let $f : S_1 \rightarrow S_2$ be any map. Then f is said to be a \mathbb{T} -model morphism if:

- (1) for each variable p and each $u \in S_1$, we have $u \in V_1(p)$ iff $f(u) \in V_2(p)$;
- (2) the map f is a *coalgebra morphism*, i.e. we have

$$\sigma_2 \circ f = \mathbb{T}f \circ \sigma_1.$$

Definition 3.4. Two pointed \mathbb{T} -models (\mathbb{S}, s) and (\mathbb{S}', s') are said to be *behaviourally equivalent*, notation: $(\mathbb{S}, s) \sim (\mathbb{S}', s')$, if s and s' can be identified by morphisms $f : \mathbb{S} \rightarrow \mathbb{T}$ and $f' : \mathbb{S}' \rightarrow \mathbb{T}$ such that $f(s) = f'(s')$ (for some \mathbb{T} -model \mathbb{T}).

It is well known that coalgebras and coalgebra morphisms for a set functor form a co-complete category, since its forgetful functor creates colimits [1]. In particular this means that coproducts of arbitrary families of \mathbb{T} -models exist, and they correspond to disjoint unions of models in the case of Kripke semantics. Concretely, we define disjoint unions of \mathbb{T} -models as follows:

Definition 3.5. Let $\{\mathbb{S}_i\}_{i \in I}$ be a family of \mathbb{T} -models, where $\mathbb{S}_i = (S_i, \sigma_i, V_i)$. Then we define the *disjoint union* $\coprod_{i \in I} \mathbb{S}_i = (S', \sigma', V')$ by first setting $S' = \coprod_{i \in I} S_i$ to be the disjoint union of the sets S_i . Let f_i denote the insertion map of S_i into S' . Then we define σ' to be the unique map with the property that $\sigma' \circ f_i = \mathbb{T}f_i \circ \sigma_i$ for all $i \in I$. For $p \in \text{Var}$ we define $V'(p) = \bigcup_{i \in I} f_i[V_i(p)]$.

Fact 3.6. For each $j \in I$, the insertion map $f_j : \mathbb{S}_j \rightarrow \coprod_{i \in I} \mathbb{S}_i$ is a \mathbb{T} -model morphism.

Kripke models are \mathbb{T} -models for \mathbb{T} equal to the (covariant) power set functor \mathcal{P} , and two pointed Kripke models are behaviourally equivalent iff they are bisimilar in the usual sense. Comprehensive lists of examples of coalgebraic modal logics for various type functors and their applications can be found in several research papers by various authors, see for example [5]. We shall review some of these examples later, when we consider applications of our main characterization results. One functor that will be of particular interest in this paper is the *monotone neighborhood* functor \mathcal{M} , usually defined as the subfunctor of $\mathcal{Q} \circ \mathcal{Q}$ given by setting $\mathcal{M}X \subseteq \mathcal{Q}QX$ to be:

$$\{N \in \mathcal{Q}QX \mid \forall Z, Z' : Z \in N \ \& \ Z \subseteq Z' \Rightarrow Z' \in N\}.$$

Following the usual terminology we shall often refer to \mathcal{M} -models as *neighborhood models*. There is a standard definition of bisimilarity for neighborhood models:

Definition 3.7. A *neighborhood bisimulation* between neighborhood models \mathbb{S}_1 and \mathbb{S}_2 is a relation $R \subseteq S_1 \times S_2$ such that, if $s_1 R s_2$ then:

- $V_1^\dagger(s_1) = V_2^\dagger(s_2)$;
- for all Z_1 in $\sigma_1(s_1)$ there is Z_2 in $\sigma_2(s_2)$ such that for all $t_2 \in Z_2$ there is $t_1 \in Z_1$ with $t_1 R t_2$;
- for all Z_2 in $\sigma_2(s_2)$ there is Z_1 in $\sigma_1(s_1)$ such that for all $t_1 \in Z_1$ there is $t_2 \in Z_2$ with $t_1 R t_2$.

Neighborhood bisimilarity coincides with behavioural equivalence for \mathcal{M} , and in the next section we shall see how neighborhood bisimulations arise from a relation lifting for \mathcal{M} .

In many cases, behavioural equivalence of coalgebras can be described more concretely by bisimulations defined via certain relation liftings. For some work specifically on this topic, see the following papers by Marti and Venema: [19, 20, 21]. As an example, bisimilarity of Kripke models can be described using what is known as the ‘‘Egli-Milner lifting’’ for the powerset functor. This lifting assigns to a relation $R \subseteq X \times Y$ the ‘‘lifted’’ relation $\overline{\mathcal{P}}(R) \subseteq \mathcal{P}X \times \mathcal{P}Y$ defined by:

$$(\alpha, \beta) \in \overline{\mathcal{P}}(R) \Leftrightarrow (\forall x \in \alpha \exists y \in \beta : xRy) \ \& \ (\forall y \in \beta \exists x \in \alpha : xRy)$$

When the functor satisfies a certain condition, *preservation of weak pullbacks*, the appropriate relation lifting can be defined directly in terms of the functor. In the special case of endofunctors on **Set**, weak pullback preservation can be described rather concretely as the following condition [9]: consider any pair of maps $f_1 : X_1 \rightarrow Y$ and $f_2 : X_2 \rightarrow Y$, and any $\alpha_1 \in \mathbb{T}X_1$, $\alpha_2 \in \mathbb{T}X_2$ with $\mathbb{T}f_1(\alpha_1) = \mathbb{T}f_2(\alpha_2)$. Let

$$R = \{(u, v) \in X_1 \times X_2 \mid f_1(u) = f_2(v)\}$$

That is, R together with its projection maps $\pi_1 : R \rightarrow X_1$ and $\pi_2 : R \rightarrow X_2$ is the pullback of the co-span f_1, f_2 in **Set**. Then \mathbb{T} preserves weak pullbacks if, in this situation, we can always find some $\gamma \in \mathbb{T}R$ with $\mathbb{T}\pi_1(\gamma) = \alpha_1$ and $\mathbb{T}\pi_2(\gamma) = \alpha_2$.

Provided that \mathbb{T} preserves weak pullbacks, we can define a relation lifting $\overline{\mathbb{T}}$ known as the *Barr extension* of \mathbb{T} [3]. This is a functor on the category **Rel** of sets with binary relations as arrows, and formally its action on a relation $R \subseteq X_1 \times X_2$ is defined by setting:

$$\overline{\mathbb{T}}R := \{(\alpha_1, \alpha_2) \in \mathbb{T}X_1 \times \mathbb{T}X_2 \mid \exists \gamma \in \mathbb{T}R : \mathbb{T}\pi_1(\gamma) = \alpha_1 \ \& \ \mathbb{T}\pi_2(\gamma) = \alpha_2\}$$

The Egli-Milner lifting is the Barr extension of the covariant powerset functor, and bisimulations that are sound and complete for behavioural equivalence can be defined from the Barr extension of a functor in precisely the same manner as standard bisimulations are defined from the Egli-Milner lifting. This will not play a big role in the present paper, but the Barr extension will make a minor appearance when we discuss some basic facts about predicate liftings. We remark that the monotone neighborhood functor is known *not* to preserve weak pullbacks, so neighborhood bisimulations can not be recovered via the Barr extension. This was in fact part of the motivation behind the work on relation liftings by Marti and Venema.

3.2. The coalgebraic μ -calculus. The modal μ -calculus is just one in a family of logical systems that may collectively be referred to as the *coalgebraic μ -calculus* [4]. Generally, the key notion that connects coalgebra with logic is that of a *predicate lifting*. Here, we recall that \mathcal{Q} denotes the contravariant powerset functor, as defined in Section 3.1.

Definition 3.8. Given a set functor \mathbb{T} , an *n -place predicate lifting* for \mathbb{T} is a natural transformation

$$\lambda : \mathcal{Q}^n \rightarrow \mathcal{Q} \circ \mathbb{T},$$

where \mathcal{Q}^n denotes the n -fold product of \mathcal{Q} with itself. A predicate lifting λ is said to be *monotone* if

$$\lambda_X(Y_1, \dots, Y_n) \subseteq \lambda_X(Z_1, \dots, Z_n),$$

whenever $Y_i \subseteq Z_i$ for each i . The *Boolean dual* λ^d of λ is defined by

$$(Z_1, \dots, Z_n) \mapsto \mathbb{T}X \setminus (\lambda_X(X \setminus Z_1, \dots, X \setminus Z_n)).$$

We shall often think of an arbitrary finite set A of size n as being identified with the set $\{1, \dots, n\}$ via some fixed one-to-one correspondence, so that we may view any natural transformation $\lambda : \mathcal{Q}^A \rightarrow \mathcal{Q} \circ \mathbb{T}$ as an n -place predicate lifting. In these cases we sometimes speak of “predicate liftings over A ” rather than “ n -place predicate liftings”, but note that this is merely a notational variation rather than a substantial generalization of the notion of predicate liftings. A predicate lifting λ over A can be represented equivalently as a subset of $\mathbb{T}\mathcal{P}A$, via an application of the Yoneda lemma (this observation is due to Schröder, [31]). Roughly speaking, predicate liftings are used to express “local” properties about the next unfolding of the transition map of a coalgebra, viewed from a given state of the coalgebra and using properties of states in the coalgebra as parameters.

Given a set functor \mathbb{T} and a set of monotone predicate liftings Λ for \mathbb{T} , the language μML_Λ of the coalgebraic μ -calculus based on Λ is defined as follows. Again we present the language in negation normal form; note that we make sure that the modal operators of the language are closed under Boolean duals.

$$\varphi ::= p \mid \neg p \mid \lambda(\varphi_1, \dots, \varphi_n) \mid \lambda^d(\varphi_1, \dots, \varphi_n) \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \eta p.\varphi$$

where $p \in \text{Var}$, $\lambda \in \Lambda$, $\eta \in \{\mu, \nu\}$, and, in $\eta p.\varphi$, no free occurrence of the variable p in φ is in the scope of a negation. If Λ consists of *all* monotone predicate liftings for \mathbb{T} , then we denote the language μML_Λ also by $\mu\text{ML}_\mathbb{T}$.

The semantics of formulas in a pointed \mathbb{T} -model (\mathbb{S}, s) , where $\mathbb{S} = (S, \sigma, V)$, is defined as follows:

- $(\mathbb{S}, s) \Vdash p$ iff $s \in V(p)$ and $(\mathbb{S}, s) \Vdash \neg p$ iff $s \notin V(p)$.
- If $\lambda \in \Lambda$ or is the dual of some lifting in Λ , we set $(\mathbb{S}, s) \Vdash \lambda(\varphi_1, \dots, \varphi_n)$ iff $\sigma(s) \in \lambda_S(\|\varphi_1\|, \dots, \|\varphi_n\|)$, where $\|\varphi_i\| = \{t \in S \mid (\mathbb{S}, t) \Vdash \varphi_i\}$.
- Standard clauses for the boolean connectives
- $(\mathbb{S}, s) \Vdash \mu p.\varphi$ iff $s \in \bigcap \{X \subseteq S \mid \varphi_p(X) \subseteq X\}$, where $\varphi_p(Z)$ denotes the truth set of the formula φ in the \mathbb{T} -model $(S, \sigma, V[p \mapsto Z])$.

It is routine to prove that all formulas in $\mu\text{ML}_\mathbb{T}$ are bisimulation invariant, i.e. they have the same truth value in every pair of behaviourally equivalent pointed \mathbb{T} -models. (We shall continue to use the term “bisimulation invariant” rather than “invariant for behavioural equivalence”, although strictly speaking bisimulation and behavioural equivalence are distinct concepts.)

Example 3.9. The powerset functor \mathcal{P} has a unary predicate lifting \diamond defined by setting, for $\alpha \in \mathcal{P}X$:

$$\alpha \in \diamond_X(Z) \Leftrightarrow \alpha \cap Z \neq \emptyset$$

The language $\mu\text{ML}_{\{\diamond\}}$ is precisely the standard modal μ -calculus μML . The dual of \diamond is denoted \square , as usual. We recall its standard definition in modal logic as $\square\psi := \neg\diamond\neg\psi$, which is in accordance with our definition of the dual of a predicate lifting.

Example 3.10. The monotone neighborhood functor comes equipped with a unary predicate lifting, which we shall also denote by the symbol \square . It is defined by setting, for $\alpha \in \mathcal{M}X$:

$$\alpha \in \square_X(Z) \Leftrightarrow Z \in \alpha$$

In fact, we can embed \mathcal{P} as a subfunctor of \mathcal{M} by mapping $\alpha \in \mathcal{P}X$ to the neighborhood structure $\{Z \subseteq X \mid \alpha \subseteq Z\}$. Clearly the two predicate liftings denoted by \square both agree on this subfunctor, so it makes sense to use the same symbol to denote both. This lifting for \mathcal{M} corresponds to the usual box modality for monotone modal logic, and so the language

$\mu\text{ML}_{\{\Box\}}$ corresponds to a fixpoint extension of monotone modal logic that we will here refer to as the *monotone μ -calculus*, denoted μMML . The logic μMML is a very natural system to study: in a sense, it is the *minimal* modal μ -calculus since the *only* constraint we put on the box modality is monotonicity, which is precisely what is needed for the semantics of the fixpoint operators to be well defined. It is a well-behaved system: it was recently shown to have uniform interpolation [22], and we shall show here that it enjoys a Janin-Walukiewicz style characterization theorem.

3.3. Coalgebraic automata. Turning to the parity automata corresponding to the language μML_Λ , we first define the *modal one-step language* 1ML_Λ .

Definition 3.11. The set $1\text{ML}_\Lambda(A)$ of *modal one-step formulas* over a set A of variables is given by the following grammar:

$$\varphi ::= \perp \mid \top \mid \lambda(\psi_1, \dots, \psi_n) \mid \lambda^d(\psi_1, \dots, \psi_n) \mid \varphi \vee \varphi \mid \varphi \wedge \varphi$$

where ψ_1, \dots, ψ_n are formulas built up from variables in A using disjunctions and conjunctions. More formally, we require that $\psi_1, \dots, \psi_n \in \text{Latt}(A)$ where $\text{Latt}(A)$ denotes the set of *lattice formulas* over A given by the grammar:

$$\perp \mid \top \mid a \mid \varphi \vee \varphi \mid \varphi \wedge \varphi$$

where a ranges over A .

Definition 3.12. Given a functor \mathbb{T} and a set of variables A , a *one-step model* over A is a triple (X, α, V) where X is any set, $\alpha \in \mathbb{T}X$ and $V : A \rightarrow \mathcal{Q}(X)$ is a valuation.

Note that we have written the valuation V as having the type $A \rightarrow \mathcal{Q}(X)$ rather than $A \rightarrow \mathcal{P}(X)$ in this definition. This notation is equally correct since the covariant and contravariant powerset functors differ only in their action on morphisms, and it is sometimes a more convenient notation since the naturality condition of predicate liftings is formulated in terms of \mathcal{Q} and not \mathcal{P} .

The semantics of formulas in the modal one-step language in a one-step model is given as follows:

- standard clauses for the boolean connectives,
- $(X, \alpha, V) \Vdash_1 \lambda(\psi_1, \dots, \psi_n)$ iff $\alpha \in \lambda_X(\|\psi_1\|_V, \dots, \|\psi_n\|_V)$

where $\|\psi_i\|_V \subseteq X$ is the (classical) truth set of the formula ψ_i under the valuation V . Formally, this is defined by $\|\perp\|_V = \emptyset$, $\|\top\|_V = X$, $\|a\|_V = V(a)$ for $a \in A$, $\|\varphi \wedge \psi\|_V = \|\varphi\|_V \cap \|\psi\|_V$ and $\|\varphi \vee \psi\|_V = \|\varphi\|_V \cup \|\psi\|_V$. For $u \in X$ and a lattice formula ψ we also write $u \Vdash_V^0 \psi$ to say that $u \in \|\psi\|_V$.

Recalling that any predicate lifting over a set of n variables A corresponds to an n -place predicate lifting, we can view any predicate lifting λ over A as a one-step formula in $1\text{ML}_{\{\lambda\}}(A)$ ³. With this mind we can write $(X, \alpha, V) \Vdash_1 \lambda$ instead of $\alpha \in \lambda_X(V)$. In this

³Or rather, if we want to be completely precise, λ corresponds to a one-step formula in $1\text{ML}_{\{\lambda'\}}(A)$ where the lifting $\lambda' : \mathcal{Q}^n \rightarrow \mathcal{Q} \circ \mathbb{T}$ is obtained by composing λ with the natural isomorphism between functors \mathcal{Q}^n and \mathcal{Q}^A induced by some fixed bijection $i : n \rightarrow A$. If we write $A = \{a_1, \dots, a_n\}$ with $a_k = f(k-1)$, this means that we have $\lambda_X(V) = \lambda'_X(V(a_1), \dots, V(a_n))$. We shall permit some abuse of notation and simply identify λ with the associated lifting λ' .

notation the naturality constraint for a predicate lifting λ over A becomes, for every $\alpha \in X$, every map $f : X \rightarrow Y$ and every $V : A \rightarrow \mathcal{Q}(Y)$:

$$(X, \alpha, \mathcal{Q}f \circ V) \Vdash_1 \lambda \Leftrightarrow (Y, \mathbb{T}f(\alpha), V) \Vdash_1 \lambda.$$

We can now define the class of automata used to characterize the coalgebraic μ -calculus.

Definition 3.13. Let P be a finite set of variables and Λ a set of predicate liftings. Then a (P -chromatic) modal Λ -automaton is a tuple (A, Δ, Ω, a_I) where A is a finite set of states with $a_I \in A$,

$$\Delta : A \times \mathcal{P}(P) \rightarrow \mathbf{1ML}_\Lambda(A)$$

is the transition map of the automaton, and $\Omega : A \rightarrow \omega$ is the parity map. The class of these automata is denoted as $Aut(\mathbf{1ML}_\Lambda)$.

Note that there are two distinct sets of “variables” involved in the automaton \mathbb{A} , and it is important to keep these apart since they have different roles: the variables P are used to provide the alphabet of the automaton (and correspond to free variables of corresponding fixpoint formulas), while the variables A are the states of the automaton (and correspond to bound variables of a corresponding fixpoint formula.)

The acceptance game for an automaton $\mathbb{A} = (A, \Delta, \Omega, a_I)$ and a \mathbb{T} -model (S, σ, V) is the parity game given by the following table:

Position	Player	Admissible moves	Priority
$(a, s) \in A \times S$	\exists	$\{U \in (\mathcal{P}S)^A \mid (S, \sigma(s), U) \Vdash_1 \Delta(a, V^\dagger(s))\}$	$\Omega(a)$
$U : A \rightarrow \mathcal{P}S$	\forall	$\{(b, t) \mid t \in U(b)\}$	0

Table 2: Acceptance game for modal automata

The loser of a finite match is the player who got stuck, and the winner of an infinite match is \exists if the greatest parity that appears infinitely often in the match is even, and the winner is \forall if this parity is odd. Note that the valuations U and V in Table 2 play fundamentally different roles: V is a fixed valuation given by the model on which the automaton is run, assigning values to the open variables of the automaton, while U is a “local” valuation assigning values to the states of the automaton, which correspond roughly to bound variables of a fixpoint formula.

Given a strategy χ for either player, a match is said to be χ -guided if it is consistent with every choice made by that player according to χ . Winning strategies are defined as usual.

Definition 3.14. The automaton \mathbb{A} *accepts* the pointed model (\mathbb{S}, s) , written $(\mathbb{S}, s) \Vdash \mathbb{A}$, if \exists has a winning strategy in the acceptance game from the starting position (a_I, s) . We say that an automaton \mathbb{A} is *equivalent* to a formula $\varphi \in \mu\mathbf{ML}_\Lambda$ if, for every pointed \mathbb{T} -model (\mathbb{S}, s) , we have that \mathbb{A} accepts (\mathbb{S}, s) iff $(\mathbb{S}, s) \Vdash \varphi$.

Fact 3.15. [8] Let \mathbb{T} be a set functor, and Λ a set of monotone predicate liftings for \mathbb{T} . Then

$$\mu\mathbf{ML}_\Lambda \equiv Aut(\mathbf{1ML}_\Lambda).$$

That is, there are effective transformations of formulas in $\mu\mathbf{ML}_\Lambda$ into equivalent automata in $Aut(\mathbf{1ML}_\Lambda)$, and vice versa.

3.4. One-step expressive completeness. Our main result in this paper is formulated in terms of the set of *all* monotone predicate liftings for a given functor, and we would argue that this is a rather natural choice. However, in some cases it is possible to choose some smaller set of liftings, in order to have a more concrete and manageable presentation of the syntax of the corresponding μ -calculus. It will then be important to choose an *expressively complete* set of predicate liftings, meaning that *any* monotone lifting for the functor can be defined by a formula in the modal one-step language. In particular, this will be required to recover the Janin-Walukiewicz theorem in its original form as a special case of our main results.

Definition 3.16. Let Λ be a set of monotone predicate liftings and λ any predicate lifting over the finite set A . We say that λ is Λ -*definable* if there is a formula $\varphi \in \mathbf{1ML}_\Lambda(A)$ such that, for any one-step model (X, α, V) we have:

$$\alpha \in \lambda_X(V) \text{ iff } (X, \alpha, V) \Vdash_1 \varphi.$$

We say that a set of monotone predicate liftings Λ is *expressively complete* if every monotone predicate lifting over any given finite set A is Λ -definable.

Our main observation in this section is that any weak pullback preserving functor that also preserves finite sets has a natural choice of expressively complete predicate liftings.

Definition 3.17. Given a functor \mathbb{T} that preserves weak pullbacks, the n -ary *Moss lifting* corresponding to $\alpha \in \mathbb{T}\{1, \dots, n\}$ is the predicate lifting $\langle \alpha \rangle : \mathcal{Q}^n \rightarrow \mathcal{Q} \circ \mathbb{T}$ such that for all sets X and $Z_1, \dots, Z_n \subseteq X$, we have:

$$\langle \alpha \rangle_X(Z_1, \dots, Z_n) = \{\beta \in \mathbb{T}X \mid (\beta, \alpha) \in \overline{\mathbb{T}}R\}$$

where $R \subseteq X \times \{1, \dots, n\}$ is defined by uRk iff $u \in Z_k$.

In particular, for any given lattice formulas $\psi_1, \dots, \psi_n \in A$, any $\beta \in \mathbb{T}\{1, \dots, n\}$ and any one-step model (X, α, V) over A , we have:

$$X, \alpha, V \Vdash_1 \langle \beta \rangle(\psi_1, \dots, \psi_n) \text{ iff } (\alpha, \mathbb{T}f(\beta)) \in \overline{\mathbb{T}}(\Vdash_V^0)$$

where $f : \{1, \dots, n\} \rightarrow \{\psi_1, \dots, \psi_n\}$ is the map defined by $k \mapsto \psi_k$.

Note that the Moss lifting $\langle \alpha \rangle$ is always a natural transformation, as required: it is obtained by composing the natural transformation $h^\alpha : \mathcal{Q}^n \rightarrow \mathbb{T} \circ \mathcal{Q}$ defined by

$$h_X^\alpha(f) := \mathbb{T}f(\alpha), \text{ for } f : n \rightarrow \mathcal{Q}X,$$

with the distributive law $\delta : \mathbb{T} \circ \mathcal{Q} \rightarrow \mathcal{Q} \circ \mathbb{T}$ defined for $\beta \in \mathbb{T}\mathcal{Q}X$ by:

$$\delta_X(\beta) := \{\gamma \in \mathbb{T}X \mid (\gamma, \beta) \in \overline{\mathbb{T}}\in_X\}$$

where \in_X is the membership relation from X to $\mathcal{Q}X$. That this distributive law is a natural transformation is a standard fact, see for example [16] where this view of Moss formulas as special predicate liftings is investigated thoroughly.

Proposition 3.18. *If \mathbb{T} preserves weak pullbacks and finite sets, then the set of all Moss liftings for \mathbb{T} (of all arities $n \in \omega$) is expressively complete.*

Proof. Fix a monotone predicate lifting λ over A . Let (X, α, V) be any one-step model over A . Let n be the number of lattice formulas over A , and let $f : \mathbf{Latt}(A) \rightarrow \{1, \dots, n\}$ be some fixed bijective enumeration of $\mathbf{Latt}(A)$, so that we can write

$$\mathbf{Latt}(A) = f^{-1}(\{1, \dots, n\}) = \{\psi_1, \dots, \psi_n\}$$

Let $\theta : X \rightarrow \mathbf{Latt}(A)$ be the map sending each $u \in X$ to $\bigwedge\{a \in A \mid u \in V(a)\}$. The *characteristic formula* of (X, α, V) , denoted $\chi(X, \alpha, V)$, is defined as:

$$\langle \mathbb{T}(f \circ \theta)(\alpha) \rangle(\psi_1, \dots, \psi_n)$$

It is an exercise in coalgebra to show that, for any one-step model (X', α', V') , we have

$$X', \alpha', V' \Vdash_1 \chi(X, \alpha, V) \text{ iff } (\alpha, \alpha') \in \overline{\mathbb{T}}(\preceq)$$

where we write $u \preceq u'$, for $u \in X$ and $u' \in X'$, if $u \in V(a)$ implies $u' \in V'(a)$ for all $a \in A$. Furthermore, there are only finitely many characteristic formulas relative to any given set of variables A , since \mathbb{T} preserves finite sets (and hence $\mathbb{T}\{1, \dots, n\}$ is finite). We claim that the formula:

$$\varphi := \bigvee\{\chi(X, \alpha, V) \mid \alpha \in \lambda_X(V)\}$$

defines the monotone lifting λ . One direction is clear: if $\alpha \in \lambda_X(V)$ then $X, \alpha, V \Vdash_1 \varphi$ since $(X, \alpha, V) \Vdash_1 \chi(X, \alpha, V)$. Conversely, suppose that $X, \alpha, V \Vdash_1 \varphi$. Then there exists some one-step model (X', α', V') such that $\alpha' \in \lambda_{X'}(V')$ and $(\alpha', \alpha) \in \overline{\mathbb{T}}(\preceq)$. Write $R = \{(u, v) \mid u \preceq v\}$, let $\pi' : R \rightarrow X'$ and $\pi : R \rightarrow X$ be the projection maps, and let $\gamma \in \mathbb{T}R$ be such that $\mathbb{T}\pi'(\gamma) = \alpha'$ and $\mathbb{T}\pi(\gamma) = \alpha$. By naturality we get $\gamma \in \lambda_R(\mathcal{Q}\pi' \circ V')$, by monotonicity we get $\gamma \in \lambda_R(\mathcal{Q}\pi \circ V)$ (since $\pi'(r) \in V'(a)$ implies $\pi(r) \in V(a)$ for all $a \in A$), and finally by naturality again we get $\gamma \in \lambda_X(V)$ as required. \square

In the special case of the powerset functor, given a set of formulas $\{\psi_1, \dots, \psi_n\}$ (viewed as a member of $\mathcal{P}(\mathbf{Latt}(A))$), the corresponding substitution instance of the n -ary Moss lifting corresponds to the usual ‘‘cover’’ formula $\nabla\{\psi_1, \dots, \psi_n\}$ which is defined in terms of \diamond and its dual by:

$$\nabla\{\psi_1, \dots, \psi_n\} := \diamond\psi_1 \wedge \dots \wedge \diamond\psi_n \wedge \square(\psi_1 \vee \dots \vee \psi_n)$$

Hence, it follows from Proposition 3.18 that the single lifting \diamond for \mathcal{P} is expressively complete.

4. COALGEBRAIC MSO

4.1. Introducing coalgebraic MSO. We now introduce coalgebraic monadic second-order logic for a set functor \mathbb{T} and a set of liftings Λ , and show how MSO can be recovered as a special case. Given a set Λ of monotone predicate liftings, we define the syntax of the monadic second-order logic \mathbf{MSO}_Λ by the following grammar:

$$\varphi ::= \mathbf{sr}(p) \mid p \subseteq q \mid \lambda(p, q_1, \dots, q_n) \mid \varphi \vee \psi \mid \neg\varphi \mid \exists p. \varphi$$

where λ is any n -place monotone predicate lifting in Λ and $p, q, q_1, \dots, q_n \in \mathit{Var}$. For Λ equal to the set of all monotone predicate liftings for \mathbb{T} , we write $\mathbf{MSO}_\Lambda = \mathbf{MSO}_\mathbb{T}$.

For the semantics, let (\mathbb{S}, s) be a pointed \mathbb{T} -model. We define the satisfaction relation $\Vdash \subseteq S \times \mathbf{MSO}_\mathbb{T}$ as follows:

- $(\mathbb{S}, s) \Vdash \mathbf{sr}(p)$ iff $V(p) = \{s\}$,
- $(\mathbb{S}, s) \Vdash p \subseteq q$ iff $V(p) \subseteq V(q)$,
- $(\mathbb{S}, s) \Vdash \lambda(p, q_1, \dots, q_n)$ iff $\sigma(v) \in \lambda_S(V(q_1), \dots, V(q_n))$ for all $v \in V(p)$,
- standard clauses for the Boolean connectives,
- $(\mathbb{S}, s) \Vdash \exists p. \varphi$ iff $(S, \sigma, V[p \mapsto Z], s) \Vdash \varphi$, some $Z \subseteq S$.

We introduce the following abbreviations:

- $p = q$ for $p \subseteq q \wedge q \subseteq p$,

- $\text{Em}(p)$ for $\forall q.(p \subseteq q)$,
- $\text{Sing}(p)$ for $\neg \text{Em}(p) \wedge \forall q(q \subseteq p \rightarrow (\text{Em}(q) \vee q = p))$,

expressing, respectively, that p and q are equal, that p denotes the empty set, and that p denotes a singleton.

Clearly, standard MSO is a notational variant of the logic $\text{MSO}_{\{\diamond\}}$, since the formula $R(p, q)$ has precisely the same satisfaction clause as $\diamond(p, q)$ - to see this, one just has to unfold the definitions. We can state the Janin-Walukiewicz theorem in a formula as:

$$\mu\text{ML}_{\{\diamond\}} \equiv \text{MSO}_{\{\diamond\}} / \sim$$

The following useful fact is fairly easy to check, so we leave its proof to the reader:

Fact 4.1. If Λ is a one-step expressively complete set of predicate liftings, then $\text{MSO}_{\Lambda} \equiv \text{MSO}_{\top}$ and $\mu\text{ML}_{\Lambda} \equiv \mu\text{ML}_{\top}$.

Since we know that \diamond is expressively complete we can now state the Janin-Walukiewicz theorem in yet another form, where we recall that \mathcal{P} is the covariant powerset functor:

$$\mu\text{ML}_{\mathcal{P}} \equiv \text{MSO}_{\mathcal{P}} / \sim$$

Generally, we let MSO_{\top} / \sim denote the fragment of MSO_{\top} that is invariant for behavioural equivalence, and we refer to it as the *bisimulation invariant fragment* of MSO_{\top} .

As mentioned in the introduction, the key question in this paper will be to compare the expressive power of coalgebraic monadic second-order logic to that of the coalgebraic μ -calculus. The following observation, of which the (routine) proof is omitted, provides the easy part of the link.

Proposition 4.2. *Let Λ be a set of monotone predicate liftings for the set functor \top . There is an inductively defined translation $(\cdot)^{\diamond}$ mapping any formula $\varphi \in \mu\text{ML}_{\Lambda}$ to a semantically equivalent formula $\varphi^{\diamond} \in \text{MSO}_{\Lambda}$.*

Considering the lifting \square for \mathcal{M} , we introduce the name *monotone monadic second-order logic* for the language $\text{MSO}_{\{\square\}}$, which we will henceforth denote by MMSO . Note that the atomic formula $\square(p, q)$ encodes a pattern of quantifier alternation of the form $\forall\exists\forall$: it says that *every* state that satisfies p has *some* neighborhood Z such that *every* state in Z satisfies q .

4.2. Automata for coalgebraic MSO. In this section we introduce automata for coalgebraic monadic second-order logic, and translate formulas of MSO_{Λ} into automata operating on tree-like coalgebras.

Since the translation of standard MSO into parity automata is valid only over trees and not over Kripke frames in general, we should expect that coalgebraic monadic second-order logic similarly translates into parity automata over a restricted class of \top -models that are in some sense “tree-like”. So we need to start by spelling out exactly what this means.

Definition 4.3. Given a set S and $\alpha \in \top S$, a subset $X \subseteq S$ is said to be a *support* for α if there is some $\beta \in \top X$ with $\top_{X,S}(\beta) = \alpha$. A *supporting Kripke frame* for a \top -coalgebra (S, σ) is a binary relation $R \subseteq S \times S$ such that, for all $u \in S$, $R(u) = \{v \mid uRv\}$ is a support for $\sigma(u)$.

Given that $X \subseteq S$ is a support for $\alpha \in \mathbb{T}S$, our convention that \mathbb{T} preserves all injectives guarantees that there is a *unique* $\alpha' \in \mathbb{T}X$ such that $\mathbb{T}\iota_{X,S}(\alpha') = \alpha$. We shall denote this α' by $\alpha|_X$. Intuitively, X is a support for α if we can know all there is to know about α just by looking at its restriction $\alpha|_X$ to X .

Definition 4.4. A \mathbb{T} -tree model is a structure (\mathbb{S}, R, u) where $\mathbb{S} = (S, \sigma, V)$ is a \mathbb{T} -model and $u \in S$, such that R is a supporting Kripke frame for the coalgebra (S, σ) , and furthermore (S, R) is a tree rooted at u (i.e. there is a unique R -path from u to w for each $w \in S$).

Our goal is to translate formulas in $\text{MSO}_{\mathbb{T}}$ to equivalent automata over \mathbb{T} -tree models. We start by introducing a very general type of automaton, originating with [35]. The motivation for taking this general perspective is to emphasize that many automata theoretic concepts and basic results apply already in this setting, and we believe the general automaton concept we introduce here has some independent theoretical interest.

Definition 4.5. Given a finite set A , a *generalized predicate lifting* over A comprises an assignment of a map

$$\varphi_X : (\mathcal{Q}X)^A \rightarrow \mathcal{Q}\mathbb{T}X.$$

to every set X . Concepts like *boolean dual* and *monotonicity* apply to these liftings in the obvious way.

The difference with respect to standard predicate liftings is that the components of a generalized predicate lifting do not need to form a natural transformation. Throughout the rest of the paper, the term “predicate lifting” will be reserved for the subclass of general predicate liftings that satisfy the naturality constraint, i.e. predicate liftings in the usual sense.⁴

Definition 4.6. A *one-step language* L for a functor \mathbb{T} consists of a collection $L(A)$ of generalized predicate liftings for \mathbb{T} , for every finite set A .

As for ordinary predicate liftings, given a generalized predicate lifting φ and a one-step model (X, α, V) we sometimes write $(X, \alpha, V) \Vdash_1 \varphi$ rather than $\alpha \in \varphi_X(V)$.

Our automata will be indexed by a (finite) set of variables involved, corresponding to the set of free variables of the $\text{MSO}_{\mathbb{T}}$ -formula.

Definition 4.7. Let $P \subseteq \text{Var}$ be a finite set of variables and let L be a one-step language for the functor \mathbb{T} . A (*P-chromatic*) *L-automaton* is a structure (A, Δ, Ω, a_I) where

- A is a finite set, with $a_I \in A$,
- $\Omega : A \rightarrow \omega$ is a parity map, and
- $\Delta : A \times \mathcal{P}(P) \rightarrow L(A)$ is the transition map of \mathbb{A} .

There are two main differences between these automata and the modal automata we have considered in Section 3.3; first, we have dropped the naturality constraint on the one-step language. Second, these automata will run on *\mathbb{T} -tree models* rather than \mathbb{T} -models. The *acceptance game* of \mathbb{A} with respect to a \mathbb{T} -tree model (T, R, σ, V, u) is given by Table 3. We say that the automaton \mathbb{A} accepts the model (T, R, σ, V, u) if \exists has a winning strategy in this game (initialized at position (a_I, u)).

⁴ In the style of abstract logic, it would make sense to require a generalized predicate lifting to be natural with respect to certain maps, the natural choice in this case being bijections. For the purpose of this paper such a restriction is not needed, however.

Position	Player	Admissible moves	Priority
$(a, s) \in A \times T$	\exists	$\{U : A \rightarrow \mathcal{P}(R(s)) \mid (R(s), \sigma(s) _{R(s)}, U) \Vdash_1 \Delta(a, V^\dagger(s))\}$	$\Omega(a)$
$U : A \rightarrow \mathcal{P}(T)$	\forall	$\{(b, t) \mid t \in U(b)\}$	0

Table 3: Acceptance game for L-automata.

4.3. Closure properties. The abstract perspective of generalized predicate liftings is useful for establishing some simple closure properties of automata, based on properties of the one-step language. The first, easy, results establish sufficient conditions for closure under union and complementation.

Proposition 4.8. *If the one-step language L is closed under disjunction, then the class of L-automata is closed under union.*

Definition 4.9. An L-automaton \mathbb{A} is said to be *monotone* if, for all $a \in A$ and all colors $c \in \mathcal{P}(P)$, the generalized lifting $\Delta(a, c)$ is monotone in each variable in A .

Proposition 4.10. *If the monotone fragment of the one-step language L is closed under boolean duals, then the class of monotone L-automata is closed under complementation.*

The most interesting property concerns closure under existential projection.

Definition 4.11. Let \mathbb{A} be a P -chromatic L-automaton, and let $p \in P$. Then a $P \setminus \{p\}$ -chromatic L-automaton \mathbb{B} is said to be an *existential projection* of \mathbb{A} if it accepts a \top -tree model (T, R, σ, V, u) precisely when there is some $Z \subseteq T$ such that \mathbb{A} accepts $(T, R, \sigma, V[p \mapsto Z], u)$.

The following notion originates with [13], but instead of relying on a particular syntactic shape of one-step formulas as in [13], we define the concepts in purely semantic terms.

Definition 4.12. A generalized predicate lifting φ over A is said to be *disjunctive* if, for every one-step model (X, α, V) such that

$$(X, \alpha, V) \Vdash_1 \varphi$$

there is a valuation $V^* : A \rightarrow \mathcal{Q}(X)$ such that

- $V^*(a) \subseteq V(a)$ for each $a \in A$,
- $V^*(a) \cap V^*(b) = \emptyset$ whenever $a \neq b$, and
- $(X, \alpha, V^*) \Vdash_1 \varphi$.

The second clause can also be stated as: $(V^*)^\dagger(u)$ is either empty or a singleton, for all $u \in X$. Call an L-automaton *non-deterministic* if every generalized lifting $\Delta(a, c)$ is disjunctive.

It is easy to see that if the language L is closed under disjunctions, then so is its disjunctive fragment. From this we obtain the following.

Proposition 4.13. *If the one-step language L is closed under disjunction, then the class of non-deterministic L-automata is closed under existential projection over \top -tree models.*

Proof. Suppose $\mathbb{A} = (A, \Delta, a_I, \Omega)$ is a non-deterministic L-automaton for the variable set P . Define the $P \setminus \{p\}$ -chromatic automaton $\exists p.\mathbb{A} = (A, \Delta^*, a_I, \Omega)$ by setting

$$\Delta^*(a, c) = \Delta(a, c) \vee \Delta(a, c \cup \{p\}).$$

Every \top -tree model accepted by \mathbb{A} is also accepted by $\exists p.\mathbb{A}$, since any admissible move for \exists at a position (a, u) in the acceptance game for \mathbb{A} is an admissible move at the same position in the acceptance game for $\exists p.\mathbb{A}$ too, by definition of the transition map Δ^* .

Conversely, suppose $\exists p.\mathbb{A}$ accepts some \mathbb{T} -tree model (S, R, σ, V, s_I) . For each winning position (a, s) in the acceptance game, let $V_{(a,s)}$ be the valuation chosen by \exists according to some given winning strategy χ . Note that we can assume that χ is a *positional* winning strategy, since $\exists p.\mathbb{A}$ is a parity automaton. Since disjunctive formulas are closed under disjunctions, the automaton $\exists p.\mathbb{A}$ is a non-deterministic automaton, and so for each winning position (a, s) there is a valuation $V_{(a,s)}^* : A \rightarrow \mathcal{P}(R(s))$, which is an admissible move for \exists , such that $V_{(a,s)}^*(b) \subseteq V_{(a,s)}(b)$ and such that for all $b_1 \neq b_2 \in A$ we have $V_{(a,s)}^*(b_1) \cap V_{(a,s)}^*(b_2) = \emptyset$. Define the strategy χ^* by letting \exists choose the valuation $V_{(a,s)}^*$ at each winning position (a, s) - this is still a winning strategy, since the valuations chosen by \exists are smaller and so no new choices for \forall are introduced. Furthermore, χ^* is clearly still a positional winning strategy.

From these facts it follows by a simple induction on the depth of the nodes in the supporting tree that the strategy χ^* is *separating*, i.e. that for every $s \in S$ there is at most one automaton state a such that (a, s) appears in a χ^* -guided match of the acceptance game. So we can define a valuation V' like V except we evaluate p to be true at all and only the states s such that

$$(R(s), \sigma(s), V_{(a_s,s)}^*) \Vdash_1 \Delta(a_s, c \cup \{p\}),$$

where a_s is a necessarily *unique* automaton state such that (a, s) appears in some χ^* -guided match, and c is the color consisting of the variables true under V at s . Then \mathbb{A} accepts (S, R, σ, V', s_I) , by the positional strategy mapping each position (a, s) to the valuation $\chi^*(a, s)$. That is, \exists can play the same strategy $\chi^*(a, s)$, now viewed as a strategy in the acceptance game for \mathbb{A} . \square

4.4. Second-order automata. We now introduce a concrete one-step language for a given set functor \mathbb{T} and a given set of monotone predicate liftings Λ , and show that \mathbf{MSO}_Λ can be translated into the corresponding class of automata. By the closure properties established in the previous section, the language needs to be closed under disjunction and negation, and furthermore we need to establish that every automaton for this language is equivalent to a non-deterministic automaton.

Definition 4.14. Let Λ be a set of monotone predicate liftings for \mathbb{T} . The set of *second-order one-step formulas* over any set of variables A and relative to the set of liftings Λ is defined by the grammar:

$$\varphi ::= a \subseteq b \mid \lambda(a_1, \dots, a_n) \mid \neg\varphi \mid \varphi \vee \varphi \mid \exists a.\varphi,$$

where $a, b, a_1, \dots, a_n \in A$ and λ is any predicate lifting in Λ .

The semantics of a second-order one-step formula in a one-step model (X, α, V) (with $V : A \rightarrow \mathcal{P}(X)$) is defined by the following clauses:

- $(X, \alpha, V) \Vdash_1 a \subseteq b$ iff $V(a) \subseteq V(b)$,
- $(X, \alpha, V) \Vdash_1 \lambda(a_1, \dots, a_n)$ iff $\alpha \in \lambda_X(V(a_1), \dots, V(a_n))$,
- standard clauses for the Boolean connectives,
- $(X, \alpha, V) \Vdash_1 \exists a.\varphi$ iff $(X, \alpha, V[a \mapsto S]) \Vdash_1 \varphi$ for some $S \subseteq X$.

Fixing an infinite set of “one-step variables” \mathbf{Var}_1 , and given a finite set A , the set of *second-order one-step sentences* over A , denoted $\mathbf{1SO}_\Lambda(A)$, is the set of one-step formulas over $A \cup \mathbf{Var}_1$, with all free variables belonging to A .

Remark 4.15. The difference between the one-step language $\mathbf{1SO}_\Lambda(A)$ and the full language \mathbf{MSO}_Λ may seem rather subtle at first sight. It is important to note that an n -place predicate lifting now corresponds to an n -place predicate in the language, *not* an $n + 1$ -place predicate as in \mathbf{MSO}_Λ . While the formula $\lambda(q_1, \dots, q_n)$ of $\mathbf{1SO}_\Lambda(A)$ expresses a proposition about the specific, fixed object α in a given one-step model, the formula $\lambda(p, q_1, \dots, q_n)$ in \mathbf{MSO}_Λ rather describes the transition map of a coalgebra as a whole: it says that the condition $\lambda(q_1, \dots, q_n)$ holds for the unfolding of each state s that satisfies p .

Any second-order one-step A -sentence φ can be regarded as a generalized predicate lifting over A , with

$$\varphi_X(V) = \{\alpha \in \mathsf{TX} \mid (X, \alpha, V) \Vdash_1 \varphi\}.$$

In particular, general concepts like monotonicity and closure under boolean duals apply to second-order one-step sentences. Also, note that second-order sentences are invariant under a natural notion of isomorphism:

Definition 4.16. An *isomorphism* between two one-step models (X_1, α_1, V_1) and (X_2, α_2, V_2) is a bijection $i : X_1 \rightarrow X_2$ such that $\mathsf{Ti}(\alpha_1) = \alpha_2$ and $V_1^\dagger(u) = V_2^\dagger(i(u))$ for each $u \in X_1$.

Proposition 4.17. *Given any set of predicate liftings Λ and a set of variables A , any two isomorphic one-step models satisfy the same formulas in the one-step language $\mathbf{1SO}_\Lambda(A)$.*

We can now introduce second-order automata; since every sentence in $\mathbf{1SO}_\Lambda(A)$ corresponds to a generalized predicate lifting, we may view $\mathbf{1SO}_\Lambda$ (the assignment of the one-step second-order A -sentences $\mathbf{1SO}_\Lambda(A)$ to every set of variables A) as a one-step language in the sense of Definition 4.6. Hence the following definition is sound:

Definition 4.18. Let Λ be a set of monotone predicate liftings for T . A *second-order Λ -automaton* is an $\mathbf{1SO}_\Lambda$ -automaton. We write $\mathit{Aut}(\mathbf{1SO}_\Lambda)$ to denote this class, and $\mathit{Aut}(\mathbf{1SO}_\mathsf{T})$ in case Λ is the set of *all* monotone predicate liftings for T .

Formulas of $\mathbf{1SO}_\mathsf{T}(A)$ are not in general monotone in the variables A since negation is present in the language (unlike the one-step language $\mathbf{1ML}_\Lambda$). However, a useful trick due to Walukiewicz allows us at any time to safely replace all one-step formulas in an automaton by monotone ones:

Proposition 4.19. *Let Λ be any set of monotone predicate liftings. Then every second-order automaton $\mathbb{A} \in \mathit{Aut}(\mathbf{1SO}_\Lambda)$ is equivalent to a monotone second-order automaton $\mathbb{A} \in \mathit{Aut}(\mathbf{1SO}_\Lambda)$.*

Proof. Enumerate A as $\{a_1, \dots, a_k\}$, and just replace each formula $\Delta(a, c)$ by

$$\exists z_1 \dots \exists z_k. z_1 \subseteq a_1 \wedge \dots \wedge z_k \subseteq a_k \wedge \Delta(a, c)[(z_i/a_i)_{i \in \{1, \dots, k\}}]$$

where $\Delta(a, c)[(z_i/a_i)_{i \in \{1, \dots, k\}}]$ is the result of substituting the variable z_i for each open variable a_i in $\Delta(a, c)$. This new formula is monotone in the variables A and the resulting automaton is equivalent to \mathbb{A} . \square

Our aim in this section is to show that every formula of \mathbf{MSO}_Λ is equivalent to a monotone second-order Λ -automaton, and we shall prove this by induction on the complexity of a formula. By the previous proposition, it suffices at each step to prove the existence of a second-order Λ -automaton equivalent to the given formula, since this automaton is then guaranteed to be equivalent to a monotone one. So the induction proceeds as follows: first, we produce second-order automata equivalent to all atomic formulas. By Proposition

4.19 this gives us monotone second-order automata for all atomic formulas. Then, for the inductive steps, we assume that we are given monotone second-order automata equivalent to φ and ψ respectively. Our aim is to produce second-order automata that are equivalent to $\varphi \vee \psi$, $\neg\varphi$ and $\exists p.\varphi$, and once this is established we can apply Proposition 4.19 again to finish the induction.

For the atomic formulas, we present only the case of formulas of the form $\lambda(p, q_1, \dots, q_n)$. First, note that by naturality of λ , we have for any given T-tree model (\mathbb{S}, R, s) that $(\mathbb{S}, R, s) \Vdash \lambda(p, q_1, \dots, q_n)$ if, and only if, for each $u \in V(p)$:

$$\sigma(s)|_{R(s)} \in \lambda_{R(s)}(V(q_1) \cap R(s), \dots, V(q_n) \cap R(s))$$

With this in mind, it is not hard to check that the following automaton $\mathbb{A} = (A, \Delta, a_I, \Omega)$ is equivalent to $\lambda(p, q_1, \dots, q_n)$ over T-tree models, where:

- $A = \{a_I, b_1, \dots, b_n\}$
- $\Omega(a_I) = \Omega(b_1) = \dots = \Omega(b_n) = 0$
- $\Delta(a_I, c) = \begin{cases} \forall z.z \subseteq a_I & \text{if } p \notin c \\ \lambda(b_1, \dots, b_n) \wedge \forall z.z \subseteq a_I & \text{if } p \in c \end{cases}$
- $\Delta(b_i, c) = \begin{cases} \perp & \text{if } q_i \notin c \\ \top & \text{if } q_i \in c \end{cases}$

Intuitively, this automaton goes on indefinitely searching the underlying tree of the model for states that satisfy p , and whenever it finds such a state u it checks whether $u \in \lambda_S(V(q_1), \dots, V(q_n))$.

For the induction, we treat the Boolean cases first: the one-step language is clearly closed under disjunction, so second-order Λ -automata are closed under union by Proposition 4.8. For negation, suppose we have a monotone second-order automaton \mathbb{A} equivalent to φ . It is not hard to show that for any monotone one-step formula we can define its Boolean dual in the one-step language, and furthermore the dual is again monotone. Hence, since \mathbb{A} was monotone, we find an automaton \mathbb{B} equivalent to $\neg\varphi$ by Proposition 4.10. To complete the translation of MSO_Λ into second-order Λ -automata we now only need to prove closure under existential projection.

4.5. A coalgebraic simulation theorem. In this section we prove that (monotone) second-order Λ -automata are closed under existential projection. The key to this result is to prove a simulation theorem, showing that every second-order automaton is equivalent to a non-deterministic one. The result then follows from Proposition 4.13.

The intuition behind the simulation theorem is the same as that behind the standard “powerset construction” for word automata: the states of the new non-deterministic automaton \mathbb{A}_n are “macro-states” representing several possible states of \mathbb{A} at once. Formally, the states of \mathbb{A}_n will be binary relations over A , and given a macro-state R , its range gives an exact description of the states in \mathbb{A} that are currently being visited simultaneously. It is safe to think of the macro-states as subsets of A , however: the only reason that we have binary relations over A as states rather than just subsets is to have a memory device so that we can keep track of traces in infinite matches. For each macro-state R and each colour c we want to be able to say that the one-step formulas corresponding to each state in the range of R hold, so we want to translate the one-step formulas over A into one-step formulas over the set of macro-states. In order to translate a formula $\Delta(a, c)$ to a new one-step formula with macro-states as variables, we have to replace the variable b in $\Delta(a, c)$ with a new variable

that acts as a stand-in for b . For this purpose we introduce a new, existentially quantified variable z_b , together with a formula stating explicitly that z_b is to represent a “disjunction” of all those macro states that contain b . Furthermore we want all the one-step formulas to be disjunctive, and for this purpose we simply add a conjunct “Disj” to each one-step formula, stating that the values of any pair of distinct variables appearing in the formula are to be disjoint. Finally, in order to turn \mathbb{A}_n into a parity automaton, we use a stream automaton to detect bad traces (see for instance [34] for the details in a more specific case).

Theorem 4.20 (Simulation). *Let Λ be a set of monotone predicate liftings for \top . For any monotone automaton $\mathbb{A} \in \text{Aut}(\mathbf{1SO}_\Lambda)$ there exists an equivalent non-deterministic $\mathbb{A}' \in \text{Aut}(\mathbf{1SO}_\Lambda)$.*

Fix an automaton $\mathbb{A} = (A, \Delta, a_I, \Omega)$. We consider the set $\mathcal{P}(A \times A)$ of binary relations over A as a set of variables. Let

$$\text{Disj} := \bigwedge_{R \neq R' \subseteq A \times A} \forall x. (x \subseteq R \wedge x \subseteq R') \rightarrow \text{Em}(x)$$

This formula says that any two distinct relations, viewed as variables, are given disjoint values. Note that the formula does *not* express the obviously contradictory condition that any two relations $R \neq R' \subseteq A \times A$ themselves are disjoint! Relations over A are here being treated as formal variables, and the formula is true in any one-step model (X, α, V) such that $V(R) \cap V(R') = \emptyset$ whenever $R \neq R'$.

Pick a fresh variable z_a for each $a \in A$. Given a 1-step formula φ , let

$$\varphi[(z_a/a)_{a \in A}]$$

be the result of substituting z_a for each free variable $a \in A$ in φ . If we enumerate the elements of A as a_1, \dots, a_k , we now define the formula $\text{Sim}(\varphi, b)$ for $b \in A$ to be

$$\exists z_{a_1} \dots \exists z_{a_k}. \bigwedge_{1 \leq i \leq k} \text{Rep}(z_{a_i}, \{R' \mid (b, a_i) \in R'\}) \wedge \varphi[(z_a/a)_{a \in A}]$$

where $\text{Rep}(z_{a_i}, \{R' \mid (b, a_i) \in R'\})$ is the formula:

$$\forall x. \text{Sing}(x) \rightarrow (x \subseteq z_{a_i} \leftrightarrow \bigvee \{x \subseteq R' \mid (b, a_i) \in R'\})$$

In informal terms, the formula $\text{Rep}(z_{a_i}, \{R' \mid (b, a_i) \in R'\})$ says that the variable z_{a_i} represents a disjunction of all the macro-states that contain a_i . The formula $\text{Sim}(\varphi, b)$ can thus be thought of as reformulating the formula φ in terms of macro-states.

Let $\mathbb{A} = (A, \Delta, a_I, \Omega)$ be any monadic Λ -automaton. We can assume w.l.o.g. that \mathbb{A} is monotone. We first construct the automaton $\mathbb{A}_n = (A_n, \Delta_n, a_I^*, F)$ with a non-parity acceptance condition $F \subseteq (A_n)^\omega$ as follows:

- $A_n = \mathcal{P}(A \times A)$
- $\Delta_n(R, c) = \text{Disj} \wedge \bigwedge_{b \in \pi_2[R]} \text{Sim}(\Delta(b, c), b)$
- $a_I^* = \{(a_I, a_I)\}$
- F is the set of streams over $\mathcal{P}(A \times A)$ with no bad traces.

Here, π_2 is the second projection of a relation R so that $\pi_2[R]$ denotes the range of R . A *trace* in a stream (R_1, R_2, R_3, \dots) over $\mathcal{P}(A \times A)$ is a stream (a_1, a_2, a_3, \dots) over A with $a_1 \in \pi_2[R_1]$ and $(a_{j-1}, a_j) \in R_j$ for $j > 1$. A trace is *bad* if the greatest number n with $\Omega(a_i) = n$ for infinitely many a_i is odd.

Proposition 4.21. *The automaton \mathbb{A}_n is disjunctive.*

Proof. Since every formula $\Delta(R, c)$ contains the formula Disj as a conjunct, any one-step model (X, α, V) that satisfies $\Delta(R, c)$ must be such that $V(R) \cap V(R') = \emptyset$ whenever $R \neq R'$. Clearly, this means that each formula $\Delta(R, c)$ is disjunctive. \square

The following lemma is routine to check, so we leave it to the reader:

Lemma 4.22. *\mathbb{A}_n accepts precisely the same \mathbb{T} -tree models as \mathbb{A} .*

Proof of Theorem 4.20. The only thing left to do at this point is to transform the automaton \mathbb{A}_n into an equivalent one that has its acceptance condition given by a parity map. We make use of the following proposition (see [34]):

Claim 4.23. The set F of streams over $\mathcal{P}(A \times A)$ that contain no bad traces w.r.t. the parity map Ω is an ω -regular stream language over the alphabet $\mathcal{P}(A \times A)$.

By Claim 4.23, there exists a deterministic parity stream automaton

$$\mathbb{Z} = (Z, \delta, z_I, \Omega_z)$$

that recognizes the language F , with $\delta : Z \times \mathcal{P}(A \times A) \rightarrow Z$. We now construct the automaton

$$\mathbb{A}_n \odot \mathbb{Z} = (A'_n, \Delta'_n, a'_I, \Omega'_n)$$

as follows:

- $A'_n = A_n \times Z$
- $a'_I = (a_I^*, z_I)$
- $\Omega'_n(R, z) = \Omega(z)$
- $\Delta'_n((R, z), c) = \Delta_n(R, s)[((R', \delta(R, z))/R')_{R' \in \mathcal{P}(A \times A)}]$

It is not difficult to check that $\mathbb{A}_n \odot \mathbb{Z}$ is equivalent to \mathbb{A}_n . Since $\mathbb{A}_n \odot \mathbb{Z}$ is clearly still a non-deterministic automaton, this ends the proof of the simulation theorem. \square

Combining Proposition 4.13 with Theorem 4.20, we obtain closure under existential projection.

Proposition 4.24. *Let Λ be a set of monotone predicate liftings for a set functor \mathbb{T} . Over \mathbb{T} -tree models, the class of second-order Λ -automata is closed under existential projection.*

We can now combine the closure properties we have established for second-order automata to give the desired translation of $\text{MSO}_{\mathbb{T}}$ into second-order automata:

Theorem 4.25. *For every formula $\varphi \in \text{MSO}_{\Lambda}$ with free variables in P , there exists a monotone P -chromatic second-order Λ -automaton $\mathbb{A}_{\varphi} \in \text{Aut}(\mathbf{1SO}_{\Lambda})$ which is equivalent to φ over \mathbb{T} -tree models.*

Proof. It suffices to show that every formula φ is equivalent to a P -chromatic Λ -automaton, since the monotonicity constraint can then be taken care of using Proposition 4.19. We prove the theorem by induction on the complexity of MSO_{Λ} -formulas. We already showed that equivalent automata are available for the atomic formulas of MSO_{Λ} . The induction step for disjunctions follows immediately from Proposition 4.8, since the one-step language $\mathbf{1SO}_{\Lambda}$ is clearly closed under disjunctions. The induction step for negations similarly follows from Proposition 4.10 once we check that the monotone fragment of the one-step language $\mathbf{1SO}_{\Lambda}$ is closed under boolean duals (here, we are again appealing to Proposition 4.19, to ensure that

it suffices to show that monotone second-order automata are closed under complementation). To see this, suppose that $\psi \in \mathbf{1SO}_\Lambda(A)$ is a monotone one-step sentence, and suppose the free variables of ψ are $a_1, \dots, a_n \in A$. For each a_i , add a fresh variable p_i , and let $\mathbf{Comp}(p_i, a_i)$ be the formula:

$$\forall x.(x \subseteq p_i \leftrightarrow \forall y.((y \subseteq x \wedge y \subseteq a_i) \rightarrow \mathbf{Em}(y)))$$

Informally, this formula says that p_i is the complement of a_i . Now, let ψ^∂ be defined as:

$$\exists p_1 \dots \exists p_n. \mathbf{Comp}(p_1, a_1) \wedge \dots \wedge \mathbf{Comp}(p_n, a_n) \wedge \neg \psi[p_1/a_1, \dots, p_n/a_n]$$

Then ψ^∂ is a monotone formula with free variables a_1, \dots, a_n , and its semantics in a one-step model is indeed that of the boolean dual of ψ .

Finally, the induction step for the existential second-order quantifier is handled as follows: suppose that \mathbb{A} is a P -chromatic second-order automaton equivalent over \mathbb{T} -tree models to the formula φ and let $p \in P$. By Theorem 4.20, we may assume without loss of generality that \mathbb{A} is disjunctive. By Proposition 4.13, there is a $P \setminus \{p\}$ -chromatic second-order automaton $\exists p.\mathbb{A}$ that defines the projection of \mathbb{A} with respect to the variable p . Hence, this automaton is equivalent to $\exists p.\varphi$ as required. \square

Corollary 4.26. *Suppose Λ is an expressively complete set of monotone predicate liftings for \mathbb{T} . Then for every formula of $\mathbf{MSO}_\mathbb{T}$, there exists an equivalent second-order Λ -automaton over \mathbb{T} -tree models.*

5. BISIMULATION INVARIANCE

5.1. Uniform translations. This section continues the program of [35], making use of the automata-theoretic translation of $\mathbf{MSO}_\mathbb{T}$ we have just established. The gist of our approach is that, in order to characterize a coalgebraic fixpoint logic $\mu\mathbf{ML}_\mathbb{T}$ as the bisimulation-invariant fragment of $\mathbf{MSO}_\mathbb{T}$, it suffices to establish a certain type of translation between the corresponding one-step languages. First we need some definitions.

Definition 5.1. Given sets X, Y , a mapping $h : X \rightarrow Y$ and a valuation $V : A \rightarrow \mathcal{Q}(Y)$, we define the valuation $V_{[h]} : A \rightarrow \mathcal{Q}(X)$ by setting $V_{[h]}(b) = h^{-1}(V(b))$ for each $b \in A$. In other words:

$$V_{[h]} := \mathcal{Q}h \circ V$$

Note that for a pair of composable maps f, g , we have $V_{[f \circ g]} = (V_{[g]})_{[f]}$. The most important concept that we take from [35] is that of a *uniform translation* (called *uniform correspondence* in [35]).

Definition 5.2. A *one-step frame* is a pair (X, α) with $\alpha \in \mathbb{T}X$. A *morphism* of one-step frames $h : (X', \alpha') \rightarrow (X, \alpha)$ is a map $h : X' \rightarrow X$ with $\mathbb{T}h(\alpha') = \alpha$. A one-step frame (X', α') together with a morphism $h : (X', \alpha') \rightarrow (X, \alpha)$ is called a *cover* of (X, α) .

We can now define the notions of uniform translations and uniform constructions:

Definition 5.3. Given a functor \mathbb{T} , a *uniform construction* for \mathbb{T} is an assignment of a cover $h_\alpha : (X_*, \alpha_*) \rightarrow (X, \alpha)$ to every one-step frame (X, α) .

We will usually denote the uniform construction consisting of an assignment of covers $h_\alpha : (X_*, \alpha_*) \rightarrow (X, \alpha)$ simply by $(-)_*$.

Remark 5.4. A good way to think of what uniform constructions will achieve, when we use them to obtain a characterization theorem later, is as methods to neutralize all the powers of distinction that the second-order one-step language has over the modal one-step language 1ML_Λ . Given two one-step models that cannot be distinguished by the modal one-step language, we want to apply suitable uniform constructions to re-shape these models so that they cannot be distinguished by the second-order one-step language either. For example, the second-order one-step language for the powerset functor is equivalent to first-order monadic logic with equality [36], which essentially adds to the modal one-step language the capability of *counting* elements of one-step models. It is precisely this capacity for counting that is the reason why formulas in this one-step language do not generally correspond to predicate liftings for the powerset functor, since conditions like “having at least two elements” is clearly not preserved by morphisms between one-step models. So we are looking for a uniform construction that prevents making any distinctions between one-step models by counting, and the solution is to create infinitely many copies of each element. This corresponds to the ω -unravelling construction used by [13].

Definition 5.5. We say that a one-step language L *admits uniform translations* if, for any finite set A and any finite set Γ of formulas in $L(A)$, there exists a uniform construction $(-)_*$ and a translation $(-)^* : \Gamma \rightarrow 1\text{ML}_\Lambda(A)$ such that for any one-step model (X, α, V) and any $\varphi \in \Gamma$, we have

$$(X, \alpha, V) \Vdash_1 \varphi^* \text{ iff } (X_*, \alpha_*, V_{[h_\alpha]}) \Vdash_1 \varphi$$

where h_α is the cover assigned to (X, α) by the construction. The pair consisting of the translation $(-)^*$ and the uniform construction $(-)_*$ will be referred to as a *uniform translation* for the set of formulas Γ .

As an example, consider the disjunctive formulas introduced by Walukiewicz,

Definition 5.6. Given a P -chromatic second-order automaton $\mathbb{A} = (A, \Delta, a_I, \Omega)$ and a uniform translation for the finite set $\Delta[A \times \mathcal{P}(P)]$, we get a corresponding modal automaton $\mathbb{A}^* = (A, \Delta^*, a_I, \Omega)$, with Δ^* given by $\Delta^*(a, c) := (\Delta(a, c))^*$. We overload the notation $(\cdot)^*$ to refer both to a uniform one-step translation and the induced translation at the level of automata.

The proof of the following result closely follows that of the main result in [35]. The only difference with [35] is that here we need an “unravelling”-like component, in order to turn an arbitrary model into a tree-like one.

Proposition 5.7. *Assume that 1SO_Λ admits uniform translations, and let \mathbb{A} be a second-order Λ -automaton. Then for each pointed \mathbb{T} -model (\mathbb{S}, s) there is a \mathbb{T} -tree model (\mathbb{T}, R, t) , with a \mathbb{T} -model morphism f from \mathbb{T} to \mathbb{S} , mapping t to s , and such that*

$$\mathbb{A} \text{ accepts } (\mathbb{T}, R, t) \text{ iff } \mathbb{A}^* \text{ accepts } (\mathbb{S}, s).$$

Proof. Consider any given pointed \mathbb{T} -model (\mathbb{S}_1, s_1) where $\mathbb{S}_1 = (S_1, \sigma_1, V_1)$. We are going to construct a \mathbb{T} -tree model (\mathbb{S}_2, R, s_2) , $\mathbb{S}_2 = (S_2, \sigma_2, V_2)$, together with a model morphism from the underlying pointed \mathbb{T} -model \mathbb{S}_2 to \mathbb{S}_1 mapping s_2 to s_1 , and such that \mathbb{A} accepts the \mathbb{T} -tree model (\mathbb{S}_2, s_2) if and only if \mathbb{A}^* accepts the pointed \mathbb{T} -model (\mathbb{S}_1, s_1) .

We construct this \mathbb{T} -tree model as follows: for each $u \in S_1$, we define an associated pair (X_u, α_u) by setting $(X_u, \alpha_u) = ((S_1)_*, \sigma_1(u)_*)$. Observe that, by the construction of these one-step models, for each $u \in S_1$, there is a mapping $\xi_u : X_u \rightarrow S_1$ such that:

$$(1) \quad \mathbb{T}(\xi_u)(\alpha_u) = \sigma_1(u)$$

(2) For each valuation $U : A \rightarrow \mathcal{Q}(S_1)$, every $u \in S_1$ and every one-step formula $\Delta(a, c)$ appearing in \mathbb{A} , we have

$$(S_1, \sigma_1(u), U) \Vdash_1 \Delta^*(a, c) \text{ iff } (X_u, \alpha_u, U_{[\xi_u]}) \Vdash_1 \Delta(a, c)$$

The map ξ_u is given by $h_{\sigma_1(u)}$. We now construct the T-tree model $(S_2, R, \sigma_2, s_2, V_2)$ as follows: first, consider the set of all non-empty finite (non-empty) tuples (v_1, \dots, v_n) of elements in

$$\{s_1\} \cup \bigcup_{u \in S_1} X_u$$

such that $v_1 = s_1$. We define, by induction, for each natural number $n > 0$ a subset M_n of this set, and a mapping $\gamma_n : M_n \rightarrow S_1$, as follows:

- Set $M_1 = \{(s_1)\}$, and define $\gamma_1(s_1) = s_1$.
- Set $M_{n+1} = \{\vec{v}w \mid \vec{v} \in M_n, w \in X_{\gamma_n(\vec{v})}\}$. Define $\gamma_{n+1}(\vec{v}w) = \xi_{\gamma_n(\vec{v})}(w)$.

Here, we write $\vec{v}w$ to denote the tuple (v_1, \dots, v_n, w) if $\vec{v} = (v_1, \dots, v_n)$. Set $S_2 = \bigcup_{n>0} M_n$, and define $\gamma = \bigcup_{n>0} \gamma_n$. Define the relation $R \subseteq S_2 \times S_2$ to be

$$\{(\vec{v}, \vec{v}w) \mid \vec{v} \in S_2, w \in X_{\gamma(\vec{v})}\}$$

Note that there is, for every $\vec{v} \in S_2$, a bijection $i_{\vec{v}} : X_{\gamma(\vec{v})} \rightarrow R(\vec{v})$ given by $w \mapsto \vec{v}w$. Note also that, for each $\vec{v} \in S_2$, we have

$$\gamma \circ \iota_{R(\vec{v}), S_2} \circ i_{\vec{v}} = \xi_{\gamma(\vec{v})}$$

With this in mind, we define the coalgebra structure σ_2 by setting

$$\sigma_2(\vec{v}) = \mathsf{T}(\iota_{R(\vec{v}), S_2} \circ i_{\vec{v}})(\alpha_{\gamma(\vec{v})})$$

Note that $\sigma_2(\vec{v})|_{R(\vec{v})} = \mathsf{T}i_{\vec{v}}(\alpha_{\gamma(\vec{v})})$. Finally, set s_2 to be the unique singleton tuple with sole element s_1 , and define the valuation V_2 by setting $V_2^\dagger(\vec{v}) = V_1^\dagger(\gamma(\vec{v}))$.

Clearly, $(S_2, R, \sigma_2, s_2, V_2)$ is a T-tree model. Denote the underlying T-model by \mathbb{S}_2 . We can then prove the following two claims:

Claim 5.8. The map γ is a T-model morphism from \mathbb{S}_2 to \mathbb{S}_1 .

Proof of Claim 5.8. The map γ clearly respects the truth values of all propositional atoms, and $\gamma(s_2) = s_1$. It suffices to show that γ is a coalgebra morphism, i.e. that $\mathsf{T}\gamma(\sigma_2(\vec{v})) = \sigma_1(\gamma(\vec{v}))$ for all \vec{v} . Pick any $\vec{v} \in S_2$. We have:

$$\begin{aligned} \mathsf{T}\gamma(\sigma_2(\vec{v})) &= \mathsf{T}\gamma \circ \mathsf{T}\iota_{R(\vec{v}), S_2} \circ \mathsf{T}(i_{\vec{v}})(\alpha_{\gamma(\vec{v})}) \\ &= \mathsf{T}(\gamma \circ \iota_{R(\vec{v}), S_2} \circ i_{\vec{v}})(\alpha_{\gamma(\vec{v})}) \\ &= \mathsf{T}(\xi_{\gamma(\vec{v})})(\alpha_{\gamma(\vec{v})}) \\ &= \sigma_1(\gamma(\vec{v})) \end{aligned}$$

as required. □

Claim 5.9. \mathbb{A} accepts $(S_2, R, \sigma_2, s_2, V_2)$ iff \mathbb{A}^* accepts \mathbb{S}_1 .

Proof of Claim 5.9: We have two parts that need to be proved here:

First part: \mathbb{A} accepts $(S_2, R, \sigma_2, s_2, V_2)$ implies \mathbb{A}^* accepts (\mathbb{S}_1, s_1) . Suppose that χ is a (positional) winning strategy in the acceptance game for \mathbb{A} and $(S_2, R, \sigma_2, s_2, V_2)$. We are going to define a strategy χ^* for \exists in the acceptance game for \mathbb{A}^* and \mathbb{S}_1 , with the property that for any χ^* -guided partial match

$$\rho = (a_1, s_1^1), \dots, (a_n, s_n^1)$$

of length n with $s_i^i \in S_1$ and $s_1^1 = s_1$, there exists a χ -guided partial match

$$\rho_* = (a_1, s_1^2), \dots, (a_n, s_n^2)$$

with $s_i^2 \in S_2$, $s_1^2 = s_2$ and $\gamma(s_i^2) = s_i^1$ for each index i , and chosen in such a way that if a χ^* -guided match ρ' is an extension of ρ , then the χ -guided match ρ'_* is an extension of the χ -guided match ρ_* .

We define the strategy χ^* by induction on the length of a partial match. For the partial match ρ consisting of the single position (a_I, s_1^1) we let ρ_* consist of the single position (a_I, s_2) , and we define the valuation $\chi^*(\rho) : A \rightarrow \mathcal{Q}(S_1)$ by setting, for each $b \in A$,

$$\chi^*(\rho)(b) = \gamma[\chi(a_I, s_2)(b)]$$

Similarly, suppose that χ^* has been defined for all matches of length less than n , and let ρ be any match of length $n > 1$. If ρ is not χ^* -guided, then we can define $\chi^*(\rho)(b) = \emptyset$ for all $b \in A$. If ρ is χ^* -guided, then write

$$\rho = (a_1, s_1^1), \dots, (a_n, s_n^1)$$

Since this match is χ^* -guided, by the induction hypothesis there is a χ -guided partial match

$$\rho_* = (a_1, s_1^2), \dots, (a_n, s_n^2)$$

with $\gamma(s_i^2) = s_i^1$ for all i . We set, for each $b \in A$,

$$\chi^*(\rho)(b) = \gamma[\chi(a_n, s_n^2)(b)]$$

Clearly, with this definition, the induction hypothesis continues to hold for all χ^* -guided matches of length $n + 1$. Furthermore, we note that, whenever ρ is a χ^* -guided match of length n , the move $\chi^*(\rho)$ is legal: since the move $\chi(a_n, s_n^2)$ must be legal, we have

$$(R(s_n^2), \sigma_2(s_n^2)|_{R(s_n^2)}, \chi(a_n, s_n^2)) \Vdash_1 \Delta(a_n, V_2^\dagger(s_n^2))$$

But since

$$\sigma_2(s_n^2)|_{R(s_n^2)} = \mathbb{T}i_{s_n^2}(\alpha_{\gamma(s_n^2)}) = \mathbb{T}i_{s_n^2}(\alpha_{s_n^1})$$

we see that $i_{s_n^2}$ is an isomorphism between the one-step models $(R[s_n^2], \sigma_2(s_n^2)|_{R(s_n^2)}, \chi(a_n, s_n^2))$ and $(X_{s_n^1}, \alpha_{s_n^1}, \chi(a_n, s_n^2)_{[i_{s_n^2}]})$, so we have

$$(X_{s_n^1}, \alpha_{s_n^1}, \chi(a_n, s_n^2)_{[i_{s_n^2}]}) \Vdash_1 \Delta(a_n, V_2^\dagger(s_n^2))$$

For all $b \in A$, we have

$$\begin{aligned} \chi(a_n, s_n^2)_{[i_{s_n^2}]}(b) &\subseteq \xi_{s_n^2}^{-1}(\xi_{s_n^2}[\chi(a_n, s_n^2)_{[i_{s_n^2}]}(b)]) \\ &= \xi_{s_n^2}^{-1}(\gamma[\chi(a_n, s_n^2)(b)]) \\ &= \xi_{s_n^2}^{-1}(\chi^*(\rho)(b)) \\ &= \chi^*(\rho)_{[\xi_{s_n^2}]}(b) \end{aligned}$$

So by monotonicity we get

$$(X_{s_n^1}, \alpha_{s_n^2}, \chi^*(\rho)_{[\xi_{s_n^2}]}) \Vdash_1 \Delta(a_n, V_2^\dagger(s_n^2))$$

Hence, since $\Delta^*(a_n, V_2^\dagger(s_n^2))$ is the uniform translation of $\Delta(a_n, V_2^\dagger(s_n^2))$, we get

$$(S_1, \sigma_1(s_n^2), \chi^*(\rho)) \Vdash_1 \Delta^*(a_n, V_2^\dagger(s_n^2))$$

as required. It is now easy to show that \exists never gets stuck in a χ^* -guided partial match, and we also see that every infinite χ^* -guided match

$$(a_I, s_1^1), (a_2, s_2^1), (a_3, s_3^1), \dots$$

corresponds to an infinite χ -guided match

$$(a_I, s_1^2), (a_2, s_2^2), (a_3, s_3^2), \dots$$

and so since \exists wins every infinite χ -guided match, she wins every infinite χ^* -guided match as well.

Second part: \mathbb{A}^* accepts (S_1, s_1) implies \mathbb{A} accepts $(S_2, R, \sigma_2, s_2, V_2)$. Let χ^* be a winning strategy for \exists in the acceptance game for \mathbb{A}^* and S_1 . We define the strategy χ as follows: given a partial match

$$\rho = (a_1, s_1^2), \dots, (a_n, s_n^2)$$

set $\chi(\rho) = \chi^*(\gamma(\rho))_{[\gamma]}$, where $\gamma(\rho)$ is the match

$$(a_1, \gamma(s_1^2)), \dots, (a_n, \gamma(s_n^2))$$

It is straightforward to check that, whenever ρ is a χ -guided match, $\gamma(\rho)$ is a χ^* -guided match, and that the move $\chi^*(\gamma(\rho))$ is legal if and only if the move $\chi(\rho)$ is legal. It follows that \exists never gets stuck in a χ -guided partial match, and that she wins every infinite χ -guided match, since an infinite χ -guided match

$$(a_I, s_1^2), (a_2, s_2^2), (a_3, s_3^2), \dots$$

corresponds to an infinite χ^* -guided match

$$(a_I, \gamma(s_1^2)), (a_2, \gamma(s_2^2)), (a_3, \gamma(s_3^2)), \dots$$

This concludes the proof. □

The lemma now follows by combining the two claims 5.8 and 5.9. □

From this, we get the following result. Here, we use the fact that every MSO_Λ -formula is equivalent to a *monotone* second-order automaton over \mathbb{T} -tree models, so that it suffices to find translations of monotone one-step formulas.

Theorem 5.10 (Auxiliary Characterization Theorem). *Let Λ be an expressively complete set of monotone predicate liftings for a set functor \mathbb{T} , and assume that the monotone fragment of the one-step language $\mathbf{1SO}_\Lambda$ admits uniform translations. Then:*

$$\text{MSO}_\Lambda / \sim \equiv \mu\text{ML}_\Lambda$$

Proof. Given a bisimulation invariant formula φ of MSO_Λ , let \mathbb{A} be an equivalent monotone second-order automaton, and suppose there is a uniform translation $(-)^*$ of the monotone fragment of the one-step language $\mathbf{1SO}_\Lambda$. Let ψ be a formula in μML_Λ equivalent to the

automaton \mathbb{A}^* . Then for any pointed \mathbb{T} -model \mathbb{S}, s , let (\mathbb{T}, R, t) be the \mathbb{T} -tree model provided by Proposition 5.7. Then we get:

$$\begin{aligned} \mathbb{S}, s \Vdash \varphi &\Leftrightarrow \mathbb{T}, t \Vdash \varphi \\ &\Leftrightarrow \mathbb{T}, R, t \Vdash \mathbb{A} \\ &\Leftrightarrow \mathbb{S}, s \Vdash \mathbb{A}^* \\ &\Leftrightarrow \mathbb{S}, s \Vdash \psi \end{aligned}$$

and the proof is done. \square

5.2. Adequate uniform constructions. The existence of uniform translations for the one-step language [35] involves two components: a translation on the syntactic side and a uniform construction on the semantic side. One of our main new observations in this paper is that we can actually forget about the syntactic translation and focus entirely on the model-theoretic side of the construction. If we can find a suitable uniform construction for the one-step models, satisfying a certain model-theoretic condition with respect to the second-order one-step language, the syntactic translation will come for free.

Definition 5.11. Let φ be any one-step formula in $1\mathbf{SO}_\Lambda(A)$. Then a uniform construction $(-)_*$ for \mathbb{T} is called *adequate* for φ if, for any pair of one-step frames (X, α) and (Y, β) , any one-step frame morphism $f : (X, \alpha) \rightarrow (Y, \beta)$ and any valuation $V : A \rightarrow \mathcal{Q}(Y)$, we have the following condition (\star) :

$$(X_*, \alpha_*, V_{[f \circ h_\alpha]}) \Vdash_1 \varphi \text{ iff } (Y_*, \beta_*, V_{[h_\beta]}) \Vdash_1 \varphi$$

where $h_\alpha : (X_*, \alpha_*) \rightarrow (X, \alpha)$ and $h_\beta : (Y_*, \beta_*) \rightarrow (Y, \beta)$ are the covering maps given by the uniform construction. We say that the construction is adequate for a set of formulas Γ if it is adequate for every member of Γ .

The following diagram clarifies condition (\star) :

$$\begin{array}{ccc} (X_*, \alpha_*, V_{[f \circ h_\alpha]}) & \xleftrightarrow{\varphi} & (Y_*, \beta_*, V_{[h_\beta]}) \\ h_\alpha \downarrow & & \downarrow h_\beta \\ (X, \alpha, V_{[f]}) & \xrightarrow{f} & (Y, \beta, V) \end{array}$$

We are now ready to prove our first main theorem, stating that if \mathbb{T} admits an adequate construction for any finite set of one-step formulas, then $\mathbf{MSO}_\mathbb{T}/\sim \equiv \mu\mathbf{ML}_\mathbb{T}$.

Proof of Theorem 1.4. By Theorem 5.10 it suffices to show that the monotone fragment of $1\mathbf{SO}_\mathbb{T}$ admits uniform translations. Let Γ be a finite set of monotone formulas of $1\mathbf{SO}_\mathbb{T}(A)$, for some finite set A , and suppose the uniform construction $(-)_*$ is adequate for Γ . Given $\varphi \in \Gamma$ we define a monotone predicate lifting λ over A by setting

$$\alpha \in \lambda_X(V) \text{ iff } (X_*, \alpha_*, V_{[h_\alpha]}) \Vdash_1 \varphi.$$

It is easy to check that this lifting is monotone, and naturality of λ is a direct consequence of the condition (\star) . If A has n elements and we list these as (a_1, \dots, a_n) , then we can view λ as a formula in $1\mathbf{ML}_\Lambda(A)$ of the form $\lambda'(a_1, \dots, a_n)$, where λ' is the n -place predicate lifting corresponding to λ with this given ordering of A . We now get a uniform translation by mapping φ to the formula $\lambda'(a_1, \dots, a_n)$. \square

Note that, if Λ is an expressively complete set of liftings for \mathbb{T} , then it suffices to find an adequate uniform construction for every finite set of formulas in $\mathbf{1SO}_\Lambda(A)$, since every formula in $\mathbf{1SO}_\mathbb{T}(A)$ is semantically equivalent to one in $\mathbf{1SO}_\Lambda(A)$. Hence, by Theorem 5.10 we get:

Corollary 5.12. *Let \mathbb{T} be any set functor, and let Λ be an expressively complete set of predicate liftings for \mathbb{T} . If \mathbb{T} admits an adequate uniform construction for every finite set of formulas of $\mathbf{1SO}_\Lambda$, then*

$$\mathbf{MSO}_\Lambda / \sim \equiv \mu\mathbf{ML}_\Lambda.$$

If \mathbb{T} preserves weak pullbacks then we can reformulate (\star) as the following condition:

Definition 5.13. Suppose \mathbb{T} preserves weak pullbacks. A *one-step $\bar{\mathbb{T}}$ -bisimulation* between one-step models (X, α, V) and (Y, β, U) is a relation $R \subseteq X \times Y$ such that:

- If xRy then $V^\dagger(x) = U^\dagger(y)$.
- $(\alpha, \beta) \in \bar{\mathbb{T}}R$.

Proposition 5.14. *Suppose that \mathbb{T} preserves weak pullbacks, and let $(-)_*$ be a uniform construction for \mathbb{T} . Then $(-)_*$ is adequate for φ if, and only if, for every pair of one-step $\bar{\mathbb{T}}$ -bisimilar one-step models (X, α, V) and (Y, β, U) we have:*

$$(X_*, \alpha_*, V_{[h_\alpha]}) \Vdash_1 \varphi \Leftrightarrow (Y_*, \beta_*, U_{[h_\beta]}) \Vdash_1 \varphi.$$

5.3. Some easy applications. We now start with some straightforward applications of the main characterization theorem.

Example 5.15. As a first application, the standard Janin-Walukiewicz characterization of the modal μ -calculus can be seen as an instance of the first main characterization result by taking $\Lambda = \{\diamond\}$ and $\mathbb{T} = \mathcal{P}$, recalling that $\mathbf{MSO} = \mathbf{MSO}_{\{\diamond\}}$. The uniform construction for \mathcal{P} , which is adequate with respect to any set of one-step formulas, is given as follows: consider a pair (X, α) with $\alpha \in \mathcal{P}(X)$. We take this to $X_* = \alpha_* = \alpha \times \omega$, and we let $h_\alpha : \alpha \times \omega \rightarrow X$ be the projection map. Proving adequacy of this construction is an exercise in model theory, and left to the reader. (The proof is a much simpler special case of the argument we will use in Section 6.2.)

Example 5.16. Consider the finitary multiset (“bags”) functor \mathcal{B} , which sends a set X to the set of mappings $f : X \rightarrow \omega$ such that the set $\{u \in X \mid f(u) = 0\}$ is cofinite. The action on morphisms is given by letting, for $f \in \mathcal{B}X$ and $h : X \rightarrow Y$, the multiset $\mathcal{B}h(f) : Y \rightarrow \omega$ be defined by $w \mapsto \sum_{h(v)=w} f(v)$. Given a pair X, α where $\alpha : X \rightarrow \omega$ has finite support, we define

$$X_* = \bigcup \{\{u\} \times \alpha(u) \mid u \in X\}.$$

Here, we follow the standard procedure and identify each finite ordinal $n \in \omega$ with the set of smaller ordinals, so that 0 is identified with \emptyset , 1 is $\{\emptyset\}$, 2 is $\{\emptyset, \{\emptyset\}\}$ and so on. The mapping $\alpha_* : X_* \rightarrow \omega$ is defined by setting $\alpha_*(w) = 1$ for all $w \in X_*$. The map $h_\alpha : X_* \rightarrow X$ is defined by $(u, i) \mapsto u$. It is easy to check that whenever two one-step frames (X, α) and (Y, β) are related by some morphism $f : (X, \alpha) \rightarrow (Y, \beta)$, the models $(X_*, \alpha_*, V_{[f \circ h_\alpha]})$ and $(Y_*, \beta_*, V_{[h_\beta]})$ are isomorphic, for any valuation $V : A \rightarrow \mathcal{Q}(Y)$. It follows that the construction is adequate for any set of one-step formulas, hence we get $\mu\mathbf{ML}_\mathcal{B} \equiv \mathbf{MSO}_\mathcal{B} / \sim$.

Example 5.17. Consider the set of all *exponential polynomial functors* [12] defined by the “grammar”:

$$\mathbb{T} ::= \mathbb{C} \mid \text{Id} \mid \mathbb{T} \times \mathbb{T} \mid \prod_{i \in I} \mathbb{T}_i \mid \mathbb{T}^{\mathbb{C}}$$

where \mathbb{C} is any constant functor for some set C , and Id is the identity functor on **Set**. These functors cover many important applications: streams, binary trees, deterministic finite automata and deterministic labelled transition systems are all examples of coalgebras for exponential polynomial functors, as is the so called *game functor* whose coalgebras provide semantics for Coalition Logic [5]. For this last instance, the “game functor” \mathcal{G} for n agents can be written in the form of an exponential polynomial functor as follows:

$$\prod_{\langle S_0, \dots, S_{n-1} \rangle \in (\mathcal{P}(\omega) \setminus \{\emptyset\})^n} \{ \langle S_0, \dots, S_{n-1} \rangle \} \times \text{Id}^{(S_0 \times \dots \times S_{n-1})}$$

Then, for a given set X , an element of $\mathcal{G}X$ will be a pair consisting of a vector $\langle S_0, \dots, S_{n-1} \rangle$ of available strategies for each player, together with an “outcome map” f assigning an element of X to each strategy profile in $S_0 \times \dots \times S_{n-1}$.

Every exponential polynomial functor admits adequate uniform constructions for all sets of one-step formulas. The proof proceeds by a straightforward induction, and similarly to the case for the bags functor it provides each exponential polynomial functor \mathbb{T} with a uniform construction such that for any one-step frame morphism $f : (X, \alpha) \rightarrow (Y, \beta)$ and any $V : A \rightarrow \mathcal{Q}(Y)$, the one-step models $(X_*, \alpha_*, V_{[f \circ h_\alpha]})$ and $(Y_*, \beta_*, V_{[h_\beta]})$ are isomorphic. Hence, we get:

Theorem 5.18. *For every exponential polynomial functor \mathbb{T} , we have $\mu\text{ML}_{\mathbb{T}} \equiv \text{MSO}_{\mathbb{T}}/\sim$.*

Proof. We provide a sketch of the inductive construction of an adequate uniform construction.

Constant functor. For the constant functor \mathbb{C} , a one-step frame is a pair (X, c) with $c \in C$. We set $X_* = \emptyset$, $c_* = c$ and $h_c : \emptyset \rightarrow X$ to be the unique inclusion of the empty set.

Identity functor. Given a one-step frame (X, u) for the identity functor, which consists of a set X and $u \in X$, we set $X_* = \{u\}$, $u_* = u$ and we set $h_u : \{u\} \rightarrow X$ to be the inclusion map sending u to itself.

Product. Suppose that \mathbb{T}_1 and \mathbb{T}_2 have associated adequate uniform constructions with the required property. Consider a one-step $\mathbb{T}_1 \times \mathbb{T}_2$ -frame $(X, (\alpha, \beta))$ with $\alpha \in \mathbb{T}_1 X$ and $\beta \in \mathbb{T}_2 X$. Let $h_1 : (X_1, \alpha_1) \rightarrow (X, \alpha)$ be the cover assigned by the uniform construction for \mathbb{T}_1 and let $h_2 : (X_2, \beta_2) \rightarrow (X, \beta)$ be the cover assigned by the uniform construction for \mathbb{T}_2 . Then we take X_* to be the disjoint union $X_1 + X_2$, and set

$$(\alpha, \beta)_* = (\mathbb{T}_1 i_1(\alpha_1), \mathbb{T}_2 i_2(\beta_2))$$

where $i_1 : X_1 \rightarrow X_1 + X_2$ and $i_2 : X_2 \rightarrow X_1 + X_2$ are the insertion maps for the co-product. Finally, we define the covering map $h_{(\alpha, \beta)} : X_1 + X_2 \rightarrow X$ be obtained by simply co-tupling

the maps h_1, h_2 , i.e. $h_{(\alpha, \beta)}$ is the map given by the universal property of the co-product applied to the diagram $X_1 \xrightarrow{h_1} X \xleftarrow{h_2} X_2$. We get:

$$\begin{aligned} (\mathbb{T}_1 \times \mathbb{T}_2)h_{(\alpha, \beta)}((\alpha, \beta)_*) &= (\mathbb{T}_1 h_{(\alpha, \beta)}(\mathbb{T}_1 i_1(\alpha_1)), \mathbb{T}_2 h_{(\alpha, \beta)}(\mathbb{T}_2 i_2(\beta_2))) \\ &= (\mathbb{T}_1(h_{(\alpha, \beta)} \circ i_1)(\alpha_1), \mathbb{T}_2(h_{(\alpha, \beta)} \circ i_2)(\beta_2)) \\ &= (\mathbb{T}_1 h_1(\alpha_1), \mathbb{T}_2 h_2(\beta_2)) \\ &= (\alpha, \beta) \end{aligned}$$

so $h_{(\alpha, \beta)}$ is indeed a covering map as required.

Exponentiation. The case of a functor \mathbb{T}^C for some constant C is handled analogously with the case of binary products, so we leave it to the reader.

Co-product. This step of the construction is actually the easiest one. Suppose that each functor \mathbb{T}_i for $i \in I$ is equipped with an adequation uniform construction, in the strong sense demanded by our induction hypothesis. Without loss of generality let's assume that $\mathbb{T}_i X \cap \mathbb{T}_j X = \emptyset$ for $i \neq j$. Let (X, α) be a one-step frame for the co-product $\coprod_{i \in I} \mathbb{T}_i$. Then there is a unique $i \in I$ with $\alpha \in \mathbb{T}_i X$, and so we define the cover $h_\alpha : (X_*, \alpha_*) \rightarrow (X, \alpha)$ merely by applying the uniform construction for \mathbb{T}_i . \square

Remark 5.19. These uniform constructions were all designed in a case-by-case fashion, and at the present time we do not know whether there is any general recipe for producing an adequate uniform construction when it exists. What the constructions mentioned so far seem to have in common is that we want to produce as many equivalent (in some sense) copies of each state in a one-step model, but this is not always sufficient. In the next section we will see a somewhat more involved construction for the monotone mu-calculus, which aims to create sufficiently many copies of each state but also, crucially, sufficiently many *pairwise disjoint* copies of all the neighborhoods. In this case we are trying to neutralize not only the capability of the second-order one-step language to count states in one-step models, but also its capability to express how certain neighborhoods are related to each other, in particular, whether they overlap or not. For example, the second-order one-step language can express the property that any two neighborhoods intersect, or there is a smallest neighborhood contained in all others etc., and the uniform construction we provide needs to trivialize all such statements. In general, applying our main result as it stands may require a bit of creativity, and we regard it as an interesting (possibly quite hard) task for future research to come up with a result that makes the task entirely mechanical. We mention some related questions in our concluding remarks section.

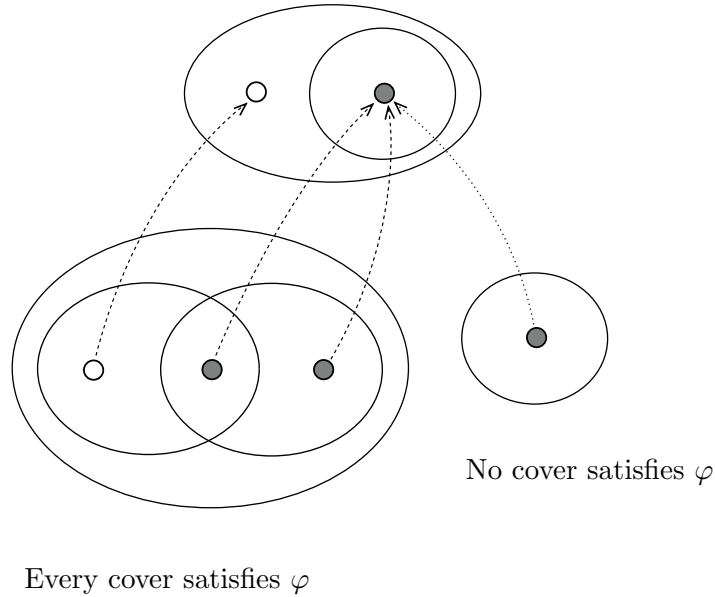
6. CHARACTERIZING THE MONOTONE μ -CALCULUS

This final section of the paper concerns the monotone neighborhood functor \mathcal{M} . There are two main difficulties here. First, the uniform construction will be less straightforward than in the previous applications, for the reasons that we explained above. Second, we need to do a bit of “pre-processing” of the monotone neighborhood functor to make the technique of uniform constructions applicable. In fact, the characterization theorem cannot possibly apply directly to the monotone μ -calculus: it turns out that there is *no* adequate uniform construction for \mathcal{M} . For these reasons, we view the results of this section as an

independent contribution of the paper, rather than just another application to illustrate the main theorem.

Proposition 6.1. *The functor \mathcal{M} does not admit an adequate uniform construction for the formula φ defined as: $\neg\text{Em}(a)$.*

Proof. The formula φ just says that the value of a is non-empty. Suppose there existed an adequate uniform construction for this formula. Consider the situation depicted in the diagram below, which shows three one-step frames together with two one-step frame morphisms, one for each of the two bottom one-step frames. The top frame consists of two points $\{w_1, w_2\}$ and has neighborhoods $\{\{w_1, w_2\}, \{w_2\}\}$, the bottom left frame has points $\{u_1, u_2, u_3\}$ and neighborhoods $\{\{u_1, u_2\}, \{u_2, u_3\}, \{u_1, u_2, u_3\}\}$. Finally the bottom right frame has a single point $\{v\}$ and this singleton as its only neighborhood. In other words, the picture shows a co-span in the category of one-step frames and one-step frame morphisms. Furthermore, consider the valuation on the topmost one-step frame which makes a true at exactly the one state w_1 , i.e. the one not belonging to the singleton neighborhood. This valuation is depicted in the diagram by representing the state where a is true by a blank circle, and the state where it is not true by a filled circle. The induced valuations on the bottom one-step frames via the frame morphisms are depicted in the same manner. With respect to these valuations, p will be true at u_1 , but false at u_2, u_3 and v .



Let us denote the top frame by (Y, β) , its valuation by V , the bottom-left frame as (X, α) , the bottom-right one as (X', α') and the corresponding frame morphisms as f and f' respectively. Now, the supposed adequate construction will assign a cover to each of the three frames, and we get valuations for each cover from the valuations depicted in the diagram. It follows from the defining condition (\star) for adequacy that φ must have the same truth value in each of these covers. But this leads to a contradiction: it is not hard to check that the cover

$h_\alpha : (X_*, \alpha_*) \rightarrow (X, \alpha)$ must be such that $(X_*, \alpha_*, V_{[f \circ h_\alpha]}) \Vdash \varphi$, because the image of the map h_α must support α and therefore include the one element of X coloured red. On the other hand, we clearly must have $(X'_*, \alpha'_*, V_{[f' \circ h_{\alpha'}]}) \not\Vdash \varphi$, which contradicts the condition (\star) . \square

6.1. The supported companion of a functor. Our solution to the problem we are facing with the monotone neighborhood functor is to perform a gentle repair, changing the functor into a closely related one that does admit an adequate uniform construction. The construction has some independent interest from a general coalgebraic perspective, so we define it in general terms:

Definition 6.2. The *supported companion* $\underline{\mathbb{T}}$ of a set functor \mathbb{T} is the sub-functor of $\mathcal{P} \times \mathbb{T}$ defined by:

$$\underline{\mathbb{T}}(X) = \{(Z, \alpha) \mid \alpha \in \mathbb{T}X \ \& \ Z \text{ supports } \alpha\}$$

It is easy to check that this is indeed a well-defined subfunctor of $\mathcal{P} \times \mathbb{T}$, since for any map $f : X \rightarrow Y$, any $\alpha \in \mathbb{T}X$ and any set $Z \subseteq X$, the image $f[Z] = \mathcal{P}f(Z)$ is a support for $\mathbb{T}f(\alpha)$ whenever Z is a support for α . The reader may also note that what we have called “ \mathbb{T} -tree models” are actually special instances of $\underline{\mathbb{T}}$ -models, so in a sense the supported companion functor has already played a role in the proof of the main theorem. We will show that this construction repairs the monotone neighborhood functor, so that the supported companion $\underline{\mathcal{M}}$ of \mathcal{M} admits an adequate uniform construction. Interestingly enough, the same construction happens to repair weak pullback preservation:

Proposition 6.3. *The functor $\underline{\mathcal{M}}$ preserves weak pullbacks.*

We leave the verification of this to the reader; the argument is similar to the reasoning in [20] used to establish the existence of a well-behaved relation lifting for \mathcal{M} .

Note that we have an obvious natural transformation $\pi : \underline{\mathbb{T}} \rightarrow \mathcal{P}$ with components given by the projection maps sending $(Z, \alpha) \in \underline{\mathbb{T}}X$ to $Z \in \mathcal{P}(X)$, and we can compose this with the predicate lifting $\square : \mathcal{Q} \rightarrow \mathcal{Q} \circ \mathcal{P}$ to obtain a predicate lifting $\underline{\square}$ for $\underline{\mathbb{T}}$ defined by:

$$\underline{\square} := \mathcal{Q}\pi \circ \square$$

We call this lifting the *support modality* for $\underline{\mathbb{T}}$, and write its dual as $\underline{\diamond}$. Furthermore, any lifting λ for \mathbb{T} corresponds to a lifting for $\underline{\mathbb{T}}$ by composing λ with the projection from $\underline{\mathbb{T}}$ to \mathbb{T} in the same manner, and we will not take care to distinguish λ as a lifting for \mathbb{T} from the corresponding lifting for $\underline{\mathbb{T}}$. A little issue that we need to address, before we can proceed to characterize the monotone μ -calculus, is just how the language μML_Λ is related to $\mu\text{ML}_{\Lambda \cup \{\underline{\square}\}}$ for a given set of liftings Λ for \mathbb{T} . The rest of this subsection will provide the answer, and give a characterization theorem for μML_Λ as the \mathbb{T} -bisimulation invariant fragment of $\mu\text{ML}_{\Lambda \cup \{\underline{\square}\}}$, where a formula is said to be \mathbb{T} -bisimulation invariant if it has the same truth value in any two pointed $\underline{\mathbb{T}}$ -models for which the corresponding \mathbb{T} -models are behaviourally equivalent. Formally we shall write $(\mathbb{S}, s) \sim (\mathbb{S}', s')$, for $\underline{\mathbb{T}}$ -models \mathbb{S} and \mathbb{S}' , to say that the respective underlying pointed \mathbb{T} -models are behaviourally equivalent. To distinguish this from actual behavioural equivalence in the sense of the companion functor $\underline{\mathbb{T}}$, we write $(\mathbb{S}, s) \simeq (\mathbb{S}', s')$ to say that these pointed models are behaviourally equivalent as $\underline{\mathbb{T}}$ -models.

We borrow a result from [8]:

Fact 6.4 ([8]). For any set of liftings Λ for any set functor \mathbb{T} , the logic μML_Λ has the finite model property.

We assume that we have at our disposal a fixed \mathbb{T} -model $\mathbb{U} = (U, \gamma, V)$ which is a disjoint union of one isomorphic copy for every *finite* pointed \mathbb{T} -model (\mathbb{S}, s) . It is not hard to see that such a model does exist: just take a disjoint union containing each \mathbb{T} -model defined on some finite subset $n = \{0, \dots, n-1\}$ of ω . Since the class of all these models clearly forms a *set*, their disjoint union is a well defined \mathbb{T} -model.

Definition 6.5. Let \mathbb{S} be any finite \mathbb{T} -model. We define the model $\mathbb{S} \oplus \mathbb{U} = (S + U, \sigma_+, V_+)$ such that, for each $s \in S$ we have

$$\sigma_+(s) = (S + U, \mathbb{T}(\iota_{S, S+U})(\alpha))$$

where $\sigma(s) = (X, \alpha)$, and for $u \in U$ we have $\sigma_+(u) = \mathbb{T}\iota_{U, S+U}(\gamma(u))$ ⁵.

Note that the underlying \mathbb{T} -model of $\mathbb{S} \oplus \mathbb{U}$ is just the co-product of the underlying \mathbb{T} -models of \mathbb{S} and \mathbb{U} respectively.

The following lemmas are immediate from the definition.

Lemma 6.6. *Let \mathbb{S} be any \mathbb{T} -model. Then for all $u \in U$, we have:*

$$(\mathbb{U}, u) \simeq (\mathbb{S} \oplus \mathbb{U}, u)$$

Lemma 6.7. *For any \mathbb{T} -model \mathbb{S} and $s \in S$: $(\mathbb{S}, s) \sim (\mathbb{S} \oplus \mathbb{U}, s)$.*

Note that Lemma 6.7 is *not* guaranteed to hold if we replace \sim by the finer equivalence relation \simeq .

It will be convenient in this section to work with a second version of the acceptance game for a modal Λ -automaton or a modal $\Lambda \cup \{\square\}$ -automaton, which we will call the *extended* acceptance game for \mathbb{A} with respect to a model \mathbb{S} , denoted $\mathcal{E}(\mathbb{A}, \mathbb{S})$. Given a modal Λ -automaton $\mathbb{A} = (A, \Delta, a_I, \Omega)$ and a model $\mathbb{S} = (S, \sigma, V)$ this game has three types of positions: pairs of the form (ψ, s) with $\psi \in \mathbf{1ML}_\Lambda(A)$, pairs of the form (ψ, s) with $\psi \in \mathbf{Latt}(A)$ (lattice formulas over A , as defined in Definition 3.11), and maps $f : \mathbf{Latt}(A) \rightarrow \mathcal{P}(S)$. Admissible moves are given in Table 4:

Position	Player	Admissible moves	Priority
$(\psi_1 \vee \psi_2, s)$	\exists	$\{(\psi_1, s), (\psi_2, s)\}$	0
$(\psi_1 \wedge \psi_2, s)$	\forall	$\{(\psi_1, s), (\psi_2, s)\}$	0
$(a, s) \in A \times S$	\exists	$\{(\Delta(a, V^\dagger(s)), s)\}$	$\Omega(a)$
$(\lambda(\psi_1, \dots, \psi_n), s)$	\exists	$\{f : \mathbf{Latt}(A) \rightarrow \mathcal{P}S \mid \sigma(s) \in \lambda(f(\psi_1), \dots, f(\psi_n))\}$	0
(\top, s)	\forall	\emptyset	0
(\perp, s)	\exists	\emptyset	0
$f : \mathbf{Latt}(A) \rightarrow \mathcal{P}(S)$	\forall	$\{(\psi, s) \mid s \in f(\psi)\}$	0

Table 4: The extended acceptance game

Here, $\lambda \in \Lambda \cup \{\square\}$. It is easy to see that the position (\top, s) is always winning for \exists , and (\perp, s) is always winning for \forall . The following result is entirely routine to prove:

⁵In this definition we have used “+” as the symbol for binary coproducts of sets.

Theorem 6.8. *Let $\mathbb{A} = (A, \Delta, a_I, \Omega)$ be a modal Λ -automaton and (\mathbb{S}, s) a pointed \mathbb{T} -model. Then \mathbb{A} accepts (\mathbb{S}, s) if, and only if, (a_I, s) is a winning position in the extended acceptance game.*

We also have the following observation:

Proposition 6.9. *Let ψ be any one-step formula in $\mathbf{1ML}_\Lambda(A)$ for a set of liftings Λ for some given set functor \mathbb{T} , and let $\mathbb{P} \in \{\forall, \exists\}$. If $(\mathbb{S}, s) \simeq (\mathbb{S}', s')$ then the position (ψ, s) is winning for \mathbb{P} in $\mathcal{E}(\mathbb{A}, \mathbb{S})$ if and only if (ψ, s') is a winning position for \mathbb{P} in $\mathcal{E}(\mathbb{A}, \mathbb{S}')$.*

This means that for a pair of \mathbb{T} -models \mathbb{S} and \mathbb{S}' such that $(\mathbb{S}, s) \simeq (\mathbb{S}', s')$, and a formula $\psi \in \mathbf{ML}_{\Lambda \cup \{\Box\}}^1(A)$, the position (ψ, s) is winning for \mathbb{P} in $\mathcal{E}(\mathbb{A}, \mathbb{S})$ if and only if (ψ, s') is a winning position for \mathbb{P} in $\mathcal{E}(\mathbb{A}, \mathbb{S}')$.

Definition 6.10. Let $\mathbb{A} = (A, \Delta, a_I, \Omega)$ be a $\Lambda \cup \{\Box\}$ -automaton and let ψ be a formula in the range of Δ . We say that ψ is \mathbb{A} -valid if for all (finite) pointed \mathbb{T} -models (\mathbb{S}, s) , the position (ψ, s) is winning in $\mathcal{E}(\mathbb{A}, \mathbb{S})$. We say that ψ is \mathbb{A} -satisfiable if there is *some* (finite) pointed \mathbb{T} -model (\mathbb{S}, s) , the position (ψ, s) is winning in $\mathcal{E}(\mathbb{A}, \mathbb{S})$.

Let Λ be any set of predicate liftings for \mathbb{T} and let \mathbb{A} be any modal $\Lambda \cup \{\Box\}$ -automaton. Making use of our fixed \mathbb{T} -model \mathbb{U} we shall define a translation $t_{\mathbb{A}} : \mathbf{1ML}_{\Lambda \cup \{\Box\}}(A) \rightarrow \mathbf{1ML}_\Lambda(A)$ by induction as follows:

- For $\lambda \in \Lambda$, we set $t_{\mathbb{A}}(\lambda(\psi_1, \dots, \psi_k)) = \lambda(\psi_1, \dots, \psi_k)$, and similarly for the dual λ^d for every lifting $\lambda \in \Lambda$.
- $t_{\mathbb{A}}(\top) = \top$ and $t_{\mathbb{A}}(\perp) = \perp$.
- $t_{\mathbb{A}}(\psi_1 \vee \psi_2) = t_{\mathbb{A}}(\psi_1) \vee t_{\mathbb{A}}(\psi_2)$ and $t_{\mathbb{A}}(\psi_1 \wedge \psi_2) = t_{\mathbb{A}}(\psi_1) \wedge t_{\mathbb{A}}(\psi_2)$.
- $t_{\mathbb{A}}(\Box\psi) = \begin{cases} \top & \text{if } \psi \text{ is } \mathbb{A}\text{-valid} \\ \perp & \text{otherwise} \end{cases}$.
- $t_{\mathbb{A}}(\Diamond\psi) = \begin{cases} \top & \text{if } \psi \text{ is } \mathbb{A}\text{-satisfiable} \\ \perp & \text{otherwise} \end{cases}$.

Note that this translation depends on the whole automaton \mathbb{A} , not just the set of variables A . So for any given set of variables A , we have one translation $t_{\mathbb{A}} : \mathbf{1ML}_{\Lambda \cup \{\Box\}}(A) \rightarrow \mathbf{1ML}_\Lambda(A)$ for each automaton \mathbb{A} with states A , and the translations will generally be different for different choices of \mathbb{A} . However, we will drop the index \mathbb{A} from now on to simplify notation. Given a modal $\Lambda \cup \{\Box\}$ -automaton $\mathbb{A} = (A, \Delta, a_I, \Omega)$, we will write $t(\mathbb{A})$ for the automaton $(A, \Delta^t, a_I, \Omega)$ where Δ^t is defined by $\Delta^t(a, c) = t(\Delta(a, c))$, where it is understood that t denotes the translation $t_{\mathbb{A}}$. Clearly $t(\mathbb{A})$ is a modal Λ -automaton.

We shall view the translation t as a map defined on the entire domain $\mathbf{1ML}_{\Lambda \cup \{\Box\}}(A) \cup \mathbf{Latt}(A)$ by setting $t(\psi) = \psi$ for $\psi \in \mathbf{Latt}(A)$.

Lemma 6.11. *For every finite pointed model (\mathbb{S}, s) , and for any modal $\Lambda \cup \{\Box\}$ -automaton \mathbb{A} , we have*

$$(\mathbb{S}, s) \Vdash t(\mathbb{A}) \text{ iff } (\mathbb{S} \oplus \mathbb{U}, s) \Vdash \mathbb{A}$$

Proof. For left to right, suppose that $(\mathbb{S}, s) \Vdash t(\mathbb{A})$. By Theorem 6.8 there is a strategy χ for \exists in the extended acceptance game for $t(\mathbb{A})$ and \mathbb{S} which is winning at (a_I, s) . Without loss of generality we may assume that χ is positional, and a winning strategy at *every* winning position in $\mathcal{E}(t(\mathbb{A}), \mathbb{S})$. Our goal is to construct a positional strategy χ' for \exists in the extended acceptance game for \mathbb{A} and $\mathbb{S} \oplus \mathbb{U}$, which prescribes a move for \exists at every position (ψ, v) belonging to \exists and with $v \in S$, such that:

- (1) χ' assigns a legitimate move to every position belonging to \exists of the form (ψ, v) with $v \in S$ such that $(t(\psi), v)$ is a winning position in $\mathcal{E}(t(\mathbb{A}), \mathbb{S})$,
- (2) every χ' -guided partial match ρ starting at (a_I, s) and ending with a position (ψ, v) satisfies one of the following two criteria:
 - a:** $v \in S$ and $t[\rho]$ is a χ -guided match in $\mathcal{E}(t(\mathbb{A}), \mathbb{S})$ (hence consists only of winning positions for \exists).
 - b:** (ψ, v) is a winning position in $\mathcal{E}(\mathbb{A}, \mathbb{S} \oplus \mathbb{U})$.

Here, given that $\rho = \pi_1 \dots \pi_n$, we define $t[\rho] = t(\pi_1) \dots t(\pi_n)$ where $t(\psi, v) = (t(\psi), v)$ if $\psi \in \mathbf{1ML}_{\mathbb{A}}(A) \cup \mathbf{Latt}(A)$, and $t(f) = f$ for a position $f : \mathbf{Latt}(A) \rightarrow \mathcal{P}(S)$. Clearly, we can build a winning strategy in $\mathcal{E}(\mathbb{A}, \mathbb{S} \oplus \mathbb{U})$ from such a strategy χ' .

We define the strategy χ' by a case distinction, given a position (ψ, v) belonging to \exists with $v \in S$. If $\psi = \alpha_1 \vee \alpha_2$ then $t(\psi) = t(\alpha_1) \vee t(\alpha_2)$, so we set $\chi'(\psi, v) = (\alpha_i, v)$ where i is chosen so that $\chi(t(\psi), v) = (t(\alpha_i), v)$. If $\psi = \lambda(\varphi_1, \dots, \varphi_n)$ then $t(\psi) = \lambda(\varphi_1, \dots, \varphi_n)$ too, so we set $\chi'(\psi, v) = \chi(t(\psi), v)$. A simple naturality argument shows that this is still a legitimate move in $\mathcal{E}(\mathbb{A}, \mathbb{S} \oplus \mathbb{U})$.

Finally, the interesting case is the one involving the support modality: at a position $(\diamond\varphi, v)$, if $(t(\diamond\varphi), v)$ is winning for \exists in $\mathcal{E}(\mathbb{A}, \mathbb{S})$ then we must have $t(\diamond\varphi) = \top$, hence φ is \mathbb{A} -satisfiable. Hence, by the construction of \mathbb{U} , and by Observation 6.9 there is some $u \in U$ such that (φ, u) is winning for \exists in $\mathcal{E}(\mathbb{A}, \mathbb{S} \oplus \mathbb{U})$. So we let the strategy χ' pick the position (φ, u) . For the case involving the dual \square , if $t(\square\varphi) = \top$ then φ is valid, and we can let \exists pick the map sending φ to $S + U$.

We also need to check that every χ' -guided partial match satisfies one of the conditions (a) or (b), and we prove this by an induction on the length of a partial match. The only interesting case is for the extension of a partial match ρ ending with a position $(\square\varphi, v)$. By the induction hypothesis on ρ , we have $v \in S$ and $(t(\square\varphi), v)$ is winning for \exists , hence we must have $t(\square\varphi) = \top$. This means that φ is \mathbb{A} -valid, so *any* move (φ, w) by \forall answering the move $S + U$ by \exists will satisfy the condition (b), i.e. (φ, w) is winning for \exists in $\mathcal{E}(\mathbb{A}, \mathbb{S} \oplus \mathbb{U})$.

For right to left, suppose that \exists has a winning strategy χ at the position (a_I, s) in the extended acceptance game for \mathbb{A} with respect to $\mathbb{S} \oplus \mathbb{U}$. We shall give \exists a winning strategy χ' at the same position in the game for $t(\mathbb{A})$ with respect to \mathbb{S} . We shall inductively associate with every χ' -guided partial match π of length k a χ -guided “shadow match” $(\psi_1, v_1), \dots, (\psi_k, v_k)$ such that π is of the form

$$(t(\psi_1), v_1), \dots, (t(\psi_k), v_k)).$$

We shall also make sure that whenever π' is an extension of π , the shadow match associated with π' is an extension of the shadow match associated with π as well. It will clearly follow that \exists wins every infinite χ' -guided match.

For the singleton match consisting of (a_I, s) we let (a_I, s) itself be the shadow match. (This is acceptable because, by convention, we have set $t(a_I) = a_I$.) For a match π of length k we define the move $\chi'(\pi)$ depending on the shape of the last position on the associated shadow match. Again we treat only the interesting cases.

If the last position on the shadow match is $(\lambda(\psi_1, \dots, \psi_m), v)$ then χ provides a map $f : \mathbf{Latt}(A) \rightarrow \mathcal{Q}(S + U)$ such that

$$\sigma_+(v) \in \lambda_{S+U}(f(\psi_1), \dots, f(\psi_m)).$$

We set $\chi'(\pi) = f'$, where $f' : \mathbf{Latt}(A) \rightarrow \mathcal{Q}(S)$ is defined by $\theta \mapsto f(\theta) \cap S$ for each $\theta \in \mathbf{Latt}(A)$. This move is legal since $v \in S$ and by naturality of λ . It is easy to see how to extend the shadow match for each response by \forall .

If the last position on the shadow match is $(\Box\psi, v)$ then since this position is winning for \exists , it must be the case that every position (ψ, w) for $w \in S + U$ is winning for \exists , hence this holds for every $w \in U$. This can only be true if ψ is \mathbb{A} -valid and so we have $t(\Box\psi) = \top$. This means we are done since (\top, v) is a winning position for \exists . Similarly, if the last position on the shadow match is $(\Diamond\psi, v)$, then there is some $w \in S + U$ such that (ψ, w) is winning for \exists , so ψ is \mathbb{A} -satisfiable. Hence $t(\Diamond\psi) = \top$, and the conclusion follows as in the previous case. \square

We can now prove our characterization theorem:

Theorem 6.12. *Let \top be any set functor and let Λ be any set of predicate liftings for \top . Then:*

$$\mu\mathbf{ML}_{\Lambda \cup \{\Box\}} / \sim \equiv \mu\mathbf{ML}_{\Lambda}.$$

Proof. Suppose a formula φ of $\mu\mathbf{ML}_{\Lambda \cup \{\Box\}}$ is invariant for \top -bisimulation. By the finite model property for $\mu\mathbf{ML}_{\Lambda \cup \{\Box\}}$ it suffices to show that φ is equivalent to a $\mu\mathbf{ML}_{\Lambda}$ -formula over finite models. Let \mathbb{A} be a modal $\Lambda \cup \{\Box\}$ -automaton equivalent to φ , and let ψ be a formula of $\mu\mathbf{ML}_{\Lambda}$ equivalent to the automaton $t(\mathbb{A})$. Consider an arbitrary finite pointed \top -model (\mathbb{S}, s) . We have:

$$\begin{aligned} (\mathbb{S}, s) \Vdash \varphi &\Leftrightarrow (\mathbb{S} \oplus \mathbb{U}, s) \Vdash \varphi && \text{(Lemma 6.7 + assumption on } \varphi\text{)} \\ &\Leftrightarrow (\mathbb{S} \oplus \mathbb{U}, s) \Vdash \mathbb{A} \\ &\Leftrightarrow (\mathbb{S}, s) \Vdash t(\mathbb{A}) && \text{(Lemma 6.11)} \\ &\Leftrightarrow (\mathbb{S}, s) \Vdash \psi \end{aligned}$$

as required. \square

6.2. An adequate uniform construction for $\underline{\mathcal{M}}$. In this section we shall apply Corollary 5.12 to derive a characterization result for the monotone μ -calculus. The first fact that we need is the following:

Proposition 6.13. *The set of liftings $\{\Box, \square\}$ is expressively complete with respect to $\underline{\mathcal{M}}$.*

Proof. Since $\underline{\mathcal{M}}$ preserves weak pullbacks, it suffices by Proposition 3.18 to show that the one-step language $\mathbf{1ML}_{\{\Box, \square\}}$ has the same expressive power as the language based on the Moss liftings for $\underline{\mathcal{M}}$. Given a finite set A , a subset $\Gamma \subseteq \mathbf{Latt}(A)$ and some $N \in \mathcal{MLatt}(A)$ that is supported by Γ , we want to define a formula γ in $\mathbf{1ML}_{\{\Box, \square\}}$ such that

$$X, (X', \alpha), V \Vdash_1 \gamma \text{ iff } ((X', \alpha), (\Gamma, N)) \in \overline{(\underline{\mathcal{M}})}(\Vdash_V^0)$$

This formula can be obtained as a conjunction $\theta_{\Gamma} \wedge \theta_N$, where θ_{Γ} is like the standard nabla-formula for the powerset functor corresponding to Γ , but expressed in terms of the support modality, and θ_N is defined as the nabla formulas for the monotone neighborhood functor introduced by Santocanale and Venema in [30]. We omit the details. \square

The remainder of this section is devoted to defining a uniform construction for the supported companion $\underline{\mathcal{M}}$ of \mathcal{M} w.r.t. a given finite set of formulas, and prove its adequacy. We proceed as follows: throughout the section we fix an arbitrary natural number k and a finite set of variables A , and define a uniform construction $(-)_*$ which will be shown to be adequate for every formula in $1\text{SO}_{\{\square, \square\}}(A)$ of quantifier depth $\leq k$. It clearly follows that there is an adequate uniform construction for every finite set Γ of formulas in $1\text{SO}_{\{\square, \square\}}$, since we can apply the result to the maximum quantifier depth of any formula in Γ .

To simplify notation a bit, we note that it suffices to define our uniform construction for what we may call *normal one-step $\underline{\mathcal{M}}$ -frames* $(X, (S, \alpha))$, satisfying the condition that $S = X$, and prove that the adequacy condition holds for all models based on such frames. It is then a straightforward matter to extend this to an adequate uniform construction defined on all $\underline{\mathcal{M}}$ -frames. We shall denote a normal one-step frame $(X, (X, \alpha))$ more simply by just (X, α) . Then for any two normal $\underline{\mathcal{M}}$ -models (X, α, V) and (X', α', V') over A , i.e. models that are based on normal one-step frames, a relation $R \subseteq X \times X'$ is a one-step $(\underline{\mathcal{M}})$ -bisimulation between these models if, and only if:

- uRu' implies $u \in V(a)$ iff $u' \in V'(a)$, for all $a \in A$;
- for all Z in α there is Z' in α' such that for all $t' \in Z'$ there is $t \in Z$ with tRt' ;
- for all Z' in α' there is Z in α such that for all $t \in Z$ there is $t' \in Z'$ with tRt' ;
- for all $u \in X$ there is $u' \in X'$ with uRu' ;
- for all $u' \in X'$ there is $u \in X$ with uRu' .

The last two conditions just say that the relation R is full on X and X' . We call such a relation a *global neighborhood bisimulation* between (X, α, V) and (X', α', V') . For any normal one-step frame (X, α) , we shall construct a cover $h_\alpha : (X_*, \alpha_*) \rightarrow (X, \alpha)$ where (X_*, α_*) is also a normal one-step frame. Since these are abbreviated notations for the one-step $\underline{\mathcal{M}}$ -frames $(X, (X, \alpha))$ and $(X_*, (X_*, \alpha_*))$, we should have $\underline{\mathcal{M}}h_\alpha(X_*, \alpha_*) = (X, \alpha)$, which means that h_α should be a surjective map from X_* to X , such that $\mathcal{M}h_\alpha(\alpha_*) = \alpha$. The intuition behind the construction is that we want to create infinitely many *disjoint* copies of each neighborhood, and furthermore we want to create sufficiently many copies of each state *within* each copy of a neighborhood.

Definition 6.14. Given a set X , and an object $\alpha \in \mathcal{M}X$, put

$$X_* := X \times 2^k \times \mathcal{P}(X) \times \omega$$

and let π_X be the projection map from X_* to X . Define $\alpha_* \in \mathcal{M}(X_*)$ by setting $Z \in \alpha_*$ for $Z \subseteq X_*$ iff $[Y, j] \subseteq Z$ for some $Y \in \alpha$ and some $j < \omega$, where

$$[Y, j] := \{(u, i, Y, j) \mid u \in Y, i < 2^k\}.$$

The sets of the form $[Y, j]$ for $Y \in \alpha$ will be called the *basic members* of α_* . The set of all elements in X_* that do not belong to any basic member will be called the *residue* of the frame (X_*, α_*) . Note that X_* is partitioned into each of the basic members along with the residue, as an extra partition cell.

Proposition 6.15. *For every given normal one-step frame (X, α) , the projection map $\pi_X : (X_*, \alpha_*) \rightarrow (X, \alpha)$ is a cover.*

Proof. Clearly the projection map π_X is always surjective since $(u, 0, \emptyset, 0) \mapsto u$ for all $u \in X$. (Note that $0 < 2^n$ for all $n \in \omega$, so we always have $(u, 0, \emptyset, 0) \in X_*$.)

We need to check that $\mathcal{M}\pi_X(\alpha_*) = \alpha$. In other words, we have to check that for all $Z \subseteq X$, we have $Z \in \alpha$ iff $\pi_X^{-1}(Z) \in \alpha_*$. For left to right, if $Z \in \alpha$ then $[Z, 0]$ is

a basic member of α_* , and clearly $[Z, 0] \subseteq \pi_X^{-1}(Z)$. Conversely, suppose $\pi_X^{-1}(Z) \in \alpha_*$. Then there is some basic member $[Y, j] \in \alpha_*$ with $[Y, j] \subseteq \pi_X^{-1}(Z)$. But then $Y \in \alpha$, and furthermore $Y \subseteq Z$: if $u \in Y$ then $(u, 0, Y, j) \in [Y, j]$, so $(u, 0, Y, j) \in \pi_X^{-1}(Z)$ meaning that $\pi_X(u, 0, Y, j) = u \in Z$. So $Z \in \alpha$ as required. \square

Our goal is to show that, for any pair of normal and globally neighborhood bisimilar models (X, α, V) and (Y, β, U) , the corresponding models $(X_*, \alpha_*, V_{[\pi_X]})$ and $(Y_*, \beta_*, U_{[\pi_Y]})$ satisfy the same formulas in $1\text{SO}_{\{\square, \diamond\}}(A)$ of quantifier depth $\leq k$. It should not be too surprising that we can prove this, but the actual proof is not entirely trivial. It suffices to check it for formulas in which the operator \square and \diamond does not occur, since over normal one-step models this operator is easily definable using the second-order quantifiers. From now on we keep the models (X, α, V) and (Y, β, U) fixed, as well as a global neighborhood bisimulation R relating these models. Throughout this section, we use the notation $(X, \alpha, V) \equiv_k (Y, \beta, U)$ to say that two one-step models satisfy the same formulas of $1\text{SO}_{\{\square, \diamond\}}(A)$ with at most k nested quantifiers.

Definition 6.16. A propositional A -type τ is a subset of A . Given a set X and a valuation $V : A \rightarrow \mathcal{Q}(X)$, the propositional A -type of $v \in X$ is defined to be $V^\dagger(v) = \{a \in A \mid v \in V(a)\}$.

Given a subset Z of X_* or Y_* , a valuation W such that $W : B \rightarrow \mathcal{Q}(X_*)$ or $W : B \rightarrow \mathcal{Q}(Y_*)$, and a natural number m , the m -signature of Z over B relative to the valuation W is the mapping $\sigma_Z : \mathcal{P}(B) \rightarrow \{0, \dots, m\}$ defined by:

$$\sigma_Z(t) := \min(|\{x \in Z \mid W^\dagger(x) = t\}|, m)$$

Definition 6.17. Let B be any set of variables containing A , and let $V_1 : B \rightarrow \mathcal{Q}(X_*)$ and $V_2 : B \rightarrow \mathcal{Q}(Y_*)$. Then for any natural number n we write

$$(X_*, \alpha_*, V_1) \approx_n (Y_*, \beta_*, V_2)$$

and say that these one-step models *match up to depth n* , if:

- (1) For every n -signature σ over variables B , the number of basic elements of signature σ in α_* and β_* respectively are either both finite and equal, or both infinite.
- (2) The residues of the two one-step models have the same n -signature.

Using the assumption that the models (X, α, V) and (Y, β, U) are globally neighborhood bisimilar, we get the following lemma.

Lemma 6.18. $(X_*, \alpha_*, V_{[\pi_X]}) \approx_{2^k} (Y_*, \beta_*, U_{[\pi_Y]})$.

Proof. To see that the residues of the two models have the same 2^k -signature, for one direction just note that if the residue of $(X_*, \alpha_*, V_{[\pi_X]})$ contains an element (u, i, Z, j) then it contains infinitely many elements of the same propositional type, namely one member (u, i, Z, p) for every $p \in \omega$. But then so will the residue of $(Y_*, \beta_*, U_{[\pi_Y]})$: just pick some v with uRv (again using the fact that R is a global neighborhood bisimulation). Since $v \notin \emptyset$, for every $p \in \omega$ the element $(v, 0, \emptyset, p)$ will be a member of the residue of $(Y_*, \beta_*, U_{[\pi_Y]})$ of the same propositional type as (u, i, Z, j) .

Now for the basic elements. First note that, for any 2^k -signature σ , α_* either contains no basic elements of signature σ , or infinitely many: if there is some basic element $[Z, j]$ of signature σ , then for any $i \neq j$, the basic element $[Z, i]$ has the same 2^k -signature as $[Z, j]$ with respect to the valuation $V_{[\pi_X]}$. The same holds for β_* with respect to the valuation

$U_{[\pi_Y]}$. Hence, it suffices to show that α_* contains a basic element of signature σ w.r.t. $V_{[\pi_X]}$ iff β_* contains a basic element of signature σ w.r.t. $U_{[\pi_Y]}$.

We consider only one direction: suppose that α_* contains a basic element $[Z, j]$ of signature σ , where $Z \in \alpha$. Then there must be $Z' \in \beta$ such that, for all $v \in Z'$, there is $u \in Z$ with uRv , since R was a global neighborhood bisimulation. Furthermore, since R is full on X , for every $w \in Z$ we can pick some $w' \in Y$ with wRw' , and put

$$Z'' = Z' \cup \{w' \mid w \in Z\}$$

By monotonicity we have $Z'' \in \beta$. Furthermore, Z and Z'' are clearly related so that the following back-and-forth conditions hold: for all $u \in Z$ there is $v \in Z''$ with uRv , and for all $v \in Z''$ there is $u \in Z$ with uRv . Since any two states related by R have the same propositional type, it follows that the same propositional types appear in $[Z, j]$ and $[Z'', j]$. But since both these sets contain at least 2^k copies of every propositional type that appears in them, it follows that $[Z, j]$ and $[Z'', j]$ have the same 2^k -signature, as required. \square

We are going to show, by induction on a natural number $m \leq k$, that if two one-step models of the form (X_*, α_*, V_1) and (Y_*, β_*, V_2) match up to depth 2^m , then they satisfy the same formulas of quantifier depth m . Together with the previous lemma, it then follows that the one-step models $(X_*, \alpha_*, V_{[\pi_X]})$ and $(Y_*, \beta_*, U_{[\pi_Y]})$ satisfy the same formulas of quantifier depth $\leq k$. For the basis case of $2^0 = 1$, we need the following result (we recall that the case of atomic formulas $\Box a$ can be safely ignored):

Lemma 6.19. *Let B be a set of variables containing A , and let $V_1 : B \rightarrow \mathcal{Q}(X_*)$ and $V_2 : B \rightarrow \mathcal{Q}(Y_*)$ be valuations such that*

$$(X_*, \alpha_*, V_1) \approx_1 (Y_*, \beta_*, V_2)$$

Then these two one-step models satisfy the same atomic formulas of the one-step language $1\text{SO}_{\{\Box\}}(B)$.

Proof. We only prove one direction for each case.

Suppose first that

$$(X_*, \alpha_*, V_1) \Vdash_1 p \subseteq q$$

where $p, q \in B$. Suppose for a contradiction that $V_2(p) \not\subseteq V_2(q)$. Then there is some $(u, i, Z, j) \in Y_*$ such that $(u, i, Z, j) \in V_2(p) \setminus V_2(q)$. If (u, i, Z, j) comes from the residue of (Y_*, β_*) then since the residue of (X_*, V_*) has the same 1-signature, it must contain some element (u', i', Z', j') of the same 1-type, and so we cannot have $V_1(p) \subseteq V_1(q)$. The case where (u, i, Z, j) comes from a basic member is similar.

Now, suppose that

$$(X_*, \alpha_*, V_1) \Vdash_1 \Box p$$

Then $V_1(p) \in \alpha_*$, so there is some basic element $[Z, j] \in \alpha_*$ with $[Z, j] \subseteq V_1(p)$. There must be some basic $[Z', j'] \in \beta_*$ of the same 1-signature over B as $[Z, j]$, and clearly it follows that $[Z', j'] \subseteq V_2(p)$ and so $V_2(p) \in \beta_*$ as required. \square

We now only need the following lemma:

Lemma 6.20. *Let B be a finite set of variables containing A , let $0 < m \leq k$ and let $V_1 : B \rightarrow \mathcal{Q}(X_*)$ and $V_2 : B \rightarrow \mathcal{Q}(Y_*)$ be valuations such that*

$$(X_*, \alpha_*, V_1) \approx_{2^m} (Y_*, \beta_*, V_2)$$

Let q be any fresh variable. Then for any valuation V'_1 over $B \cup \{q\}$ extending V_1 with some value for q , there exists a valuation V'_2 over $B \cup \{q\}$ extending V_2 , such that

$$(X_*, \alpha_*, V'_1) \approx_{2^{(m-1)}} (Y_*, \beta_*, V'_2)$$

and vice versa.

Proof. We only prove one direction since the other direction can be proved by a symmetric argument.

Let V'_1 be given. By the hypothesis, for any 2^m -signature σ over the variables B , the number of basic elements of signature σ in α_* and β_* relative to V_1 and V_2 are either both finite and the same, or both infinite. Let $\sigma_1, \dots, \sigma_k$ be a list of all the distinct 2^m -signatures over B such that the set of basic elements of α_* and β_* of signature σ_i , with $1 \leq i \leq k$, is non-empty but finite, and let $\sigma_{k+1}, \dots, \sigma_l$ be a list of all the 2^m -signatures such that, for $k+1 \leq i \leq l$, there are infinitely many basic elements of α_* and of β_* of signature σ_i . Then, for each $i \in \{1, \dots, l\}$, let $\alpha_*[\sigma_i]$ denote the set of basic elements in α_* of signature σ_i , and similarly let $\beta_*[\sigma_i]$ denote the set of basic elements of β_* of signature σ_i . Then $\alpha_*[\sigma_1], \dots, \alpha_*[\sigma_l]$ is a partition of the set of basic elements of α_* into non-empty cells, and similarly $\beta_*[\sigma_1], \dots, \beta_*[\sigma_l]$ is a partition of the set of basic elements of β_* .

Given the extended valuation V'_1 on X_* defined on variables $B \cup \{q\}$, we similarly let $\tau_1, \dots, \tau_{k'}$ be a list of all the 2^{m-1} -signatures over $B \cup \{q\}$ such that, for $1 \leq i \leq k'$, the set of basic elements of α_* of 2^{m-1} -signature τ_i is non-empty but finite. We let $\tau_{k'+1}, \dots, \tau_{l'}$ be a list of all the 2^{m-1} -signatures over $B \cup \{q\}$ such that, for each i with $k'+1 \leq i \leq l'$, the set of basic elements of α_* of 2^{m-1} -signature τ_i is infinite. Let $\alpha_*[\tau_i]$ denote the set of basic elements of α_* of 2^{m-1} -signature τ_i , so that the collection $\alpha_*[\tau_1], \dots, \alpha_*[\tau_{l'}]$ constitutes a second partition of the set of basic elements of α_* . It will be useful to introduce the abbreviation D_1 for the finite set $\alpha_*[\sigma_1] \cup \dots \cup \alpha_*[\sigma_k]$, and the abbreviation D_2 for the finite set $\alpha_*[\tau_1] \cup \dots \cup \alpha_*[\tau_{k'}]$.

For each i with $1 \leq i \leq k$, there is a bijection between the set $\alpha_*[\sigma_i]$ and $\beta_*[\sigma_i]$, and we can paste all these bijections together into a bijective map

$$f : \alpha_*[\sigma_1] \cup \dots \cup \alpha_*[\sigma_k] \rightarrow \beta_*[\sigma_1] \cup \dots \cup \beta_*[\sigma_k]$$

Since every basic element of α_* not in D_1 belongs to a 2^m -signature of which there are infinitely many basic elements in β_* , and since $D_1 \cup D_2$ is finite, it is easy to see that we can extend the map f to a map g which is an injection from the set $D_1 \cup D_2$ into the set of basic elements of β_* , such that for each basic element $[Z, j]$ in $D_1 \cup D_2$, $[Z, j]$ and $g([Z, j])$ have the same 2^m -signature over B , and such that $g \upharpoonright D_1 = f$. Each basic element of β_* not in the image of g must then be of one of the 2^m -signatures $\sigma_{k+1}, \dots, \sigma_l$, and so we can partition the set of basic elements of β_* outside the image of g into the cells $\beta_*[\sigma_{k+1}] \setminus g[D_2], \dots, \beta_*[\sigma_l] \setminus g[D_2]$. For each i with $k+1 \leq i \leq l$, let $\gamma_1^i, \dots, \gamma_r^i$ list all infinite sets of the form $\alpha_*[\sigma_i] \cap \alpha_*[\tau_j]$ for $k'+1 \leq j \leq l'$. The list $\gamma_1^i, \dots, \gamma_r^i$ must be non-empty, and so since the set $\beta_*[\sigma_i] \setminus g[D_2]$ is also infinite, we may partition it into r many infinite cells and list these as $\delta_1^i, \dots, \delta_r^i$. Now, for each basic element $[Z, j]$ of β_* , we define a map $W_{[Z, j]}$ from $B \cup \{q\}$ to $\mathcal{P}([Z, j])$ by a case distinction as follows:

Case 1: $[Z, j] = g([Z', j'])$ for some $[Z', j'] \in D_1 \cup D_2$. Then $[Z, j]$ and $[Z', j']$ have the same 2^m -signature over B . Using this fact we define the valuation $W_{[Z, j]}$ so that, for each $p \in B$, we have $W_{[Z, j]}(p) = V_2(p) \cap [Z, j]$, and so that $[Z', j']$ and $[Z, j]$ have the same 2^{m-1} -signature over $B \cup \{q\}$ with respect to the valuations V_1' and $W_{[Z, j]}$. We show how to assign the value of the variable q : for each propositional type t over $B \cup \{q\}$, there are three different possible cases to consider. If $[Z', j']$ has $m' < 2^{m-1}$ elements of type $t \cup \{q\}$ over $B \cup \{q\}$, then pick m' many elements of $[Z, j]$ of type t and mark them by q . This is possible since $m' < 2^{m-1} \leq 2^m$ and $[Z', j']$ and $[Z, j]$ have the same 2^m -signature. If there are $m' < 2^{m-1}$ elements of $[Z', j']$ of type t over $B \cup \{q\}$, then pick m' elements of $[Z, j]$ of type t over B , and mark all the other elements of $[Z, j]$ of type t by q . Finally, if there are at least 2^{m-1} elements of $[Z', j']$ of type $t \cup \{q\}$ over $B \cup \{q\}$ and at least 2^{m-1} elements of $[Z', j']$ of type t over $B \cup \{q\}$, then all in all there must be at least 2^m elements of $[Z', j']$ of type t over B , and so there must be at least 2^m elements of $[Z, j]$ of type t over B . Pick 2^{m-1} of these and mark them by q . Finally, let $W_{[Z, j]}(q)$ be the set of elements of $[Z, j]$ marked by q .

Case 2: $[Z, j]$ is not in the image of g . Then there must be some $i \in \{k' + 1, \dots, l'\}$ such that $[Z, j] \in \beta_*[\sigma_{k+1}] \setminus g[D_2]$, and this set is partitioned into $\delta_1^i, \dots, \delta_r^i$. Let $[Z, j] \in \delta_r^i$, and pick some arbitrary element $[Z', j']$ of the set γ_j^i . Then $[Z', j']$ and $[Z, j]$ have the same 2^m -signature over B and we can proceed as in Case 1.

We define the valuation V_2' so that the intersection of $V_2'(q)$ with the union of all the basic members of β_* equals the union of the sets $W_{[Z, j]}(q)$ for $[Z, j]$ a basic element in β_* , and so that the residue of (Y_*, β_*) has the same $2^{(m-1)}$ -signature as the residue of (X_*, α_*) with respect to the valuations V_1' and V_2' . This can be done using the same reasoning as in the two previous cases. We now need to check that

$$(X_*, \alpha_*, V_1') \approx^{2^{(m-1)}} (Y_*, \beta_*, V_2')$$

First, suppose there are infinitely many basic elements of α_* of some 2^{m-1} signature τ_j , meaning that $k' \leq j \leq l'$. Then since the set $\alpha_*[\tau_j]$ is infinite, D_1 is finite and

$$\alpha_*[\tau_j] = (D_1 \cap \alpha_*[\tau_j]) \cup (\alpha_*[\sigma_{k+1}] \cap \alpha_*[\tau_j]) \cup \dots \cup (\alpha_*[\sigma_l] \cap \alpha_*[\tau_j])$$

there must be some $i \in \{k + 1, \dots, l\}$ such that the set $\alpha_*[\sigma_i] \cap \alpha_*[\tau_j]$ is infinite. This means that $\alpha_*[\sigma_i] \cap \alpha_*[\tau_j]$ appears in the list $\gamma_1^i, \dots, \gamma_r^i$, and so we see that all elements of some member of the list $\delta_1^i, \dots, \delta_r^i$ will have the 2^{m-1} -signature τ_j . Since each member of this list is infinite, we see that there must be infinitely many basic elements of β_* of signature τ_j .

Conversely, suppose there are infinitely many basic elements of β_* of 2^{m-1} -signature τ_j over $B \cup \{q\}$. Then since the image of g is finite, some of these elements must be outside the image of g , which means that for some $i \in \{k + 1, \dots, l\}$, some member of the list $\delta_1^i, \dots, \delta_r^i$ will consist of elements of signature τ_j . This means that some member of the list $\gamma_1^i, \dots, \gamma_r^i$ will consist of elements of signature τ_j , and since each member of this list is infinite we see that α_* has infinitely many basic elements of 2^{m-1} -signature τ_j over $B \cup \{q\}$.

Finally, suppose that there are finitely many basic elements of α_* and β_* of 2^{m-1} -signature τ_j . We check that the mapping g restricts to a bijection between the basic elements of α_* and β_* of this signature. First, g is injective and maps basic elements of α_* of signature τ_j to basic elements of β_* of signature τ_j . It only remains to show that (the restriction of) g is surjective, i.e. each basic element $[Z, r]$ of signature τ_j is equal to $g([Z', r'])$ for some $[Z', r']$. But suppose $[Z, r]$ is not in the image of g ; then it is in one of the members of the list $\delta_1^i, \dots, \delta_r^i$ for some i , and since each of these members is an infinite set of basic elements

of the same signature, we see that there are infinitely many basic elements of β_* of signature τ_j , contrary to our assumption. Hence, the proof is done. \square

Theorem 6.21. *The monotone μ -calculus is the neighborhood bisimulation invariant fragment of monotone monadic second-order logic. In a formula:*

$$\text{MMSO}/\sim \equiv \mu\text{MML}$$

Proof. It suffices to prove that $\text{MSO}_{\{\square, \square\}}/\sim$ is equal to $\mu\text{ML}_{\{\square, \square\}}$, and then apply Theorem 6.12. For this, by Theorem 1.4 it suffices in turn to prove that the construction $(-)_*$ is adequate for all formulas in $\text{ISO}_{\{\square, \square\}}(A)$ of quantifier depth $\leq k$. We can prove this by combining the last three lemmas, using Ehrenfeucht-Fraïssé games for the one-step language. Lemmas 6.18 and 6.20 provide a recipe for how “Duplicator” can survive k steps of the game comparing the models $(X_*, \alpha_*, V_{[\pi_X]})$ and $(Y_*, \beta_*, U_{[\pi_Y]})$, and 6.19 guarantees that the valuations constructed at the end of the game will satisfy the same atomic formulas. Working out the full argument is entirely standard, so we leave the details to the reader. \square

7. FUTURE WORK

We conclude by mentioning some questions for future research:

- (1) Is there a good *categorical* and “logic free” characterization of those set functors \mathbb{T} that admit an adequate (or weakly adequate) uniform construction, for instance, in terms of \mathbb{T} preserving certain limits or colimits? Trying to answer this question would involve a deeper study of how the model theory of one-step languages is related to categorical properties of the type functor involved. Related to this, an anonymous referee pointed out to us that the machinery of one-step frames, covers and uniform constructions are all taking place in the *category of elements* associated with the set functor \mathbb{T} , and properties of the functor \mathbb{T} are closely related to properties of the corresponding category of elements [2]. This could very well be a fruitful direction to investigate further.
- (2) Can we improve our work on the supported companion functor, to the effect that every set functor \mathbb{T} has a companion \mathbb{T}' that admits an adequate uniform construction? Relating this to the previous question, we would like to understand *why* the supported companion to \mathcal{M} admits an adequate uniform construction, but not \mathcal{M} itself. The construction achieves two things, in general: first, it ensures that every $\alpha \in \underline{\mathbb{T}}X$ has a unique smallest support (even when X is infinite), often called its *base*. Second, and in our view more importantly, we get that the map $\text{base}_X : \underline{\mathbb{T}}X \rightarrow \mathcal{P}X$ defined by sending each α to its base is a natural transformation. Conditions for a functor under which this holds has been isolated by Gumm in [9]. Are these conditions related to the existence of adequate uniform constructions?
- (3) It would be interesting to further explore the relation between $\text{MSO}_{\mathbb{T}}$ and the first-order logic of Litak & alii [17] for \mathbb{T} -coalgebras. For instance, an interesting question would be whether (on \mathbb{T} -tree models) $\text{MSO}_{\mathbb{T}}$ is equivalent to some extension of this first-order language with fixpoint operators.
- (4) Finally, there is the question of finding sufficient *and necessary* conditions for a Janin-Walukiewicz theorem to hold. Related to this question, we have not been able to produce an example of a functor for which the Janin-Walukiewicz theorem does *not* hold (the question of whether such an example can be found was raised to us by an anonymous referee). Given all that we can say for sure, it could be the case that a

Janin-Walukiewicz theorem simply holds for *every* set functor, but we conjecture that at least some conditions on the functor are required.

Acknowledgement. We wish to thank the anonymous referees for their very thorough and constructive feedback, which led to many substantial improvements on the original draft of this paper.

REFERENCES

- [1] M. Barr, “Terminal coalgebras in well-founded set theory”. *Theoretical Computer Science*, 114: 299–315, 1993.
- [2] M. Barr and C. Wells, “Category theory for computing science”, Prentice Hall New York, 1990.
- [3] A. Carboni, G.M. Kelly, and R.J. Wood, “A 2-categorical approach to change of base and geometric morphisms I”. *Cahiers de Topologie et Géométrie Différentielle Catégoriques*, 32: 47-95, 1991.
- [4] C. Cirstea, C. Kupke, and D. Pattinson, “EXPTIME tableaux for the coalgebraic μ -calculus,” *Lecture Notes in Computer Science* 5771 : 179–193, 2009.
- [5] C. Cirstea, A. Kurz, D. Pattinson, L. Schröder, and Y. Venema, “Modal logics are coalgebraic,” *The Computer Journal*, 54: 524–538, 2011.
- [6] A.E. Emerson and C.S. Jutla, “Tree automata, mu-calculus and determinacy”, Proceedings of *32nd Annual Symposium on Foundations of Computer Science*, pp.368–377, 1991.
- [7] S. Enqvist, F. Seifan and Y. Venema, “Monadic second-order logic and bisimulation invariance for coalgebras”. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2015)*, pp.353–365, 2015.
- [8] G. Fontaine, R. Leal, and Y. Venema, “Automata for coalgebras: an approach using predicate liftings,” in *Proc. 37th International Colloquium on Automata, Languages, and Programming (ICALP 2010)*, *Lecture Notes in Computer Science* 6199: 381–392, 2010.
- [9] H.P. Gumm, “From T-coalgebras to filter structures and transition systems”, in *Proc. International Conference on Algebra and Coalgebra in Computer Science (CALCO 2005)*, 194–212, 2005.
- [10] H. Hansen and C. Kupke, “A coalgebraic perspective on monotone modal logic,” *Electronic Notes in Theoretical Computer Science*, 106 : 121 – 143, 2004.
- [11] J.E. Hopcroft, R. Motwani and J.D. Ullman, “Introduction to automata theory, languages, and computation”, *ACM SIGACT News*, 32 : 60–65, 2001.
- [12] B. Jacobs, “Introduction to coalgebra. Towards mathematics of states and observations,” Cambridge University Press, 2016.
- [13] D. Janin and I. Walukiewicz, “On the expressive completeness of the propositional μ -calculus w.r.t. monadic second-order logic,” *Lecture Notes in Computer Science* 1119: 263–277, 1996.
- [14] K. Kapulkin, A. Kurz and J. Velebil, “Expressiveness of Positive Coalgebraic Logic”. *Advances in Modal Logic* 9: 368–385, 2012.
- [15] C. Kupke and Y. Venema, “Coalgebraic automata theory: basic results,” *Logical Methods in Computer Science* 4: 1–43, 2008.
- [16] R. A. Leal, “Predicate Liftings Versus Nabla Modalities”, *Electronic Notes in Theoretical Computer Science* 203: 195–220, 2008.
- [17] T. Litak, D. Pattinson, K. Sano, and L. Schröder, “Coalgebraic predicate logic,” *Lecture Notes in Computer Science* 7392: 299–312, 2012.
- [18] D.A. Martin, “Borel Determinacy”, *Annals of Mathematics* 102 : 363–371, 1975.
- [19] J. Marti, “Relation liftings in coalgebraic modal logic”. MSc thesis, Institute for Logic, Language and Computation, 2011.
- [20] J. Marti and Y. Venema, “Lax extensions of coalgebra functors”, Proceedings of the International Workshop on Coalgebraic Methods in Computer Science (CMCS), *Lecture Notes in Computer Science* 7399: 150 – 16, 2012.
- [21] J. Marti and Y. Venema, “Lax extensions of coalgebra functors and their logic”, *Journal of Computer and System Sciences* 81.5 (2015): 880-900.

- [22] J. Marti, F. Seifan and Y. Venema, “Uniform Interpolation for Coalgebraic Fixpoint Logic”, in Proceedings of the 6th Conference on Algebra and Coalgebra in Computer Science (CALCO), Leibniz International Proceedings in Informatics, 35: 238–252, 2015.
- [23] L. Moss, “Coalgebraic logic,” *Annals of Pure and Applied Logic*, 96 : 277–317, 1999. (Erratum published in *Ann.P.Appl.Log.* 99:241–259, 1999.)
- [24] A.W. Mostowski, “Games with forbidden positions”, Technical Report 78, Uniwersytet Gdański, Instytut Matematyki, 1991.
- [25] D. Muller and P. Schupp, “Simulating alternating tree automata by nondeterministic automata,” *Theoretical Computer Science*, 141: 69–107, 1995.
- [26] R. Parikh, “The Logic of Games and Its Applications”, in *Selected Papers of the International Conference on “Foundations of Computation Theory” on Topics in the Theory of Computation*, 111–139, 1985.
- [27] D. Pattinson, “Coalgebraic modal logic: Soundness, completeness and decidability of local consequence,” *Theoretical Computer Science*, 309: 177–193, 2003.
- [28] M. Rabin, “Decidability of second-order theories and automata on infinite trees,” *Transactions of the American Mathematical Society*, 141: 1–35, 1969.
- [29] J. Rutten, “Universal coalgebra: A theory of systems,” *Theoretical Computer Science*, 249: 3–80, 2000.
- [30] L. Santocanale and Y. Venema, “Uniform Interpolation for Monotone Modal Logic,” *Advances in Modal Logic* 8 : 350–370, 2010.
- [31] L. Schröder, “Expressivity of coalgebraic modal logic: the limits and beyond,” *Theoretical Computer Science*, 390: 230–247, 2008.
- [32] V. Trnková, “Some properties of set functors”, *Commentationes Mathematicae Universitatis Carolinae*, 10:323–352, 1969.
- [33] Y. Venema, “Automata and fixed point logic: a coalgebraic perspective,” *Information and Computation*, 204: 637–678, 2006.
- [34] Y. Venema, “Lectures on the modal μ -calculus,” 2012, Lecture Notes, ILLC, University of Amsterdam.
- [35] Y. Venema, “Expressiveness modulo bisimilarity: a coalgebraic perspective,” in *Johan van Benthem on Logic and Information Dynamics*, Outstanding Contributions to Logic 5: 33–65 2014.
- [36] I. Walukiewicz, “Monadic second order logic on tree-like structures,” *Theoretical Computer Science* 275: 311–346, 2002.
- [37] W. Zielonka, “Infinite games on finitely coloured graphs with applications to automata on infinite trees”, *Theoretical Computer Science* 200 :135–183, 1998.