Deciding the First Levels of the Modal μ Alternation Hierarchy by Formula Construction

Karoliina Lehtinen and Sandra Quickert

Laboratory for Foundations of Computer Science, University of Edinburgh
10 Crichton Street, Edinburgh, UK
M.K.Lehtinen@sms.ed.ac.uk, squicke1@inf.ed.ac.uk

Abstract

We construct, for any sentence Ψ of the modal μ calculus (L_μ), a derived sentence Ψ^{ML} in the modal fragment ML of L_μ and a sentence $\Psi^{\Pi_1^\mu}$ in the fragment Π_1^μ of L_μ without least fixpoints such that Ψ is equivalent to a formula in ML or Π_1^μ if and only if it is equivalent to Ψ^{ML} or $\Psi^{\Pi_1^\mu}$ respectively. The formula $\Psi^{\Sigma_1^\mu}$ such that Ψ is equivalent to $\Psi^{\Sigma_1^\mu}$ if and only if Ψ is semantically in the greatest-fixpoint free fragment Σ_1^μ is obtained by duality to $\Psi^{\Pi_1^\mu}$. This yields a new proof of decidability of the first levels of the modal μ alternation hierarchy. The blow-up incurred by turning Ψ into the modal formula Ψ^{ML} is shown to be necessary: there are ML formulas that can be expressed sub-exponentially more efficiently with the use of fixpoints. For Π_1^μ and Σ_1^μ however, as long as formulas are in guarded disjunctive form, the transformation into a syntactically Π_1^μ or Σ_1^μ does not increase the size of the formula.

1998 ACM Subject Classification F.4.1 Mathematical Logic

Keywords and phrases modal μ calculus, fixpoint logic, alternation hierarchy

Digital Object Identifier 10.4230/LIPIcs.CSL.2015.457

1 Introduction

The modal μ calculus (L_μ), a logic expressing properties of labelled transition systems, was first introduced by Kozen in 1983 [5]. Its popularity is due to its simple but productive syntax and appealing decidability: deciding satisfiability is EXPTIME-complete; model checking is in NP and conjectured to be in P.

Syntactically, L_μ consists simply of a propositional modal logic augmented with its namesake least fixpoint operator μ and the dual greatest fixpoint operator ν. Both the expressivity and complexity of the logic stem from the alternating usage of μ and ν: the more alternations are allowed, the richer the fragment of L_μ but the more difficult its model-checking. Indeed, the alternation hierarchy, consisting of L_μ fragments for which the number of alternations is fixed is strict [10, 1]. For each fixed alternation-depth, the model-checking problem is of polynomial complexity, but for whole of L_μ the best current algorithms still have complexity exponential in a function of the alternation depth.

It is therefore of both practical and theoretical interest to reduce, whenever possible, the number of alternations used to express a property. Even though the problem must be at least EXPTIME-hard, in practice model checking is likely to benefit from the one-time cost of reducing a formula to its simplest form, especially since the size of the formula is unlikely to dominate the runtime complexity of the model checking. However, only properties expressible in modal logic or with a single type of fixpoint operator are currently known to be recognisable. In general, for a given Φ, finding an equivalent Ψ with smallest alternation depth is one of the main open problems surrounding the modal μ calculus.

24th EACSL Annual Conference on Computer Science Logic (CSL 2015).
Editor: Stephan Kreutzer; pp. 457–471
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Here we focus on the lowest levels of the alternation hierarchy, which are known to be decidable. The class *ML* of properties expressible in modal logic, L_μ without fixpoints, resides at the base of the alternation hierarchy. These are properties which dictate some behaviour in the initial fragment of a structure, up to fixed depth. Otto [13] showed that properties of this class are recognisable via a reduction to S2S, the monadic second order logic over binary trees. Küsters and Wilke showed in [8] that the problem of deciding whether a property of L_μ can be expressed with only least fixpoints, or, by duality, only greatest fixpoints, is EXPTIME-complete. Their proof first constructs a bottom-up tree automaton of which the states correspond to sets of subformulas based on the L_μ formula. Roughly speaking, the bottom-up automaton accepts a structure if it has an initial fragment such that every completion admits a valid assignment of automaton states to its nodes. This automaton is closed under bisimulation if and only if it is Σ_1^μ definable and equivalent to the original formula.

Both Otto's, and Küsters and Wilke's results focus on deciding whether a property is expressible with a formula of the lower class, but they pay little heed to the target formula. The Σ_1^μ-formula is described in the technical report [9] as part of the proof of decidability of the first alternation level. Unfortunately the transformation can incur a double-exponential blow-up in the size of the formula. From a model-checking point of view, this is problematic: not only is the transformation into a Σ_1^μ-formula non trivial, but the transformation does not reduce the complexity of the model-checking procedure. The formula is also quite complex and does not necessarily resemble the original formula, so it is difficult to follow how the redundant fixpoints were eliminated. Otto does not describe the target *ML* formula at all but it seems that if one can be extracted from the decision procedure for *ML*, it will also be based on a power-set construction around subformulas of the original formula.

This paper puts the focus on the relation between a formula and its equivalent formulas in lower alternation classes. It describes $\Psi^{ML}, \Psi^{\Pi_1^\mu}$ and $\Psi^{\Sigma_1^\mu}$, formulas based on, and syntactically close to Ψ such that Ψ is semantically equivalent to Ψ^C if Ψ is semantically in the class C. We show that the required transformations into a *ML*, Π_1^μ or Σ_1^μ formula are conceptually very simple and easily implementable. The formula Ψ^{ML} is perhaps as one could anticipate: if Ψ is semantically in *ML*, then there is some m such that Ψ is equivalent to the formula obtained by approximating all fixpoints to their m^{th} stage and truncating the resulting formula at modal depth m. As it turns out, m needs to be at most exponential in the length of the formula. Interestingly, the potential blow-up in the size of the formula is not accidental: there are properties which are semantically modal but can be expressed with much shorter Σ_1^μ-formulas than *ML*-formulas. We show that in this sense, Σ_1^μ is at least sub-exponentially more concise than *ML*. There is a clear trade-off between syntactic complexity and formula length. From the model checking point of view, this means that if a formula has high modal depth, it may be wise to retain some fixpoint operators which will keep the size of the formula down. In contrast to Ψ^{ML}, the most interesting aspect of $\Psi^{\Pi_1^\mu}$ is perhaps its simplicity. As long as Ψ is given in disjunctive form, $\Psi^{\Pi_1^\mu}$ and its dual $\Psi^{\Sigma_1^\mu}$ are at most as large as Ψ: for disjunctive formulas, Π_1^μ and Σ_1^μ are perfectly concise in the sense that using further alternations to express the same property does not reduce the size of the formula. This is significant in that the transformation from Ψ to $\Psi^{\Pi_1^\mu}$ results in a genuinely simpler formula instead of a formula in which alternations are eliminated at the cost of conciseness. The exponential complexity of the resulting decision procedure which compares Ψ with $\Psi^{\Pi_1^\mu}$ is also optimal. The transformation itself is also noteworthy: it consists roughly speaking of replacing every μ-operator with either $\perp$ or a ν-operator. In other words, in L_μ, any satisfiable μ-subformula is either necessary or interchangeable with the identical ν-subformula.

The key to the transformation into $\Psi^{\Pi_1^\mu}$ is the use of disjunctive form. Disjunctive form, introduced in [4], is a syntactic constraint on conjunctions and universal modalities. It has been used in the context of tableau methods to decide satisfiability for example but as this paper shows, it is also a promising tool for syntactic manipulations.

2 Preliminaries

▶ **Definition 1** (Modal μ). Given a set of atomic propositions $Prop = \{P, Q, ...\}$ and a set of fixpoint variables $Var = \{X, Y, ...\}$, the syntax of L_μ is given by:

$$\phi := P \mid X \mid \neg P \mid \phi \wedge \phi \mid \phi \vee \phi \mid \Diamond\phi \mid \Box\phi \mid \mu X.\phi \mid \nu X.\phi \mid \bot \mid \top$$

This definition only allows formulas in positive from: negation is only applied to propositional variables. Positivity does not restrict the expressivity of the logic. A formula is guarded if every fixpoint variable is within its binding in the scope of a modality. As is well documented in the literature [11, 7] every L_μ formula is equivalent to a formula in guarded form. Without loss of expressivity, we therefore restrict ourselves to L_μ in guarded positive form. For the sake of clarity, we only consider the uni-modal L_μ but expect the multi-modal case, as defined in [2] for example, to behave broadly speaking similarly.

Notation

If $\phi(X)$ is a formula, we write $\phi(\psi)$ for the formula ϕ where every occurrence of the variable X is replaced with ψ. For readability, if ϕ is the binding formula of the fixpoint variable X as in $\mu X.\phi$, then $\phi(\psi)$ is also ϕ with X substituted by ψ.

Formulas of L_μ are evaluated on transition systems, referred to as structures, represented by potentially infinite trees annotated with propositions.

▶ **Definition 2** (Structures). A structure $\mathcal{M} = (S, s_0, R, P)$ consists of a set of states S, rooted at some initial state $s_0 \in S$, and a successor relation $R \subseteq S \times S$ between the states. Every state s is associated with a set of propositions $P(s) \subseteq Prop$ which it is said to satisfy. In this document it is sufficient for us to consider finitely branching structures, so we require that nodes only have finitely many successors. It is well-known that any structure can be represented as a potentially infinite tree. To ease the manipulation of structures, we adopt this representation.

For clarity and conciseness, we give the semantics directly in terms of parity games – the equivalence between these and the usual semantics is a standard result. For a presentation of the standard semantics of L_μ and a proof of the equivalence to the above, see for example Bradfield and Stirling [2].

▶ **Definition 3** (Parity games). A parity game is a potentially infinite two-player game on a graph $\mathcal{G} = (V_0, V_1, E, v_I, \Omega)$ of which the vertices consist of two disjoint sets, V_0 and V_1 belonging to the players Even and Odd respectively, and are annotated with positive integer priorities bounded by some maximal priority q, via $\Omega : V_0 \cup V_1 \to \{0, 1, ..., q\}$. Player Even and her opponent, player Odd, move a token along the edges $E \in (V_0 \cup V_1) \times (V_0 \cup V_1)$ of the graph starting from an initial position $v_I \in V_0 \cup V_1$, each choosing the next position when the token is on a vertex in their partition. Some positions p might have no successors in which case they are winning for the player of the parity of $\Omega(p)$. A play consists of the potentially infinite sequence of vertices visited by the token. For finite plays, the last visited parity decides the winner of the play. For infinite plays, the parity of the lowest priority visited infinitely often decides the winner of the game: Even wins if the lowest priority visited infinitely often is even; otherwise Odd wins. Note that in the literature, the highest priority is sometimes used, equivalently, as the most significant priority.

The winner of a parity game is defined in terms of existence of winning strategies. Strategies in general can depend on the history of the game, but in the case of parity games positional strategies which depend on the current position alone are sufficient, so we define strategies as mappings from position to position.

▶ **Definition 4** (Positional Strategies)**.** A positional strategy σ for one of the players in a parity game $\mathcal{G}$ is a mapping from the Player's positions V_0 or V_1 in the game to a valid successor position. A play respects a Player's strategy σ if the successor positions in the play belonging to the Player are those dictated by σ. If σ is Even's strategy and τ is Odd's strategy, then there is a unique play $\sigma \times \tau$ respecting both strategies. The winner of the parity game at a position is the player who has a strategy σ, said to be a winning strategy, such that they win $\sigma \times \tau$ from that position for any counter-strategy τ. A strategy σ is said to reach a position if there is a counter-strategy τ such that the position is along the play $\sigma \times \tau$.

Parity games are positionally determined: for every position either Even or Odd has a winning positional strategy [3]. This means that strategies gain nothing from looking at the whole play rather than just the current position. As a consequence, we may take a strategy to be memoryless: it maps each position of a player to a successor.

For any modal μ formula ϕ and a structure $\mathcal{M}$ we define a parity game $\mathcal{M} \times \phi$, constructed in polynomial time, and say that $\mathcal{M}$ satisfies ϕ, written $\mathcal{M} \models \phi$, if and only if Even has a winning strategy in $\mathcal{M} \times \phi$.

▶ **Definition 5** (Model-checking parity game)**.** For any formula ϕ of L_μ, taken to be in positive form, and a model $\mathcal{M}$, define a parity game $\mathcal{M} \times \phi$ with positions (s, ψ) where s is a state of $\mathcal{M}$ and ψ is a subformula of ϕ. The initial position is (s_0, ϕ) where s_0 is the root of $\mathcal{M}$. Positions (s, ψ) where ψ is a disjunction or a formula starting with an existential modality $\Diamond$ belong to Even while conjunctions and formulas starting with a universal modality $\Box$ belong to Odd. Other positions have at most one successor so their owner is irrelevant; let them be Even's. There are edges from $(s, \psi \vee \psi')$ and $(s, \psi \wedge \psi')$ to both (s, ψ) and (s, ψ'); from $(s, \mu X.\phi)$ and $(s, \nu X.\phi)$ to (s, ϕ); from (s, X) to $(s, \nu X.\psi)$ if X is bound by ν, or $(s, \mu X.\psi)$ if it is bound by μ; finally, from $(s, \Diamond\psi)$ and $(s, \Box\psi)$ to every (s', ψ) where (s, s') is an edge in the model $\mathcal{M}$. Positions (s, P), $(s, \neg P), (s, \top)$ and $(s, \bot)$ have no successors. The parity function assigns an even priority to $(s, \top)$ and also to (s, P) if P satisfies s in $\mathcal{M}$ and to $(s, \neg P)$ if s does not satisfy P in $\mathcal{M}$; otherwise (s, P) and $(s, \neg P)$ receive odd priorities, along with $(s, \bot)$. Fixpoint variables are given distinct priorities such that ν-bound variables receive even priorities while μ-bound variables receive odd priorities. Furthermore, whenever X has priority i, Y has priority j and $i > j$, X must not appear free in the formula ψ binding Y in $\mu Y.\psi$ or $\nu Y.\psi$. In other words, inner fixpoints receive higher, less significant priorities while outer fixpoint receive low priorities. Other nodes receive a priority max which is larger than any of the priorities assigned to fixpoint nodes. This ensures that these will never be the lowest priority seen infinitely often.

We now use parity games to define the semantics of L_μ.

▶ **Definition 6** (Satisfaction relation)**.** A structure $\mathcal{M}$, rooted at s_0 is said to satisfy a formula Ψ of L_μ, written $\mathcal{M} \models \Psi$ if and only if the Even player has a winning strategy from (s_0, Ψ)

in $\mathcal{M} \times \Psi$. $\mathcal{M}$ satisfies a subformula ϕ of Ψ if it satisfies the formula ϕ where free fixpoint variables X are recursively replaced with their fixpoint binding $\mu X.\phi$ or $\nu X.\phi$ from Ψ. This is the case if and only if Even has a winning strategy from (s, ϕ) in $\mathcal{M} \times \Psi$.

Formulas are semantically equivalent if they are satisfied by exactly the same structures.

▶ **Definition 7** (Modal Logic, Π_1^μ and Σ_1^μ)**.** *ML* is the class of properties of structures expressible in modal logic, that is to say in L_μ without any fixpoint operators. A formula without fixpoint operators is said to be modal and has a modal depth which is the greatest number of nested modal operators in it. Π_1^μ is the class of properties expressible by a formula in positive form without using the least fixpoint operator μ and Σ_1^μ is the class of properties expressible by a formula in positive form without using the greatest fixpoint operator ν. If a formula does not contain μ, ν or both it is said to be *syntactically* in Π_1^μ, Σ_1^μ or *ML* respectively.

Beyond *ML*, Π_1^μ and Σ_1^μ, the syntactic complexity of formulas is measured by the number of alternations between least and greatest fixpoint operators. This is the alternation depth of a formula and corresponds to the number of priorities needed in the model-checking parity game. A precise definition of the alternation depth is given for example in [2]. The fragments of L_μ with bounded alternation depth form the alternation hierarchy, which is known to be strict: for each level, there are formulas which cannot be expressed with a formula of lower alternation depth [1, 10].

If a formula is equivalent to a formula syntactically in some alternation level, it is said to be semantically in that alternation level. Thus a formula of high syntactic alternation level may be of low semantic alternation level.

Our concern is to decide whether a formula is semantically in one of Π_1^μ, Σ_1^μ or *ML* and produce an equivalent formula syntactically in the appropriate alternation level.

3 The formula for *ML*

The modal fragment of L_μ, or *ML*, was shown by Otto to be decidable: for any formula of L_μ, we can decide whether there is an equivalent modal formula [13]. This section proposes a proof of decidability by formula construction: given a guarded formula Ψ, it presents a formula Ψ^{ML} in *ML* which Ψ is equivalent to if and only if Ψ is semantically a modal formula. The crux of the argument is that a semantically modal formula Ψ can only reach depth $2^{2|\Psi|}$ in any structure and therefore, if a formula is equivalent to a modal formula, it is sufficient to first approximate all fixpoints to the $2^{2|\Psi|}$-th stage of induction and then truncate the formula at modal depth $2^{2|\Psi|}$. At first sight this might seem like a wasteful solution since the size of the formula increases as it is unfolded. However, Example 20 shows that there are formulas which cannot be expressed in modal logic without at least a sub-exponential increase in formula size. This proves that approximating the fixponts to their $2^{2|\Psi|th}$ stage of induction is hardly excessive.

The first two definitions fix the notation for measuring the depth of a structure and approximating fixpoints. Note that all structures are represented as trees since we are interested in how far into a structure a formula can reach and this is easier to do when reasoning about trees rather than graphs.

▶ **Definition 8** (Rank and depth)**.** The rank of a state without successors is 0. The rank of a state with finitely many successors is $h + 1$ where h is the maximal rank amongst the state's successors. The depth of the root of a structure is 0; otherwise the depth of a state is one greater than the depth of its parent.

The rank of a finite tree is the rank of its root and corresponds to the length of the longest path in the tree.

The next definition formalises the notion of simultaneously evaluating all fixpoints to their n^{th} stage of induction.

▶ **Definition 9** (Approximants). Let $\mu X^0.\phi = \bot$ and $\mu X^n.\phi(X) = \phi(\mu X^{n-1}.\phi(X))$; let $\nu X^0.\phi = \top$ and $\nu X^n.\phi(X) = \phi(\nu X^{n-1}.\phi(X))$.
Then, using the notation $\phi[a/b]$ to mean ϕ where all instances of b are substituted with a, and *fixpoints*(ϕ) is the set of fixpoint variables in ϕ, let $\phi_n = \phi[\nu X^n/\nu X; \mu X^n/\mu X]_{\forall X \in \mathit{fixpoints}(\phi)}$, the formula ϕ where every fixpoint μX or νX is substituted with its n^{th} approximation μX^n or νX^n. See, e.g., [2].

We can then show easily enough that for trees of bounded height n, there is never any need to go beyond the n^{th} stage of induction. The intuition is as follows: without change in semantics, we can unfold all the fixpoints in ϕ all of n times each to obtain a formula equivalent to ϕ which only differs syntactically from ϕ_n at modal depth greater than n. On trees of bounded height n, the model-checking parity game will never reach this point for either formula. Therefore, both games for ϕ and ϕ_n must agree.

▶ **Lemma 10.** *If ϕ is guarded, ϕ and ϕ_n agree on trees of rank bounded by n.*

Proof. The game for ϕ_n is similar to the game for ϕ except for the additional rule that each priority p has a counter attached to it which counts how many times p occurs in the play without a smaller priority occurring in between. If a counter for an odd priority reaches $n+1$, Even loses immediately; if an even priority counter reaches $n+1$, then Odd loses immediately. If $\mathcal{M}$ is a tree of bounded rank, that is to say with a longest path of length no more than n, and ϕ is guarded, all plays in $\mathcal{M} \times \phi$ visit at most n states and do not visit a position (s, ψ) more than once. Therefore no priority is seen more than n times: no counter can reach $n+1$, so Even wins $\phi_n \times \mathcal{M}$ if and only if she can win $\phi \times \mathcal{M}$. Hence, on $\mathcal{M}$ a tree of rank at most n, ϕ and ϕ_n must agree. Write $\mathcal{M} \models \phi$ if and only if $\mathcal{M} \models \phi_n$. ◀

The formula ϕ_n is a modal formula but it may have modal depth greater than n, for example if a fixpoint is guarded by more than one modality or if it has interacting fixpoints. We will therefore define a truncating operation which reduces the modal depth of a formula to n.

▶ **Definition 11.** Let the formula ϕ^n be the one obtained from a modal formula ϕ by replacing subformulas $\Box\psi$ of modal depth n or larger with $\top$ and $\Diamond\psi$ of modal depth n or larger with $\bot$.

▶ **Lemma 12.** *Let Φ be guarded. Then $\mathcal{M} \models \phi^n$ iff $\mathcal{M}^n \models \phi$ where $\mathcal{M}^n$ is the infinite tree of $\mathcal{M}$ truncated at depth n. That is to say, ϕ^n is true in $\mathcal{M}$ iff the initial three of height n of $\mathcal{M}$ satisfies ϕ.*

Proof. First note that in the model checking parity game of modal formulas, a state at depth n can only be reached at a subformula that is itself at modal depth n. Assuming $\mathcal{M} \models \phi^n$, Even has a winning strategy σ in $\mathcal{M} \times \phi^n$ to prove it. This game is identical to $\mathcal{M} \times \phi$ until a position s at depth n is reached at $\bot$ or $\top$ instead of $\Diamond\psi$ or $\Box\psi$ respectively. If Even can win, her strategy cannot reach any position $(s, \bot)$. The game $\mathcal{M}^n \times \phi$ is also identical to $\mathcal{M} \times \phi$ until a position at depth n is reached. The strategy σ is winning in $\mathcal{M}^n \times \phi$ since it can avoid $(s, \Diamond\psi)$ positions where s it at depth n and positions $(s, \Box\psi)$ are automatically winning for states s at depth n since they have no successors in $\mathcal{M}^n$.
Conversely, assume Even has a winning strategy σ in $\mathcal{M}^n \times \phi$. She can use this strategy

in $\mathcal{M} \models \phi^n$ until it reaches positions (s, ψ) where s is at depth n. Since these are leaves, her winning strategy does not reach any state $(s, \Diamond\psi)$ where s is at depth n. In $\mathcal{M} \models \phi^n$ her strategy σ therefore only reaches final positions (s, P) and $(s, \top)$ where s is at depth n, which are winning for her. A strategy is therefore winning in $\mathcal{M} \times \phi^n$ if and only if it is winning in $\mathcal{M}^n \times \phi$ and therefore $\mathcal{M} \models \phi^n$ iff $\mathcal{M}^n \models \phi$. ◀

▶ **Example 13.** Consider the modal formula $\phi = A_0 \wedge \Diamond A_1 \wedge \Box(\Diamond\Diamond A_2 \vee \Box A_3) \wedge \Box\Box\Box A_4$. Then $\phi^3 = A_0 \wedge \Diamond A_1 \wedge \Box(\Diamond\bot \vee \Box A_3) \wedge \Box\Box\top$ is true in $\mathcal{M}$ iff its initial tree of height 3 satisfies ϕ. Similarly $\phi^2 = A_0 \wedge \Diamond A_1 \wedge \Box(\bot \vee \top) \wedge \Box\top$ is true in $\mathcal{M}$ if its initial tree of height 2 satisfies ϕ. Finally $\phi^1 = A_0 \wedge \bot \wedge \top = \bot$, which is reasonable since the root of $\mathcal{M}$ cannot satisfy $\Diamond A_1$ without having any successors.

We use the following lemma, which Otto also uses in [13]. The intuition is that if a formula is equivalent to a *ML*-formula of modal depth m, then what happens beyond depth m in a structure can have no effect on whether the formula holds in this structure or not. The semantic modal depth of a semantically *ML* formula is the least modal depth of any equivalent *ML* formula. Note that the syntactic alternation-depth of a formula is irrelevant to its semantic modal depth but only semantically *ML* formulas have a finite modal depth.

▶ **Lemma 14.** *If ϕ is guarded and of semantic modal depth m, then $\mathcal{M} \models \phi$ iff $\mathcal{M}^m \models \phi$ where $\mathcal{M}^m$ is the infinite tree of $\mathcal{M}$ truncated at depth m.*

Proof. If ϕ is guarded and of semantic modal depth m, there are formulas ψ equivalent to ϕ of syntactic modal depth m. The model checking parity game for a modal formula has no infinite paths in it. Furthermore for a formula of modal depth m, a play can visit at most m distinct states. As a result, in the games $\mathcal{M} \times \psi$, only positions containing states no deeper than m are reachable: $\mathcal{M} \times \psi$ and $\mathcal{M}^m \times \psi$ are identical. Since ψ is equivalent to ϕ, $\mathcal{M} \models \phi$ iff $\mathcal{M}^m \models \phi$. ◀

We can now show that if ϕ is of semantic modal depth m, then it is equivalent to the formula ϕ^m_m where fixpoints are first approximated to the m^{th} stage of induction as detailed in Definition 9 and then truncated at modal depth m as per Definition 11.

▶ **Theorem 15.** *If ϕ is guarded and of semantic modal depth m, then ϕ is equivalent to ϕ^m_m.*

Proof. Let ϕ be a guarded L_μ formula equivalent to a modal formula of modal depth m. Then $\mathcal{M} \models \phi$ iff $\mathcal{M}^m \models \phi$. However, ϕ agrees with ϕ_m on all trees of height at most m. Therefore the following are equivalent:
(1) $\mathcal{M} \models \phi$
(2) $\mathcal{M}^m \models \phi$
(3) $\mathcal{M}^m \models \phi_m$
(4) $\mathcal{M} \models \phi^m_m$
The conditions (1) and (2) are equivalent since ϕ is semantically modal of depth m, as per Lemma 14. Then (2) and (3) are equivalent since ϕ and ϕ_m have the same truth-value on $\mathcal{M}^m$, from Lemma 10. Finally, (3) and (4) are equivalent by definition of ϕ^m_m. ◀

Next we aim to show that m can be calculated from ϕ, using an argument similar to the one used by Otto [13]. The argument relies on labelling the states of structures with the subformulas of ϕ it satisfies and noting that the successors of a state can freely change as long as the set of successor-labels remains the same without affecting the formulas the state satisfies. The crux of the argument is that if two structures only differ at very high depth, but one satisfies ϕ and the other one does not, then the state labels must repeat themselves before the point at which the structures differ. Then we can duplicate a portion of the branch leading to the difference in order to create structures which are differentiated even deeper but still only one of them satisfies ϕ. This shows that if ϕ is modal, its modal depth cannot be deeper that the point at which the state labels need to start repeating themselves. $2^{2|\phi|}+1$ is an upper bound for that point.

The next lemma uses the fact that in an infinite tree, any subtree rooted at s can be replaced with a distinct subtree rooted at s' without affecting the subformulas of Ψ satisfied above depth s as long as the subtrees rooted at s and s' agree on all subformulas of Ψ. For a proof, see for example [6]. This should be clear from the notion that whether a state satisfies a subformula of Φ depends only on the propositional variables that state satisfies and the subformulas satisfied by its successor states.

▶ **Definition 16.** Let $\mathcal{M} = (M, i_M, E_M, P_M)$ be an infinite tree and t be a state of $\mathcal{M}$, and let $\Psi \in L_\mu$. We denote by $\alpha^\Psi_M(t)$ the set of subformulas of Ψ satisfied by the state t in $\mathcal{M}$.

▶ **Lemma 17** (Consistent labelling). *Let there be two disjoint trees, $\mathcal{M} = (M, i_M, E_M, P_M)$ and $\mathcal{M}' = (M', i_{M'}, E_{M'}, P_{M'})$, and a sentence $\Psi \in L_\mu$. Let s and s' be states of $\mathcal{M}$ and $\mathcal{M}'$ respectively, such that $\alpha^\Psi_M(s) = \alpha^\Psi_{M'}(s')$, and let v be the predecessor of s in M.*

Replace the edge e from v to s within $\mathcal{M}$ by a new edge e' from v to s' to obtain a new model $\mathcal{N}$ built from parts of $\mathcal{M}$ and $\mathcal{M}'$. More precisely, $\mathcal{N} = (N, i_N, E_N, P_N)$ with

- *$N = M \setminus \{u \in M \mid u$ extends or is equal to $s\} \cup \{u \in M' \mid u$ extends or is equal to $s'\}$ as set of states;*
- *$i_N = i_M$ as initial node;*
- *$E_N = (N \times N \upharpoonright E_M) \cup (N \times N \upharpoonright E_{M'}) \cup \{e'\}$ as set of edges, where $\upharpoonright$ denotes restriction; and*
- *$P_N = (P_M \upharpoonright N \cup P_{M'}) \upharpoonright N$ as propositional variables.*

Then, since $N \subseteq M \uplus M'$ (where $\uplus$ denotes disjoint union), the labelling $(\alpha^\Psi_M \cup \alpha^\Psi_{M'}) \upharpoonright N$ is defined on all states of $\mathcal{N}$ as well. Moreover, for all $s \in N$ we have $(\mathcal{N}, s) \models \phi$ if and only if $\phi \in ((\alpha^\Psi_M \cup \alpha^\Psi_{M'}) \upharpoonright N)(s)$, meaning that $(\alpha^\Psi_M \cup \alpha^\Psi_{M'}) \upharpoonright N$ is identical to α^Ψ_N.

▶ **Lemma 18.** *Let ϕ be guarded and semantically modal, i.e. ϕ is equivalent to a formula in ML. Then the semantic modal depth m of ϕ is bounded above by $2^{2|\phi|}+1$.*

Proof. Assume $m > 2^{2|\phi|}+1$ to be the semantic modal depth of ϕ. Then there exists a tree $\mathcal{M}$ of height $2^{2|\phi|}+1$ which is the prefix of two models $\mathcal{M}_1$ and $\mathcal{M}_2$ such that $\mathcal{M}_1 \models \phi$ and $\mathcal{M}_2 \not\models \phi$. That is to say, for every state s of $\mathcal{M}$, there are states s_1 and s_2 in $\mathcal{M}_1$ and $\mathcal{M}_2$ respectively such that s, s_1 and s_2 agree on propositions and for all inner nodes of $\mathcal{M}$, s' is a successor of s if and only if s'_1 is a successor of s_1, if and only if s'_2 is a successor of s_2. If d is maximal such that $\mathcal{M}_1$ and $\mathcal{M}_2$ agree up to depth d, write $agree(\mathcal{M}_1, \mathcal{M}_2) = d$. To start with, $agree(\mathcal{M}_1, \mathcal{M}_2) > 2^{2|\phi|}$ since $\mathcal{M}_1$ and $\mathcal{M}_2$ agree on their prefix $\mathcal{M}$ of rank $2^{2|\phi|}+1$.

Label every state s of $\mathcal{M}$ with a set $\alpha^\phi_{M_1}(s)$ consisting of subformulas of ϕ which are true in $\mathcal{M}_1$ and a set $\alpha^\phi_{M_2}(s)$ consisting of subformulas of ϕ which are true in $\mathcal{M}_2$. For each branch of $\mathcal{M}$, that is to say a path from the root of $\mathcal{M}$, if the branch is longer than $2^{2|\phi|}$, there are two states a, b in $\mathcal{M}$ along the branch such that $\alpha^\phi_{M_1}(a) = \alpha^\phi_{M_1}(b)$ and $\alpha^\phi_{M_2}(a) = \alpha^\phi_{M_2}(b)$. For each branch i, choose b^i to be the first state on a branch which has an ancestor a^i such that $\alpha^\phi_{M_1}(a^i) = \alpha^\phi_{M_1}(b^i)$ and $\alpha^\phi_{M_2}(a^i) = \alpha^\phi_{M_2}(b^i)$. Note that for any pair of branches i and j, either $b^i = b^j$ or b^i and b^j are not reachable from one another.

For each branch i and its states a^i and b^i, let ${a^i}'_1$ be the root of a distinct copy of the subtree in $\mathcal{M}_1$ rooted at a^i. Similarly, let ${a^i}'_2$ be the root of a distinct copy of the subtree rooted at a^i in $\mathcal{M}_2$. Let $\mathcal{M}'_1$ be obtained from $\mathcal{M}_1$ where for each branch i, the state b^i is replaced with ${a^i}'_1$ and its induced subtree; let $\mathcal{M}_2'$ be obtained from $\mathcal{M}_2$ where b^i is replaced with ${a^i}'_2$ and its induced subtree. Note that these transformations do not affect each other: recall that each b^i is on a distinct branch and is replaced with a subtree of the original structure. Since $\alpha^{\phi}_{M_1}({a^i}'_1) = \alpha^{\phi}_{M_1}(b^i)$ and $\alpha^{\phi}_{M_2}({a^i}'_2) = \alpha^{\phi}_{M_2}(b^i)$, all states preserve their labels and we know that $\mathcal{M}_1' \models \phi$ and $\mathcal{M}_2' \not\models \phi$, from Lemma 17.

We now show that if $\mathcal{M}_1$ and $\mathcal{M}_2$ agree up to depth d, then $\mathcal{M}_1'$ and $\mathcal{M}_2'$ agree up to depth $d+1$. Let i be a branch in $\mathcal{M}$ of length d such that i is extended differently in models $\mathcal{M}_1$ and $\mathcal{M}_2$. Since $depth(b^i) > depth(a^i)$ the models $\mathcal{M}_1'$ and $\mathcal{M}_2'$ agree along all extensions of i to depth $d - depth(a^i) + depth(b^i) > d$. That is to say $\mathcal{M}_1'$ and $\mathcal{M}_2'$ agree at least up to $d+1$. This establishes $agree(\mathcal{M}_1', \mathcal{M}_2') > agree(\mathcal{M}_1, \mathcal{M}_2)$. In $z = m - d$ many steps, we will reach models $\mathcal{M}_1^z$ and $\mathcal{M}_2^z$ such that $agree(\mathcal{M}_1^z, \mathcal{M}_2^z) \geq m$ but $\mathcal{M}_1^z \models \phi$ and $\mathcal{M}_2^z \not\models \phi$. This contradicts m being the modal depth of ϕ. ◀

▶ **Corollary 19.** *Whether a guarded formula ϕ is equivalent to a modal formula can be decided by testing whether ϕ is equivalent to $\phi^{2^{|\phi|}+1}_{2^{|\phi|}+1}$.*

Proof. From the previous lemma, if a formula ϕ is modal, its semantic modal depth m is no greater than $2^{2|\phi|} + 1$. If $\phi \neq \phi^{2^{|\phi|}+1}_{2^{|\phi|}+1}$, then ϕ must disagree with $\phi^{2^{|\phi|}+1}_{2^{|\phi|}+1}$ on some structure $\mathcal{M}$. However, the two model checking games $\mathcal{M} \times \phi$ and $\mathcal{M} \times \phi^{2^{|\phi|}+1}_{2^{|\phi|}+1}$ are identical on plays which do not reach states deeper than $2^{|\phi|} + 1$, which, since $m \leq 2^{|\phi|} + 1$, contradicts the fact that the modal depth of ϕ is m. ◀

The most surprising aspect of this result is perhaps the exponential modal depth. This is not due to the authors' laziness: formulas with fixpoints can indeed be at least sub-exponentially more compact than the equivalent modal formulas. The following exhibits syntactically Σ^{μ}_1-formulas but semantically modal formulas with sub-exponential modal depth. The idea of these formulas is to require a series of propositional variables to occur at different frequencies until they all occur at the same time. The modal depth of the formula is then the least common multiple of the frequencies.

▶ **Example 20.** There is a family of formulas $\Phi_n \in \Sigma^{\mu}_1$ which are semantically modal but have modal depth $\Omega(2^{\sqrt{n}})$ in the length of Φ_n.

Proof. Write $\Box^n$ for $\Box ... \Box$ repeated n times. The formula $\mu X.(A \wedge (\Box^n X \vee B)$ states that A occurs every n^{th} state on any path until B also occurs at a state whose depth is a multiple of n. By combining such formulas we can write $[\Box^a \mu X.A \wedge (\Box^a X \vee (B \wedge C))] \wedge [\Box^b \mu X.B \wedge (\Box^b X \vee (A \wedge C))] \wedge [\Box^c \mu X.C \wedge (\Box^c X \vee (A \wedge B))]$ which sets the frequencies at which A, B and C are seen until they are seen simultaneously. This formula is modal since if it is true, at the latest at depth $a \times b \times c$, all of A, B and C are seen simultaneously. More precisely, its modal depth is the least common multiple of a, b and c. Generalising this, for a fixed n, let $\psi_d = \mu X.\Box^d(P_d \wedge X) \vee (\bigwedge_{i \leq n} P_i)$ be the formula stating that the proposition P_d occurs at frequency d until all propositions P_i for $i \leq n$ occur at the same time, at a depth multiple of d. Now, let $\Phi_n = \bigwedge_{d \leq n} \psi_d$. The modal depth of Φ_n is the least common multiple of the integers up to n, written $lcm(n)$. For sufficiently large n, $lcm(n) > 2^n$ [12] so the formula Φ_n is of length $O(n^2)$ and has modal depth $\Omega(2^n)$ which proves the correctness of the example. ◀

4 The formulas for Π^{μ}_1 and Σ^{μ}_1

The previous section addressed how to eliminate accidental complexity from semantically modal formulas. This section studies the same question for Π^{μ}_1, the class of properties expressible without least fixpoint operators, and its dual, Σ_1^μ. Küsters and Wilke [8] showed that it is decidable whether a formula is equivalent to a Π_1^μ formula; this section constructs the desired formula, yielding an alternative decision procedure for $\Pi_1^\mu$and Σ_1^μ. We first formalise the idea that if a property is in Σ_1^μ, then some finite initial tree is always sufficient to show that a structure satisfies the property. We then introduce disjunctive form. The final subsection shows how unnecessary fixpoints can be eliminated syntactically from formulas in disjunctive form by using the fact that Σ_1^μ formulas have finite proofs.

4.1 Properties in Π_1^μ have finite counter-proofs

In this section we characterise properties in Σ_1^μ and Π_1^μ as properties with finite proofs and counter-proofs respectively. Informally, μ-formulas express finite behaviour such as reachability – proofs of such properties are finite: once the desired state is reached, the rest of the structure is irrelevant. Dually, ν-formulas express infinite behaviour and if a structure fails to display infinite behaviour, the state at which it fails must be finitely reachable.

▶ **Lemma 21.** *Let $\mathcal{M}$ be a structure with finite branching such that $\mathcal{M} \not\models \Psi$. If $\Psi \in \Pi_1^\mu$ then there is some n such that for any structure $\mathcal{M}'$, if $\mathcal{M}'$ agrees with $\mathcal{M}$ up to depth n, then $\mathcal{M}' \not\models \Psi$.*

Proof. Assume Ψ is semantically in Π_1^μ and Φ is the equivalent formula with no least fixpoints. Since $\mathcal{M} \not\models \Phi$, Even has a winning strategy in $\mathcal{M} \times \neg\Phi$. Note that $\neg\Phi$ is a formula without greatest fixpoints. That means that Even has a strategy σ winning in $\mathcal{M} \times \neg\Phi$ which only agrees with finite plays. Let n be the depth of the furthest state in $\mathcal{M}$ which σ reaches – since $\mathcal{M}$ has finite branching, there is such an n. Note that agreement between $\mathcal{M}$ and $\mathcal{M}'$ up to $n+1$ requires any leaves at depth less than n in $\mathcal{M}$ to remain leaves in $\mathcal{M}'$. Now for any $\mathcal{M}'$ which agrees with $\mathcal{M}$ up to $n+1$, the strategy σ is still winning for Even so $\mathcal{M} \not\models \Psi$. ◀

4.2 Disjunctive form

Disjunctive form was introduced in [4] as a syntactic restriction to universal branching. It had been used for example to show the completeness of Kozen's axiomatisation [14]. Here we show that disjunctive forms are also a tool for simplifying syntactic manipulations. Informally, the idea of disjunctive form is to push conjunctions into the leaves and allow player Odd to make exactly one choice per state.

▶ **Definition 22.** *(Disjunctive formulas)* The set of disjunctive form formulas of (unimodal) L_μ is the smallest set $\mathcal{F}$ satisfying:

- Propositional variables and their negations, fixpoint variables and $\top$ and $\bot$ are in $\mathcal{F}$;
- If $\psi \in \mathcal{F}$ and $\phi \in \mathcal{F}$, then $\psi \vee \phi \in \mathcal{F}$;
- If $\mathcal{A}$ is a set of literals and $\mathcal{B} \subseteq \mathcal{F}$ ($\mathcal{B}$ is finite), then $\bigwedge \mathcal{A} \wedge \rightarrow\mathcal{B}$ where $\rightarrow\mathcal{B}$ is short for $(\bigwedge_{\psi\in\mathcal{B}} \Diamond\psi) \wedge \Box \bigvee_{\psi\in\mathcal{B}} \psi$ – that is to say, every formula in $\mathcal{B}$ is realised by at least one successor and every successor realises at least one of the formulas in $\mathcal{B}$;
- $\mu X.\psi$ and $\nu X.\psi$ are in $\mathcal{F}$ as long as $\psi \in \mathcal{F}$ and X only appears positively and never in a conjunction $X \wedge \alpha$ where α is another formula.

The last constraint ascertains that if $\mu X.\phi(X)$ is in disjunctive form, then $\phi(\mu X.\phi(X))$ is also in disjunctive form.

Every formula is known to be equivalent to an effectively computable formula in disjunctive form [14]. The transformation preserves guardedness.

We can now prove our key lemma about formulas in disjunctive form which exploits the restriction imposed on player Odd's choices. With an arbitrary L_μ formula, once Even has fixed a strategy, Odd may be able to choose to play to a state s at various different formulas. For example, from $(s, \Box\psi \wedge \Box\phi)$ Odd can choose to play any successor of s at either ϕ or ψ. However, with some minor assumption about the structure, once a formula is in disjunctive form and Even has fixed her strategy, Odd is much more restricted in his choices: if he chooses to play to a state s, he can only choose to play a formula fixed by Even's strategy or a literal.

First we define well behaved strategies and models in which whenever a state s is required to satisfy a modal formula $\to\mathcal{B}$, s has a distinct successor for each formula in $\mathcal{B}$. Such a well-behaved model can easily be derived from any model by duplicating the successor states of s as necessary. A well-behaved model will allow the even player to use a well-behaved strategy which chooses a distinct successor for each of the formulas Odd can choose at $\to\mathcal{B}$.

▶ **Definition 23** (Well behaved models)**.** A model $\mathcal{M}$ of Ψ is well behaved with respect to an Even's winning strategy σ in $\mathcal{M} \times \Psi$ if for each position $(s, \to\mathcal{B})$ reachable with σ, s has distinct successors s_ϕ such that $s_\phi \models \phi$ for each $\phi \in \mathcal{B}$ and σ plays s_ϕ if Odd picks ϕ from s and ϕ if Odd picks s_ϕ. Note that s may have more than one successor satisfying ϕ but σ chooses only one such successor, s_ϕ to play to whenever Odd chooses to play ϕ. A model is well-behaved if it is well-behaved with respect to some winning strategy.

Every model $\mathcal{M}$ of Ψ is bisimilar to a well behaved model of Ψ obtained by duplicating successor states as necessary. A strategy is said to be well-behaved if the model is well-behaved with respect to that strategy. Next we define the tree of Odd's playable positions induced by Even's strategy. This tree consists of the choices which Odd is left with once Even has fixed her strategy.

▶ **Definition 24** (Odd's position tree)**.** If σ is a strategy for Even, we consider the tree made out of positions belonging to Odd which are reachable by plays respecting σ. One step in the Odd's position tree corresponds to one move by Odd followed by as many moves dictated by σ as necessary to reach the next position belonging to Odd. Note that since σ is Even's strategy, Odd's position tree does not have any disjunctive positions any more, only conjunctions: all the non-leaf positions of this tree are of the form $(s, \bigwedge\mathcal{A} \wedge \to\mathcal{B})$ for some set of literals $\mathcal{A}$ and a set of formulas $\mathcal{B}$. The leaves are of the form (s, A) where A is a literal.

The following lemma shows how the syntactic constraints of disjunctive form simplify the strategies in the parity game. It is the key to our proof of decidability. It states that once Even has fixed her strategy, Odd can only reach a state s at a single formula $\bigwedge\mathcal{A} \wedge \to\mathcal{B}$. This will allow us to replace s with the root of any structure which satisfies the same formula, while preserving Even's winning strategy.

▶ **Lemma 25.** *If Ψ is in disjunctive form and $\mathcal{M}$ is the tree-representation of a well-behaved model with respect to a strategy σ, then each state of $\mathcal{M}$ appears in Odd's position tree for σ at most once at a non-leaf position.*

Proof. For a state s to appear twice at such a formula, (s, ϕ_0) and (s, ϕ_1) must be two positions of the tree with a last common ancestor, (t, ψ) where ψ has to be $\to\mathcal{B}$ for some $\mathcal{B}$ containing ϕ_0 and ϕ_1 . However, since σ is well behaved, if Odd chooses either $\phi \in \mathcal{B}$ or the successor t_ϕ, the game goes to (t_ϕ, ϕ) so each successor t_ϕ only appears in one successor position of (t, ψ). Furthermore, each non-indexed successor of t also appears in only one successor position of (t, ψ). This can therefore not be the last common ancestor of (s, ϕ_0) and (s, ϕ_1). ◀

4.3 The formula $\Psi^{\Pi_1^\mu}$

This section proves the main theorem on the constructive decidability of Π_1^μ: any semantically Π_1^μ formula in disjunctive form can be transformed into an equivalent syntactically Π_1^μ formula by changing every occurrence of μ into either ν or $\bot$.

To show this, we first select for each μ-subformula $\mu X.\phi$ in Ψ, a structure $\mathcal{M}$ such that whether $\mathcal{M}$ satisfies Ψ or not depends on a restricted set of states satisfying the formula $\mu X.\phi$, as shown in Lemma 26. We then show in Lemma 27 that if a μ-subformula $\mu X.\phi$ of Ψ is satisfiable but cannot be replaced with the corresponding ν-formula $\nu X.\phi$, then for any n we can build a twin structure for $\mathcal{M}$ agreeing with $\mathcal{M}$ up to n but disagreeing on Ψ. This implies that Ψ is not a Π_1^μ formula, as Lemma 21 shows that Π_1^μ formulas have finite counter-proofs. This leaves us with two scenarios: either the μ-subformula is unsatisfiable, in which case it can be replaced by $\bot$, using Lemma 28, or the μ-formula can be replaced with the corresponding ν-formula. In either case, we can turn any semantically Π_1^μ formula into a syntactically Π_1^μ formula by replacing μ-subformulas with either $\bot$ or the dual ν-formula.

Notation. Let $\Psi(\psi)$ be a formula in disjunctive form which contains a subformula ψ. We will write $\Psi(\psi')$ for the formula in which ψ is substituted with the formula ψ'. With this notation, we will use formulas related to Ψ in order to specify structures where the players' strategies must exhibit some desired behaviours. For example, if for some structure $\mathcal{M}$, Odd can win $\mathcal{M} \times \Psi(\bot)$, then Even can only win $\mathcal{M} \times \Psi(\mu X.\phi)$ by playing eventually to $\mu X.\phi$.

In the following lemma we show that for a structure to satisfy $\neg\Psi(\phi) \wedge \Psi(\top)$ means that Odd can win the game for $\Psi(\phi)$, but only by playing to a position (s, ϕ) for some s in a set S. Then, if states of S are substituted with new substructures, Odd may only win if he can win from one of the new substructures at ϕ.

▶ **Lemma 26.** *Let Ψ be a formula in disjunctive guarded form with a subformula ϕ. If $\mathcal{M}$ is a structure such that $\mathcal{M} \models \neg\Psi(\phi) \wedge \Psi(\top)$ and $\mathcal{M}$ is well-behaved for $\Psi(\top)$, then there is a non-empty set of states S in $\mathcal{M}$ such that in $\mathcal{M} \times \Psi(\phi)$ each of Odd's winning strategies reaches (s, ϕ) for some $s \in S$– that is to say, for each of Odd's winning strategies τ there is a counter strategy σ such that (s, ϕ) is on the play $\tau \times \sigma$ for some $s \in S$. Furthermore, if every state s_i of S is replaced with some state t_i, yielding a new model $\mathcal{M}'$, Odd only wins in $\mathcal{M}' \times \Psi(\phi)$ if Odd wins from (t_i, ϕ) in the same game for some t_i.*

Proof. If Even wins $\mathcal{M} \times \Psi(\top)$ but Odd wins $\mathcal{M} \times \Psi(\phi)$, then Odd cannot win $\mathcal{M} \times \Psi(\phi)$ with a strategy which avoids ϕ, otherwise the same strategy would be winning in $\mathcal{M} \times \Psi(\top)$. Let τ be one of Odd's winning strategies in $\mathcal{M} \times \Psi(\phi)$ and let σ be Even's well-behaved winning strategy in $\mathcal{M} \times \Psi(\top)$. Since $\mathcal{M} \times \Psi(\top)$ is identical to $\mathcal{M} \times \Psi(\phi)$ until a play reaches ϕ, the strategy σ is also an initial strategy in $\mathcal{M} \times \Psi(\phi)$, defined until ϕ is reached. The play $\tau \times \sigma$ must reach ϕ because otherwise it would have to be winning for Even due to it being identical to a play respecting her winning strategy in $\mathcal{M} \times \Psi(\top)$. Let s_τ be the first state at which the play $\tau \times \sigma$ reaches ϕ in $\mathcal{M} \times \Psi(\phi)$. Then $S = \{s_\tau | \tau \text{ is a winning strategy for Odd }\}$ is the set such that in $\mathcal{M} \times \Psi(\phi)$ each of Odd's winning strategies reaches (s, ϕ) for some $s \in S$.

For the second part of the lemma, first observe that if Even wins from (t_i, ϕ) for all i, then Odd cannot use any of his winning strategies from $\mathcal{M} \times \Psi(\phi)$ to win in $\mathcal{M}' \times \Psi(\phi)$ since if Even initially plays according to σ, the play reaches (t_i, ϕ) from where Even has a winning strategy. As a result, Odd cannot avoid all t_i without losing. From Lemma 25 we know that each t_i is only seen at position (t_i, ϕ) so not only can Odd not avoid all t_i, Odd cannot avoid all (t_i, ϕ) without losing. Hence, if Odd loses from (t_i, ϕ) for all i, Odd loses in $\mathcal{M}' \times \Psi(\phi)$. ◀

We can now prove the main result: to obtain the syntactically Π_1^μ formula equivalent to a semantically Π_1^μ formula in guarded disjunctive form, it is sufficient to replace each least fixpoint with either $\bot$ or a greatest fixpoint. The crux is to show that each μ-binding in a semantically Π_1^μ formula can either be replaced by $\bot$ or ν. The following lemma identifies two cases. The first is that the subformula $\mu X.\phi$ is unsatisfiable in the sense that there is no structure $\mathcal{T}$ from the root of which Even can win at $\mu X.\phi$ in $\mathcal{T} \times \Psi(\mu X.\phi)$. Then it can be replaced with $\bot$. In the other case, $\mu X.\phi$ can be replaced with $\nu X.\phi$.

▶ **Lemma 27.** *If $\Psi(\mu X.\phi)$, a guarded formula in disjunctive form with a subformula $\mu X.\phi$, is semantically in Π_1^μ, then either there is no structure $\mathcal{T}$ such that Even wins from $(r_0, \mu X.\phi)$ in $\mathcal{T} \times \Psi(\mu X.\phi)$ where r_0 is the root of $\mathcal{T}$, or $\Psi(\mu X.\phi)$=$\Psi(\nu X.\phi)$.*

Proof. Assume that $\Psi(\mu X.\phi) \neq \Psi(\nu X.\phi)$ and that there is a structure $\mathcal{T}$ such that Even wins from $(r_0, \mu X.\phi)$ in $\mathcal{T} \times \Psi(\mu X.\phi)$ where r_0 is the root of $\mathcal{T}$. Since $\Psi(\mu X.\phi)$ implies $\Psi(\nu X.\phi)$ but not the other way around, then there is a structure $\mathcal{M}$ such that $\mathcal{M} \models \neg\Psi(\mu X.\phi) \wedge \Psi(\nu X.\phi)$ and $\mathcal{M}$ is well-behaved with respect to $\Psi(\nu X.\phi)$. Recall that we require for $\mathcal{M}$ to be finitely branching. We will show that for any n there is a structure $\mathcal{M}'$ which agrees with $\mathcal{M}$ up to depth n but which satisfies $\Psi(\mu X.\phi)$. Using Lemma 21, this will contradict $\Psi(\mu X.\phi) \in \Pi_1^\mu$.

For any n, we can write $\Psi(\mu X.\phi)$ as $\Psi(\overbrace{\phi...\phi}^{n}(\mu X.\phi))$, where we drop some brackets for readability, so $\phi\phi(X)$ should be understood as $\phi(\phi(X))$. Then, the structure $\mathcal{M}$ satisfies $\neg\Psi(\overbrace{\phi...\phi}^{n}(\mu X.\phi)) \wedge \Psi(\overbrace{\phi...\phi}^{n}(\top))$ since $\overbrace{\phi...\phi}^{n}(\top)$ is implied by $\nu X.\phi$. Furthermore, $\mathcal{M}$ is well-behaved for $\Psi(\overbrace{\phi...\phi}^{n}(\top))$. From Lemma 26 we know that there is a set S of states in $\mathcal{M}$ such that for each $s_i \in S$, Even loses from $(s, \neg\mu X.\phi)$ and if each $s_i \in S$ is replaced with r_0, the root of $\mathcal{T}$, to yield a new model $\mathcal{M}'$, then Ψ holds in $\mathcal{M}'$. Furthermore, since X is guarded in $\mu X.\phi$, a play can only reach $\mu X.\phi$ from $\overbrace{\phi...\phi}^{n}(\mu X.\phi)$ at depth at least n: each $s_i \in S$ is at least at depth n therefore $\mathcal{M}'$ agrees with $\mathcal{M}$ up to depth n. We have built for any n, a structure that agrees with $\mathcal{M}$, a counter-model of Ψ, up to n but satisfies Ψ. This contradicts the assumption that Ψ is in Π_1^μ, and has finite counter-proofs using Lemma 21. ◀

It now suffices to show that if a subformula $\mu X.\phi$ is unsatisfiable in the sense that there is no structure $\mathcal{T}$ from the root of which Even can win at $\mu X.\phi$ in $\mathcal{T} \times \Psi(\mu X.\phi)$, then $\mu X.\phi$ can be replaced with $\bot$. This should be intuitively justified by the idea that in no structure can Even win by playing to $\mu X.\phi$, so it is no worse for her to have $\bot$ instead.

▶ **Lemma 28.** *If there is no structure $\mathcal{T}$ rooted at t_0 such that Even wins from $(t_0, \mu X.\phi)$ in $\mathcal{T} \times \Psi(\mu X.\phi)$, then $\Psi(\mu X.\phi) = \Psi(\bot)$.*

Proof. If Even wins $\mathcal{M} \times \Psi(\mu X.\phi)$ but there is no $\mathcal{T}$ rooted at t_0 such that Even wins from $(t_0, \mu X.\phi)$, then Even's winning strategy cannot reach any position $(s, \mu X.\phi)$. Then the same strategy can be used in $\mathcal{M} \times \Psi(\bot)$ to avoid any position $(s, \bot)$. Since these two games are identical up until $\mu X.\phi$ or $\bot$ is reached, Even also wins in $\mathcal{M} \times \Psi(\bot)$. This shows $\Psi(\mu X.\phi) \implies \Psi(\bot)$. The other direction is trivial since $\bot \implies \mu X.\phi$ and L_μ is monotone. ◀

▶ **Theorem 29.** *If Ψ is a formula in guarded disjunctive form and semantically in Π_1^μ, then either $\Psi = \Psi[\bot/\mu X.\phi]$ or $\Psi[\nu X.\phi/\mu X.\phi]$ for any subformula $\mu X.\phi$ of Ψ.*

Proof. If there is no structure $\mathcal{T}$ such that Even wins from $(r_0, \mu X.\phi)$ in $\mathcal{T} \times \Psi(\mu X.\phi)$ where r_0 is the root of $\mathcal{T}$, then from the previous lemma, $\Psi = \Psi[\bot/\mu X.\phi]$. If there is such a structure, then from Lemma 27 we know that $\Psi = \Psi[\nu X.\phi/\mu X.\phi]$. ◀

▶ **Corollary 30.** Π_1^μ *and by duality* Σ_1^μ *are decidable.*

Proof. Any formula Ψ of L_μ can be turned into a guarded formula in disjunctive guarded form. Then, if Ψ is semantically in Π_1^μ, every occurrence of $\mu X.\phi$ can be eliminated either by replacing it with $\bot$ or $\nu X.\phi$. Hence to decide whether a formula is semantically in Π_1^μ, it is sufficient to decide whether it is equivalent to the formula where each $\mu X.\phi$ formula reachable by Even in the game for $\Psi(\mu X.\phi)$ is replaced with $\nu X.\phi$.

By duality, to decide whether a formula is semantically in Σ_1^μ it is sufficient to decide whether its negation is in Π_1^μ. If this is the case, the Π_1^μ formula can be syntactically negated to yield a formula in Σ_1^μ. ◀

5 Conclusion

We have defined syntactic transformations from L_μ into *ML*, Π_1^μ and Σ_1^μ which preserve meaning for formulas which are semantically, but not yet syntactically in the target class. A straight-forward corollary of this result is an alternative decision procedure for the low levels of the alternation hierarchy: to decide whether a L_μ formula is in Π_1^μ,Σ_1^μ or *ML*, it suffices to check whether it is equivalent to its projection into that class.

For the modal fragment of L_μ, the transformation we describe incurs a potentially exponential blow-up in the size of the formula – as such, it may be more concise to represent a formula with some fixpoints. This blow-up is however necessary since fixpoint formulas which are semantically modal can have at least subexponential modal depth.

For Π_1^μ on the other hand, assuming formulas are in guarded disjunctive form, the target formula is no larger than the original one. The transformation into guarded disjunctive form itself can incur an exponential blow-up, not least because the transformation involves distributing conjunctions over disjunctions, causing duplication.

This result is of both practical and theoretical interest since the complexity of model checking depends on the *syntactic* alternation depth of a formula, rather than the semantic one. Thus for formulas that are semantically in a low alternation class, this transformation can potentially turn an exponential model-checking procedure into a polynomial one by eliminating the exponent. By providing a concise formula in the semantic alternation class of a formula, our method provides an appealing pre-processing step for model checking.

References

1 J.C. Bradfield. The modal mu-calculus alternation hierarchy is strict. In Ugo Montanari and Vladimiro Sassone, editors, *CONCUR'96: Concurrency Theory*, volume 1119 of *Lecture Notes in Computer Science*, pages 233–246. Springer Berlin Heidelberg, 1996.

2 J.C. Bradfield and C. Stirling. Modal mu-calculi. *Handbook of modal logic*, 3:721–756, 2007.

3 E. A. Emerson and C. S. Jutla. Tree automata, mu-calculus and determinacy. In *Foundations of Computer Science, 1991. Proceedings., 32nd Annual Symposium on*, pages 368–377. IEEE, 1991.

4 D. Janin and I. Walukiewicz. Automata for the modal μ-calculus and related results. In *Proc. MFCS'95 LNCS 969*, pages 552–562, 1995.

5 D. Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27(3):333–354, 1983. Special Issue Ninth International Colloquium on Automata, Languages and Programming (ICALP) Aarhus, Summer 1982.

6 D. Kozen. A finite model theorem for the propositional μ-calculus. *Studia Logica*, 47(3):233–241, 1988.

7 O. Kupferman, M. Y. Vardi, and P. Wolper. An automata-theoretic approach to branching-time model checking. *Journal of the ACM (JACM)*, 47(2):312–360, 2000.

8 R. Küsters and T. Wilke. Deciding the first level of the μ-calculus alternation hierarchy. In *Proc. FSTTCS 2002, LNCS 2556*, page 241–252, 2002.

9 R. Küsters and T. Wilke. Deciding the First Level of the μ-calculus Alternation Hierarchy. Technical Report 0209, Institut für Informatik und Praktische Mathematik, CAU Kiel, Germany, 2002.

10 G. Lenzi. A hierarchy theorem for the μ-calculus. In Friedhelm Meyer and Burkhard Monien, editors, *Automata, Languages and Programming*, volume 1099 of *Lecture Notes in Computer Science*, pages 87–97. Springer Berlin Heidelberg, 1996.

11 R. Mateescu. Local model-checking of modal mu-calculus on acyclic labeled transition systems. In Joost-Pieter Katoen and Perdita Stevens, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, volume 2280 of *Lecture Notes in Computer Science*, pages 281–295. Springer Berlin Heidelberg, 2002.

12 M. Nair. A new method in elementary prime number theory. *Journal of the London Mathematical Society*, 2(3):385–391, 1982.

13 M. Otto. Eliminating recursion in the μ-calculus. In *STACS 99*, pages 531–540. Springer, 1999.

14 I. Walukiewicz. Completeness of Kozen's Axiomatisation of the Propositional μ-Calculus. *Information and Computation*, 157(1–2):142–182, 2000.