On the size of disjunctive formulas in the μ-calculus

Clemens Kupke*
University of Strathclyde
Glasgow, Scotland
clemens.kupke@strath.ac.uk

Johannes Marti† Yde Venema
ILLC, University of Amsterdam
Amsterdam, The Netherlands
johannes.marti@gmail.com y.venema@uva.nl

A key result in the theory of the modal μ-calculus is the disjunctive normal form theorem by Janin & Walukiewicz, stating that every μ-calculus formula is semantically equivalent to a so-called disjunctive formula. These disjunctive formulas have good computational properties and play a pivotal role in the theory of the modal μ-calculus. It is therefore an interesting question what the best normalisation procedure is for rewriting a formula into an equivalent disjunctive formula of minimal size. The best constructions that are known from the literature are automata-theoretic in nature and consist of a guarded transformation, i.e., the constructing of an equivalent guarded alternating automaton from a μ-calculus formula, followed by a Simulation Theorem stating that any such alternating automaton can be transformed into an equivalent non-deterministic one. Both of these transformations are exponential constructions, making the best normalisation procedure doubly exponential. Our key contribution presented here shows that the two parts of the normalisation procedure can be integrated, leading to a procedure that is single-exponential in the closure size of the formula.

1 Introduction

The modal μ-calculus [1] is a general modal logic enriched with fixpoint operators that allow to reason about the ongoing, possibly infinite behaviour of a system. The generality and complexity of the modal μ-calculus calls for research into fragments of the logic. On the one hand, this concerns fragments tailor-made for certain application domains such as the temporal logics LTL, CTL or dynamic logics [20]. On the other hand, one focuses on fragments of the μ-calculus that are either semantically or syntactically well-behaved and where a better understanding increases our knowledge about the full μ-calculus. One key fragment of the latter kind is formed by the so-called disjunctive formulas [13] . These are formulas, where the use of conjunctions is strictly limited to conjunctions of propositional atoms and the formula $\top$ (thought of as the empty conjunction). We state the formal definition here - the meaning of the ∇-operator will be discussed later.

Definition 1.1 *The set μDML of disjunctive μ-calculus formulas is given by the following grammar:*

$$\varphi ::= \bot \mid \top \mid \bigwedge L \wedge \nabla\Phi \mid (\varphi_1 \vee \varphi_2) \mid \mu p.\varphi \mid \nu p.\varphi$$

where L is a finite set of literals (i.e., propositional variables or their negations), Φ is a finite set of formulas in μDML, and p is a propositional variable. Furthermore we require that in a formula $\eta p.\varphi$ all occurrences of p in φ are positive, guarded and not in the context of a conjunction $p \wedge \psi$.

*Partially supported by Leverhulme grant RPG-2020-232.

†The research of this author has been made possible by a grant from the Dutch Research Council NWO, project nr. 617.001.857.

D. Bresolin and P. Ganty (Eds.): 12th International Symposium
on Games, Automata, Logics, and Formal Verification (GandALF 2021)
EPTCS 346, 2021, pp. 291–307, doi:10.4204/EPTCS.346.19

While disjunctive formulas correspond to a proper syntactic fragment of the μ-calculus, it is a somewhat surprising fact that each formula of the μ-calculus is semantically equivalent to a disjunctive one. This has many applications, e.g., satisfiability checking of a disjunctive formula can be carried out efficiently [13] (being ExpTime-complete in for arbitrary formulas [7, 8]) and disjunctive formulas facilitate the computation of uniform interpolants [4, 5]. Furthermore, disjunctive formulas play a pivotal role in the completeness theory of the modal μ-calculus [23, 9]. Finally, disjunctive formulas also provide insights for characterising important semantic fragments such as the continuous, additive and monotone fragments of the modal μ-calculus [10].

Recipes to rewrite a given arbitrary μ-calculus formula into an equivalent disjunctive one are well-known [13]. The size of the resulting disjunctive formula is crucial, in particular, in connection with satisfiability and uniform interpolation. The construction of a disjunctive formula usually proceeds in two stages: first a given μ-calculus formula is transformed into an equivalent, possibly alternating modal automaton, which is then transformed into an equivalent non-deterministic "disjunctive" modal automaton. The latter can be easily translated back into a disjunctive formula. We will argue that the outlined two-stage construction will inevitably lead to a double-exponential blow-up in the size of the formula. This is, because the first move from formulas to automata involves bringing the formula into a guarded format, i.e., a form where each fixpoint variable is in the scope of at least one modality. That guarding is problematic has been observed in the work by Bruse et al. [2] - we will argue that it necessarily entails an exponential blow-up in the size of the structures involved.

This sets the stage for our main result: a procedure that directly turns an arbitrary, possibly unguarded formula in the modal μ-calculus into an equivalent disjunctive automaton of exponential size. The latter can be turned easily into a disjunctive formula, which leads to our main theorem.

Theorem 1.2 *For any μ-calculus formula φ we can construct an equivalent disjunctive μ-calulus formula φ^d of size $2^{\mathcal{O}(n^2 k\cdot\log(nk))}$ and alternation depth $\mathcal{O}(n\cdot k)$, where $n = |\varphi|$ and where k is the alternation depth of φ.*

In the above theorem, the size of a formula refers to the size of its Fisher-Ladner closure, which has been shown in [2] to provide the tightest measure of formula size. For a discussion and comparison of different size measures see [15] where we also propose so-called "parity formulas" as a versatile tool to study the complexity of formula constructions. Parity formulas are a graph-shaped variant of μ-calculus formulas which is closely related to Wilke's alternating automata [24] and hierarchical equation systems [6]. While we stated the above theorem with reference to standard formulas, we will work throughout the paper with parity formulas instead. At the same time we will explain why Thm. 1.2 is a consequence of our work.

The outline of our paper is thus as follows: we will first introduce the necessary terminology for parity formulas, modal automata and their respective disjunctive variants. After that, in Section 3, we will discuss why guarding a parity formula can lead to an exponential blow-up. We then demonstrate that turning an arbitrary formula into an equivalent modal automaton is at least as costly as guarding a formula which means that the earlier mentioned two-stage method of constructing a disjunctive formula will in general lead to a double-exponential blow-up. Section 5 contains the central result of this paper, a construction that turns any given parity formula into an equivalent disjunctive modal automaton. This will provide a proof of Thm. 1.2.

Related Work. In addition to the already mentioned papers we would like to highlight a few more closely related lines of research. In spirit, our construction is related to the work by Friedmann & Lange on tableaux for unguarded μ-calculus formulas [11] but the cited paper is not concerned with disjunctive normal forms. Similarly, our automata theoretic result could be obtainable from a more general result in [21] - a key definition in that paper, however, appears to make an implicit assumption on the names of fixpoint variables ("cleanness") whereas our results from [15] demonstrate that cleaning a formula can lead to an exponential blow-up in (closure) size. In addition, it is not clear how to extract a disjunctive formula from the purely automata-theoretic constructions in [21]. Finally, the work by Lehtinen [17] studies how the alternation depth of a formula relates to the alternation depth of an equivalent disjunctive formula. While it turns out that the difference in alternation depth can be arbitrarily big, we note that this is not in conflict with the bound in our theorem, as we refer to a particular disjunctive equivalent as opposed to an arbitrary one.

Acknowledgements. We would like to thank the anonymous referees for valuable comments that helped to improve this paper.

2 Preliminaries

2.1 The μ-calculus and parity formulas

We will now recall the standard syntax of the modal μ-calculus and its reformulation in terms of parity formulas. It will be convenient for us to assume that μ-calculus formulas are in so-called negation normal form. We assume an infinite set of propositional variables, and define a *literal* to be either a propositional variable p or its negation $\overline{p}$.

Definition 2.1 *The formulas of the modal μ-calculus $\mu\mathtt{ML}$ are given by the following grammar:*

$$\varphi \ ::= \ \ell \mid \bot \mid \top \mid (\varphi_1 \vee \varphi_2) \mid (\varphi_1 \wedge \varphi_2) \mid \Diamond\varphi \mid \Box\varphi \mid \mu p.\varphi \mid \nu p.\varphi,$$

where ℓ is a literal, p is a propositional variable, and the formation of the formulas $\mu p.\varphi$ and $\nu p.\varphi$ is subject to the constraint that φ is positive *in p, i.e., there are no occurrences of $\overline{p}$ in φ. With $|\varphi|$ we denote the size of a formula measure in the number of distinct formulas in its Fisher-Ladner closure.*

We often restrict attention to formulas of which the free variables belong to some fixed finite set P; these are interpreted over Kripke models over P (in the following referred to as models), i.e., triples $\mathbb{S} = (S, R, \mathrm{Val})$ where S is a set of points, R is a binary relation and $\mathrm{Val} : \mathsf{P} \to \wp S$. We sometimes refer to the propositional type $c_s := \{p \in \mathsf{P} \mid s \in \mathrm{Val}(p)\}$ of $s \in S$ as the *colour* of s. A *pointed model* is a model $\mathbb{S}$ together with a designated point $s_I \in S$. Finally, it will be convenient to extend Val to all literals by putting $\mathrm{Val}(\overline{p}) = S \setminus \mathrm{Val}(p)$.

In this paper, we will not work with μ-calculus formulas in their usual shape, but with formulas represented as graphs, so-called "parity formulas". Parity formulas will facilitate discussing the complexity of our constructions. In addition, the fact that parity formulas resemble automata will simplify our proofs, as key constructions in our paper turn formulas into automata and vice versa. While parity formulas were introduced in [15] they are closely related to alternating automata [24] and hierarchical equation systems, see for instance [6]. A detailed discussion of the connections can be found in [15, Section 5]. Before giving the definition it will be useful to fix some terminology for directed graphs (which we will also apply to structures such as parity formulas that possess a directed graph structure). For binary relations $R \subseteq X \times X$ and $x \in X$ we will use the notation $R[x]$ to denote the set $\{x' \in X \mid (x, x') \in R\}$.

Definition 2.2 *Let (V, E) be a directed graph. A path π through (V, E) is a finite, non-empty sequence $\pi = v_0 \dots v_n \in V^*$ such that $v_{i+1} \in E[v_i]$ for all $i \in \{0, \dots, n-1\}$. We denote by first(π) and last(π) the first and last element of the path π, respectively. Concretely, for the above path we have first$(\pi) = v_0$ and last$(\pi) = v_n$. A path π with first$(\pi) =$ last(π) is called a cycle if it consists of at least two nodes.*

Position		Player	Admissible moves
(v,s)	with $L(v) = l$ and $s \in \mathrm{Val}(l)$	$\forall$	$\varnothing$
(v,s)	with $L(v) = l$ and $s \notin \mathrm{Val}(l)$	$\exists$	$\varnothing$
(v,s)	with $L(v) = \varepsilon$	-	$E[v] \times \{s\}$
(v,s)	with $L(v) = \vee$	$\exists$	$E[v] \times \{s\}$
(v,s)	with $L(v) = \wedge$	$\forall$	$E[v] \times \{s\}$
(v,s)	with $L(v) = \Diamond$	$\exists$	$E[v] \times R[s]$
(v,s)	with $L(v) = \Box$	$\forall$	$E[v] \times R[s]$

Table 1: The evaluation game $\mathscr{E}(\mathbb{G},\mathbb{S})$.

Given the set P of proposition letters, we let $\mathtt{Lit}(\mathsf{P})$ and $\mathtt{At}(\mathsf{P}) := \mathtt{Lit}(\mathsf{P}) \cup \{\bot, \top\}$ denote the set of *literals* and *atomic formulas over* P, respectively.

Definition 2.3 *Let* P *be a finite set of proposition letters. A* parity formula over P *is a quintuple* $\mathbb{G} = (V, E, L, \Omega, v_I)$*, where*

- (V,E) *is a finite, directed graph, with* $|E[v]| \leq 2$ *for every vertex* v*;*
- $L : V \to \mathtt{At}(\mathsf{P}) \cup \{\wedge, \vee, \Diamond, \Box, \varepsilon\}$ *is a labelling function;*
- $\Omega : V \stackrel{\circ}{\to} \omega$ *is a partial map, the* priority *map of* $\mathbb{G}$*; and*
- v_I *is a vertex in* V*, referred to as the* initial *node of* $\mathbb{G}$*;*

such that

1. $|E[v]| = 0$ *if* $L(v) \in \mathtt{At}(\mathsf{P})$*, and* $|E[v]| = 1$ *if* $L(v) \in \{\Diamond, \Box\} \cup \{\varepsilon\}$*;*
2. *every cycle of* (V,E) *contains at least one node in* $\mathsf{Dom}(\Omega)$.

A node $v \in V$ *is called* atomic *if it is either constant or literal,* boolean *if* $L(v) \in \{\wedge, \vee\}$*, and* modal *if* $L(v) \in \{\Diamond, \Box\}$*. We denote by* V_a*,* V_b *and* V_m *the collections of atomic, boolean and modal nodes, respectively. The elements of* $\mathsf{Dom}(\Omega)$ *will be called* states*. The* size *of a parity formula* $\mathbb{G} = (V, E, L, \Omega, v_I)$ *is defined as its number of nodes:* $|\mathbb{G}| := |V|$.

Definition 2.4 *Let* $\mathbb{S} = (S, R, \mathrm{Val})$ *be a Kripke model for a set* P *of proposition letters, and let* $\mathbb{G} = (V, E, L, \Omega, v_I)$ *be a parity formula over* P*. We define the* evaluation game $\mathscr{E}(\mathbb{G},\mathbb{S})$ *as the parity game* (G, E, Ω') *of which the board consists of the set* $V \times S$*, the priority map* $\Omega' : V \times S \to \omega$ *is given by*

$$\Omega'(v,s) := \begin{cases} \Omega(v) & \text{if } v \in \mathsf{Dom}(\Omega) \\ 0 & \text{otherwise}, \end{cases}$$

and the game graph is given in Table 1. Here all possible game positions are listed in the left column, the owner of a position is either $\forall$ *or* $\exists$[1] *and the set of possible moves is specified in the right column. As usual, finite plays of the game are lost by the player who owns the last position of the play from which no more move is possible ("the player who gets stuck loses"). An infinite play is won by* $\exists$ *if the maximum priority occurring infinitely often along the play is even, and by* $\forall$ *if it is odd.*

The parity formula $\mathbb{G}$ holds *at a point* s *if the pair* (v_I, s) *is winning for* $\exists$ *in the evaluation game.*

[1] Note that we do not need to assign a player to positions that admit a single move only.

A central complexity measure for both parity formulas and modal automata will be the so-called *index*. We define the index of a parity formula as the size of the range of its priority function Ω. We will rely on the following result from [15] that ensures that throughout the paper we are able to work on parity formulas instead of formulas in standard syntax.

Proposition 2.5 *There is an algorithm that constructs for any formula $\varphi \in \mu\mathtt{ML}$ an equivalent parity formula $\mathbb{G}_\varphi$ such that $|\mathbb{G}_\varphi| = |\varphi|$ and such that the index of $\mathbb{G}_\varphi$ is smaller or equal to the alternation depth of φ. Conversely, there is an algorithm that constructs for a given parity formula $\mathbb{G}$ an equivalent formula $\varphi_\mathbb{G} \in \mu\mathtt{ML}$ such that $|\varphi_\mathbb{G}| \leq 2 \cdot |\mathbb{G}|$ and such that the alternation depth of $\varphi_\mathbb{G}$ is smaller or equal to the index of $\mathbb{G}$.*

2.2 Modal Automata

Intuitively, modal automata correspond to parity formulas in a certain normal form - the precise connection will be discussed in Section 4 below. Modal automata are based on the *modal one-step language*. This language consists of very simple modal formulas, built up from a collection A of *variables*, which represent the states of the automaton and correspond to the fixpoint variables of a formula.

Definition 2.6 *Given a set A. The set $\mathtt{1ML}(A)$ of* modal one-step formulas *over A is inductively given as follows:*

$$\alpha ::= \bot \mid \top \mid \Diamond a \mid \Box a \mid \alpha \wedge \alpha \mid \alpha \vee \alpha,$$

where $a \in A$. We let $Sfor(\alpha)$ denote the collection of subformulas of a one-step formula α.

Definition 2.7 *Let P be a finite set of propositional variables. A* modal P-automaton $\mathbb{A}$ *is a quadruple (A, Δ, Ω, a_I) where A is a non-empty finite set of* states*, of which $a_I \in A$ is the* initial *state, $\Omega : A \to \omega$ is the* priority map*, and the* transition map $\Delta : A \times \wp\mathsf{P} \to \mathtt{1ML}(A)$ *maps pair of states and colors to one-step formulas.*

The size of a modal automaton is defined as follows.

Definition 2.8 *Let $\mathbb{A} = (A, \Delta, \Omega, a_I)$ be a modal automaton. We define its* state size $|\mathbb{A}|^s := |A|$*, its* size *as*

$$|\mathbb{A}| := \left|\bigcup\{Sfor(\alpha) \mid \alpha \in \mathsf{Ran}(\Delta)\}\right| + |\mathbb{A}|^s,$$

and its index *as $ind(\mathbb{A}) := |\mathsf{Ran}(\Omega)|$.*

Modal automata operate on pointed models, acceptance is defined via parity graph games.

Definition 2.9 *Let $\mathbb{A} = (A, \Delta, \Omega, a_I)$ be a modal automaton and let $(\mathbb{S}, s_I)$ be a model. The acceptance game $\mathscr{A}(\mathbb{A}, \mathbb{S})$ of $\mathbb{A}$ has the game board displayed in Table 2. A pointed model $(\mathbb{S}, s_I)$ is accepted by $\mathbb{A}$ if $\exists$ has a winning strategy at position (a_I, s_I) in $\mathscr{A}(\mathbb{A}, \mathbb{S})$.*

2.3 Disjunctive Formulas & Automata

In this section we introduce disjunctive formulas and their automata-theoretic pendant, so-called disjunctive automata. Disjunctive formulas can be best characterised in a modal language that has one “cover modality” ∇ that takes a finite set of formulas as its argument. Given such a set Φ, one may think of the formula $\nabla\Phi$ as the abbreviation

$$\nabla\Phi \equiv \bigwedge_{\varphi\in\Phi} \Diamond\varphi \wedge \Box \bigvee_{\varphi\in\Phi} \varphi.$$

Position	Player	Admissible moves
$(a,s) \in A \times S$	–	$\{(\Delta(a,c_s),s)\}$ with $c_s = \{p \in Q \mid s \in \mathrm{Val}(p)\}$
$(\bot,s) \in \mathtt{1ML}(A) \times S$	$\exists$	$\emptyset$
$(\top,s) \in \mathtt{1ML}(A) \times S$	$\forall$	$\emptyset$
$(\alpha_1 \vee \alpha_2, s)$	$\exists$	$\{(\alpha_1,s),(\alpha_2,s)\}$
$(\alpha_1 \wedge \alpha_2, s)$	$\forall$	$\{(\alpha_1,s),(\alpha_2,s)\}$
$(\Diamond a, s)$	$\exists$	$\{(a,s') \mid s' \in R[s]\}$
$(\Box a, s)$	$\forall$	$\{(a,s') \mid s' \in R[s]\}$

Table 2: The game board of the accepance game of a modal automaton.

Position		Player	Admissible moves
(v,s)	with $L(v) = \varepsilon$	-	$E[v] \times \{s\}$
(v,s)	with $L(v) = \top$	$\forall$	$\varnothing$
(v,s)	with $L(v) = \vee$	$\exists$	$E[v] \times \{s\}$
(v,s)	with $L(v) = \wedge_l$ and $s \notin \mathrm{Val}(l)$	$\exists$	$\varnothing$
(v,s)	with $L(v) = \wedge_l$ and $s \in \mathrm{Val}(l)$	$\forall$	$E[v] \times \{s\}$
(v,s)	with $L(v) = \nabla$	$\exists$	$\{Z \subseteq V \times S \mid (E[v],R[v]) \in \overline{Z}\}$
Z	$\subseteq V \times S$	$\forall$	$\{(v',s') \mid (v',s') \in Z\}$

Table 3: The evaluation game $\mathscr{E}(\mathbb{G},\mathbb{S})$.

It is called the "cover modality" since, intuitively, the formula $\nabla\Phi$ holds at a point s if the set of successors of s and the set of elements of Φ cover each other, in a sense that can be made precise using the notion of *relation lifting*. For a relation $Z \subseteq X \times Y$ we define its lifting $\overline{Z} \subseteq \wp X \times \wp Y$ by putting

$$(U,V) \in \overline{Z} \quad \text{if} \quad \forall x \in U.\exists y \in V.(x,y) \in Z \text{ and } \forall y \in V.\exists x \in U.(x,y) \in Z.$$

It is then easy to verify that the formula $\nabla\Phi$ holds at a point s if the pair $(R[s],\Phi)$ belongs to the lifting of the truth relation between points and formulas. The operator ∇ is well-known from the literature, cf. e.g. [13, 22, 17].

Definition 2.10 *Let* P *be a finite set of propositional variables. A* disjunctive parity formula *over* P *is a quintuple* $\mathbb{G} = (V,E,L,\Omega,v_I)$ *such that*

- $L : V \to \{\wedge_l \mid l \in \mathtt{Lit}(\mathsf{P})\} \cup \{\nabla, \vee, \top, \varepsilon\}$ *is a labelling function;*
- *for all* $v \in V$ *with* $L(v) \neq \nabla$ *we have* $|E[v]| \leq 2$

and such that all other conditions of the definition of parity formulas in Def. 2.3 are met. The board of the evaluation game $\mathscr{E}(\mathbb{G},\mathbb{S})$ *of a disjunctive parity formula on a model* $\mathbb{S}$ *is displayed in Table 3.*

Intuitively, a node v labelled with ∇ represents the formula $\nabla\{\varphi_w \mid w \in E[v]\}$, where for each $w \in E[v]$ we write φ_w for the formula represented by w. Furthermore $\wedge_l$ is intended to be a unary operator that represents the conjunction of its argument with l. In other words, disjunctive formulas are formulas that contain conjunctions only in the form of conjunctions with literals and in the form of $\top$ that can be thought of as the empty conjunction. Disjunctive formulas have their automata-theoretic pendant, the so-call disjunctive modal automata – the so-called μ-*automata* of Janin & Walukiewicz [13]. These are defined by restricting the shape of transition conditions.

Definition 2.11 *Given a finite set A, we define the set* $\mathtt{1DML}(A)$ *of* disjunctive modal one-step formulas *over A via the following grammar:*

$$\alpha ::= \bot \mid \top \mid \nabla B \mid \alpha \vee \alpha,$$

where $B \subseteq A$. *A disjunctive modal* P*-automaton is a tuple* $\mathbb{A} = (A, \Delta, \Omega, a_I)$ *such that* $\Delta : A \times \wp \mathsf{P} \to \mathtt{1DML}(A)$. *The acceptance game* $\mathscr{A}(\mathbb{A}, \mathbb{S})$ *on a model* $\mathbb{S}$ *is defined as for general modal automata with the difference that the rule for* $\wedge$ *no longer applies and that the rules for* $\Box$ *and* $\Diamond$ *are replaced by*

Position	*Player*	*Admissible moves*
$(\nabla B, s)$	$\exists$	$\{Z \subseteq A \times S \mid (B, R[s]) \in \overline{Z}\}$
$Z \subseteq A \times S$	$\forall$	$\{(v', s') \mid (v', s') \in Z\}$

3 Guarding Revisited

Existing approaches for turning a μ-calculus formula into a modal automaton rely on the assumption that the input formula is guarded. As [2] have shown this assumption is problematic because existing algorithms for guarding formulas, which have long been thought to be polynomial, are in fact exponential. In this section we discuss two results on the complexity of guarding formulas. We do this in the setting of parity formulas.

Definition 3.1 *A path* $\pi = v_0 v_1 \cdots v_n$ *through* $\mathbb{G}$ *is* unguarded *if* $n \geq 1$, $v_0, v_n \in \mathsf{Dom}(\Omega)$ *while there is no i, with* $0 < i \leq n$, *such that* v_i *is a modal node. A parity formula is* guarded *if it has no unguarded cycles, and* strongly guarded *if it has no unguarded paths.*

Adapting the well-known construction for guarding formulas one can show that it is possible to guard parity formulas, with an exponential blow-up in the number of states [15].

Theorem 3.2 *There is an algorithm that transforms a parity formula* $\mathbb{G} = (V, E, L, \Omega, v_I)$ *into a strongly guarded parity formula* $\mathbb{G}^g$ *such that*

1) $\mathbb{G}^g \equiv \mathbb{G}$;
2) $|\mathbb{G}^g| \leq 2^{1+|\mathsf{Dom}(\Omega)|} \cdot |\mathbb{G}|$;
3) $ind(\mathbb{G}^g) \leq ind(\mathbb{G})$.

It is unclear whether this result can be improved such that the number of states of $\mathbb{G}^g$ is polynomial in the number of states in $\mathbb{G}$. The results in section 4 of Bruse, Friedmann & Lange [2] show that certain guarded transformation procedures are as hard[2] as solving parity games. Theorem 3.3 below can be seen as our parity-formula version of this observation. Our proof is in fact simpler because we can exploit the close connection between parity games and parity formulas and thus do not need the product construction from [14] that is used for the results from [2].

Theorem 3.3 *If there is a procedure that runs in polynomial time and transforms a parity formula* $\mathbb{G}$ *to a guarded parity formula* $\mathbb{G}^\gamma$ *with* $\mathbb{G}^\gamma \equiv \mathbb{G}$ *then solving parity games is in* PTIME.

For a proof of Theorem 3.3 we refer to [15], where we also discuss in some detail the relation with other results in [2].

[2] It is an open question whether parity games can be solved in polynomial time. Despite considerable efforts no polynomial algorithm has been found so far. In the recent literature, however, various quasi-polynomial algorithms have been given, following the breakthrough work of Calude et alii [3].

4 Modal Automata and Strongly Guarded Parity Formulas

In this section we establish a close connection between modal automata and parity formulas. We will first see that a modal automaton can be turned into a strongly guarded parity formula that is of linear size if we ignore propositional variables. In particular this shows that turning a parity formula into an equivalent modal automaton is at least as hard as the guarding construction from the previous section (hardness of the latter does not depend on formulas containing propositional variables). We close by showing how to turn a parity formula into an equivalent modal automaton of exponential size. Collectively the results from this section will show that constructing a disjunctive modal automaton from a parity formula by first turning the latter into an equivalent modal automaton to which we then apply a known "non-determinisation" construction would yield a doubly exponential blow-up (in closure size). This sets the stage for our main result in the next section where a new construction that is single exponential in closure size from parity automata to disjunctive modal automata is provided.

Theorem 4.1 *There is an algorithm that constructs, given a modal automaton $\mathbb{A}$, a strongly guarded parity formula $\mathbb{G}$ such that*
1) $\mathbb{G} \equiv \mathbb{A}$;
2) $|\mathbb{G}| \leq 2^{|\mathsf{P}|} \cdot |\mathbb{A}|$;
3) $ind(\mathbb{G}) \leq ind(\mathbb{A})$.
4) If $\mathbb{A}$ is disjunctive then so is $\mathbb{G}$.

Proof. We only sketch the construction. For each $a \in A$ we let $\Delta'(a) = \bigvee_{c \in \wp \mathsf{P}}(\bigwedge_{p \in c} p \wedge \bigwedge_{p \notin c} \overline{p} \wedge \Delta(a,c))$. Given a modal automaton $\mathbb{A} = (A, \Delta, \Omega, a_I)$ we define the set of nodes of $\mathbb{G}$ by

$$V = A \cup \bigcup \{Sfor(\alpha) \mid \alpha \in \mathsf{Ran}(\Delta')\}$$

To defined the edge relation E of $\mathbb{G}$ we put $E[a] = \Delta'(a)$ for all $a \in A$. For all other elements of V we let E be the "immediate subformula" relation, e.g. $E[\alpha_1 \wedge \alpha_2] = \{\alpha_1, \alpha_2\}$, $E[\Diamond a] = \{a\}$, etc. The priority map $\Omega_{\mathbb{G}}$ of $\mathbb{G}$ assigns to each element $a \in A$ the priority $\Omega(a)$. The initial state v_I of $\mathbb{G}$ is defined as $v_I = a_I$. Finally, the map L assigns to each element $a \in A$ the label ε and for each formula α the label consists of the top-most operator of α. It is easy to see that $\mathbb{G}$ thus defined satisfies conditions 2) and 3). The proof of condition 1) is following a standard argument and is omitted here. Concerning 4) it suffices to note that for a disjunctive $a \in A$ we put

$$\Delta'(a) = \bigvee_{c \in \wp \mathsf{P}} (\wedge_{l_1} (\cdots \wedge_{l_m} (\Delta(a,c))))$$

where $l_1, \ldots, l_n$ is an enumeration of all propositional variables in c and the negation of all variables not in c. Otherwise $\mathbb{G}$ is defined as in the non-disjunctive case and the result is a disjunctive formula. QED

Theorem 4.2 *There is an algorithm that constructs, given a parity formula $\mathbb{G}$, a modal automaton $\mathbb{A}$ such that*
1) $\mathbb{A} \equiv \mathbb{G}$;
2) $|\mathbb{A}|^s \leq 2^{1+|\mathsf{Dom}(\Omega)|} \cdot |\mathbb{G}|$ and $|\mathbb{A}| \leq 2^{|\mathsf{P}|+1+|\mathsf{Dom}(\Omega)|} \cdot |\mathbb{G}|$;
3) $ind(\mathbb{A}) \leq ind(\mathbb{G})$.

Proof. By Theorem 3.2 we may effectively construct from $\mathbb{G}$ an equivalent, strongly guarded parity formula $\mathbb{H} = (V, E, L, \Omega, v_I)$ such that $|\mathbb{H}| \leq 2^{1+|\mathsf{Dom}(\Omega)|} \cdot |\mathbb{G}|$ and $ind(\mathbb{H}) \leq ind(\mathbb{G})$. As shown in [15], we may additionally assume that in $\mathbb{H}$, every predecessor of a node $v \in \mathsf{Dom}(\Omega)$ is a modal node. The state space of the modal automaton $\mathbb{A}$ will be given as the set $A := E[V_m] \cup \{v_I\}$, so that by the assumption on $\mathbb{H}$ every state of $\mathbb{H}$ is a state of $\mathbb{A}$. In addition, v_I and possible other successors of modal nodes are states of $\mathbb{A}$ as well. We can then simply define $\Omega_{\mathbb{A}} := \Omega$, and take v_I as the initial state of $\mathbb{A}$. It remains to define the transition function Δ of $\mathbb{A}$.

Our first step will be to associate, with each node $v \in V$, a formula $\alpha(v)$, which belongs to the collection $\mathtt{1AML}(\mathsf{P}, A)$ of *alternative* one-step formulas given by the following grammar:

$$\alpha ::= \bot \mid \top \mid p \mid \overline{p} \mid \Diamond a \mid \Box a \mid \alpha \wedge \alpha \mid \alpha \vee \alpha,$$

where $p \in \mathsf{P}$ and $a \in A$. As the formula $\mathbb{H}$ is strongly guarded and by the additional property that $E^{-1}[\mathsf{Dom}(\Omega)] \subseteq V_m$, there is a unique map $\alpha : V \to \mathtt{1AML}(\mathsf{P}, A)$ which satisfies the following conditions:

$$\alpha(v) = \begin{cases} L(v) & \text{if } v \in V_l \\ \heartsuit a \quad \text{where } L(v) = \heartsuit \text{ and } E[v] = \{a\} & \text{if } v \in V_m \\ \bigodot \{\alpha(u) \mid u \in E[v]\} \text{ where } L(v) = \odot & \text{if } v \in V_b \\ \alpha(u) \quad \text{where } E[v] = \{u\} & \text{if } v \in V_\varepsilon \end{cases}$$

We can now define the transition map $\Delta : A \times \wp(\mathsf{P}) \to \mathtt{1ML}(A)$ as follows. For each state $a \in A$ and color $c \in \wp(\mathsf{P})$ we define the formula $\Delta(a, c)$ as $\Delta(a, c) := \alpha(a)[\sigma_c]$, where the substitution $\sigma_c : \mathtt{1AML}(\mathsf{P}, A) \to \mathtt{1ML}(A)$ is given by putting

$$\sigma_c(p) := \begin{cases} \top & \text{if } p \in c \\ \bot & \text{if } p \notin c. \end{cases}$$

This completes the definition of the automaton $\mathbb{A}$. It is easy to see that $|\mathbb{A}|^s \leq |V| \leq 2^{1+|\mathsf{Dom}(\Omega)|} \cdot |\mathbb{G}|$, that $|\mathbb{A}| \leq 2^{|\mathsf{P}|} \times V \leq 2^{|\mathsf{P}|+1+|\mathsf{Dom}(\Omega)|} \cdot |\mathbb{G}|$ and that $ind(\mathbb{A}) = ind(\mathbb{H}) \leq ind(\mathbb{G})$, which proves the items 2) and 3) of the theorem. The equivalence of $\mathbb{A}$ and $\mathbb{H}$ (and thus, of $\mathbb{A}$ and $\mathbb{G}$) can be proved by a routine argument. QED

5 The Simulation Theorem

The main result of this section, and the main technical contribution of the paper, is the following theorem.

Theorem 5.1 *Let $\mathbb{G}$ be a parity formula of size n and index k with propositional variables contained in P with $|\mathsf{P}| = l$. Then we can effectively construct a disjunctive modal automaton*[3] $\mathbb{A} = (A, \Theta, \Omega, a_I)$ *with* $|A| \leq 2^{n^2 k}$, $|\mathbb{A}| \leq n 2^{n^2 k + l + n}$ *and* $\mathbb{G} \equiv \mathbb{A}$.

Convention. In the remainder of this section we fix a parity formula $\mathbb{G}$ with $\mathbb{G} = (V, E, L, \Omega, v_I)$ with $|V| = n$ and $|\mathsf{Ran}(\Omega)| = k$. It will be convenient to make the following assumptions on $\mathbb{G}$: (i) Ω is total, (ii) $L^{-1}(\varepsilon) = \varnothing$, and (iii) $E[v] \neq \varnothing$ if $L(v) \in \{\wedge, \vee\}$. We leave it for the reader to convince themselves that this is without loss of generality. Furthermore we define $V_\vee := L^{-1}(\vee)$, etc. We now turn to the proof of Theorem 5.1.

[3] For the time being this will be an automaton with a regular acceptance condition. We will transform this into an automaton with a parity condition later.

5.1 Macrostates

We shall construct the simulating automaton via a powerset construction. That is, for the states of $\mathbb{A}$ we will in principle take subsets of $\mathbb{G}$. However, in order to handle infinite matches correctly we need to store some more information in these states: A state of $\mathbb{A}$ will be a *macrostate* over $\mathbb{G}$, that is, a ternary relation $m \subseteq V \times \mathsf{Ran}(\Omega) \times V$, representing various pieces of information. Basically, each triple $(u,k,v) \in m$ represents the projection on $\mathbb{G}$ of a partial play of the evaluation game of $\mathbb{G}$ which starts at u, ends at v, and reaches k as its highest priority (after u). More precisely, the triple (u,k,v) represents a path π through $\mathbb{G}$ with $\mathit{first}(\pi) = u$, $\mathit{last}(\pi) = v$, and such that k is the highest priority reached on π after u. Consequently, one single match of the acceptance game of $\mathbb{A}$ on a pointed Kripke structure will represent a certain bundle of matches of evaluation game of $\mathbb{G}$.

Before defining and discussing macrostates formally, we need some auxiliary terminology.

Definition 5.2 *A subset $U \subseteq V$ is* inconsistent *if there is $u \in U$ with $L(u) = \bot$, or if there are nodes $u, v \in U$ with $L(u) = p$ and $L(v) = \overline{p}$ for some $p \in \mathsf{P}$. Given a color $c \in \wp\mathsf{P}$ we say that U is* compatible *with c if $L(u) \neq \bot$, $L(u) = p$ implies $p \in c$, and $L(u) = \overline{p}$ implies $p \notin c$, for all $u \in U$ and $p \in \mathsf{P}$.*

Definition 5.3 *We define the set M_Ω of* macrostates *of $\mathbb{G}$ by putting $M_\Omega := \wp(V \times \mathsf{Ran}(\Omega) \times V)$. The* range $\mathsf{Ran}(m)$ *of a macrostate $m \in M_\Omega$ is the set of all $v \in V$ such that $(u,k,v) \in m$ for some $u \in V$ and $k \in \mathsf{Ran}(\Omega)$. With $m, m' \in M_\Omega$, we define the* composition *$m\,;m' \in M_\Omega$ as the set of triples $(v,k,v'') \in V \times \mathsf{Ran}(\Omega) \times V$ for which we can find $(v,k',v') \in m$ and $(v',k'',v'') \in m'$ such that $k = \max(k',k'')$. For a subset $U \subseteq V$, we define $\Delta_U := \{(u,0,u) \mid u \in U\}$; where $v \in V$, we abbreviate $\Delta_v := \Delta_{\{v\}}$.*

A macrostate m is called consistent*, respectively* compatible *with a color $c \in \wp\mathsf{P}$, if $\mathsf{Ran}(m) \subseteq V$ satisfies the mentioned property w.r.t. c, in the sense of Definition 5.2.*

Given a stream $\alpha = (m_i)_{i\in\omega}$ of macrostates, we say that a stream $(v_i,k_i)_{i\in\omega} \in (V \times \mathsf{Ran}(\Omega))^\omega$ is a trace *on α if $(v_i,k_i,v_{i+1}) \in m_i$, for all $i \in \omega$. Such a trace is* good *(*bad*, respectively) if the maximum number k occurring as k_i for infinitely many i is even (odd, respectively). We let NBT_Ω denote the collection of M_Ω-streams that do not carry a bad trace.*

Proposition 5.4 *The set NBT_Ω is an ω-regular language over M_Ω. Concretely, there is a deterministic parity automaton $\mathbb{P}$ over M_Ω such that $L_\omega(\mathbb{P}) = \mathrm{NBT}_\Omega$ and $\mathbb{P}$ has size $2^{\mathcal{O}(nk\cdot\log(nk))}$ and index $\mathcal{O}(nk)$.*

Proof. We first observe that there is a non-deterministic parity word automaton $\mathbb{W} = (Q, q_I, \Delta_{\mathbb{W}}, \Omega_{\mathbb{W}})$ that accepts the language $\mathscr{L}_{\text{bad}} := (M_\Omega)^\omega \setminus \mathrm{NBT}_\Omega$, i.e., all infinite streams of macrostates that do contain a bad trace. To define $\mathbb{W}$ we put $Q := V$, $q_I := v_I$, $\Delta_{\mathbb{W}}(v,m) := \{u \in V \mid \exists k.(v,k,u) \in m\}$ and $\Omega_{\mathbb{W}}(v) := \Omega(v) + 1$ for all $v \in V$ and all $m \in M_\Omega$. It is easy to verify $\mathbb{W}$ is a parity automaton with n states and index k that accepts $\mathscr{L}_{\text{bad}}$. Standard constructions can be used to first transform $\mathbb{W}$ into an equivalent non-deterministic Büchi automaton $\mathbb{W}'$ of size $\mathcal{O}(n \cdot k)$ (cf. e.g. [12]) which can be in turn transformed into an equivalent deterministic parity word automaton $\mathbb{W}''$ that meets the size bounds of the proposition (cf. [18, 19]). The automaton $\mathbb{P}$ is now constructed as the deterministic parity word automaton that accepts the complement of the language of $\mathbb{W}''$ by adding 1 to all the priorities of states in $\mathbb{W}''$. QED

5.2 Local strategies & the disjunctive modal automaton

In our approach of dealing with the possible unguardedness of the input parity formula, the key concept is that of a (positional) *local strategy* for $\exists$. A local strategy represents a complete set of choices of $\exists$ for all disjunction nodes in $\mathbb{G}$. Intuitively, one may think of a local strategy as some part of a positional strategy of $\exists$ where we stay put at a point in the model. More precisely, a local strategy χ induces, in the evaluation game of $\mathbb{G}$, for any point in the model and any vertex in $\mathbb{G}$, a unique (partial) play that does not leave the mentioned point and stops whenever a modal vertex in $\mathbb{G}$ is met. The projections of these matches will be called *stationary plays*. Formally we define these notion, together with some related concepts that we will discuss in a moment, as follows.

Definition 5.5 *A* local strategy *on $\mathbb{G}$ is a map $\chi : V_\vee \to V$ such that $\chi(v) \in E[v]$, for all $v \in V_\vee$. The collection of local strategies on $\mathbb{G}$ is denoted by $LS_{\mathbb{G}}$.*

Now fix such a local strategy χ. Given a vertex $v \in V$ we define the set $SP_\chi(v)$ of stationary χ-plays *from v as the smallest set $S \subseteq V^*$ such that*
(1a) if $v \in V_\vee$ then $v \cdot \chi(v) \in S$;
(1b) if $v \in V_\wedge$ then $v \cdot w \in S$, for each $w \in E[v]$;
(2a) if $\pi \in S$ and $u := last(\pi) \in V_\vee$, then $\pi \cdot \chi(u) \in S$; and
(2b) if $\pi \in S$ and $u := last(\pi) \in V_\wedge$, then $\pi \cdot w \in S$, for each $w \in E[u]$.
Given $\pi = vv_1 \cdots v_k \in SP_\chi(v)$, define $\widetilde{\Omega}(\pi) := \max\{\Omega(v_i) \mid 1 \leq i \leq k\}$. Via these stationary plays we define the following macrostates:

$$\begin{array}{rcl} e^-_\chi & := & \big\{(v,n,u) \quad \mid v \in V_b \text{ and } u = last(\pi) \text{ for some } \pi \in SP_\chi(v) \text{ with } n = \widetilde{\Omega}(\pi)\big\} \\ e_\chi & := & e^-_\chi \cup \Delta_V \end{array}$$

We say that on a macrostate m, χ is locally compatible *with a color $c \in \wp\mathsf{P}$ if (i) $\mathsf{Ran}(m\,;e_\chi)$ is compatible with c and (ii) the stream $m\,;(e^-_\chi)^\omega$ does not contain any bad trace.*

Here are some intuitions about these notions. First, note that $SP_\chi(v)$ only contains paths of length > 1, corresponding to matches where actually a move is made at v. The macrostate e^-_χ represents the relevant information about all finite stationary χ-plays; hence, the collection of all *infinite* stationary χ-plays corresponds to the set of all traces over the M_Ω-stream $(e^-_\chi)^\omega$.

Now fix a macro-state m. The macrostate $m\,;e^-_\chi$ represents the version of m that absorbs all continuations of matches in m with one of these finite stationary χ-plays; thus the set $m \cup (m\,;e^-_\chi)$ represents all triples in m that are possibly continued with such a play. For technical reasons it is convenient to define this set using the macro-state e_χ, in the sense that $m\,;e_\chi = m\,;(e^-_\chi \cup \Delta_V) = (m\,;e^-_\chi) \cup (m\,;\Delta_V) = (m\,;e^-_\chi) \cup m$. The stream $m\,;(e^-_\chi)^\omega$ represents all *infinite* plays that start with an m-play, and then continue with an infinite stationary χ-play. To say that there is a bad trace on such a stream means that for some $v \in \mathsf{Ran}(m)$, $\exists$'s *opponent* $\forall$ has a strategy such that, played against χ, the resulting match (path through the graph of $\mathbb{G}$) is bad, in the sense of the highest priority met infinitely often being odd. To say that, relative to m, χ is locally compatible with a color c indicates, roughly, that after any m-match, if the local point of the Kripke model has color c, then it is safe for $\exists$ (until the next modal vertex in $\mathbb{G}$ is encountered) to continue by playing χ. Finally, it may also be useful to observe that for all m and χ we find $\mathsf{Ran}(m) \cap V_p \subseteq \mathsf{Ran}(m\,;e_\chi) \cap V_p$ and $\mathsf{Ran}(m) \cap V_m \subseteq \mathsf{Ran}(m\,;e_\chi) \cap V_m$.

We turn to the definition of the disjunctive modal automaton $\mathbb{A}_{\mathbb{G}}$. For the definition of its transition map Θ, note a macrostate m may be thought of as representing the conjunction of states in $\mathsf{Ran}(m)$ that are visited by "parallel" plays of the evaluation game for $\mathbb{G}$. The transition function Θ will implement the intuition that a modal step needs to satisfy the "demands" posed by the modal nodes in $\mathsf{Ran}(m)$. These demands are formulated separately for all box nodes and for each individual diamond node.

Definition 5.6 *Let $m \in M_\Omega$ be some macrostate, and let $x \in \mathsf{Ran}(m) \cap V_\Diamond$. Then we define*

$$\begin{array}{rcl} d_\Box(m) & := & \{(u,\Omega(v),v) \mid u \in \mathsf{Ran}(m) \cap V_\Box \text{ and } v \in E[u]\}, \\ d_x(m) & := & \{(x,\Omega(v),v) \mid v \in E[x]\} \cup d_\Box(m). \end{array}$$

The macrostates $d_\Box(m)$ and $d_x(m)$ correspond to, respectively, the *universal* and *existential* requirements made by the vertices in the range of m.

Definition 5.7 *Let $\mathbb{G} = (V, E, L, \Omega, v_I)$ be a parity formula. We define the automaton $\mathbb{A}_\mathbb{G} = (A, \Theta, Acc, m_I)$ as follows. To start with, its carrier is the collection of all macrostates: $A := M_\Omega$ and its initial state is given as $m_I = \Delta_{v_I} = \{(v_I, 0, v_I)\}$. The transition map Θ is given as follows. For a macrostate $m \in M_\Omega$ and a local strategy χ we define $A_{m,\chi} \subseteq M_\Omega$:*

$$A_{m,\chi} \quad := \quad \{e_\chi \,;\, d_\Box(m\,;\,e_\chi)\} \cup \{e_\chi \,;\, d_x(m\,;\,e_\chi) \mid x \in \mathsf{Ran}(m\,;\,e_\chi) \cap V_\Diamond\},$$

and for a macrostate $m \in M_\Omega$ and a color $c \in \wp\mathsf{P}$ we put:

$$\Theta(m,c) := \bigvee \big\{\theta(m,.c,\chi) \mid \chi \in LS_\mathbb{G} \text{ is locally compatible with } c \text{ on } m\big\},$$

where

$$\theta(m,c,\chi) := \begin{cases} \nabla A_{m,\chi} & \text{if } \mathsf{Ran}(m\,;\,e_\chi) \cap V_\Diamond \neq \varnothing \\ \nabla A_{m,\chi} \vee \nabla\varnothing & \text{if } \mathsf{Ran}(m\,;\,e_\chi) \cap V_\Diamond = \varnothing. \end{cases}$$

Finally, for its acceptance condition Acc we take the ω-regular language NBT_Ω, i.e, an infinite play of $\mathscr{A}$ will be winning for $\exists$ if the corresponding stream of macrostates is in NBT_Ω.

To get some intuitions: to define $\Theta(m,c)$, we nondeterministically *guess* a local strategy χ that is compatible with c on m – this guess is reflected by the disjunction in the definition of $\Theta(m,c)$. For each such χ, we absorb its stationary plays into m and turn to the set of modal nodes in the range of the resulting macro-state $m\,;\,e_\chi$. We gather the universal and existential requirements of $m\,;\,e_\chi$ into an appropriate collection $A_{m,\chi}$ of "next" macro-states. This set $A_{m,\chi}$ is then to be *covered* by the collection of successors of the point in the Kripke model under inspection, as encoded by the formula $\nabla A_{m,\chi}$. In the special case where $m\,;\,e_\chi$ makes no existential demands (i.e., if $\mathsf{Ran}(m\,;\,e_\chi) \cap V_\Diamond = \varnothing$), we add the disjunct $\nabla\varnothing \equiv \Box\bot$, allowing for the possibility that the point has no successors at all. To see how this all works out precisely, the reader is advised to look at the proof of Proposition 5.8 below.

5.3 Proof of Theorem 5.1

Turning to the proof of the main theorem, we first show that the disjunctive modal automaton from Definition 5.7 is equivalent to the parity formula $\mathbb{G}$ that we started with. After that we prove Theorem 5.1 by showing that $\mathbb{A}_\mathbb{G}$ is within the desired size bounds.

Proposition 5.8 *For any parity formula $\mathbb{G}$, we have $\mathbb{G} \equiv \mathbb{A}$, where $\mathbb{A} := \mathbb{A}_\mathbb{G}$ is given as in Definition 5.7.*

Proof. We show that

$$(v_I, s_I) \in \mathrm{Win}_\exists(\mathscr{E}(\mathbb{G},\mathbb{S})) \text{ iff } (m_I, s_I) \in \mathrm{Win}_\exists(\mathscr{A}(\mathbb{A},\mathbb{S})), \tag{1}$$

where $(\mathbb{S}, s_I)$ is an arbitrary but fixed pointed model, $m_I := \Delta_{v_I}$, and we write $\mathrm{Win}_\exists$ for the set of winning positions for $\exists$ in a game. In the sequel we will abbreviate $\mathscr{E} := \mathscr{E}(\mathbb{G},\mathbb{S})$ and $\mathscr{A} := \mathscr{A}(\mathbb{A},\mathbb{S})$.

For the direction from left to right of (1), fix a positional winning strategy f for $\exists$ in $\mathscr{E}$. For any point $s \in S$, we may associate a map $\chi_s : V_\vee \to V$ with f, as follows. Given a vertex $u \in V_\vee$, if $(u,s) \in \mathrm{Win}_\exists(\mathscr{E})$ then χ_s maps u to the element $v \in E[u]$ such that $(v,s) = f(u,s)$; if $(u,s) \notin \mathrm{Win}_\exists(\mathscr{E})$ we define $\chi_s(u)$ to be an arbitrary element in $E[u]$. Clearly χ_s is a local strategy in the sense of Definition 5.5, and it is not hard to prove that χ_s is locally compatible with the color $\mathrm{Val}(s)$ of s, on any $m \in M_\Omega$ such that $(v,s) \in \mathrm{Win}_\exists(\mathscr{E})$ for any $v \in \mathsf{Ran}(m)$. We define the following (positional) strategy f' for $\exists$ in $\mathscr{A}$. Let

Σ be a partial $\mathscr{A}$-match with $last(\Sigma) = (m,s)$. In case $\mathsf{Ran}(m) \times \{s\} \not\subseteq \mathrm{Win}_\exists(\mathscr{E})$ then $\exists$ plays randomly (one may show that this will never happen). If, on the other hand, $\mathsf{Ran}(m) \times \{s\} \subseteq \mathrm{Win}_\exists(\mathscr{E})$, then we already saw that χ_s is locally compatible with $\mathrm{Val}(s)$ on m. Compute $e := e_{\chi_s}$ and recall that with each element $(u,k,v) \in e$ we may associate a partial f-guided match $\pi : (u,s)\cdots(v,s)$, which is stationary at s and such that k is the highest priority met on π after u: $k = \widetilde{\Omega}(\pi)$. It is then not hard to see that, for an arbitrary element $w \in \mathsf{Ran}(m\,;e)$ we have $(w,s) \in \mathrm{Win}_\exists(\mathscr{E})$. In particular, if such a w belongs to $V_\Diamond$, then $\exists$'s winning strategy f in $\mathscr{E}$ selects a successor $t_w \in R[s]$. This assignment, $V_\Diamond \ni w \mapsto t_w$, determines $\exists$'s strategy f' in $\mathscr{A}$. That is, at any partial match Σ with $last(\Sigma) = (m,s)$, let $\exists$ play

$$f'(m,s) := \Big\{\big(e\,;d_\Box(m\,;e),t\big) \mid t \in R[s]\Big\} \cup \Big\{\big(e\,;d_w(m\,;e),t_w\big) \mid w \in V_\Diamond \cap \mathsf{Ran}(m\,;e)\Big\}. \qquad (2)$$

It is easy to see that this move is legitimate. To show that f' is actually winning for $\exists$ consider an arbitrary f'-guided partial match $\Sigma = (m_0,s_0)\cdots(m_n,s_n)$ with $(m_0,s_0) = (m_I,s_I)$. Via an inductive proof one can show that $\mathsf{Ran}(m_i) \subseteq \{v \in V \mid (v,s_i) \in \mathrm{Win}_\exists(\mathscr{E})\}$, for each $i \leq m$. Therefore $\exists$'s moves in Σ are indeed legitimately guided as in (2) and she wins every finite f'-guided match of $\mathscr{A}$ starting at (m_I,s_I).

To see that $\exists$ also wins the infinite f'-guided matches, let $\Sigma = (m_i,s_i)_{i\in\omega}$ be an arbitrary such match, and consider an arbitrary trace $(v_i,k_i)_{i\in\omega}$ on Σ. Let k be the maximum number occurring as k_i for infinitely many i; it then suffices to show that k is even. For every $i < \omega$ one may associate f-guided partial $\mathscr{E}$-matches $\sigma_i = (u^i_0,s_i)(u^i_1,s_i)\cdots(u^i_{n_i},s_i)$ and $\sigma^+_i := \sigma_i \cdot (u^{i+1}_0,s_{i+1})$ such that $v_i = u^i_0$, $v_{i+1} = u^{i+1}_0$ and $k = \widetilde{\Omega}(\sigma^+_i)$. Putting these partial plays together, with the trace $(v_i,k_i)_{i\in\omega}$ we have thus associated a (full) infinite f-guided $\mathscr{E}$-match $\sigma = \sigma_0\sigma_1\cdots$, such that $\sigma_i = (u^i_0,s_i)(u^i_1,s_i)\cdots(u^i_{n_i},s_i)$ and $k_i = \max\big\{\Omega(\sigma^i_1),\ldots,\Omega(\sigma^i_{n_i}),\Omega(\sigma^{i+1}_0)\big\}$. This means that k is the highest priority that occurs infinitely often in σ, and since σ is guided by $\exists$'s winning strategy f, k must be even indeed.

For the other direction of (1), fix a winning strategy h for $\exists$ in $\mathscr{A}$. W.l.o.g. we may assume that $\mathbb{S}$ is an ω-expanded tree[4] with root s_I, so that with each $s \in S$ we may associate a *unique* state m_s such that (m_s,s) can be reached during an h-guided match of $\mathscr{A}$ starting from (m_I,s_I). By definition of $\mathbb{A}_\mathbb{G}$ and $\mathscr{A}$ then, with each $s \in S$ we may also associate a local strategy χ_s, which is locally compatible with the color $\mathrm{Val}(s)$ on m_s and such that $\exists$'s strategy h is aimed at satisfying the one-step formula $\nabla A_{m_s,\chi_s}$.

To define $\exists$'s strategy h' in $\mathscr{E}$, consider an arbitrary finite $\mathscr{E}$-match σ. It is not hard to see that σ admits a unique *modal decomposition* $\sigma = \sigma_0\cdots\sigma_l$, where for all $i < l$, $last(\sigma_i)$ is the unique modal position in σ_i, and σ_l either contains no modal positions, or a unique one at $last(\sigma_l)$. This means that we may present each σ_i as $\sigma_i = (v^i_0,s_i)(v^i_1,s_i)\cdots(v^i_{n_i},s_i)$ for some fixed point s_i in $\mathbb{S}$. The key idea underlying the definition of h' is that with every h'-guided finite match σ, with $\sigma = \sigma_0\cdots\sigma_l$ as above, we associate an h-guided $\mathscr{A}$-match $\Sigma_\sigma = (m_0,s_0)\cdots(m_l,s_l)$ satisfying the condition (†) given below:

(†1) for each $i \leq l$, $j \leq n_i$, we have $v^i_j \in \mathsf{Ran}(m_i\,;e_{\chi_{s_i}})$.

(†2) for each $i \leq l$, $m_i = m_{s_i}$, and the sequence $v^i_0\cdots v^i_{n_i}$ is χ_{s_i}-guided; for each pair j,k with $j < k \leq n_i$, we have $(v^i_j,N^i_{j,k},v^i_k) \in e^-_{\chi_{s_i}}$, where $N^i_{j,k} = \max\{\Omega(v^i_r) \mid j < r \leq k\}$;

(†3) for each $i < l$, with $w := v^i_{n_i}$, if $w \in V_\Diamond$, then $m_{i+1} = e_{\chi_{s_i}}\,;d_w(m_i\,;e_{\chi_{s_i}})$ and $(v^i_0,M_i,v^{i+1}_0) \in m_{i+1}$, where $M_i = \max\{\Omega(v^i_1),\ldots,\Omega(v^i_{n_i}),\Omega(v^{i+1}_0)\}$.

(†4) for each $i < l$, with $w := v^i_{n_i}$, if $w \in V_\Box$, then $m_{i+1} = e_{\chi_{s_i}}\,;d_\Box(m_i\,;e_{\chi_{s_i}})$ and $(v^i_0,M_i,v^{i+1}_0) \in m_{i+1}$, where $M_i = \max\{\Omega(v^i_1),\ldots,\Omega(v^i_{n_i}),\Omega(v^{i+1}_0)\}$.

[4]A pointed model $(\mathbb{S},s_I)$ is *ω-expanded* if R is the parent-child relation of a tree (S,R) which has s_I as its root, and is such that every non-root node s in $\mathbb{S}$ has at least ω many bisimilar siblings. It is not hard to see that every pointed model can be unravelled to a bisimilar model that is ω-expanded.

Based on this connection, we define the following strategy h' for $\exists$ in $\mathscr{E}$; we show at the same time that, playing h', $\exists$ can maintain the condition (†) and wins all finite matches. Consider a partial h'-guided match σ, modally decomposed as $\sigma = \sigma_0 \cdots \sigma_l$ as above, where $\sigma_l = (v_0^l, s_l) \cdots (v_k^l, s_l)$, and let $\Sigma_\sigma = (m_0, s_0) \cdots (m_l, s_l)$ be an associated $\mathscr{A}$-match satisfying (†). We distinguish cases, writing $v := v_k^l$ and $\chi := \chi_{s_l}$ for brevity.

- If v is a propositional node, we need to show that σ is won by $\exists$. This is immediate in case $L(v) = \top$, so assume that $L(v) = \bot$ or $L(v) \in \{p, \overline{p}\}$ for some proposition letter p. We only treat the case where $L(v) = p$, the other cases being similar. By (†2) we have $m_l = m_{s_l}$, so that $\chi = \chi_{s_l}$ is locally compatible with the color $\mathrm{Val}(s_l)$ on $m_l \,;\, e_\chi$. But then (†1) implies that $L(v) \in \mathrm{Val}(s)$.
- If $v \in V_\vee$, define $h'(\sigma) := (\chi(v), s_l)$. It is easy to see that $\sigma \cdot (\chi(v), s_l)$ and Σ_σ are related by (†).
- If $v \in V_\wedge$, suppose that $\forall$ picks a conjunct u of v. Then $\sigma \cdot (u, s_l)$ and Σ_σ are related by (†).
- If $v \in V_\Diamond$, first define $m_{l+1} := e_\chi \,;\, d_v(m_l \,;\, e_\chi)$. Note that, since $v \in \mathsf{Ran}(m_l \,;\, e_\chi)$ by the inductive hypothesis (†1), we find $m_{l+1} \in A_{m_l,\chi}$. Furthermore, recall that by our assumption on h, $\exists$'s move at position (m_l, s_l) in $\mathscr{A}$ is aimed at satisfying the one-step formula $\nabla A_{m_l,\chi}$, and so this move must contain a pair of the form (m_{l+1}, t) for some $t \in R[s_l]$. Now define $h'(\sigma) := (u, t)$, where u is the (unique) element of $E[v]$. The modal decomposition of $\sigma' := \sigma \cdot (u, t)$ is then $\sigma' = \sigma_1 \cdots \sigma_m \sigma_{m+1}$, where $\sigma_{m+1} = (u, t)$. (That is, in the terminology of (†) we have $v_0^{l+1} = u$ and $s_{l+1} = t$.)

  We now check that σ' and $\Sigma' := \Sigma \cdot (m_{l+1}, t)$ are related by (†). For (†1), it suffices to show that $u \in \mathsf{Ran}(m_{l+1} \,;\, e_{\chi_t})$. But this is immediate by $(v, \Omega(u), u) \in d_v(m_l \,;\, e_\chi) = \Delta_V \,;\, d_v(m_l \,;\, e_\chi) \subseteq e_\chi \,;\, d_v(m_l \,;\, e_\chi) = m_{l+1} \subseteq m_{l+1} \,;\, \Delta_V \subseteq m_{l+1} \,;\, e_{\chi_t}$. For (†2) it suffices to show that $m_{l+1} = m_t$ but this holds by construction. For (†3) we likewise have $m_{l+1} = e_\chi \,;\, d_v(m_l \,;\, e_\chi)$ by construction. We already saw that $(v, \Omega(u), u) \in d_v(m_l \,;\, e_\chi)$, and we have $(v_0^l, N_{0,k}^l, v_k^l) \in e_\chi^-$ by the induction hypothesis (†2). From this, and the observation that $M_l = \max(N_{0,k}^l, \Omega(u))$ we obtain $(v_0^l, M_l, u) \in e_\chi^- \,;\, d_v(m_l \,;\, e_\chi) \subseteq e_\chi \,;\, d_v(m_l \,;\, e_\chi) = m_{l+1}$. Finally, for (†4) there is nothing to prove.
- The case where $v \in V_\Box$ is similar to the previous one, so we skip some details. Let $u \in E[v]$ be the unique successor of v in $\mathbb{G}$, and suppose that in our $\mathscr{E}$-match, $\forall$ picks a successor t of s_l; that is, we now look at the continuation $\sigma' := \sigma \cdot (u, t)$ of the $\mathscr{E}$-match σ. Consider $\exists$'s move in $\mathscr{A}$ at position (m_l, s_l), which makes the one-step formula $\nabla A_{m_l,\chi}$ true, and thus contains a pair (m, t) for some $m \in A_{m_l,\chi}$. Now define $m_{l+1} := m$, and let $\Sigma' := (m, t)$; this is an h-guided continuation of Σ.

  To verify that σ' and Σ' satisfy (†), first note that by definition of the set $A_{m_l,\chi}$, m must be of the form $e_\chi \,;\, d_x(m_l \,;\, e_\chi)$, where either $x = \Box$ or $x \in \mathsf{Ran}(m_l \,;\, e_\chi)$. Based on this observation, checking the conditions (†1), (†2) and (†4) are similar to the respective conditions (†1), (†2) and (†3) in the previous case (with the only difference that we now must also take the possibility that $x = \Box$ into account). Finally, condition (†3) needs no check since it is not applicable.

To see that h' is winning for $\exists$ consider a full h'-guided match, and distinguish cases. For finite matches one can check that $\exists$ never gets stuck. In case σ is an infinite h'-guided match, we make a further distinction as to whether the number of modal positions that σ passes through is finite or infinite. If σ passes through infinitely many modal positions, there is a unique way of decomposing σ as $\sigma = \sigma_0 \sigma_1 \cdots$, where $(last(\sigma_i))_{i \in \omega}$ is the sequence of (all) modal positions in σ. By construction there is an associated *infinite* h-guided $\mathscr{A}$-match $\Sigma_\sigma = (m_i, s_i)_{i \in \omega}$ related to σ via the condition (†). It is now possible to prove that $\exists$ is the winner of σ by using that h is winning for $\exists$ in $\mathscr{A}$ (details omitted due to space limitations). If σ only passes finitely many modal positions, we may represent $\sigma = \sigma_0 \cdots \sigma_l$, where each σ_i with $i < l$ is finite, σ_l is infinite, and $(last(\sigma_i))_{i<l}$ is the sequence of all modal positions in σ. We only consider the subcase where $l > 0$. Let $\Sigma_\sigma = (m_0, s_0) \cdots (m_l, s_l)$ be the h-guided $\mathscr{A}$-match that we have associated with σ (or, to be more precise, with the initial segments of σ that are long enough to have passed the last modal node of σ). Observe that since Σ_σ is h-guided, the position (m_l, s_l) must be winning for $\exists$, and that by (†2) the macrostate m_l is the unique state m in A such that (m_l, s_l) is met during an h-guided $\mathscr{A}$-match. Write $\sigma_m = (u_0, s)(u_1, s)(u_2, s)\cdots$; that is, we write $u_j := v_j^l$ and $s := s_l$. The sequence $u_0u_1u_2\cdots$ is a trace on the stream $m_l\,;(e^-_{\chi_s})^\omega$; but then, by the compatibility of χ_s with m_l on $\mathrm{Val}(s)$, $u_0u_1u_2\cdots$ must be a *good* trace. Since $(u_0, s)(u_1, s)\cdots$ is a tail of σ, this means that σ is won by $\exists$, as required. QED

Proof of Theorem 5.1. The equivalence part of the disjunctive modal automaton $\mathbb{A}_{\mathbb{G},v_I}$ to $\mathbb{G}$ was proved in Proposition 5.8. It remains to check the sizes of the components of the automaton $\mathbb{A}$. But this is fairly straightforward. To start with, from the definition of $\mathbb{A}$ we have $M_\Omega = \wp(V \times \mathsf{Ran}(\Omega) \times V)$ it immediately follows that $|A| \leq 2^{|V\times\mathsf{Ran}(\Omega)\times V|} = 2^{n^2k}$. To compute the size of Θ, first observe that the number of local strategies is equal to $2^{|A_\vee|}$, and that for each macrostate m, local strategy χ and color $c \in \wp\mathsf{P}$ we find $|A_{m,\chi}| \leq |V_\Diamond| + 1 \leq n$. From this it is immediate that for each formula $\Theta(m,c)$ we have $|\Theta(m,c)| \leq 2^{|A_\vee|}\cdot(|V_\Diamond|+1) \leq n2^n$. Finally, the table of Θ has $|A|\cdot 2^{|\mathsf{P}|} \leq 2^{n^2k}\cdot 2^l = 2^{n^2k+l}$ entries, so that its total size is bounded by $n2^n \cdot 2^{n^2k+l} = n2^{n^2k+l+n}$, as stated by the theorem. QED

Corollary 5.9 *The disjunctive modal automaton $\mathbb{A}_{\mathbb{G},v_I}$ can be turned into an equivalent disjunctive parity automaton $\mathbb{A}$ with index $\mathscr{O}(n\cdot k)$ and size $2^{\mathscr{O}(n^2k\cdot\log(nk))}$.*

Proof. A standard construction, the so-called wreath product, can be used to turn the automaton $\mathbb{A}_{\mathbb{G},v_I}$ together with the automaton $\mathbb{P} = (P, \delta, \Omega_P, p_I)$ from Proposition 5.4 into a parity automaton (cf. e.g. [16, Definition 4.3]). The transition map of the resulting automaton $\mathbb{A}$ will have the same size as the one of $\mathbb{A}_\mathbb{G}$, the set of states is given as the product $M_\Omega \times P$ and the index of $\mathbb{A}$ is equal to the index of $\mathbb{P}$. Hence the parity automaton has size $2^{\mathscr{O}(nk\cdot\log(nk))}\cdot n2^{n^2k+l+n} = 2^{\mathscr{O}(n^2k\cdot\log(nk))}$, and index $\mathscr{O}(nk)$. QED

We finish with the main result of our paper: there is an algorithm turning a parity formula into an equivalent disjunctive one in exponential size of the original formula. Due to our size preserving translations from parity formulas to μ-calculus formulas in the standard syntax, the result carries directly over to formulas in standard syntax if we measure the size of this formula in terms of its closure.

Corollary 5.10 *For any parity formula $\mathbb{G}$ we can construct an equivalent disjunctive parity formula $\mathbb{G}^d$ with $|\mathbb{G}^d| \leq 2^{\mathscr{O}(n^2k\cdot\log(nk))}$ and with index $\mathscr{O}(n\cdot k)$. Here $n = |\mathbb{G}|$ and k is the index of $\mathbb{G}$.*

The corollary is an immediate consequence of Corollary 5.9 and Theorem 4.1. An application of Prop. 2.5 shows that the corollary implies Thm. 1.2.

6 Conclusions

We have presented an algorithm that constructs for a given arbitrary formula in the modal μ-calculus an equivalent disjunctive formula with a single exponential blow-up when measuring the size of a formula in closure size. While the complexity of this construction is likely to be optimal, it is an interesting question for future work whether or not the construction can be optimised to obtain "nice" disjunctive formulas. In particular, the move from modal automata to parity formulas potentially adds a large number of unnecessary disjuncts. Obtaining a nicer formula could be relevant for computing uniform interpolants. Another nagging question is the exact repercussion of our work for satisfiability checking. While satisfiability checking for disjunctive formulas is linear in subformula size, our formulas are measured in closure size, which is potentially an exponential smaller. It has to be checked whether one can use our result for ExpTime satisfiability checking when the input formula is measured in closure size.

References

[1] J. Bradfield & C. Stirling (2006): *Modal μ-calculi*. In J. van Benthem, P. Blackburn & F. Wolter, editors: *Handbook of Modal Logic*, Elsevier, pp. 721–756.

[2] F. Bruse, O. Friedmann & M. Lange (2015): *On guarded transformation in the modal μ-calculus*. *Logic Journal of the IGPL* 23(2), pp. 194–216, doi:10.1093/jigpal/jzu030.

[3] C.S. Calude, S. Jain, B. Khoussainov, W. Li & F. Stephan (2017): *Deciding parity games in quasipolynomial time*. In H. Hatami, P. McKenzie & V. King, editors: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, (STOC 2017)*, ACM, pp. 252–263, doi:10.1145/3055399.3055409.

[4] G. D'Agostino & M. Hollenberg: *Logical questions concerning the μ-calculus*. *Journal of Symbolic Logic* 65, pp. 310–332, doi:10.1080/11663081.1991.10510772.

[5] G. D'Agostino & G. Lenzi (2006): *On modal mu-calculus with explicit interpolants*. *Journal of Applied Logic* 4(3), pp. 256–278, doi:10.1016/j.jal.2005.06.008.

[6] S. Demri, V. Goranko & M. Lange (2016): *Temporal Logics in Computer Science: Finite-State Systems*. Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, doi:10.1017/CBO9781139236119.

[7] E.A. Emerson & C.S. Jutla (1988): *The complexity of tree automata and logics of programs (extended abstract)*. In: *Proceedings of the 29th Symposium on the Foundations of Computer Science*, IEEE Computer Society Press, pp. 328–337.

[8] E.A. Emerson & C.S. Jutla (1991): *Tree automata, mu-calculus and determinacy (extended abstract)*. In: *Proceedings of the 32nd Symposium on the Foundations of Computer Science*, IEEE Computer Society Press, pp. 368–377.

[9] S. Enqvist, F. Seifan & Y. Venema (2018): *Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics*. *Theoretical Computer Science* 727, pp. 37–100, doi:10.1016/j.tcs.2018.03.001.

[10] G. Fontaine & Y. Venema (2018): *Some model theory for the modal mu-calculus: syntactic characterizations of semantic properties*. *Logical Methods in Computer Science* 14(1).

[11] O. Friedmann & M. Lange (2013): *Deciding the unguarded modal μ-calculus*. *Journal of Applied Non-Classical Logics* 23(4), pp. 353–371, doi:10.1080/11663081.2013.861181.

[12] E. Grädel, W. Thomas & T. Wilke, editors (2002): *Automata, Logic, and Infinite Games*. *LNCS* 2500, Springer.

[13] D. Janin & I. Walukiewicz (1995): *Automata for the modal μ-calculus and related results*. In: *Proceedings of the Twentieth International Symposium on Mathematical Foundations of Computer Science, MFCS'95*, *LNCS* 969, Springer, pp. 552–562.

[14] O. Kupferman & M.Y. Vardi (2005): *From Linear Time to Branching Time*. *ACM Transactions on Computational Logic* 6(2), pp. 273–294, doi:10.1145/1055686.1055689.

[15] C. Kupke, J. Marti & Y. Venema (2020): *Size matters in the modal μ-calculus*. *CoRR* abs/2010.14430.

[16] C. Kupke & Y. Venema (2008): *Coalgebraic Automata Theory: Basic Results*. *Logical Methods in Computer Science* 4(4), doi:10.2168/LMCS-4(4:10)2008.

[17] K. Lehtinen (2015): *Disjunctive form and the modal μ alternation hierarchy*. In: *Proceedings of FICS 2015*, *EPTCS* 191, pp. 117–131, doi:10.4204/EPTCS.191.11.

[18] N. Piterman (2007): *From Nondeterministic Büchi and Streett Automata to Deterministic Parity Automata*. *LMCS* 3(3), doi:10.2168/LMCS-3(3:5)2007.

[19] S. Schewe (2009): *Tighter Bounds for the Determinisation of Büchi Automata*. In Luca de Alfaro, editor: *Proceedings of FOSSACS 2009*, *LNCS* 5504, Springer, pp. 167–181, doi:10.1007/978-3-642-00596-1_13.

[20] C. Stirling: *Modal and Temporal Properties of Processes*. Texts in Computer Science, Springer-Verlag, doi:10.1007/978-1-4757-3550-5.

[21] M. Y. Vardi (1998): *Reasoning about the Past with Two-Way Automata*. In: *Proceedings of ICALP'98*, *LNCS* 1443, Springer, pp. 628–641, doi:10.1007/BFb0055090.

[22] Y. Venema (2006): *Automata and Fixed Point Logic: a Coalgebraic Perspective*. *Information and Computation* 204, pp. 637–678, doi:10.1016/j.ic.2005.06.003.

[23] I. Walukiewicz (2000): *Completeness of Kozen's axiomatisation of the propositional μ-calculus*. *Information and Computation* 157, pp. 142–182, doi:10.1006/inco.1999.2836.

[24] T. Wilke (2001): *Alternating tree automata, parity games, and modal μ-calculus*. *Bulletin of the Belgian Mathematical Society* 8, pp. 359–391.