The descriptive complexity of L_μ model-checking parity games

M. K. Lehtinen
University of Kiel
Kiel, Germany
University of Edinburgh
Edinburgh, United Kingdom
kleh@informatik.uni-kiel.de

This paper revisits the well-established relationship between the modal μ calculus L_μ and parity games to show that it is even more robust than previously known. It addresses the question of whether the descriptive complexity of L_μ model-checking games, previously known to depend on the syntactic complexity of a formula, depends in fact on its semantic complexity. It shows that up to formulas of semantic complexity Σ_2^μ, the descriptive complexity of their model-checking games coincides exactly with their semantic complexity. Beyond Σ_2^μ, the descriptive complexity of the model-checking parity games of a formula Ψ is shown to be an upper bound on the semantic complexity of Ψ; whether it is also a lower bound remains an open problem.

1 Introduction

The modal μ-calculus [9], written L_μ, is a verification logic consisting of a simple modal logic augmented with a least fixpoint μ and its dual ν. It is a prime example of the intersection between logic and games. In its model-checking games, the antagonism between the verifying player Even and her opponent Odd describes the duality between conjunctions and disjunctions, and between least and greatest fixpoints. The complexity of the winning condition – a parity condition over a set of priorities – corresponds to the complexity of the formula, as measured by its index, *i.e.* the number of alternations between least and greatest fixpoints.

The index of a formula is a robust measure of a formula's complexity: for each fixed index, the model-checking problem is in P, but no fixed index suffices to express all L_μ-expressible properties [1]. The decidability of the *semantic index* of a formula, that is to say the least index of any equivalent formula, is a long-standing open problem. Despite the efforts put towards solving it and its automata-theoretic counterpart, the Mostowski–Rabin index problem for parity automata, only the low levels of the alternation hierarchy are currently known to be decidable. Formulas semantically in *ML*, the fragment without fixpoints, and those in Π_1^μ and Σ_1^μ, the fragments with only one type of fixpoint, were first shown to be decidable by Otto [20] and Küsters and Wilke [12] respectively. Alternative formula-focused decidability proofs for both were derived in [13]. Some results in automata theory [4] can be interpreted within the L_μ context to yield another decidability result: given a formula in the Π_2^μ alternation class, it is decidable whether it is equivalent on ranked trees to a formula in the Σ_2^μ alternation class. An alternative proof with a topological account was given in [21]. This result extends onto unranked trees: Σ_2^μ is decidable for Π_2^μ [15].

This paper revisits one of the fundamental concepts of L_μ literature in light of the index-problem: the relationship between L_μ formulas and parity games. Parity games can be used to describe the semantics

Submitted to:
GandALF 2017

of L_μ and are therefore a recurring topic in L_μ literature. This relationship is complexity-preserving, in the following sense: On the one hand, the model-checking games for L_μ formulas of index I are parity games with priorities from I; on the other, a L_μ formula of index I suffices to describe the winning regions of such games. As a result, the descriptive complexity of the model-checking games of a formula Ψ is bounded by the index of Ψ. This paper asks whether this complexity-preserving relationship can be extended to also account for the *semantic* complexity of a formula: can the bounds on the complexity of model-checking games be tightened to the semantic index of Ψ?

In its first part (Section 3), this paper argues that if the winning regions of the model-checking games for Ψ are described by Φ – written Φ interprets Ψ – then Ψ is semantically as simple as Φ, *i.e.* equivalent to a formula of the same index as Φ. In other words, the descriptive complexity of the model-checking games of a formula is an upper bound on the semantic complexity of the formula. Besides completing our understanding of the L_μ– parity game correspondance, this result opens up a new route to simplyfing L_μ formulas: for some formulas (such as Example 15) it is easy to show that they only generate model-checking games of low descriptive complexity; it follows that such formulas are semantically simple.

The second part of the paper (Section 4) considers the converse: if a formula is semantically of index I, then is it interpreted by a formula of index I? That is to say, do semantically simple formulas generate equally simple model-checking games? This is shown to be the case for the first semantic levels of the alternation hierarchy. In some cases, the input formula must be transformed into disjunctive form [8] first, a normal form for L_μ:

- If a L_μ formula is equivalent to a modal formula, it is interpreted by a modal formula;
- If a disjunctive L_μ formula is equivalent to a Π^μ_1 formula, it is interpreted by a Π^μ_1 formula;
- If a disjunctive L_μ formula is equivalent to a Σ^μ_2 formula, it is interpreted by a Σ^μ_2 formula;

These theorems are based on characterising the formulas semantically in ML, Π^μ_1 and Σ^μ_2 as those for which the model-checking games coincide with variations of parity games that characterise the class. For these (semantic) alternation-classes at least, the descriptive complexity of the model-checking games of a formula Ψ corresponds exactly to the semantic index of Ψ, rather than its syntactic index.

Section 4.4 considers the general case and shows that if a *co-disjunctive* formula is equivalent to a *disjunctive* formula of index I, then it is also interpreted by a disjunctive formula of index I: the complexity of an equivalent disjunctive formula bounds the descriptive complexity of the model-checking games generated by the co-disjunctive form of an input formula. I conjecture that the same holds for disjunctive input formulas, but the currently available techniques do not suffice to prove this.

In short, this paper connects two important concepts of L_μ theory: the index-problem and the relationship between L_μ and parity games. It shows that the latter can be extended in a natural way to account for the former: the descriptive complexity of the model-checking games of Ψ is an upper bound on the semantic complexity of Ψ. For several non-trivial cases, and in particular beyond the alternation classes currently known to be decidable, the converse also holds: a formula of semantic index $I \in \{\{\}, \{0\}, \{1,0\}\}$ has model-checking games of descriptive complexity exactly I. In the general case, the index of an equivalent disjunctive formula bounds the descriptive complexity of a co-disjunctive formula. Overall, the L_μ–parity game relationship is even more robust than previously thought.

The practical consequence of this insight is that some formulas can be simplified by analysing the model-checking games they generate (see Example 15). On the theoretical side, the study of the descriptive complexity of the model-checking games is a novel approach to the L_μ-index problem. Although the three decision procedures for ML, Π^μ_1 and Σ^μ_2 (for Π^μ_2 formulas) each uses a completely different strategy, they can all be reformulated as ways to describe the winning regions of formulas semantically in the target class. Extending the interpretation theorems of this paper seems likely to lead to characterisations of higher L_μ alternation classes, providing a stepping stone towards further decidability results. In particular, a generalisation of the Σ_2^μ interpretation theorem would yield a reduction of the decidability of the index problem to a boundedness question, in the style of what Colcombet *et. al.* achieved for non-deterministic automata [5].

The descriptive complexity of parity games has been considered in [6] where the authors ask which formalisms can describe the winning regions of classes of parity games without a fixed index. They show that guarded second order logic suffices in the general case while least fixed point logic does not. On finite game arenas, the winning regions are definable in least fixed point logic if and only if solving parity games is in P. In contrast, the winning regions of the classes of parity games considered here are trivially L_μ expressible, and the focus is on exactly which fragment of L_μ is necessary.

The following section fixes the notation and terminology used throughout. Some familiarity with L_μ and parity games is assumed – see for example [3] for an introduction.

2 Preliminairies

2.1 L_μ syntax

Fix countably infinite sets $Prop = \{P, Q, \ldots\}$ of propositional variables, and $Var = \{X, Y, \ldots\}$ of fixpoint variables. For the sake of clarity and conciseness the scope of this paper is restricted to the unimodal L_μ; however, the results presented here extend easily to the multi-modal case.

Definition 1. *(L_μ)* The syntax of unimodal L_μ is given by:

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

Conjunctions take precedence over disjunctions. The scope of fixpoint bindings extends as far as possible to the right while the scope of modalities extends as little as possible to the right. For example, $\mu X.\Diamond X \wedge C \vee B$ is parsed as $\mu X.(((\Diamond X) \wedge C) \vee B)$. The syntax of L_μ presented here only allows negation to be applied to propositional variables.

A fixpoint variable X which does not appear in the scope of $\mu X.\phi$ or $\nu X.\phi$ is said to be *free*, and the set of free fixpoint variables of a formula ψ is written $Free_\psi$. A sentence is a formula without free fixpoint variables. A fixpoint μ or ν *binds* a variable $X \in Free_\phi$ in $\mu X.\phi$ or $\nu X.\phi$ respectively. The formula ϕ is called the binding formula of X, and will be written ϕ_X. For notational purposes, the binding formula of X will be treated as an immediate subformula of X. The *parse tree* of a sentence is the tree with the subformulas of Ψ for nodes, rooted at the formula itself, and where the children of a node consist of the parse-trees of its immediate subformulas. In other words it consists of the formula, written as a tree, with back-edges from fixpoint variables to their bindings.

A formula is *guarded* if every fixpoint variable is in the scope of a modality within its binding. Every L_μ formula is equivalent to a formula in guarded form [18, 11]. Without loss of expressivity, all L_μ formulas are assumed to be in guarded form.

Disjunctive form is a normal form for L_μ formulas [8] which imposes some additional structure onto the L_μ syntax. The dual definition yields co-disjunctive formulas. Every L_μ can be turned into a disjunctive formula, although the transformation may not preserve complexity [14]. Unlike in general L_μ model-checking games, in the model-checking games of disjunctive formulas, the verifying player has strategies which only agree with one play per branch [13].

2.2 Alternations

The definition of alternation depth, here referred to as index to match the automata-theoretic terminology, is meant to only capture alternations which generate algorithmic complexity, matching the definition given in Niwiński 1986 [19]. The presentation here emphasizes the relationship of a formula's index to the priorities in parity games. A thorough discussion on how to define alternation depth, and comparison of definitions used in the literature can be found in Bradfield and Stirling 2007 [3].

Definition 2. *(Priority assignment and index)* A priority assignment for a sentence Ψ is a mapping $\Omega : Var_\Psi \to \{m,...,q-1,q\}$, where $m \in \{0,1\}$ and q is a positive integer, of priorities to the fixpoint variables Var_Ψ of Ψ such that μ-bound variables receive odd priorities while ν-bound variables receive even priorities. A priority assignment is order preserving if whenever X is free in the formula binding Y, $\Omega(X) \geq \Omega(Y)$ holds.

A formula has index I if it has an order preserving priority assignment with co-domain I. A formula has semantic index I if it is equivalent to a formula of index I.

Definition 3. *(Alternation hierarchy)* The base of the alternation hierarchy is *ML*, the modal fragment of L_μ, consisting of formulas without any fixpoints. Formulas with only ν-bound fixpoints, or only μ-bound fixpoints respectively, have index $\{0\}$, or $\{1\}$, corresponding to the alternation classes Π^μ_1, or Σ^μ_1. The classes Π^μ_i and Σ^μ_i for positive even i correspond to formulas with indices $\{1,...,i\}$ and $\{0,...,i-1\}$ respectively, while for odd i they correspond to formulas with indices $\{0,...,i-1\}$ and $\{1,...,i\}$ respectively.

Example 4. The formula $\mu X.\nu Y.(\Box Y \wedge \mu Z.\Box(X \vee Z))$ accepts the order-preserving assignment $\Omega(X) = 1, \Omega(Y) = 0$ and $\Omega(Z) = 1$, and has therefore index $\{0,1\}$ and is in Σ^μ_2. It is however equivalent to $\mu X.\Box X$ (see Section 2.3) and is therefore semantically in Σ^μ_1.

Theorem 5. *[1, 16] The alternation hierarchy is strict: for each index $I = \{m,...,q\}$, where $m \in \{0,1\}$, there are formulas that are not equivalent to a formula with smaller index.*

Deciding an index I means deciding, given an arbitrary L_μ formula, whether it is equivalent to any formula of index I.

2.3 L_μ semantics and parity games

We define the semantics of L_μ in terms of its model-checking games, known as parity games. The correspondence with the classical semantics is standard [3].

L_μ formulas operate on regular unranked trees. Since we consider only unimodal L_μ, these trees do not have edge-labels.

Definition 6. *(Trees)* An unranked tree $\mathfrak{T} = (S,E,r,P)$, rooted at $r \in S$ consists of:

- a set of states S,
- a successor relation $E \subseteq S \times S$,
- a labelling $P : S \to \mathscr{P}(Prop_\mathfrak{T})$ mapping states to subsets of a finite set $Prop_\mathfrak{T} \subset Prop$.

For every state $s \in S$, the set of its ancestors, $\{w \in S \mid \exists w_1,\dots,w_k.\ (w,w_1) \in E, (w_1,w_2) \in E, \dots (w_k,s) \in E\}$, is finite and well-ordered with respect to the transitive closure of E.

The scope of this paper is restricted to regular trees with finite but unbounded branching, which can be finitely represented as trees with back edges. Since L_μ enjoys a finite model property [10], the index problem on trees with infinite branching reduces to the index problem on trees with finite branching.

Definition 7. *(Parity game)* A parity game arena is $A = (V, E, v_\iota, \Omega)$ consisting of: a set V of states partitioned into those belonging to Even, V_e and those belonging to Odd, V_o; an edge relation $E \subseteq V \times V$; an initial node v_ι; and a priority assignment $\Omega : V \to I$. The co-domain I of Ω is said to be the index of A. I is a prefix of the natural numbers, starting at either 0 or 1.

A play in a parity game is a potentially infinite sequence of successive positions starting with v_ι. A play is finite if it ends in a position without successors; then the owner of the position loses. For infinite plays, the winner depends on the highest priority seen infinitely often, called the dominant priority. The winner is the player of the parity of the dominant priority.

A positional (or memoryless) strategy σ for Even (and similarly for Odd) consists of a choice $\sigma(v)$ of successor at the nodes v in V_e (or V_o). A play $\pi = v_0, v_1, ...$, where $v_0 = v_\iota$, agrees with a strategy σ if, at every position $v_i \in V_e$ (or V_o) along the play, $v_{i+1} = \sigma(v_i)$.

A pair of strategies, one for each player, induces a unique play. A strategy for one of the players is winning if every play that agrees with it is winning for the player. Parity games are positionally determined [7, 17]: positional strategies suffice and exactly one of the players has a winning strategy from every position. A parity game $G(A)$ on arena A is said to be winning for the player with a winnning strategy.

We define the semantics of L_μ in terms of winning strategies in parity games: For an unranked tree $\mathfrak{T}$ and formula Ψ, we say that $\mathfrak{T}$ satisfies Ψ, written $\mathfrak{T} \models \Psi$ if and only if Even has a winning strategy in the model-checking parity game $G(\mathfrak{T} \times \Psi)$ defined below.

Definition 8. *(Model-checking parity game)* Let Ψ be a sentence of L_μ and Ω_Ψ a priority assignment with co-domain I on the fixpoint variables of Ψ. Then, for any unranked tree $\mathfrak{T}$, define the model-checking parity game arena $\mathfrak{T} \times \Psi$ as follows: $\mathfrak{T} \times \Psi = (V, E, v_\iota, \Omega)$ where:

- V is the set of states (s, ϕ), where s is a state of $\mathfrak{T}$ and ϕ is a subformula of Ψ.
- V_o consists of positions $(s, \phi \wedge \psi)$ and $(s, \Box\phi)$ while $V_e = V \setminus V_o$;
- Positions (s, C), where C is a literal, are terminal and if $C \in P(s)$, belong to Odd; else it belongs to Even;
- There is an edge from $(s, \phi \vee \psi)$ to (s, ϕ) and (s, ψ);
  an edge from $(s, \phi \wedge \psi)$ to (s, ϕ) and (s, ψ);
  an edge from (s, X), $(s, \mu X.\phi_X)$ and $(s, \nu X.\phi_X)$ to (s, ϕ_X);
  an edge from $(s, \Diamond\phi)$ and $(s, \Box\phi)$ to (s', ϕ) for every successor s' of s;
- v_ι the initial position is (r, Ψ), where r is the root of $\mathfrak{T}$;
- Ω assigns $\Omega_\Psi(X)$ to positions $s \times X$; Ω assigns the minimal priority of I to all other positions.

Definition 9. *(Semantics of L_μ via games)* If Ω_Ψ is an order-preserving priority assignment for Ψ, then for all $\mathfrak{T}$, Even has a winning strategy in the parity game $G(\mathfrak{T} \times \Psi)$ if and only if $\mathfrak{T} \models \Psi$. Furthermore, Even has a winning strategy in $G(\mathfrak{T} \times \Psi)$ from each position (s, ϕ) of $\mathfrak{T} \times \Psi$, where s is a state of $\mathfrak{T}$ and ϕ is a subformula of Ψ, if and only if $s \models \phi$.

Definition 10. The model-checking parity games generated by a formula Ψ are encoded as trees by assigning propositional variables E_i to positions belonging to Even with priority i and O_i to positions belonging to Odd of priority i. Encoding more data about the provenance of positions will be useful in the second half of this paper: we will mark modal positions, that is to say positions $s \times \Box\phi$ and $s \times \Diamond\phi$ with a propositional variable M.

The winning regions of parity games can be described by L_μ formulas [7], shown to be complete for their index (*i.e.* not equivalent to any formula of lower index)[2].

For an index $I = \{m, ..., q\}$ with $m \in \{0,1\}$, define the formula:

$$\mathsf{Parity}_I = \gamma_q X_q \dots \gamma_m X_m . \bigvee_{i \in I} (E_i \wedge \Diamond X_i) \vee (O_i \wedge \Box X_i)$$

where γ_i is μ for odd i and ν for even i.

This formula *describes the winning regions* of parity games of index I in the sense that if A is a parity game arena of index I, encoded as a tree, then $A \models \mathsf{Parity}_I$ if and only if Even wins $G(A)$.

Definition 11. Φ *interpret* Ψ if for all trees $\mathfrak{T}$, it is the case that $\mathfrak{T} \models \Psi$ if and only if $\mathfrak{T} \times \Psi \models \Phi$.

Theorem 12. *[7, 22] For all Ψ of index I, Parity_I interprets Ψ.*

3 The descriptive complexity of parity games

In the last section we reviewed the classic result that to interpret a formula, a formula of the same alternation depth is sufficient: the syntactic complexity of Ψ is an upper bound on the descriptive complexity of its model-checking game. Here we show that if a formula in an alternation class C interprets Ψ, then the formula Ψ is equivalent to one in C. In other words, the semantic complexity of a L_μ formula is a lower bound to the descriptive complexity of its model-checking games.

Theorem 13. *Let Ψ be a formula of L_μ. If for some formula Win and all structures $\mathfrak{T}$ we have $\mathfrak{T} \times \Psi \models \mathsf{Win}$ if and only if $\mathfrak{T} \models \Psi$, that is to say Win interprets Ψ, then Ψ is equivalent to a formula which has the same alternation depth as Win.*

The proof of this theorem depends on a natural, product-like operation on formulas which, if Win interprets Ψ, yields a formula $\Psi \times \mathsf{Win}$, equivalent to Ψ, with the alternation depth of Win. The formula $\Psi \times \mathsf{Win}$ is built with the intention that the parity game arena $(\mathfrak{T} \times \Psi) \times \mathsf{Win}$ is the same as $\mathfrak{T} \times (\Psi \times \mathsf{Win})$. The choice of overloading the notation $\times$ is meant to emphasize this associativity. Note however the type of the objects in these statements: $\Psi \times \mathsf{Win}$ is a formula if Ψ and Win are both formulas while $\mathfrak{T} \times \Psi$ is a tree if $\mathfrak{T}$ is a tree.

Definition 14. *($\Psi \times \mathsf{Win}$)* Let Ψ be a formula of L_μ with index J and priority assignment Ω_Ψ, and let Win be a formula over propositional variables $\mathscr{P}_{\mathsf{Win}}$ with priority assignment Ω_{Win} with co-domain I with minimal element m. Define $\Psi \times \mathsf{Win}$, using a fresh set of fixpoint variables $W_{\phi \times X}$ where ϕ ranges over subformulas of Ψ and X ranges over the fixpoint variables of Win, as follows:

$$P \times E_m = \neg P \text{ where } P \text{ is a literal;}$$
$$P \times O_m = P \text{ where } P \text{ is a literal;}$$
$$\phi \times P = \top \text{ if for all states } s \text{ of any tree, the position } (s,\phi) \text{ satisfies } P \in \mathscr{P}_{\mathsf{Win}}.$$

In particular:

$$\phi \wedge \psi \times O_m = \top;$$
$$\phi \vee \psi \times E_m = \top;$$
$$\Diamond\phi \times E_m = \top;$$
$$\Box\phi \times O_m = \top;$$
$$\top \times O_m = \bot \times E_m = \top;$$
$$X \times E_{\Omega_\Psi(X)} = \top;$$
$$\phi \times P = \bot \text{ for } P \in \mathscr{P}_{\mathsf{Win}} \text{ otherwise.}$$
$$\Box\phi \times \Box\psi = \Diamond\phi \times \Box\psi = \Box(\phi \times \psi);$$
$$\phi \times \Box\psi = \bigwedge_{\phi' \in im(\phi)} \phi' \times \psi \text{ where } im(\phi) \text{ is the set of immediate subformulas of } \phi;$$
$$\Box\phi \times \Diamond\psi = \Diamond\phi \times \Diamond\psi = \Diamond(\phi \times \psi);$$
$$\phi \times \Diamond\psi = \bigvee_{\phi' \in im(\phi)} \phi' \times \psi \text{ where } im(\phi) \text{ is the set of immediate subformulas of } \phi;$$

If reached computing a subformula of $\mu W_{\phi\times X}.(\phi \times \psi_X)$ (or $\nu W_{\phi\times X}.(\phi \times \psi_X)$) but not a subformula $\nu W_{\phi'\times Z}.\phi_{W_{\phi'\times Z}}$ (or $\mu W_{\phi'\times Z}.\phi_{W_{\phi'\times Z}}$) thereof with $\Omega(Z) > \Omega(X)$ then:

$$\phi \times X = W_{\phi\times X}, \text{ else:}$$
$$= \mu W_{\phi\times X}.(\phi \times \phi_X) \text{ if } X \text{ is a } \mu \text{ variable;}$$
$$= \nu W_{\phi\times X}.(\phi \times \phi_X) \text{ if } X \text{ is a } \nu \text{ variable.}$$
$$\phi \times \mu X.\psi_X = \mu W_{\phi\times X}.\phi \times \phi_X;$$
$$\phi \times \nu X.\psi_X = \nu W_{\phi\times X}.\phi \times \phi_X;$$
$$\phi \times \psi \wedge \psi' = (\phi \times \psi) \wedge (\phi \times \psi');$$
$$\phi \times \psi \vee \psi' = (\phi \times \psi) \vee (\phi \times \psi').$$

where $\Omega_{\Psi\times\mathsf{Win}}(W_{X\times\phi}) = \Omega_{\mathsf{Win}}(X)$. Superfluous fixpoint bindings which do not bind any free variables can then be removed. The construction terminates as computations of $\phi \times X$ will eventually result in the variable $W_{\phi\times X}$ being bound.

Proof. (Theorem 13) For the correctness of this construction, it is sufficient to compare rule by rule the parity games $(\mathfrak{T} \times \Psi) \times \mathsf{Win}$ and $\mathfrak{T} \times (\Psi \times \mathsf{Win})$.

Case $P \times E_m = \neg P$: On one hand, the position $((s,P),E_m)$ is winning for Even exactly when (s,P) belongs to Even: when $s \models \neg P$. On the other hand, the position $(s,\neg P)$ is also winning for Even exactly when $s \models \neg P$.

Case $P \times O_m = P$: On one hand the position $((s,P),O_m)$ is winning for Even exactly when (s,P) belongs to Odd: when $s \models P$. On the other hand the position (s,P) is also winning for Even exactly when $s \models P$.

Case $\phi \times P$: For $\phi \notin \mathscr{P}_{\mathsf{Win}}$, whether (s,ϕ) satisfies P can only depend on ϕ and not s. If for all states s of any tree, (s,ϕ) satisfies P, then $((s,\phi),P)$ is winning for Even, as is $(s,\top)$. Otherwise $((s,\phi),P)$ is winning for Odd, as is $(s,\bot)$.

Case $\Box\phi \times \Box\psi = \Box(\phi \times \psi)$: $((s,\Box\phi),\Box\psi)$ and $((s,\Diamond\phi),\Box\psi)$ are both positions belonging to Odd with successors $((s',\phi),\psi)$ for s' a successor of s. On the other hand $(s,\Box(\phi,\psi))$ has successors $(s',\phi \times \psi)$ for s' a successor of s. The case for $\Box\phi \times \Diamond\phi = \Diamond\phi \times \Diamond\psi = \Diamond(\phi \times \psi)$ is similar.

Case $\phi \times \Box\psi = \bigwedge_{\phi' \in im(\phi)} \phi' \times \psi$: The position $((s,\phi),\Box\psi)$ belongs to Odd and has successors $((s,\phi'),\psi)$ where ϕ' is an immediate subformula of ϕ. On the other hand, $(s, bigwedge_{\phi' \in im(\phi)} \phi' \times \psi)$ also belongs to Odd and has successors $(s,\phi' \times \psi)$ where ϕ' is an immediate subformula of ϕ. The case for $\phi \times \Diamond\psi) = \bigvee_{\phi' \in im(\phi)} \phi' \times \psi$ is similar.

Case for fixpoints: $((s,\phi),X)$ has a unique successor $((s,\phi),\phi_X))$ where ϕ_X is the formula binding X and is of the priority of X. On the other hand $(s,W_{\phi\times X})$ and $(s,\gamma W_{\phi\times X})$ for $\gamma \in \{\mu,\nu\}$ have a unique successor $(s,\phi \times \phi_X)$, and is of the priority of $W_{\phi\times X}$ which is the same as the priority of X.

Case for $\phi \times \psi \wedge \psi'$: The position $((s,\phi),\psi \wedge \psi')$ belongs to Odd and has successors $((s,\phi),\psi)$ and $((s,\phi),\psi')$. The position $(s,(\phi \times \psi) \wedge (\phi \times \psi')$ also belongs to Odd and has successors $(s,\phi \times \psi)$ and $(s,\phi \times \psi')$. The case for $(\phi \times \psi \vee \psi')$ is similar.

The formula $\Psi \times \mathsf{Win}$ inherits its priority assignment from Win, and the condition on the introduction of fixpoint variables guarantees that the priority assignment is order-preserving. Then we have that $\mathfrak{T} \models \Psi \Leftrightarrow \mathfrak{T} \times \Psi \models \mathsf{Win} \Leftrightarrow \mathfrak{T} \models \Psi \times \mathsf{Win}$: the formula $\Psi \times \mathsf{Win}$, of index I, is equivalent to Ψ. $\Box$

This concludes the argument that if Ψ can be interpreted by a formula Φ with index I, then Ψ is semantically of index I. Since $\Psi \times \Phi$ is an effective transformation that turns Ψ into a formula which is syntactically of index I, a formula can be simplified whenever a suitable interpreting formula is found. The following example shows how this can be used to argue that a formula is semantically simple without recourse to the formula's semantics.

Example 15. Consider this formula

$$\Psi = \mu X.\nu Y.\mu Z.(A \wedge \Diamond Y) \vee (B \wedge \Diamond(Z \wedge \phi)) \vee (C \wedge \Box X)$$

where $\phi \in \Pi_2^\mu$ and does not have variables X,Y,Z. We use Theorem 13 to argue that this formula, with seemingly opaque semantics, is semantically Π_2^μ.

The first thing to note is that a winning strategy σ for Even in a model-checking game of Ψ has the following property: a play that agrees with σ eventually reaches a point from which no play that agrees with σ sees a position $s \times X$. We can describe the winning regions of such games with the following formula:

$$\mathsf{Interpretor} = \mu X. \bigvee_{i \in \{1,2,3\}} (E_i \wedge \Diamond X \vee O_i \wedge \Box X) \vee \nu X_2.\mu X_1. \bigvee_{i \in \{2,1\}} (E_i \wedge \Diamond X_i \vee O_i \wedge \Box X_i)$$

The intuition here is that the first part of the Interpretor formula describes Even's strategy until she can guarantee to no longer a position of priority 3; the second part of the formula is just the usual $\mathsf{Parity}_{2,1}$ formula. A strategy in the model-checking game for this formula translates directly into the model-checking game of $\mathsf{Parity}_{3,2,1}$. On the other hand, if Even has a winning strategy in a model-checking game $\mathfrak{T} \times \mathsf{Parity}_{3,2,1}$ with the above property, then she also has a winning strategy in $\mathfrak{T} \times \mathsf{Interpretor}$: once a play reaches the point from where her strategy no longer sees priority 3, she moves to the second part of the formula.

The Π_2^μ formula $\mathsf{Interpretor}$ therefore interprets Ψ. From Theorem 13, Ψ is semantically Π_2^μ. For the curious reader who wants to gain an intuition of how the transformation operates, the simplified formula is discussed in the Appendix.

4 Interpretation theorems

So far, we have seen that if a formula Ψ is interpreted by a formula of index I, then Ψ is itself semantically of index I. This extension of the L_μ-parity game relationship is natural and not too surprising. This section considers the converse: when is it the case that a formula semantically of index I can be interpreted by a formula of syntactic index I? The formula in Example 15 shows that this may happen for syntactic reasons. Here we discuss whether it holds in general. This conjecture can be studied with respect to any index I:

Conjecture 16. *If Ψ is semantically of index I, then there is a formula Φ of syntactic index I such that for all structures $\mathfrak{T}$, $\mathfrak{T} \times \Psi \models \Phi$ if and only if $\mathfrak{T} \models \Psi$.*

If this conjecture is found to be true in general, then the descriptive complexity of the model-checking games of a L_μ formula Ψ is exactly the complexity of the simplest formula equivalent to Ψ.

Say that the class of formulas of index I admits an *interpretation theorem* if the above conjecture holds for I. This section considers this conjecture in light of what we know about deciding ML, Π_1^μ and Σ_2^μ. It shows that these classes admit interpretation theorems:

- ML has an interpretation theorem;
- Π_1^μ has an interpretation theorem for disjunctive L_μ;
- Σ_2^μ has an interpretation theorem for disjunctive L_μ;

It then considers what can be said in the general case, and shows that disjunctive L_μ alternation classes all have an interpretation theorem for input in co-disjuntive form.

4.1 Interpretation theorem for *ML*

This section shows that conjecture 16 holds for semantically modal formulas:

Theorem 17. (Interpretation Theorem for *ML*)

Let Ψ be a semantically modal formula, then there exists (effectively) a modal formula Φ such that Φ interprets Ψ.

It is known that if a guarded formula Ψ is equivalent to a modal formula of modal depth m, then it is equivalent to the formula obtained by approximating all fixpoints further than m and then truncating the formula at a modal depth of m [13]. More precisely, a formula Ψ equivalent to a modal formula of modal depth $m > 0$ is equivalent to Ψ^n, defined as follows:

Definition 18.

$$(\Box\phi)^0 = \Box\top$$
$$(\Diamond\phi)^0 = \Diamond\bot$$
$$\phi^n = \phi[(\Box\psi')^{n-1}/\Box\psi; (\Diamond\psi')^{n-1}/\Diamond\psi]$$
$$\text{where } \psi' = \psi[\phi_X/X]_{\forall X \in Free(\psi)}.$$

That is to say, fixpoints are unfolded until a formula nested within m modalities is reached. Then, it is replaced with either $\top$ or $\bot$, according to the last modality. Note that the recursive case is only applied to formulas prefixed with a modality and guarded formulas without modalities are already modal.

Definition 19. A n-bounded parity game is a parity game, augmented with a counter, of which some positions $M \in V$ are marked. M must be such that every infinite path goes through M. Then, a play of this game is a play of the parity game, except that the counter, initially at n, is decremented whenever a position in M is seen. If the play reaches M at counter value 0, then the owner of the position loses the game.

If the longest path between two positions in M is no longer than p, the winning regions of these games are described by the modal formula $\mathsf{Bounded}_{p,m}$ defined inductively as follows:

$$\begin{aligned}
\mathsf{Bounded}_{0,b} =& \bot \\
\mathsf{Bounded}_{a,0} =& (E_i \wedge \neg M \wedge \Diamond \mathsf{Bounded}_{a-1,0}) \vee \\
& (E_i \wedge M \wedge \bot) \vee \\
& (O_i \wedge \neg M \wedge \Box \mathsf{Bounded}_{a-1,0}) \vee \\
& (O_i \wedge M \wedge \top) \\
\mathsf{Bounded}_{a,b} =& (E_i \wedge \neg M \wedge \Diamond \mathsf{Bounded}_{a-1,b}) \vee \\
& (E_i \wedge M \wedge \Diamond \mathsf{Bounded}_{p,b-1}) \vee \\
& (O_i \wedge \neg M \wedge \Box \mathsf{Bounded}_{a-1,b}) \vee \\
& (O_i \wedge M \wedge \Box \mathsf{Bounded}_{p,b-1})
\end{aligned}$$

Proof. (of Theorem 17) Let Ψ be equivalent to a modal formula of modal depth m. Let p be the longest path in the parse-tree of Ψ (where ϕ_X is taken to be the child of X) without modalities. Since Ψ is taken to be guarded, p is finite. Then the model-checking games of Ψ are m-bounded parity games where M is the set of positions $s \times \phi$ where ϕ is a modal formula: a winning strategy for either player in $\mathfrak{T} \times \Psi^m$ informs a winning strategy in the m-bounded parity game on $\mathfrak{T} \times \Psi$. Then, noting that the longest path between two positions in M in any model-checking parity game of Ψ is of length at most p, $\mathsf{Bounded}_{p,m}$ interprets Ψ. □

Hence the modal fragment of L_μ has an effective interpretation theorem: any semantically modal formula can be interpreted by a syntactically modal formula. This means that all semantically modal formulas, no matter how high their syntactic index, generate a class of model-checking games with modal descriptive complexity. The interpreting formula depends on the size of Ψ as well as its semantic modal depth, but is computable from the formula.

4.2 Interpretation theorem for Π_1^μ

The conjecture 16 also holds for *disjunctive formulas* that are semantically Π_1^μ.

Theorem 20. (Interpretation theorem for Π_1^μ)
Let Ψ be a disjunctive formula which is semantically in Π_1^μ. Then there exists (effectively) a Π_1^μ formula Φ such that Φ interprets Ψ.

The decision procedure for Π_1^μ in [13] shows that in disjunctive semantically Π_1^μ formulas, every satisfiable μ-subformula can be replaced with the same subformula bound by ν, while unsatisfiable ones can be replaced by $\bot$. This can be translated into the formula describing the winning regions of the model-checking games of the formula, simply by substituting the least fixpoint in the parity game formula corresponding to a subformula $\mu X.\phi$ with either $\bot$ or a greatest fixpoint, depending on whether the subformula is unsatisfiable.

For example, the appropriate formula describing the winning regions of a disjunctive formula in which all μ-bound subformulas are satisfiable would simply be:

$$\nu Y.(E_e \wedge \Diamond Y) \vee (E_o \wedge \Diamond Y) \vee (O_e \wedge \Box Y) \vee (O_o \wedge \Box Y)$$

Where E_e (E_o) is the disjunction of E_i for even (odd) i; O_e (O_o) is the disjunction of O_i for even (odd) i. For μ-bound subformulas which are unsatisfiable, it is enough to turn the clause $E_i \wedge \Diamond X_i$ corresponding to the μ-variable X_i in question into $E_i \wedge \bot$.

Remark 21. Note that the encoding of the model-checking game arena can incorporate data about the provenance of a node, rather than just its priority, for instance by using additional propositional variables E_X to indicate positions which stem from the fixpoint variable X. To describe the winning regions of games encoded in this manner, we use $E_X \wedge \Diamond Y_{\Omega(X)}$ instead of $E_i \wedge \Diamond Y_i$.

Recall that for modal formulas, the interpreting formula depends only on the modal rank of the formula and the length of the longest path without fixpoints in the formula. Since both of these can be bound in relation to the size of the formula, all semantically modal L_μ formulas of the same size can be interpreted by the same formula. This cannot be said for semantically Π^μ_1 formulas: the interpreting formulas depend on which fixpoint subformulas are unsatisfiable and which are interchangeable with ν. It remains an open question whether a *uniform* interpreting formula could be devised for all semantically Π^μ_1 formulas in disjunctive form. It is also open whether the restriction to disjunctive formulas can be lifted.

4.3 Interpretation theorem for Σ^μ_2

The previous sections establish effective interpretation theorems for Π^μ_1, restricted to disjunctive formulas, and for the modal fragment of L_μ. This section shifts the focus onto an alternation class not known to be decidable: Σ^μ_2.

Theorem 22. (Interpretation Theorem for Σ^μ_2)
If a disjunctive formula Ψ is semantically in Σ^μ_2, then it is interpreted by a Σ^μ_2 formula.

The proof argues that first part of the decidability proof of Σ^μ_2 for Π^μ_2 presented in [15] translates into an interpretation theorem for Σ^μ_2. Similarly to the modal case, a variation on parity games yields the model-checking games for semantically Σ^μ_2 formulas. Their winning regions are described by a Σ^μ_2 formula. Here we briefly recall the n-challenge game, and the result that its winning regions are described by a Σ^μ_2 formula.

Definition 23. Fix a formula Ψ in disjunctive form, of index $\{q,...,0\}$. Let $I = \{q,...,0\}$ if q is even and $\{q+1,q,...,0\}$ otherwise. Write I_e for the even priorities in I.

The n-challenge game consists of a normal parity game augmented with a set of challenges, one for each even priority i. A challenge can either be *open* or *met* and has a counter c_i attached to it. Each counter is initialised to n, and decremented when the corresponding challenge is opened. The Odd player can at any point open challenges of which the counter is non-zero, but he must do so in decreasing order: an i-challenge can only be opened if every j-challenge for $j > i$ is opened. When a play encounters the priority j while the j-challenge is open, the challenge is said to be met. All i-challenges for $i < j$ are then *reset*. This means that the counters c_i are set back to n and marked *met*.

A play of this game is a play in a parity game, augmented with the challenge and counter configuration at each step. A play with dominant priority d is winning for Even if either d is even or if every opened $d+1$ challenge is eventually met or reset.

Proof. *(Theorem 22)* It is known from [15] that disjunctive Ψ is equivalent to a Σ_2^μ formula if an only if there is an n such that for all trees $\mathfrak{T}$, Even wins the n-challenge game on $\mathfrak{T} \times \Psi$ if and only if she wins the parity game. The winning regions of the n-challenge games on parity game arenas of index I are described by a Σ_2^μ formula $\mathsf{Challenge}_I^n$. Hence, if Ψ is a disjunctive formula semantically in Σ_2^μ, then there is an n such that for all $\mathfrak{T}$, $\mathfrak{T} \models \Psi$ if and only if $\mathfrak{T} \times \Psi \models \mathsf{Challenge}_I^n$. That is to say, $\mathsf{Challenge}_I^n$ interprets Ψ. □

This concludes the argument that for formulas up to semantic index $\{1,0\}$, the descriptive complexity of the model-checking games coincides exactly with the semantic complexity of the formula. The dependence of these theorems disjunctive form highlights the importance of this normal form for the index problem.

4.4 General interpretation theorem for disjunctive L_μ

This section generalises the argument used for Σ_2^μ. With some concessions over the structure of target formulas, that is to say only considering disjunctive alternation classes, this yields a general interpretation theorem for co-disjunctive formulas.

We generalise the n-challenge game by considering a set of challenges (one per target priority) per input priority rather than a single challenge. Fix J to be the *input* index, *i.e.* the index of the input formula, while I is the *target* index, *i.e.* the index of the target alternation class. Then, for the pair J,I, we defines a parameterised challenge game such that an input formula Ψ in co-disjuntive form of index J is interpreted by the disjunctive formula (of index I) describing the winning regions of these games if and only if Ψ is equivalent to a disjunctive formula of index I.

As usual, this game is played on a parity game arena with priorities from J. Unlike for Σ_2^μ, the challenging player is now Even. This is because the target formulas are in disjunctive form, which means Even has strategies which only see one play per branch – this is necessary as the target index is no longer restricted to $\{0,1\}$. The input formula on the other hand is in co-disjunctive form, to yield such strategies for Odd.

Definition 24. *(Generalised challenge games)* A *challenge configuration* $(\bar{a},\bar{c})$ consists of $|I_o| \times |J_o|$ challenges and counters, where I_o and J_o are the odd priorities in I and J respectively. Write $a_{i,j} = met$ or *open* for $i \in I$ and $j \in J$ to indicate whether the i-level challenge on j is open. Each challenge $a_{i,j}$ is attached to a positive integer counter value $c_{i,j}$, bounded by n.

Given a configuration $(\bar{a},\bar{c})$, the least j for which $a_{i,j}$ is open (for any i) is the *priority* of the challenge configuration, and the highest level i at which $a_{i,j}$ is open is its *level*. A valid challenge configuration respects the following constraint: if $a_{i,j} = open$ then $a_{i,k} = open$ for all $k > j$. That is to say, challenges are opened in decreasing order.

The *game configuration* consists of a position in the parity game and a valid challenge configuration. The progress of the game can be divided into two rounds: in the first round Even opens or resets challenges, while the second round is a step in the parity game. In the first round her possible actions are:

- To *k-reset* for any odd $k \in I_o$, setting $a_{i,j} := met$ and $c_{i,j} := n$ for all $i \le k$ and all $j \in J$;
- To open at any level i, challenges up to any p, as long as the counters allow it: for all $j \ge p$ such that $a_{i,j} = met$, set $a_{i,j} := open$ and $c_{i,j} := c_{i,j} - 1$ if $c_{i,j} > 0$.

Then, in the second round, the player whose turn it is in the parity game picks a successor position. If the underlying parity game ends in a terminal state, then the winner of the underlying parity game immediately wins the challenge game, too. The challenge configuration is updated according to the priority p of this new position:

- $a_{i,j} := met$ for all $j \leq p$, all $i \in I_o$;
- $c_{i,j} := n$ for all $j < p$, all $i \in I_o$.

If $c_{i,p} = 0$ for some i, then the game ends immediately with a win for Odd.

A play is a potentially infinite sequence of game configurations: an underlying parity game play augmented with challenge configurations. The initial challenge-configuration is $a_{i,j} = met$ and $c_{i,j} = n$ for all i, j. The dominant priority of the parity game play is also the dominant priority of the challenge game play. An infinite play with dominant priority d is winning for Odd if: d is odd, or all $a_{i,d+1}$ challenges are in the *met* state infinitely often.

Lemma 25. *The winning regions of a generalised n-challenge game for J, I are described by a disjunctive L_μ formula with index I.*

The constructionof the formula $\mathsf{GChallenge}^n_{J,I}$, detailed the Appendix, is similar in spirit to the one described in [15], although slightly more involved, given the additional input priorities to account for.

Lemma 26. *A co-disjunctive L_μ formula Ψ of index J is interpreted by the I-index formula* $\mathsf{GChallenge}^n_{J,I}$ *if it is equivalent to a disjunctive formula of index I.*

Again, the proof, in the Appendix, follows the same structure as the proof of the Σ^μ_2 case. In brief, it assumes that Ψ of index J is equivalent to some disjunctive Φ of index I and that for all m, there is a structure $\mathfrak{T}$, such that Even wins the parity game $\mathfrak{T} \times \Psi$ but Odd wins the generalised m-challenge game on $\mathfrak{T} \times \Psi$. It then takes a sufficiently large m, and uses a winning strategy σ for Even in $\mathfrak{T} \times \Phi$ to define a challenging strategy γ for her in the generalised m-challenge game on $\mathfrak{T} \times \Psi$. In the Σ^μ_2 case the challenges were issues when the higher priority had been seen on all plays since the last challenge. This time, different even priorities cause Even to issue challenges at different levels while different odd priorities cause Even to reset at different levels. Then Odd's winning strategy τ in the m-challenge game on $\mathfrak{T} \times \Psi$ is used to add back edges to $\mathfrak{T}$, turning it into a new structure $\mathfrak{T}'$ which preserves Even's winning strategy σ in $\mathfrak{T}' \times \Phi$ while turning τ_γ into a winning strategy in $\mathfrak{T}' \times \Psi$. This contradicts the equivalence of Φ and Ψ. The main technical difference is that we no longer invoke König's lemma, and instead uses the fact that target formulas are in disjunctive form to find suitable times for Even to challenge and reset.

This gives us a general interpretation theorem for all disjunctive Σ^μ alternation classes, for co-disjunctive input formulas: the complexity of an equivalent disjunctive formula is an upper bound on the descriptive complexity of the model-checking games generated by the co-disjunctive form of Ψ.

Comparison with automata-theoretic results

The readers familiar with the automata-theoretic efforts to tackle the Rabin–Mostowski index problem may find some parallels between this last interpretation theorem and the reduction of the decidability of the Rabin–Mostowski hierarchy of non-deterministic automata to a boundedness question [5]. In both cases, the core construction gives a way of mapping input priorities to the target priorities. In [5], it is a question of guessing mappings between input and target priorities a bounded number of times. Here it is achieved through the parameterised challenges. The proof in [5] relies on *guidable* automata – a concept for which the L_μ-theoretic conterpart is conspicuously absent.

I believe this comparison to be a cause for cautious optimism when looking for further interpretation theorems: indeed, the automata-theoretic result is not burdened by the same constraints to disjunctive and co-disjuntive form as Theorem 26; perhaps these constraints can be lifted in the L_μ framework, too. Perhaps the mapping construction in [5] could yield a family of interpreting formulas, for which the interpretation theorem does not place restrictions on the syntax of the input formula.

The transferability of the techniques is however not a trivial question. Disjunctive L_μ was designed to be the L_μ counter-part of non-deterministic automata. However, although the model-checking problems for non-deterministic automata and disjunctive L_μ reduce to each other by encoding unranked trees as ranked trees, this reduction is not necessarily index-preserving. Furthermore, it seems straightforward to show that the index-problem for non-deterministic automata corresponds exactly to the index problem of disjunctive L_μ *restricted to existential modalities*; it seems less likely that the index problem of this restricted fragment of disjunctive L_μ would suffice to decide the index problem for disjunctive L_μ. Therefore, how the disjunctive L_μ index problem and the non-deterministic parity automata index problem relate to each other is not yet settled.

5 Discussion

We have shown that for formulas equivalent to formulas of index $I \in \{\{\}, \{0\}, \{0,1\}\}$, the descriptive complexity of their medel-checking games coincides exactly with I. Beyond these alternation classes, the descriptive complexity of the model-checking games of Ψ is an upper bound on the semantic complexity of Ψ. Furthermore, for a co-disjunctive formula equivalent to a disjunctive formula of index I, the descriptive complexity of its model-checking games is bounded by I.

This links the index-problem to the parity game–L_μ relationship, which is shown to be even more robust than previously accounted for. As a result, the index-problem can be approached by studying the descriptive complexity of model-checking games. On one hand, for practical purposes formulas can be simplified by looking at the model-checking games they generate. On the other hand, interpretation theorems are a stepping stone on the pursuit of further decision procedures.

Besides finding a truly general interpretation theorem which would show that the descriptive complexity of model-checking games coincides exactly with the semantic complexity of formulas, I leave the reader with a couple of related open questions.

For low alternation classes ML, Π_1^μ and Σ_2^μ, some of the existing decision procedures can be recast as interpretation theorems. For both Π_1^μ and Σ_2^μ, the interpretation theorem only applies for disjunctive input formula. This once again highlights the importance of disjunctive form for the index problem: the descriptive complexity of the model-checking games of disjunctive formulas appears to be lower than for general L_μ formulas. Whether this is truly the case, remains an open question.

The generality of the last interpretation theorem comes at a cost. In particular, it is not constructive, in the sense that if a co-disjunctive formula is indeed equivalent to a disjunctive formula of index I, even if we find the interpreting disjunctive formula of index I, this only allows us to construct a formula of index I, which may not itself be disjunctive. I conjecture that a similar theorem could be shown for disjunctive input formulas. As discussed in the previous section, some automata-theoretic techniques [5] could perhaps be used to address some of the limits of the current techniques.

References

[1] Julian C. Bradfield. The modal mu-calculus alternation hierarchy is strict. *Theoretical Computer Science*, 195(2):133–153, 1998.

[2] Julian C Bradfield. Fixpoint alternation: Arithmetic, transition systems, and the binary tree. *RAIRO-Theoretical Informatics and Applications*, 33(4-5):341–356, 1999.

[3] Julian C. Bradfield and Colin Stirling. Modal mu-calculi. *Handbook of modal logic*, 3:721–756, 2007.

[4] Thomas Colcombet, Denis Kuperberg, Christof Löding, and Michael Vanden Boom. Deciding the weak definability of Büchi definable tree languages. In *LIPIcs-Leibniz International Proceedings in Informatics*, volume 23. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2013.

[5] Thomas Colcombet and Christof Löding. The non-deterministic Mostowski hierarchy and distance-parity automata. In *International Colloquium on Automata, Languages, and Programming*, pages 398–409. Springer, 2008.

[6] Anuj Dawar and Erich Grädel. The descriptive complexity of parity games. In *International Workshop on Computer Science Logic*, pages 354–368. Springer, 2008.

[7] E Allen Emerson and Charanjit S Jutla. Tree automata, mu-calculus and determinacy. In *Foundations of Computer Science, 1991. Proceedings., 32nd Annual Symposium on*, pages 368–377. IEEE, 1991.

[8] David Janin and Igor Walukiewicz. *Automata for the modal μ-calculus and related results*. Springer, 1995.

[9] Dexter 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.

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

[11] Orna Kupferman, Moshe Y Vardi, and Pierre Wolper. An automata-theoretic approach to branching-time model checking. *Journal of the ACM (JACM)*, 47(2):312–360, 2000.

[12] Ralf Küsters and Thomas Wilke. Deciding the First Level of the μ-calculus Alternation Hierarchy. (0209), 2002.

[13] Karoliina Lehtinen and Sandra Quickert. Deciding the first levels of the modal μ alternation hierarchy by formula construction. In *LIPIcs-Leibniz International Proceedings in Informatics*, volume 41. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.

[14] M Karoliina Lehtinen. Disjunctive form and the modal μ alternation hierarchy. *FICS15, The 10th International Workshop on Fixed Points in Computer Science, EPTCS 191*, page 117, 2015.

[15] M Karoliina Lehtinen and Sandra Quickert. Σ_2^μ is decidable for Π_2^μ. *Accepted to Computation in Europe 2017, preprint available: arXiv:1703.03239*, 2017.

[16] Giacomo 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.

[17] Donald A Martin. Borel determinacy. *Annals of Mathematics*, pages 363–371, 1975.

[18] Radu 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.

[19] Damian Niwiński. *On fixed-point clones*, pages 464–473. Springer Berlin Heidelberg, Berlin, Heidelberg, 1986.

[20] Martin Otto. *Eliminating Recursion in the μ-Calculus*, pages 531–540. Springer Berlin Heidelberg, Berlin, Heidelberg, 1999.

[21] Michal Skrzypczak and Igor Walukiewicz. Deciding the Topological Complexity of Büchi Languages. In Yuval Rabani Ioannis Chatzigiannakis, Michael Mitzenmacher and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, volume 55 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 99:1–99:13, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.

[22] Igor Walukiewicz. Monadic second-order logic on tree-like structures. *Theoretical Computer Science*, 275(1):311–346, 2002.

A The simplified formula for Example 15

Example 15 argued that the following formula is semantically Π^μ_2:

$$\Psi = \mu X.\nu Y.\mu Z.A \wedge \Diamond Y \vee B \wedge \Diamond(Z \wedge \phi) \vee C \wedge \Box X$$

One can obtain the simplified formula by applying the construction of Theorem 13. Unfortunately, the mechanics of the transformation may come in the way of clarity, and obscure the intuition of what the transformation achieves. Here is a tidied-up version of the transformed formula which shows how the interpreting formula restructures the input formula into the target formula.

Recall the interpreting formula:

$$\begin{aligned}\mathsf{Interpretor} = \mu X. \bigvee_{i\in\{1,2,3\}} (E_i \wedge \Diamond X \vee O_i \wedge \Box X) \vee \\ \nu X_2.\mu X_1. \bigvee_{i\in\{2,1\}} (E_i \wedge \Diamond X_i \vee O_i \wedge \Box X_i)\end{aligned}$$

I claim, without proof, that Ψ is equivalent to the following formula Ψ':

$$\begin{aligned}\mu X.A \wedge \Diamond X \vee B \wedge \Diamond(X \wedge \phi) \vee C \wedge \Box X \vee \\ \nu Y.\mu Z.A \wedge \Diamond Y \vee B \wedge \Diamond(Z \wedge \phi) \vee C \wedge \Box \bot\end{aligned}$$

Notice that the structural transformation from $\mathsf{Parity}_{\{1,2,3\}}$ to $\mathsf{Interpretor}$ is the same as the one from Ψ to Ψ': the formula is first copied with all fixpoints replaced by a unique μ-bound one; then the disjunction is taken with another copy of the original formula, this time replacing the most significant μ-variable with $\bot$.

B The formula for generalised challenge games: Proof of Lemma 25

For clarity, describe the formula $\mathsf{GChallenge}^n_{J,I}$ as an alternating parity automaton on unranked trees.

Let A be the automaton of $\mathsf{Parity}_\mathsf{J}$, with priority assignment Ω_A. For each valid challenge configuration $(\bar{a},\bar{c})$, of priority p and level k, let $A(\bar{a},\bar{c})$ be a copy of A with the following modifications:

- $\Omega_{(\bar{a},\bar{c})}(q) = k$ (odd) if $\Omega_A(q) < p-1$;
- $\Omega_{(\bar{a},\bar{c})}(q) = k+1$ (even) if $\Omega_A(q) \geq p-1$;

The components $A(\bar{a},\bar{c})$ are joined to form one automaton by adding transitions as follows:

- For a state q in $A(\bar{a},\bar{c})$ with $\Omega_A(q) = d \geq p$ and some i such that $a_{i,p} = open$ and $c_{i,p} = 0$, there is a unique transition to $\bot$, indicating an immediate win for Odd; else
- A state q in $A(\bar{a},\bar{c})$ with $\Omega_A(q) = d \geq p$ has a transition to the copy of q in $A(\bar{a}',\bar{c}')$ where: $a'_{i,j} = met$ for all i and $j \leq d$, $a'_{i,j} = a_{i,j}$ otherwise, and $c'_{i,j} = n$ for all i and $j < d$, and $c'_{i,j} = c_{i,j}$ otherwise; this indicates all challenges up to d being met.
- All other states q have a transition $\delta(q) = \bigvee Q_{open} \cup Q_{reset}$ where Q_{open} and Q_{reset} are as follows. Q_{open} is the set of states corresponding to q in the components $A(\bar{a}',\bar{c}')$ such that if $a_{i,j} = open$, then $a'_{i,j} = open$ and whenever $a_{i,j} = met$ but $a'_{i,j} = open$, then $c'_{i,j} = c_{i,j} - 1 \geq 0$; this corresponds to Even opening some challenges.

  Q_{reset} is the set of states corresponding to q in the components $A(\bar{a}',\bar{c}')$ such that for some $k \in I_o$, $a'_{i,j} = met$ and $c'_{i,j} = n$ for all $i \leq k$ and $a'_{i,j} = a_{i,j}$ and $c'_{i,j} = c_{i,j}$ otherwise; This corresponds to Even resetting counters up to level k.

A play of this automaton can be read as a play of the automaton A augmented with the challenge configurations $(\bar{a},\bar{c})$ corresponding to what component a state is visited in. The original priorities of states, as defined by Ω_A, correspond to the priorities in the underlying game. The moves between the components describe how the challenge configuration evolves as the Even player sets and resets challenges and as these are met. It suffices to check that the winning conditions of the challenge game are respected by the new parity assignment.

A play that reaches a terminal position corresponds to either a finite play in the parity game, or a challenge on p being met when a counter $a_{i,p}$ is at 0. We then need to consider three cases.

First consider infinite plays, dominated in the underlying parity game by an even priority d, where some $d+1$ challenge $a_{i,d+1}$ is eventually always *open*. Such a play should, according to the rules of the challenge game, be winning for Even. In the automaton of $\mathsf{GChallenge}^{\mathsf{n}}_{\mathsf{J},\mathsf{I}}$, the play eventually gets stuck in automaton components at some level k, corresponding to the most significant level $d+1$ is opened but not met at. The play visits a component with level k and priority $d+1$ infinitely often, at least whenever d is seen. The play will therefore see the even priority $k+1$ infinitely often. This will be the dominant priority since it is the largest priority within components of level k. Such a play is therefore winning for Even, as it should be.

Now consider infinite plays, dominated in the underlying parity game by an even priority d, where all $d+1$ challenges are always eventually *met*. Such a play should be winning for Odd according to the challenge game specifications. We consider the case that $d+1$ is opened finitely many times first. Then, the game eventually gets stuck in components of priority $d+3$ or more. In such components priorities d and lower receive Odd priorities, so Odd wins this game.

If $d+1$ is opened infinitely often at some maximal level k, Even must be resetting level k or higher infinitely often. The play therefore visits a component of level $k+2$ or higher with priority $d+3$ infinitely often, where it must see the odd priority k. The play can see no higher even priority (available only when $a_{i,j}$ is open for $i \geq k+2$ when a priority equal or higher to $j-1$ is seen) infinitely often since the $d+1$ challenge is not opened at a more significant level infinitely often. The play is therefore winning for Odd, as required.

Finally, if an infinite play is dominated by an odd priority d, all d challenges are always eventually *met*. Again, we can consider both cases where a d challenge is opened finite and infinitely often. If such challenges are only opened finitely often, the game gets stuck in components of priority $d+2$ where all nodes of original priorities d and lower receive an odd priority. If a d challenge is opened infinitely often at the maximal level k, then Even must be resetting level k or higher infinitely often, which means that the play visits a node of priority k infinitely often and, as above, cannot visit a higher even priority.

As this formula is based on Parity_J, a disjunctive formula, and Even's additional choices are disjunctions, this formula is in disjunctive form.

C Proof of Lemma 26

This proof uses the property that if Ψ is in disjunctive form, then in the model-checking games for Ψ we can restrict ourselves to strategies for Even which see one play per branch. See for example [13]. Dually, co-disjunctive formulas yield such strategies for Odd.

Proof. (Lemma 26) The proof structure is similar to the less general version in [15] – here it just incorporates the additional challenges and resets. Furthermore,we use the properties of disjuntive formulas rather than König's Lemma.

Assume that Ψ is equivalent to some disjunctive Φ of index I and that for all m, in particular a fixed one larger than $(|\Phi|+|\Psi|)^2$, there is a structure $\mathfrak{T}$, such that Even wins the parity game $\mathfrak{T}\times\Psi$ but Odd wins the generalised m-challenge game on $\mathfrak{T}\times\Psi$. Without loss of generality, take $\mathfrak{T}$ to be finitely branching.

Part I. Let σ be Even's winning strategy in the parity game $\mathfrak{T}\times\Phi$ which only sees one play per branch. For a branch b of $\mathfrak{T}$, on which σ reaches a node v, indicate by $next_k(b,v)$, for even $k\in I$, the next node along b at which k or a higher even priority is seen. If k or higher is not seen after v, leave it undefined. For each odd $k\in I$ define R_k the set of nodes at which σ sees the odd priority k.

Now consider the generalised m-challenge game on the arena $\mathfrak{T}\times\Psi$. Let Even's challenging strategy γ be: to open all challenges at the start of the game, and whenever its counter is reset; if challenges for a priority j is met at v, and none of its counters $c_{i,j}$ is at 0, open the next challenge $a_{i,j}$ when the play reaches a node $next_{i-1}(b,v)$ for any branch b, unless the counter is reset before then. Even resets level k upon reaching a position in R_k.

Part II. Odd wins the generalised m-challenge game on $\mathfrak{T}\times\Psi$, so let τ be his winning strategy. Recall that τ_γ is an Odd's strategy for Ψ up to the point where an m^{th} challenge is met, and undefined thereafter. Since Ψ is co-disjunctive, we can adjust $\mathfrak{T}$ into a bisimilar structure in which the pure parity game strategy τ_γ is well-behaved wherever it is defined – it reaches each position of $\mathfrak{T}$ at either one subformula, or none.

The strategy τ_γ is winning in the challenge game against any strategy for Even which uses the challenging strategy γ. Even always eventually opens a challenge on every priority, at some level, and only resets level k infinitely often on a branch if all challenges at level $k+2$ are *open* infinitely often. Hence the only way for her to lose is to either lose finitely in the underlying parity game or for the play to reach a position of priority p when for some i, $c_{i,p}=0$. Thus, every play in the challenge game that agrees with τ_γ and the challenging strategy γ is finite.

Since τ_γ is well-behaved, each branch carries at most one play of the parity game. For every branch b with a play on it, the play either ends in a win for Odd in the underlying parity game, or in a long streak in which the highest priority seen is some odd p, and it is seen at least m times, corresponding to every instance of Odd meeting a p-challenge, set on some level k. As long as m is sufficiently large, on every such branch there are two nodes v and its descendant w, at which Even opens challenge $a_{k,p}$, which agree on the set of subformulas that τ_γ reaches there in $\mathfrak{T}\times\Psi$ and that σ reaches there in $\mathfrak{T}\times\Phi$. Now consider the structure $\mathfrak{T}'$, which is as $\mathfrak{T}$ except that the predecessor of each w-node has an edge to v instead. The strategies τ_γ and σ transfer in the obvious way to $\mathfrak{T}'$.

Part III. Finally, let us show that σ is winning in $\mathfrak{T}'\times\Phi$ and that τ_γ is winning in $\mathfrak{T}'\times\Psi$. Starting with τ_γ, consider plays that do not go through back edges infinitely often. They must finite and winning for Odd, as in the challenge game on $\mathfrak{T} \times \Psi$. Any play in $\mathfrak{T} \times \Psi$ that agrees with τ_γ which sees both v and w is dominated by an odd priority between v and w. Then, as the w and v agree on which subformula σ_γ reaches them at, an odd priority dominates any play that goes through back edges in $\mathfrak{T}' \times \Psi$ infinitely many times. The strategy τ_γ is therefore winning in $\mathfrak{T}' \times \Psi$.

Now onto σ in $\mathfrak{T}' \times \Phi$. If a branch is unchanged by the transformation, then any play on it is still winning for σ. If a branch that σ plays on has been changed, then consider in $\mathfrak{T}$ the two nodes v and w at which the transformation is done. These both are nodes at which according to γ, Even opens a $a_{k,p}$ challenge for the k and p for which Even runs out of challenges in the challenge game on $\mathfrak{T}$. Therefore, from the definition of $next_{k-1}$ and γ, σ sees $k-1$ between them and no odd priority k or larger, as witnessed by the lack of reset of k-counters. Since v and w agree on which subformula σ reaches them at, any play in $\mathfrak{T}' \times \Phi$ which goes through a back-edge infinitely often is dominated by the even priority $k-1$.

This contradicts the equivalence of Ψ and Φ. Therefore, if Ψ is semantically equivalent to a disjunctive formula of index I, then for all structures $\mathfrak{T}$ the m-challenge game and the parity game on $\mathfrak{T} \times \Psi$ have the same winner for $m > (|\Phi| + |\Psi|)^2$, *i.e*, Ψ is interpreted by $\mathsf{GChallenge}^m_{J,I}$ □