



Formula size games for modal logic and mu-calculus

Citation

Hella, L. T., & Vilander, M. S. (2019). Formula size games for modal logic and mu-calculus. *JOURNAL OF LOGIC AND COMPUTATION*, [exz025]. <https://doi.org/10.1093/logcom/exz025>

Year

2019

Version

Early version (pre-print)

Link to publication

[TUTCRIS Portal \(http://www.tut.fi/tutcris\)](http://www.tut.fi/tutcris)

Published in

JOURNAL OF LOGIC AND COMPUTATION

DOI

[10.1093/logcom/exz025](https://doi.org/10.1093/logcom/exz025)

Take down policy

If you believe that this document breaches copyright, please contact cris.tau@tuni.fi, and we will remove access to the work immediately and investigate your claim.

Formula size games for modal logic and μ -calculus

Lauri Hella
Tampere University
FI-33014
Tampere University
lauri.hella@tuni.fi
+358503053204

Miikka Vilander
Tampere University
FI-33014
Tampere University
miikka.vilander@tuni.fi
+358405677359

February 5, 2019

Abstract

We propose a new version of formula size game for modal logic. The game characterizes the equivalence of pointed Kripke-models up to formulas of given numbers of modal operators and binary connectives. Our game is similar to the well-known Adler-Immerman game. However, due to a crucial difference in the definition of positions of the game, its winning condition is simpler, and the second player does not have a trivial optimal strategy. Thus, unlike the Adler-Immerman game, our game is a genuine two-person game. We illustrate the use of the game by proving a non-elementary succinctness gap between bisimulation invariant first-order logic FO and (basic) modal logic ML. We also present a version of the game for the modal μ -calculus L_μ and show that FO is also non-elementarily more succinct than L_μ .

Keywords: Succinctness, formula size game, modal logic, modal μ -calculus, bisimulation invariant first-order logic.

1 Introduction

Logical languages are often compared in terms of expressiveness and computational complexity. The authors of [13] argue that another important semantic aspect of a logical language is the size of formulas needed for expressing properties of structures. If two logics L and L' are equivalent in terms of expressivity, one of them may be able to express interesting properties much more succinctly than the other. According to the standard terminology, for a given function f on natural numbers, L is said to be f times more succinct than L' if there is a sequence $(\varphi_n)_{n \in \mathbb{N}}$, of L -formulas such that for any sequence $(\psi_n)_{n \in \mathbb{N}}$ of equivalent L' -formulas, the size of ψ_n is at least $f(m_n)$, where m_n is the size of φ_n .

The succinctness of various modal and temporal logics has been an active area of research for the last couple of decades; see e.g. [32, 21, 5, 1, 22, 20] for earlier work on the topic and [9, 28, 30, 31, 24, 7] for recent work. Typical results in the area state an exponential succinctness gap between two equally expressive logics. Often such a gap is reflected in the complexity of the logics in question. For example, Etessami, Vardi and Wilke proved in [5] that the two-variable fragment FO^2 of first-order logic and unary-TL (a weak version of temporal logic) have the same expressive power over ω -words, but FO^2 is exponentially more succinct than unary-TL. Furthermore, the complexity of satisfiability for FO^2 is NEXPTIME-complete, while the complexity of unary-TL is in NP [26]. However, being more succinct does not always imply higher complexity: for example, public announcement logic PAL is exponentially more succinct than epistemic logic EL, but the complexity of satisfiability is the same for both of them [20].

The most commonly used methods for proving succinctness results are formula size games and extended syntax trees. Formula size games were first introduced by Adler and Immerman in

[1] for branching-time temporal logic CTL. The method of extended syntax trees was originally formulated by Grohe and Schweikardt in [14] for first-order logic. The notion of extended syntax tree was actually inspired by the Adler-Immerman game, and in a certain sense these two methods are equivalent: an extended syntax tree can be interpreted as a winning strategy for one of the players of the corresponding formula size game. Both of these methods have been adapted to a large number of modal languages, including epistemic logic [8], multimodal logics with union and intersection operators on modalities [29] and modal logic with contingency operators [30].

The basic idea of the Adler-Immerman game is that one of the players, S (spoiler), tries to show that two sets of pointed models \mathbb{A} and \mathbb{B} can be separated by a formula of size n , while the other player, D (duplicator), aims to show that no formula of size at most n suffices for this. The moves that S makes in the game reflect directly the logical operators in a formula that is supposed to separate the sets \mathbb{A} and \mathbb{B} . Any pair (σ, δ) of strategies for the players S and D produces a finite game tree $T_{\sigma, \delta}$, and S wins this play if the size of $T_{\sigma, \delta}$ is at most n . The strategy σ is a winning strategy for S if using it, S wins every play of the game. If this is the case, then there is a formula of size at most n that separates the sets, and this formula can actually be read from the strategy σ .

A peculiar feature of the Adler-Immerman game is that the second player, duplicator, can be completely eliminated from it. This is because D has an optimal strategy δ_{\max} , which is to always choose the maximal allowed answer; this strategy guarantees that the size of the tree $T_{\sigma, \delta}$ is as large as possible. Thus, in this sense the Adler-Immerman game is not a genuine two-person game, but rather a one-person game. Extended syntax trees, on the other hand, do away with the game aspect entirely.

In the present paper, we propose another type of formula size game for modal logic. Our game is a natural adaptation of the game introduced by Hella and Väänänen [17] for propositional logic and first-order logic. The basic setting in our game is the same as in the Adler-Immerman game: there are two players, S and D, and two sets of structures that S claims can be separated by a formula of some given size. The crucial difference is that in our game we define positions to be tuples $(k, \mathbb{A}, \mathbb{B})$ instead of just pairs (\mathbb{A}, \mathbb{B}) of sets of structures, where k is a parameter referring to the number of modal operators and binary connectives in a formula. In each move S has to decrease the parameter k . The game ends when the players reach a position $(k^*, \mathbb{A}^*, \mathbb{B}^*)$ such that either there is a literal separating \mathbb{A}^* and \mathbb{B}^* , or S cannot make any moves because $k^* = 0$. In the former case, S wins the play; otherwise D wins.

Thus, in contrast to the Adler-Immerman game, to determine the winner in our game it suffices to consider a single “leaf-node” $(k^*, \mathbb{A}^*, \mathbb{B}^*)$ of the game tree. This also means that our game is a real two-person game: the final position $(k^*, \mathbb{A}^*, \mathbb{B}^*)$ of a play depends on the moves of D, and there is no simple optimal strategy for D that could be used for eliminating the role of D in the game.

We believe that our game is more intuitive and thus, in some cases it may be easier to use than the Adler-Immerman game. On the other hand, it should be remarked that the two games are essentially equivalent: The moves corresponding to connectives and modal operators are the same in both games (when restricting to the sets \mathbb{A} and \mathbb{B} in a position $(k, \mathbb{A}, \mathbb{B})$). Hence, in principle, it is possible to translate a winning strategy in one of the games to a corresponding winning strategy in the other.

Additionally, we introduce a formula size game for modal μ -calculus. This game is obtained by adapting the formula size game of modal logic to the setting with fixed point operators μ and ν . A new challenge in defining such a game is that if S uses a fixed point ηX ($\eta \in \{\mu, \nu\}$) as the logical operator in his move, and later uses the corresponding variable X , then in the next round, the game has to return to the subformula that follows ηX . This means that the play may become infinite, and defining the correct winning condition for infinite plays is complicated. We solve this problem by adding ordinal clocks to the pointed Kripke models in the sets \mathbb{A} and \mathbb{B} . The idea is that the the ordinals corresponding to a fixed point variable X decrease each time

the game returns to an earlier formula from a position with label X . This, in conjunction with keeping the sets \mathbb{A} and \mathbb{B} always finite, guarantees that every play of the game is finite. The idea of using ordinal clocks is also used in [16] to define finite semantic games for L_μ .

We illustrate the use of our games by proving two non-elementary succinctness gaps, one between first-order logic FO and (basic) modal logic ML and the other between FO and modal μ -calculus L_μ . More precisely, we define a bisimulation invariant property of pointed Kripke-models by a first-order formula of linear size, and show that this property cannot be defined by any ML- or L_μ -formula of size less than the exponential tower of height $n - 1$. Furthermore, we show that the same property of pointed Kripke-models is already definable by a formula of size $\mathcal{O}(2^n)$ in a version ML^2 of 2-dimensional modal logic. Hence the same non-elementary succinctness result holds for ML^2 over ML.

A similar gap between FO and temporal logic follows from a construction in the PhD thesis [27] of Stockmeyer. He proved that the satisfiability problem of FO over words is of non-elementary complexity. Etessami and Wilke [6] observed that from Stockmeyer's proof it is possible to extract FO-formulas of size $\mathcal{O}(n)$ whose smallest models are words of length non-elementary in n . On the other hand, it is well known that any satisfiable formula of temporal logic has a model of size $\mathcal{O}(2^n)$, where n is the size of the formula. Another result related to ours can be found in [25], where Otto shows that FO is exponentially more succinct than ML by relating the modal depth of the ML-formula to the quantifier rank of the FO-formula. In contrast to this, our proof relies entirely on the number of disjunctions and conjunctions in the ML-formula.

For modal μ -calculus, the literature regarding succinctness is scarcer. In [15] Grohe and Schweikardt show several succinctness gaps between monadic second-order logics, many with fixed points. They use automata-theoretic techniques and cite a non-elementary succinctness gap between MSO and L_μ as well-known¹.

The structure of the paper is as follows. In Section 2 we present the logics used in the paper, fix some notation and define our notion of formula size. In Section 3 we present the formula size game for ML and show some basic results for it. Section 4 is dedicated to the non-elementary succinctness gap between FO and ML and all necessary definitions and lemmas to prove it. In Section 5 we define the formula size game for L_μ and show basic results. Finally, in Section 6 we show the non-elementary succinctness of FO over L_μ . Section 7 is the conclusion.

The work on modal logic was previously published in the conference paper [18]. This version has some minor changes to the modal logic part and the sections on modal μ -calculus are completely new.

2 Preliminaries

In this section we fix some notation, define the syntax and semantics of basic modal logic and modal μ -calculus, and define our notions of formula size. For a detailed account on the logics used in the paper, we refer to the textbook [2] of Blackburn, de Rijke and Venema for basic modal logic and [3] for the modal μ -calculus.

Basic modal logic and first-order logic

Let $Prop$ be an infinite set of proposition symbols and let $\Phi \subseteq Prop$. Let $\mathcal{M} = (W, R, V)$, where W is a set, $R \subseteq W \times W$ and $V : \Phi \rightarrow \mathcal{P}(W)$, and let $w \in W$. The structure (\mathcal{M}, w) is called a *pointed Kripke-model for Φ* .

Let (\mathcal{M}, w) be a pointed Kripke-model. We use the notation

$$\Box(\mathcal{M}, w) := \{(\mathcal{M}, v) \mid v \in W, wR^{\mathcal{M}}v\}.$$

¹We were unable to find a source in the literature for this result but we are reasonably convinced that a non-elementary gap already between FO and L_μ is a new result.

If \mathbb{A} is a set of pointed Kripke-models, we use the notation

$$\Box\mathbb{A} := \bigcup_{(\mathcal{M}, w) \in \mathbb{A}} \Box(\mathcal{M}, w).$$

Furthermore, if f is a function $f : \mathbb{A} \rightarrow \Box\mathbb{A}$ such that $f(\mathcal{M}, w) \in \Box(\mathcal{M}, w)$ for every $(\mathcal{M}, w) \in \mathbb{A}$, then we use the notation

$$\Diamond_f\mathbb{A} := f(\mathbb{A}).$$

Intuitively $\Box(\mathcal{M}, w)$ is the set of all successor models of (\mathcal{M}, w) , $\Box\mathbb{A}$ is the collection of all successor models of all models $(\mathcal{M}, w) \in \mathbb{A}$ and $\Diamond_f\mathbb{A}$ consists of one successor for each model in \mathbb{A} , where the successors are given by the function f . We now define the syntax and semantics of basic modal logic for pointed models.

Let $\Phi \subseteq Prop$. The set of formulas of $ML(\Phi)$ is generated by the following grammar

$$\varphi := \top \mid \perp \mid p \mid \neg p \mid (\varphi \wedge \psi) \mid (\varphi \vee \psi) \mid \Diamond\varphi \mid \Box\varphi,$$

where $p \in \Phi$.

As is apparent from the definition of the syntax, we assume that all ML-formulas are in negation normal form. This is useful for the formula size game that we introduce in the next section.

The satisfaction relation $(\mathcal{M}, w) \models \varphi$ between pointed Kripke-models (\mathcal{M}, w) and $ML(\Phi)$ -formulas φ is defined as follows:

- (1) $(\mathcal{M}, w) \models \top$ for all (\mathcal{M}, w) , and $(\mathcal{M}, w) \not\models \perp$ for all (\mathcal{M}, w) ,
- (2) $(\mathcal{M}, w) \models p \Leftrightarrow w \in V(p)$, and $(\mathcal{M}, w) \models \neg p \Leftrightarrow w \notin V(p)$,
- (3) $(\mathcal{M}, w) \models (\varphi \wedge \psi) \Leftrightarrow (\mathcal{M}, w) \models \varphi$ and $(\mathcal{M}, w) \models \psi$,
- (4) $(\mathcal{M}, w) \models (\varphi \vee \psi) \Leftrightarrow (\mathcal{M}, w) \models \varphi$ or $(\mathcal{M}, w) \models \psi$,
- (5) $(\mathcal{M}, w) \models \Diamond\varphi \Leftrightarrow$ there is $(\mathcal{M}, v) \in \Box(\mathcal{M}, w)$ such that $(\mathcal{M}, v) \models \varphi$,
- (6) $(\mathcal{M}, w) \models \Box\varphi \Leftrightarrow$ for every $(\mathcal{M}, v) \in \Box(\mathcal{M}, w)$ it holds that $(\mathcal{M}, v) \models \varphi$.

Furthermore, if \mathbb{A} is a class of pointed Kripke-models, then

$$\mathbb{A} \models \varphi \Leftrightarrow (\mathcal{A}, w) \models \varphi \text{ for every } (\mathcal{A}, w) \in \mathbb{A}.$$

For the sake of convenience we also use the notation

$$\mathbb{A} \models \neg\varphi \Leftrightarrow (\mathcal{A}, w) \not\models \varphi \text{ for every } (\mathcal{A}, w) \in \mathbb{A}.$$

Note that this is only a notational convention as $\neg\varphi$ is not in negation normal form and as such is generally not a formula in our syntax.

In Section 4, we consider the case $\Phi = \emptyset$. In this case the only available literals are the constants \top and \perp , which are always true or false respectively.

The syntax and semantics for first-order logic are defined in the standard way. Each ML-formula φ defines a class $\text{Mod}(\varphi)$ of pointed Kripke-models:

$$\text{Mod}(\varphi) := \{(\mathcal{M}, w) \mid (\mathcal{M}, w) \models \varphi\}.$$

In the same way, any FO-formula $\psi(x)$ in the vocabulary consisting of the accessibility relation symbol R and unary relation symbols U_p for $p \in \Phi$ defines a class $\text{Mod}(\psi)$ of pointed Kripke-models:

$$\text{Mod}(\psi) := \{(\mathcal{M}, w) \mid \mathcal{M} \models \psi[w/x]\}.$$

The formulas $\varphi \in ML$ and $\psi(x) \in FO$ are *equivalent* if $\text{Mod}(\varphi) = \text{Mod}(\psi)$.

The well-known link between ML and FO is the following theorem.

Theorem 2.1 (van Benthem Characterization Theorem). *A first-order formula $\psi(x)$ is equivalent to some formula in ML if and only if $\text{Mod}(\psi)$ is bisimulation invariant.*

If a property of pointed Kripke-models is n -bisimulation invariant for some $n \in \mathbb{N}$, then it is also bisimulation invariant. Thus, FO-definability and n -bisimulation invariance imply ML-definability for any property of pointed Kripke-models. We will use this version of van Benthem's characterization in Section 4.1 for showing that a certain property is ML-definable. For the sake of easier reading, we give here the definition of n -bisimulation.

Definition 2.2. Let (\mathcal{M}, w) and (\mathcal{M}', w') be pointed Φ -models. We say that (\mathcal{M}, w) and (\mathcal{M}', w') are n -bisimilar, $(\mathcal{M}, w) \simeq_n (\mathcal{M}', w')$, if there are binary relations $Z_n \subseteq \dots \subseteq Z_0$ such that for every $0 \leq i \leq n-1$ we have

- (1) $(\mathcal{M}, w)Z_n(\mathcal{M}', w')$,
- (2) if $(\mathcal{M}, v)Z_0(\mathcal{M}', v')$, then $(\mathcal{M}, v) \models p \Leftrightarrow (\mathcal{M}', v') \models p$ for each $p \in \Phi$,
- (3) if $(\mathcal{M}, v)Z_{i+1}(\mathcal{M}', v')$ and $(\mathcal{M}, u) \in \square(\mathcal{M}, v)$ then there is $(\mathcal{M}', u') \in \square(\mathcal{M}', v')$ such that $(\mathcal{M}, u)Z_i(\mathcal{M}', u')$,
- (4) if $(\mathcal{M}, v)Z_{i+1}(\mathcal{M}', v')$ and $(\mathcal{M}', u') \in \square(\mathcal{M}', v')$ then there is $(\mathcal{M}, u) \in \square(\mathcal{M}, v)$ such that $(\mathcal{M}, u)Z_i(\mathcal{M}', u')$.

It is well known that if Φ is finite, two pointed Φ -models are n -bisimilar if and only if they are equivalent with respect to ML(Φ)-formulas of modal depth at most n .

Modal μ -calculus

Let $\Phi \subseteq \text{Prop}$ and let Var be an infinite set of variables. The syntax of the modal μ -calculus $L_\mu(\Phi)$ is given by the grammar:

$$\varphi ::= \top \mid \perp \mid p \mid \neg p \mid (\varphi \vee \psi) \mid (\varphi \wedge \psi) \mid \diamond\varphi \mid \square\varphi \mid X \mid \mu X.\varphi \mid \nu X.\varphi,$$

where $p \in \Phi$ and $X \in \text{Var}$. Note that all formulas are again in negation normal form. We additionally assume for simplicity that variables of different fixed points are distinct.

Truth of formulas of $L_\mu(\Phi)$ is, like ML, evaluated on pointed Kripke models (\mathcal{M}, w) , where $\mathcal{M} = (W, R, V)$. Let $\varphi \in L_\mu(\Phi)$ and let $\rho : \text{Var} \rightarrow \mathcal{P}(W)$ be a valuation of variables. We define truth relation $(\mathcal{M}, w) \models_\rho \varphi$ between pointed models and $L_\mu(\Phi)$ -formulas. Let $\|\varphi\|_\rho := \{w \in W \mid (\mathcal{M}, w) \models_\rho \varphi\}$ and let $\Gamma_{\varphi, \rho} : \mathcal{P}(W) \rightarrow \mathcal{P}(W)$ be an operator which maps W' to $\|\varphi\|_{\rho[W'/X]}$. The notation LFP stands for least fixed point of an operator and GFP for greatest fixed point. Since variables only occur positively in fixed point formulas, $\Gamma_{\varphi, \rho}$ is a monotone operator. By the Tarski-Knaster Theorem, the least and greatest fixed points of such a monotone operator always exist. The recursive definition of \models_ρ is as follows:

- $(\mathcal{M}, w) \models_\rho p \Leftrightarrow w \in V(p)$,
- $(\mathcal{M}, w) \models_\rho X \Leftrightarrow w \in \rho(X)$,
- $(\mathcal{M}, w) \models_\rho (\varphi \vee \psi) \Leftrightarrow (\mathcal{M}, w) \models_\rho \varphi$ or $(\mathcal{M}, w) \models_\rho \psi$,
- $(\mathcal{M}, w) \models_\rho (\varphi \wedge \psi) \Leftrightarrow (\mathcal{M}, w) \models_\rho \varphi$ and $(\mathcal{M}, w) \models_\rho \psi$,
- $(\mathcal{M}, w) \models_\rho \diamond\varphi \Leftrightarrow$ there is $(\mathcal{M}, v) \in \square(\mathcal{M}, w)$ such that $(\mathcal{M}, v) \models_\rho \varphi$,
- $(\mathcal{M}, w) \models_\rho \square\varphi \Leftrightarrow$ for every $(\mathcal{M}, v) \in \square(\mathcal{M}, w)$ it holds that $(\mathcal{M}, v) \models_\rho \varphi$,
- $(\mathcal{M}, w) \models_\rho \mu X.\varphi \Leftrightarrow w \in \text{LFP}(\Gamma_{\varphi, \rho})$,
- $(\mathcal{M}, w) \models_\rho \nu X.\varphi \Leftrightarrow w \in \text{GFP}(\Gamma_{\varphi, \rho})$.

Formula size

We define notions of formula size for ML, L_μ and FO. Note that many different notions are called formula size in the literature and our notion is close to the length of the formula as a string rather than, say, the DAG-size² of it.

Definition 2.3. The *size* of a formula $\varphi \in \text{ML}$, denoted $\text{sz}(\varphi)$, is defined recursively as follows:

- (1) If φ is a literal, then $\text{sz}(\varphi) = 1$.
- (2) If $\varphi = \psi \vee \vartheta$ or $\varphi = \psi \wedge \vartheta$, then $\text{sz}(\varphi) = \text{sz}(\psi) + \text{sz}(\vartheta) + 1$.
- (3) If $\varphi = \diamond\psi$ or $\varphi = \square\psi$, then $\text{sz}(\varphi) = \text{sz}(\psi) + 1$.

Definition 2.4. The *size* of a formula $\varphi \in L_\mu$, denoted $\text{sz}(\varphi)$, is defined recursively as follows:

- (1) $\text{sz}(l) = \text{sz}(X) = 1$, where l is a literal and X is a variable,
- (2) $\text{sz}(\varphi \vee \psi) = \text{sz}(\varphi \wedge \psi) = \text{sz}(\varphi) + \text{sz}(\psi) + 1$,
- (3) $\text{sz}(\diamond\varphi) = \text{sz}(\square\varphi) = \text{sz}(\mu X.\varphi) = \text{sz}(\nu X.\varphi) = \text{sz}(\varphi) + 1$.

The size of a formula is essentially its length as a string. Note however, that we do not count negations as we view them as parts of literals. Another aspect worth mentioning is the size of descriptions of proposition symbols. If we have an infinite set of propositions, the size of the encoding of each symbol in a fixed size vocabulary necessarily grows logarithmically. Here we consider all propositions to be of size one.

Similarly we define formula size for FO to be the number of binary connectives, quantifiers and literals in the formula. In general this could lead to an arbitrarily large difference between formula size and actual string length. For example if f is a unary function symbol, then atomic formulas of the form $f(x) = x$, $f(f(x)) = x$ and so on, all have size 1. In this paper however, we only consider formulas with one binary relation so this is not an issue.

Definition 2.5. The *size* of a formula $\varphi \in \text{FO}$, denoted by $\text{sz}(\varphi)$, is defined recursively as follows:

- (1) If φ is a literal, then $\text{sz}(\varphi) = 1$.
- (2) If $\varphi = \neg\psi$, then $\text{sz}(\varphi) = \text{sz}(\psi)$.
- (3) If $\varphi = \psi \vee \vartheta$ or $\varphi = \psi \wedge \vartheta$, then $\text{sz}(\varphi) = \text{sz}(\psi) + \text{sz}(\vartheta) + 1$.
- (4) If $\varphi = \exists x\psi$ or $\varphi = \forall x\psi$, then $\text{sz}(\varphi) = \text{sz}(\psi) + 1$.

To refer to some rather large formula sizes we need the exponential tower function.

Definition 2.6. We define the function $\text{twr} : \mathbb{N} \rightarrow \mathbb{N}$ recursively as follows:

$$\begin{aligned} \text{twr}(0) &= 1 \\ \text{twr}(n+1) &= 2^{\text{twr}(n)}. \end{aligned}$$

We will also use in the sequel the binary logarithm function, denoted by \log .

²The DAG-size of a formula φ is the number of edges of the syntactic structure of φ in the form of a directed acyclic graph. Thus since the fan-out in the DAG is at most two, the DAG-size is at most two times the number of subformulas of φ .

Separating classes by formulas

The definitions of the formula size games in sections 3 and 5 are based on the notion of separating classes of pointed Kripke-models by formulas. Recall that by the notation $\mathbb{B} \models \neg\varphi$ we mean that for every model $(\mathcal{B}, w) \in \mathbb{B}$, we have $(\mathcal{B}, w) \not\models \varphi$. As formulas of ML are also in L_μ , we only define the following for L_μ and FO.

Definition 2.7. Let \mathbb{A} and \mathbb{B} be classes of pointed Kripke-models.

- (a) We say that a formula $\varphi \in L_\mu$ *separates* \mathbb{A} from \mathbb{B} if $\mathbb{A} \models \varphi$ and $\mathbb{B} \models \neg\varphi$.
- (b) Similarly, a formula $\psi(x) \in \text{FO}$ separates \mathbb{A} from \mathbb{B} if for all $(\mathcal{M}, w) \in \mathbb{A}$, $\mathcal{M} \models \psi[w/x]$ and for all $(\mathcal{M}, w) \in \mathbb{B}$, $\mathcal{M} \models \neg\psi[w/x]$.

In other words, a formula $\varphi \in L_\mu$ separates \mathbb{A} from \mathbb{B} if $\mathbb{A} \subseteq \text{Mod}(\varphi)$ and $\mathbb{B} \subseteq \overline{\text{Mod}(\varphi)}$, where $\overline{\text{Mod}(\varphi)}$ is the complement of $\text{Mod}(\varphi)$.

3 The formula size game for ML

As in the Adler-Immerman game, the basic idea in our formula size game is that there are two players, S (Samson) and D (Delilah), who play on a pair (\mathbb{A}, \mathbb{B}) of two sets of pointed Kripke-models. The aim of S is to show that \mathbb{A} and \mathbb{B} can be separated by a formula with size at most k , while D tries to refute this. The moves of S reflect the connectives and modal operators of a formula that is supposed to separate the sets.

The crucial difference between our game and the Adler-Immerman game is that we define positions in the game to be tuples $(k, \mathbb{A}, \mathbb{B})$ instead of just pairs (\mathbb{A}, \mathbb{B}) . As in the A-I game, D chooses for connective moves, which branch she would like to see played next. However, our game never returns to the branch not chosen, so D has a genuine choice to make. The winning condition of our game is based on a natural property of single positions instead of the size of the entire game tree.

We give now the precise definition of our game.

Definition 3.1. Let \mathbb{A}_0 and \mathbb{B}_0 be sets of pointed Φ -Kripke-models and let $k_0 \in \mathbb{N}$. *The formula size game between the sets \mathbb{A}_0 and \mathbb{B}_0 , denoted $\text{FS}_{k_0}^\Phi(\mathbb{A}_0, \mathbb{B}_0)$, has two players, S and D. The number k_0 is the *resource parameter* of the game. The starting position of the game is $(k_0, \mathbb{A}_0, \mathbb{B}_0)$. Let the position after n moves be $(k, \mathbb{A}, \mathbb{B})$. If $k = 0$, D wins the game. If $k > 0$, S has the following five moves to choose from:*

- *\vee -move:* First, S chooses natural numbers k_1 and k_2 and sets \mathbb{A}_1 and \mathbb{A}_2 such that $k_1 + k_2 + 1 = k$ and $\mathbb{A}_1 \cup \mathbb{A}_2 = \mathbb{A}$. Then D decides whether the game continues from the position $(k_1, \mathbb{A}_1, \mathbb{B})$ or the position $(k_2, \mathbb{A}_2, \mathbb{B})$.
- *\wedge -move:* First, S chooses natural numbers k_1 and k_2 and sets \mathbb{B}_1 and \mathbb{B}_2 such that $k_1 + k_2 + 1 = k$ and $\mathbb{B}_1 \cup \mathbb{B}_2 = \mathbb{B}$. Then D decides whether the game continues from the position $(k_1, \mathbb{A}, \mathbb{B}_1)$ or the position $(k_2, \mathbb{A}, \mathbb{B}_2)$.
- *\diamond -move:* S chooses a function $f : \mathbb{A} \rightarrow \square\mathbb{A}$ such that $f(\mathcal{A}, w) \in \square(\mathcal{A}, w)$ for all $(\mathcal{A}, w) \in \mathbb{A}$ and the game continues from the position $(k - 1, \diamond_f\mathbb{A}, \square\mathbb{B})$.
- *\square -move:* S chooses a function $g : \mathbb{B} \rightarrow \square\mathbb{B}$ such that $g(\mathcal{B}, w) \in \square(\mathcal{B}, w)$ for all $(\mathcal{B}, w) \in \mathbb{B}$ and the game continues from the position $(k - 1, \square\mathbb{A}, \diamond_g\mathbb{B})$.
- *Lit-move:* S chooses a literal $l \in \text{Lit}(\Phi)$. If l separates the sets \mathbb{A} and \mathbb{B} , S wins. Otherwise D wins.

Since D wins if k runs out, the parameter k can be thought of as a resource of S that she spends on connectives and literals. In addition if there is a model $(\mathcal{M}, w) \in \mathbb{A}$ (or \mathbb{B}) for which $\Box(\mathcal{M}, w) = \emptyset$, then S cannot make a \Diamond - (or \Box -)move.

We prove that the formula size game indeed characterizes the separation of two sets of pointed Kripke-models by a formula of a given size.

Theorem 3.2. *Let \mathbb{A} and \mathbb{B} be sets of pointed Φ -models and let k be natural number. Then the following conditions are equivalent:*

(win) $_k$ *S has a winning strategy in the game $\text{FS}_k^\Phi(\mathbb{A}, \mathbb{B})$.*

(sep) $_k$ *There is a formula $\varphi \in \text{ML}(\Phi)$ such that $\text{sz}(\varphi) \leq k$ and the formula φ separates \mathbb{A} from \mathbb{B} .*

Proof. The proof proceeds by induction on the number k . First assume $k = 1$. If S makes any non-literal move, D wins since $k = 0$ in the following position. So the only possibility for a winning strategy is a literal move. There is a winning literal move if and only if there is a literal which separates \mathbb{A}_0 from \mathbb{B}_0 . Thus (win) $_1 \Leftrightarrow$ (sep) $_1$.

Suppose then that $k > 1$ and (win) $_l \Leftrightarrow$ (sep) $_l$ for all $l < k$. Assume first that (win) $_k$ holds. Consider the following cases according to the first move in the winning strategy of S. For \vee - and \wedge -moves we use the index i to always mean $i \in \{1, 2\}$.

- (a) Assume the first move of the winning strategy is a literal move and φ is the literal chosen by S. Then φ separates \mathbb{A} and \mathbb{B} and $\text{sz}(\varphi) = 1$ so (sep) $_k$ trivially holds.
- (b) Assume that the first move of the winning strategy of S is a \vee -move choosing numbers $k_1, k_2 \in \mathbb{N}$ such that $k_1 + k_2 + 1 = k$, and sets $\mathbb{A}_1, \mathbb{A}_2 \subseteq \mathbb{A}$ such that $\mathbb{A}_1 \cup \mathbb{A}_2 = \mathbb{A}$. Since this move is given by a winning strategy, S has a winning strategy for both possible continuations of the game, $(k_1, \mathbb{A}_1, \mathbb{B})$ and $(k_2, \mathbb{A}_2, \mathbb{B})$. Since $k_i < k$, by induction hypothesis there is a formula ψ_i such that $\text{sz}(\psi_i) \leq k_i$ and ψ_i separates \mathbb{A}_i from \mathbb{B} . Thus $\mathbb{A}_i \models \psi_i$ so $\mathbb{A} \models \psi_1 \vee \psi_2$. On the other hand $\mathbb{B} \models \neg\psi_1$ and $\mathbb{B} \models \neg\psi_2$ so $\mathbb{B} \models \neg(\psi_1 \vee \psi_2)$. Therefore the formula $\psi_1 \vee \psi_2$ separates \mathbb{A} from \mathbb{B} . In addition $\text{sz}(\psi_1 \vee \psi_2) = \text{sz}(\psi_1) + \text{sz}(\psi_2) + 1 \leq k_1 + k_2 + 1 = k$ so (sep) $_k$ holds.
- (c) The case in which the first move of the winning strategy of S is a right splitting move is proved in the same way as the previous one, with the roles of \mathbb{A} and \mathbb{B} switched, and disjunction replaced by conjunction.
- (d) Assume that the first move of the winning strategy of S is a \Diamond -move choosing a function $f : \mathbb{A} \rightarrow \Box\mathbb{A}$ such that $f(\mathcal{A}, w) \in \Box(\mathcal{A}, w)$ for all $(\mathcal{A}, w) \in \mathbb{A}$. The game continues from the position $(k - 1, \Diamond_f\mathbb{A}, \Box\mathbb{B})$ and S has a winning strategy from this position. By induction hypothesis there is a formula ψ such that $\text{sz}(\psi) \leq k - 1$ and ψ separates $\Diamond_f\mathbb{A}$ from $\Box\mathbb{B}$. Now for every $(\mathcal{A}, w) \in \mathbb{A}$ we have $f(\mathcal{A}, w) \in \Box(\mathcal{A}, w)$ and $f(\mathcal{A}, w) \models \psi$. Therefore $\mathbb{A} \models \Diamond\psi$. On the other hand $\Box\mathbb{B} \models \neg\psi$ so for every $(\mathcal{B}, w) \in \mathbb{B}$ and every $(\mathcal{B}, v) \in \Box(\mathcal{B}, w)$ we have $(\mathcal{B}, v) \not\models \psi$. Thus $\mathbb{B} \models \neg\Diamond\psi$. So the formula $\Diamond\psi$ separates \mathbb{A} from \mathbb{B} and since $\text{sz}(\Diamond\psi) = \text{sz}(\psi) + 1 \leq k$, (sep) $_k$ holds.
- (e) The case in which the first move of the winning strategy of S is a right successor move is similar to the case of left successor move. It suffices to switch the classes \mathbb{A} and \mathbb{B} , and replace \Diamond with \Box .

Now assume (sep) $_k$ holds, and φ is the formula separating \mathbb{A} from \mathbb{B} . We obtain a winning strategy of S for the game $\text{FS}_k^\Phi(\mathbb{A}, \mathbb{B})$ using φ as follows:

- (a) If φ is a literal, S wins the game by making the corresponding literal move.

- (b) Assume that $\varphi = \psi_1 \vee \psi_2$. Let $\mathbb{A}_i := \{(\mathcal{A}, w) \in \mathbb{A} \mid (\mathcal{A}, w) \models \psi_i\}$. Since $\mathbb{A} \models \varphi$ we have $\mathbb{A}_1 \cup \mathbb{A}_2 = \mathbb{A}$. In addition, since $\mathbb{B} \models \neg\varphi$, we have $\mathbb{B} \models \neg\psi_i$. Thus ψ_i separates \mathbb{A}_i from \mathbb{B} . Since $\text{sz}(\psi_1) + \text{sz}(\psi_2) + 1 = \text{sz}(\varphi) \leq k$, there are $k_1, k_2 \in \mathbb{N}$ such that $k_1 + k_2 + 1 = k$ and $\text{sz}(\psi_i) \leq k_i$. By induction hypothesis S has winning strategies for the games $\text{FS}_{k_i}^\Phi(\mathbb{A}_i, \mathbb{B})$. Since $k \geq \text{sz}(\varphi) \geq 1$, S can start the game $\text{FS}_k^\Phi(\mathbb{A}, \mathbb{B})$ with a \vee -move choosing the numbers k_1 and k_2 and the sets \mathbb{A}_1 and \mathbb{A}_2 . Then S wins the game by following the winning strategy for whichever position D chooses.
- (c) Assume that $\varphi = \psi_1 \wedge \psi_2$. Let $\mathbb{B}_1 := \{(\mathcal{B}, w) \in \mathbb{B} \mid (\mathcal{B}, w) \not\models \psi_1\}$ and $\mathbb{B}_2 := \{(\mathcal{B}, w) \in \mathbb{B} \mid (\mathcal{B}, w) \not\models \psi_2\}$. Since $\mathbb{B} \models \neg\varphi$, we have $\mathbb{B}_1 \cup \mathbb{B}_2 = \mathbb{B}$. In addition, since $\mathbb{A} \models \varphi$, we have $\mathbb{A} \models \psi_1$ and $\mathbb{A} \models \psi_2$. Thus ψ_1 separates \mathbb{A} from \mathbb{B}_1 while ψ_2 separates \mathbb{A} from \mathbb{B}_2 . As in the previous case, there are $k_1, k_2 \in \mathbb{N}$ such that $k_1 + k_2 = k$, $\text{sz}(\psi_1) \leq k_1$ and $\text{sz}(\psi_2) \leq k_2$. By induction hypothesis S has a winning strategy for the games $\text{FS}_{k_1}^\Phi(\mathbb{A}, \mathbb{B}_1)$ and $\text{FS}_{k_2}^\Phi(\mathbb{A}, \mathbb{B}_2)$. S wins the game $\text{FS}_k^\Phi(\mathbb{A}, \mathbb{B})$ by starting with a \wedge -move choosing the numbers k_1 , and k_2 and the sets \mathbb{B}_1 and \mathbb{B}_2 and proceeding according to the winning strategies for the games $\text{FS}_{k_1}^\Phi(\mathbb{A}, \mathbb{B}_1)$ and $\text{FS}_{k_2}^\Phi(\mathbb{A}, \mathbb{B}_2)$.
- (d) Assume that $\varphi = \diamond\psi$. Since $\mathbb{A} \models \varphi$, for every $(\mathcal{A}, w) \in \mathbb{A}$ there is $(\mathcal{A}, v_w) \in \square(\mathcal{A}, w)$ such that $(\mathcal{A}, v_w) \models \psi$. We define the function $f : \mathbb{A} \rightarrow \square\mathbb{A}$ by $f(\mathcal{A}, w) = (\mathcal{A}, v_w)$. Clearly $\diamond_f\mathbb{A} \models \psi$. On the other hand $\mathbb{B} \models \neg\varphi$ so for each $(\mathcal{B}, w) \in \mathbb{B}$ and each $(\mathcal{B}, v) \in \square(\mathcal{B}, w)$ we have $(\mathcal{B}, v) \not\models \psi$. Therefore $\square\mathbb{B} \models \neg\psi$ and the formula ψ separates $\diamond_f\mathbb{A}$ from $\square\mathbb{B}$. Moreover, $\text{sz}(\psi) = \text{sz}(\varphi) - 1 \leq k - 1$ so by induction hypothesis S has a winning strategy for the game $\text{FS}_{k-1}^\Phi(\diamond_f\mathbb{A}, \square\mathbb{B})$. Since $k \geq \text{sz}(\varphi) \geq 1$, S can start the game $\text{FS}_k^\Phi(\mathbb{A}, \mathbb{B})$ with a \diamond -move choosing the function f . Then S wins the game by following the winning strategy for the game $\text{FS}_{k-1}^\Phi(\diamond_f\mathbb{A}, \square\mathbb{B})$.
- (e) Assume finally that $\varphi = \square\psi$. Since $\mathbb{A} \models \varphi$, as in the previous case we obtain $\square\mathbb{A} \models \psi$. On the other hand, since $\mathbb{B} \models \neg\varphi$, for every $(\mathcal{B}, w) \in \mathbb{B}$ there is $(\mathcal{B}, v_w) \in \square(\mathcal{B}, w)$ such that $(\mathcal{B}, v_w) \not\models \psi$. We define the function $g : \mathbb{B} \rightarrow \square\mathbb{B}$ by $g(\mathcal{B}, w) = (\mathcal{B}, v_w)$. Clearly $\diamond_g\mathbb{B} \models \neg\psi$ so the formula ψ separates the sets $\square\mathbb{A}$ and $\diamond_g\mathbb{B}$. By induction hypothesis S has a winning strategy for the game $\text{FS}_{k-1}^\Phi(\square\mathbb{A}, \diamond_g\mathbb{B})$. S wins the game $\text{FS}_k^\Phi(\mathbb{A}, \mathbb{B})$ by starting with a \square -move choosing the function g and proceeding according to the winning strategy of the game $\text{FS}_{k-1}^\Phi(\square\mathbb{A}, \diamond_g\mathbb{B})$.

□

Remark 3.3. In this form, the game $\text{FS}_k^\Phi(\mathbb{A}, \mathbb{B})$ tracks the size of the separating formula but with slight modifications it could track different things such as the number or nesting depth of specific operators. See e.g. the conference paper [18] where the game counts propositional connectives and modal operators with two separate parameters.

Note that in Theorem 3.2 we allow the set of proposition symbols Φ to be infinite. This is in contrast with other similar games, such as the bisimulation game and the n -bisimulation game. For an example let $\Phi = \{p_i \mid i \in \mathbb{N}\}$ and $W = \{w\} \cup \{w_i \mid i \in \mathbb{N}\}$. Furthermore let (\mathcal{A}, w) be a pointed model, where $\text{dom}(\mathcal{A}) = W$, $R^{\mathcal{A}} = \{(w, w_i) \mid i \in \mathbb{N}\}$ and $V^{\mathcal{A}}(p_i) = \{w_j \mid j \geq i\}$ for each $i \in \mathbb{N}$. Let (\mathcal{B}, w) be the same model with the addition of a point $w_{\mathbb{N}}$ in which all propositions are true. In other words $\text{dom}(\mathcal{B}) = W \cup \{w_{\mathbb{N}}\}$, $R^{\mathcal{B}} = R^{\mathcal{A}} \cup \{(w, w_{\mathbb{N}})\}$ and $V^{\mathcal{B}}(p_i) = V^{\mathcal{A}}(p_i) \cup \{w_{\mathbb{N}}\}$ for each $i \in \mathbb{N}$.

We see that by moving to $w_{\mathbb{N}}$, S wins the (n) -bisimulation game between the models (\mathcal{A}, w) and (\mathcal{B}, w) , even though the models satisfy exactly the same ML-formulas.

We prove next that k -bisimilarity implies that D has winning strategy in the formula size game with resource parameter k . This simple observation is used in the next section, when we apply the game FS_k^Φ for proving a succinctness result for FO over ML.

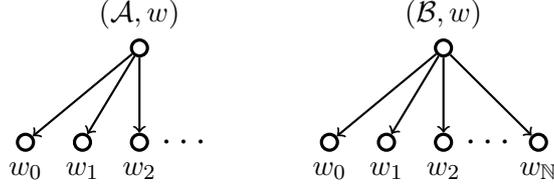


Figure 1: The pointed models (\mathcal{A}, w) and (\mathcal{B}, w) .

Theorem 3.4. *Let \mathbb{A} and \mathbb{B} be sets of pointed models and let $k \in \mathbb{N}$. If there are $(k-1)$ -bisimilar pointed models $(\mathcal{A}, w) \in \mathbb{A}$ and $(\mathcal{B}, v) \in \mathbb{B}$, then D has a winning strategy for the game $\text{FS}_k^\Phi(\mathbb{A}, \mathbb{B})$.*

Proof. The proof proceeds by induction on the number $k \in \mathbb{N}$. If $k = 1$ and $(\mathcal{A}, w) \in \mathbb{A}$ and $(\mathcal{B}, v) \in \mathbb{B}$ are 0-bisimilar and thus satisfy the same literals. Thus there is no literal $\varphi \in \text{ML}$ that separates the sets \mathbb{A} and \mathbb{B} . Thus any literal move by S leads to D winning. In addition, any non-literal move leads to a following position with $k = 0$ so D wins the game $\text{FS}_1^\Phi(\mathbb{A}, \mathbb{B})$.

Assume that $k > 1$ and $(\mathcal{A}, w) \in \mathbb{A}$ and $(\mathcal{B}, v) \in \mathbb{B}$ are $(k-1)$ -bisimilar. We consider the cases of the first move of S in the game $\text{FS}_k^\Phi(\mathbb{A}, \mathbb{B})$.

If S makes a literal move, D will win as in the basic step.

If S starts with a \vee -move choosing the numbers k_1 and k_2 and the sets \mathbb{A}_1 and \mathbb{A}_2 , then since $\mathbb{A}_1 \cup \mathbb{A}_2 = \mathbb{A}$, D can choose the next position $(k_i, \mathbb{A}_i, \mathbb{B})$, in such a way that $(\mathcal{A}, w) \in \mathbb{A}_i$. Then we have $k_i < k$ so by induction hypothesis D has a winning strategy for the game $\text{FS}_{k_i}^\Phi(\mathbb{A}_i, \mathbb{B})$. The case of a \wedge -move is similar.

If S starts with a \diamond -move choosing a function $f : \mathbb{A} \rightarrow \square\mathbb{A}$, then since (\mathcal{A}, w) and (\mathcal{B}, v) are $(k-1)$ -bisimilar, there is a pointed model $(\mathcal{B}, v') \in \square(\mathcal{B}, v)$ that is $(k-2)$ -bisimilar with the pointed model $f(\mathcal{A}, w)$. By induction hypothesis D has a winning strategy in $\text{FS}_{k-1}^\Phi(\diamond_f\mathbb{A}, \square\mathbb{B})$. The case of a \square -move is similar. \square

4 Succinctness of FO over ML

In this section, we illustrate the use of the formula size game FS_k^Φ by proving a non-elementary succinctness gap between bisimulation invariant first-order logic and modal logic. We also show that this gap is already present between a limited 2-dimensional modal logic ML^2 and basic modal logic.

A similar gap between FO and linear temporal logic LTL has already been established in the literature. In his PhD thesis [27], Stockmeyer proved that the satisfiability problem of FO over words is of non-elementary complexity. He reduced the problem of nonemptiness of star-free regular expressions to this satisfiability problem. Etessami and Wilke pointed out in [5] that careful examination of Stockmeyer's proof yields FO sentences with size $\mathcal{O}(n)$ such that the minimal words satisfying these sentences have length non-elementary in n^3 . Since all satisfiable formulas of LTL have a satisfying model at most exponential in the size of the formula, a non-elementary succinctness gap between FO and LTL is obtained.

4.1 A property of pointed models

For the remainder of this section we consider only the case where the set Φ of propositional symbols is empty. This makes all points in Kripke-models propositionally equivalent so the only formulas available for the win condition of S in the game FS_k^Φ are \perp and \top . Thus S can only win with a literal move from position $(k, \mathbb{A}, \mathbb{B})$ if either $\mathbb{A} = \emptyset$ and $\mathbb{B} \neq \emptyset$, or $\mathbb{A} \neq \emptyset$ and $\mathbb{B} = \emptyset$.

We will use the following two classes in our application of the formula size game FS_k^Φ :

- \mathbb{A}_n is the class of all pointed models (\mathcal{A}, w) such that for all $(\mathcal{A}, u), (\mathcal{A}, v) \in \square(\mathcal{A}, w)$, the models (\mathcal{A}, u) and (\mathcal{A}, v) are n -bisimilar.

- \mathbb{B}_n is the complement of \mathbb{A}_n .

Lemma 4.1. *For each $n \in \mathbb{N}$ there is a formula $\varphi_n(x) \in \text{FO}$ that separates the classes \mathbb{A}_n and \mathbb{B}_n such that the size of $\varphi_n(x)$ is linear with respect to n , i.e., $\text{sz}(\varphi_n) = \mathcal{O}(n)$.*

Proof. We first define formulas $\psi_n(x, y) \in \text{FO}$ such that $(\mathcal{M}, u) \simeq_n (\mathcal{M}, v)$ if and only if $\mathcal{M} \models \psi_n[u/x, v/y]$. The formulas $\psi_n(x, y)$ are defined recursively as follows:

$$\begin{aligned} \psi_1(x, y) &:= \exists s R(x, s) \leftrightarrow \exists t R(y, t) \\ \psi_{n+1}(x, y) &:= \forall s \exists t \left((R(x, s) \rightarrow R(y, t)) \wedge (R(y, s) \rightarrow R(x, t)) \right. \\ &\quad \left. \wedge (R(x, s) \vee R(y, s) \rightarrow \psi_n(s, t)) \right) \end{aligned}$$

Clearly these formulas express n -bisimilarity as intended. When we interpret the equivalence and implications as shorthand in the standard way, we get the sizes $\text{sz}(\psi_1) = 11$ and $\text{sz}(\psi_{n+1}) = \text{sz}(\psi_n) + 14$. Thus $\text{sz}(\psi_n) = 14n - 3$.

Now we can define the formulas φ_n :

$$\varphi_n(x) := \forall y \forall z (R(x, y) \wedge R(x, z) \rightarrow \psi_n(y, z)).$$

Clearly for every $(\mathcal{A}, w) \in \mathbb{A}_n$ we have $\mathcal{A} \models \varphi_n[w/x]$ and for every $(\mathcal{B}, v) \in \mathbb{B}_n$ we have $\mathcal{B} \models \neg \varphi_n[w/x]$ so the formula φ_n separates the classes \mathbb{A}_n and \mathbb{B}_n . Furthermore, $\text{sz}(\varphi_n) = \text{sz}(\psi_n) + 6 = 14n + 3$ so the size of φ_n is linear³ with respect to n . \square

Lemma 4.2. *For each $n \in \mathbb{N}$, the formula φ_n is $(n + 1)$ -bisimulation invariant.*

Proof. Let (\mathcal{A}, w) and (\mathcal{B}, v) be $(n + 1)$ -bisimilar pointed models. Assume that $\mathcal{A} \models \varphi_n[w/x]$. If $(\mathcal{B}, v_1), (\mathcal{B}, v_2) \in \square(\mathcal{B}, v)$, by $(n + 1)$ -bisimilarity there are $(\mathcal{A}, w_1), (\mathcal{A}, w_2) \in \square(\mathcal{A}, w)$ such that $(\mathcal{A}, w_1) \simeq_n (\mathcal{B}, v_1)$ and $(\mathcal{A}, w_2) \simeq_n (\mathcal{B}, v_2)$. Since $\mathcal{A} \models \varphi_n[w/x]$, we have $(\mathcal{B}, v_1) \simeq_n (\mathcal{A}, w_1) \simeq_n (\mathcal{A}, w_2) \simeq_n (\mathcal{B}, v_2)$ so $\mathcal{B} \models \psi_n[v_1/x, v_2/y]$. Thus, we see that $\mathcal{B} \models \varphi_n[v/x]$. \square

It follows now from van Benthem's characterization theorem that each φ_n is equivalent to some ML-formula. Thus, we get the following corollary.

Corollary 4.3. *For each $n \in \mathbb{N}$, there is a formula $\vartheta_n \in \text{ML}$ that separates the classes \mathbb{A}_n and \mathbb{B}_n .*

4.2 Set theoretic construction of pointed models

We have shown that the classes \mathbb{A}_n and \mathbb{B}_n can be separated both in ML and in FO. Furthermore the size of the FO-formula is linear with respect to n . It only remains to ask: what is the size of the smallest ML-formula that separates the classes \mathbb{A}_n and \mathbb{B}_n ? To answer this we will need suitable subsets of \mathbb{A}_n and \mathbb{B}_n to play the formula size game on.

Definition 4.4. Let $n \in \mathbb{N}$. *The finite levels of the cumulative hierarchy* are defined recursively as follows:

$$\begin{aligned} \mathbb{V}_0 &= \emptyset \\ \mathbb{V}_{n+1} &= \mathcal{P}(\mathbb{V}_n) \end{aligned}$$

For every $n \in \mathbb{N}$, \mathbb{V}_n is a *transitive set*, i.e., for every $a \in \mathbb{V}_n$ and every $b \in a$ it holds that $b \in \mathbb{V}_n$. Thus it is reasonable to define a model $\mathcal{F}_n = (\mathbb{V}_n, R_n)$, where for all $a, b \in \mathbb{V}_n$ it holds that $(a, b) \in R_n \Leftrightarrow b \in a$.

For every point $a \in \mathbb{V}_n$ we denote by (\mathcal{M}_a, a) the pointed model, where \mathcal{M}_a is the submodel of \mathcal{F}_n generated by the point a .

³**Acknowledgement.** We are grateful to Martin Lück for coming up with a linear size formula $\psi_n(x, y)$ to replace our previous one that was of exponential size.

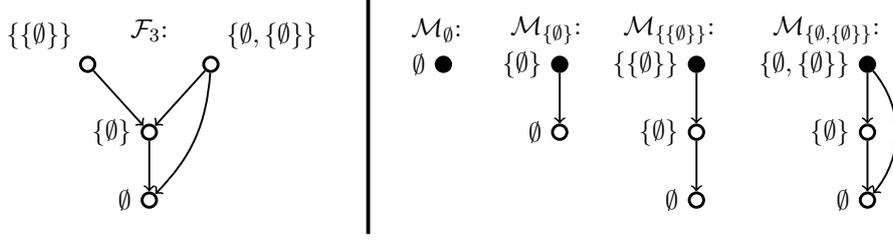


Figure 2: The model \mathcal{F}_3 and its generated submodels

Lemma 4.5. *Let $n \in \mathbb{N}$ and $a, b \in \mathbb{V}_{n+1}$. If $a \neq b$, then $(\mathcal{M}_a, a) \not\equiv_n (\mathcal{M}_b, b)$.*

Proof. We prove the claim by induction on n . The basic step $n = 0$ is trivial since \mathbb{V}_1 only has one element. For the induction step, assume that $a, b \in \mathbb{V}_{n+1}$ and $a \neq b$. Assume further for contradiction that $(\mathcal{M}_a, a) \equiv_n (\mathcal{M}_b, b)$. Since $a \neq b$, by symmetry we can assume that there is $x \in a$ such that $x \notin b$. By n -bisimilarity there is $y \in b$ such that (\mathcal{M}_x, x) and (\mathcal{M}_y, y) are $(n - 1)$ -bisimilar. Since $x \in a \in \mathbb{V}_{n+1}$ and $y \in b \in \mathbb{V}_{n+1}$, we have $x, y \in \mathbb{V}_n$. By induction hypothesis we obtain $x = y$. This is a contradiction, since $x \notin b$ and $y \in b$. \square

If \mathbb{A} is a set of pointed models, the pointed model $\Delta\mathbb{A}$ is formed by taking all the pointed models of \mathbb{A} and connecting a new root to their distinguished points as illustrated in Figure 3. To make sure that $(\Delta\mathbb{A}, w)$ is bisimilar with (\mathcal{A}, v) for any $(\mathcal{A}, v) \in \Delta\mathbb{A}$, we require that the models in \mathbb{A} are compatible in possible intersections. The precise definition is the following.

Let \mathbb{A} be a set of pointed models such that for all $(\mathcal{A}, v), (\mathcal{A}', v') \in \mathbb{A}$ it holds that $R^{\mathcal{A}} \upharpoonright (\text{dom}(\mathcal{A}) \cap \text{dom}(\mathcal{A}')) = R^{\mathcal{A}'} \upharpoonright (\text{dom}(\mathcal{A}) \cap \text{dom}(\mathcal{A}'))$ and let $w \notin \text{dom}(\mathcal{A})$ for all $(\mathcal{A}, v) \in \mathbb{A}$. We use the notation $\Delta\mathbb{A} := (\mathcal{M}, w)$, where

$$\text{dom}(\mathcal{M}) = \{w\} \cup \bigcup \{\text{dom}(\mathcal{A}) \mid (\mathcal{A}, v) \in \mathbb{A}\}, \text{ and}$$

$$R^{\mathcal{M}} = \{(w, v) \mid (\mathcal{A}, v) \in \mathbb{A}\} \cup \bigcup \{R^{\mathcal{A}} \mid (\mathcal{A}, v) \in \mathbb{A}\}.$$

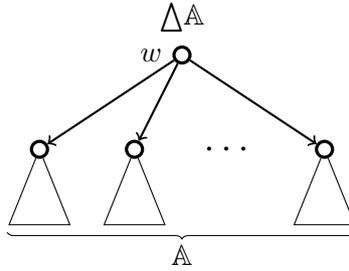


Figure 3: The pointed model $\Delta\mathbb{A}$

For each $n \in \mathbb{N}$ we define the following sets of pointed models:

$$\mathbb{C}_n := \{\Delta\{(\mathcal{M}_a, a)\} \mid a \in \mathbb{V}_{n+1}\}$$

$$\mathbb{D}_n := \{\Delta\{(\mathcal{M}_a, a), (\mathcal{M}_b, b)\} \mid a, b \in \mathbb{V}_{n+1}, a \neq b\}.$$

In other words the pointed models in \mathbb{C}_n have a single successor from level $n + 1$ of the cumulative hierarchy, whereas the pointed models in \mathbb{D}_n have two different successors from the same set. Therefore clearly $\mathbb{C}_n \subseteq \mathbb{A}_n$ and by Lemma 4.5 also $\mathbb{D}_n \subseteq \mathbb{B}_n$. In the next subsection we will use these sets in the formula size game.

It is well known that the cardinality of \mathbb{V}_n is the exponential tower of $n - 1$. Thus, the cardinality of \mathbb{C}_n is $\text{twr}(n)$.

Lemma 4.6. *If $n \in \mathbb{N}$, we have $|\mathbb{C}_n| = |\mathbb{V}_{n+1}| = \text{twr}(n)$.* \square

4.3 Graph colorings and winning strategies in FS_k^Φ

Our aim is to prove that any ML-formula ϑ_n separating the sets \mathbb{C}_n and \mathbb{D}_n is of size at least $\text{twr}(n-1)$. To do this, we make use of a surprising connection between the chromatic numbers of certain graphs related to pairs of the form (\mathbb{V}, \mathbb{E}) , where $\mathbb{V} \subseteq \mathbb{C}_n$ and $\mathbb{E} \subseteq \mathbb{D}_n$, and existence of a winning strategy for D in the game $\text{FS}_k^\Phi(\mathbb{V}, \mathbb{E})$.

Let $n \in \mathbb{N}$, $\emptyset \neq \mathbb{V} \subseteq \mathbb{C}_n$ and $\mathbb{E} \subseteq \mathbb{D}_n$. Then $\mathcal{G}(\mathbb{V}, \mathbb{E})$ denotes the graph (V, E) , where

$$V = \square\mathbb{V} \text{ and} \\ E = \{((\mathcal{M}, w), (\mathcal{M}', w')) \in V \times V \mid \Delta\{(\mathcal{M}, w), (\mathcal{M}', w')\} \in \mathbb{E}\}.$$

That is, since models on the left all have exactly one successor, and ones on the right have exactly two successors from the same basic set, we can take the graph where these successors are nodes and the pairs on the right define the edges. Note that a pair on the right only produces an edge if both elements of the pair are present on the left.

Definition 4.7. Let $\mathcal{G} = (V, E)$ be a graph and let C be a set. A function $\chi : V \rightarrow C$ is a *coloring* of the graph \mathcal{G} if for all $u, v \in V$ it holds that if $(u, v) \in E$, then $\chi(u) \neq \chi(v)$. If the set C has k elements, then χ is called a *k-coloring* of \mathcal{G} .

The *chromatic number* of \mathcal{G} , denoted by $\chi(\mathcal{G})$, is the smallest number $k \in \mathbb{N}$ for which there is a *k-coloring* of \mathcal{G} .

When playing the formula size game $\text{FS}_k^\Phi(\mathbb{V}, \mathbb{E})$, connective moves correspond with dividing either the vertex set or the edge set of the graph $\mathcal{G}(\mathbb{V}, \mathbb{E})$ into two parts, forming two new graphs. In the next lemma we get simple arithmetic estimates for the behaviour of chromatic numbers in such divisions. In the case of a vertex set split, if the two new graphs are colored with separate colors, combining these colorings yields a coloring of the whole graph. For an edge split, the full graph is colored with pairs of colors given by the two new colorings. If two vertices are adjacent in the full graph, at least one of the new colorings will color them with a different color and the pairs of colors will be different.

Lemma 4.8. *Let $\mathcal{G} = (V, E)$ be a graph.*

1. *Let $V_1, V_2 \subseteq V$ be nonempty such that $V_1 \cup V_2 = V$ and let $\mathcal{G}_1 = (V_1, E \upharpoonright V_1)$ and $\mathcal{G}_2 = (V_2, E \upharpoonright V_2)$. Then we have $\chi(\mathcal{G}) \leq \chi(\mathcal{G}_1) + \chi(\mathcal{G}_2)$.*
2. *Let $E_1, E_2 \subseteq E$ such that $E_1 \cup E_2 = E$ and let $\mathcal{G}_1 = (V, E_1)$ and $\mathcal{G}_2 = (V, E_2)$. Then $\chi(\mathcal{G}) \leq \chi(\mathcal{G}_1)\chi(\mathcal{G}_2)$.*

Proof. 1. Let V_1, V_2, \mathcal{G}_1 and \mathcal{G}_2 be as in the claim and let $k_1 = \chi(\mathcal{G}_1)$ and $k_2 = \chi(\mathcal{G}_2)$. Let $\chi_1 : V_1 \rightarrow \{1, \dots, k_1\}$ be a k_1 -coloring of the graph \mathcal{G}_1 and let $\chi_2 : V_2 \rightarrow \{k_1+1, \dots, k_1+k_2\}$ be a k_2 -coloring of the graph \mathcal{G}_2 . Then it is straightforward to show that $\chi = \chi_1 \cup (\chi_2 \upharpoonright (V_2 \setminus V_1))$ is a $k_1 + k_2$ -coloring of the graph \mathcal{G} , whence $\chi(\mathcal{G}) \leq k_1 + k_2 = \chi(\mathcal{G}_1) + \chi(\mathcal{G}_2)$.

2. Let $\chi_1 : V \rightarrow \{1, \dots, k_1\}$ and $\chi_2 : V \rightarrow \{1, \dots, k_2\}$ be colorings of the graphs \mathcal{G}_1 and \mathcal{G}_2 , respectively. Then it is easy to verify that the map $\chi : V \rightarrow \{1, \dots, k_1\} \times \{1, \dots, k_2\}$ defined by $\chi(v) = (\chi_1(v), \chi_2(v))$ is a coloring of \mathcal{G} . Thus we obtain $\chi(\mathcal{G}) \leq |\{1, \dots, k_1\} \times \{1, \dots, k_2\}| = \chi(\mathcal{G}_1)\chi(\mathcal{G}_2)$. □

For the condition D maintains to win the game, we use the logarithm of the chromatic number of $\mathcal{G}(\mathbb{V}, \mathbb{E})$ as it behaves nicely with both kinds of splittings. Note that to achieve non-elementary formula size, it suffices to consider the number of binary connectives required before any modal moves can be made.

Lemma 4.9. *Assume $\emptyset \neq \mathbb{V} \subseteq \mathbb{C}_n$ and $\mathbb{E} \subseteq \mathbb{D}_n$ for some $n \in \mathbb{N}$ and let $k \in \mathbb{N}$. If $k \leq \log(\chi(\mathcal{G}(\mathbb{V}, \mathbb{E})))$, then D has a winning strategy in the game $\text{FS}_k^\Phi(\mathbb{V}, \mathbb{E})$.*

Proof. Let $n, k \in \mathbb{N}$ and assume that $\emptyset \neq \mathbb{V} \subseteq \mathbb{C}_n$, $\mathbb{E} \subseteq \mathbb{D}_n$ and $k \leq \log(\chi(\mathcal{G}(\mathbb{V}, \mathbb{E})))$. We prove the claim by induction on k .

If $k = 0$, then D wins the game.

If $k = 1$, any non-literal move of S leads to D winning. Since $\mathbb{V}, \mathbb{E} \neq \emptyset$ and all models are propositionally equivalent, D will also win if S makes a literal move.

Assume then that $k > 1$. If S starts the game with a literal move, then D wins as described above.

Assume that S begins the game with a \diamond - or \square -move. Since $\chi(\mathcal{G}(\mathbb{V}, \mathbb{E})) \geq 2$, there are pointed models $(\mathcal{M}, w), (\mathcal{M}', w') \in V$ such that $((\mathcal{M}, w), (\mathcal{M}', w')) \in E$. Thus $\Delta\{(\mathcal{M}, w)\}, \Delta\{(\mathcal{M}', w')\} \in \mathbb{V}$ and $\Delta\{(\mathcal{M}, w), (\mathcal{M}', w')\} \in \mathbb{E}$. In the following position $(k-1, \mathbb{V}', \mathbb{E}')$ it holds that $(\mathcal{M}, w) \in \mathbb{V}' \cap \mathbb{E}'$ or $(\mathcal{M}', w') \in \mathbb{V}' \cap \mathbb{E}'$. Thus the same pointed model is present on both sides of the game and by Theorem 3.4, D has a winning strategy for the game $\text{FS}_{k-1}^\Phi(\mathbb{V}', \mathbb{E}')$.

Assume that S begins the game with a \vee -move choosing the numbers $k_1, k_2 \in \mathbb{N}$ and the sets $\mathbb{V}_1, \mathbb{V}_2 \subseteq \mathbb{V}$. Consider the graphs $\mathcal{G}(\mathbb{V}, \mathbb{E}) = (V, E)$ and $\mathcal{G}(\mathbb{V}_i, \mathbb{E}) = (V_i, E_i)$. Since $\mathbb{V}_1 \cup \mathbb{V}_2 = \mathbb{V}$, we have $V_1 \cup V_2 = V$. In addition, by the definition of the graphs $\mathcal{G}(\mathbb{V}, \mathbb{E})$ and $\mathcal{G}(\mathbb{V}_i, \mathbb{E})$ we see that $E_i = E \upharpoonright V_i$. Thus by Lemma 4.8, we obtain $\chi(\mathcal{G}(\mathbb{V}, \mathbb{E})) \leq \chi(\mathcal{G}(\mathbb{V}_1, \mathbb{E})) + \chi(\mathcal{G}(\mathbb{V}_2, \mathbb{E}))$. It must hold that $k_1 \leq \log(\chi(\mathcal{G}(\mathbb{V}_1, \mathbb{E})))$ or $k_2 \leq \log(\chi(\mathcal{G}(\mathbb{V}_2, \mathbb{E})))$, since otherwise we would have

$$\begin{aligned} k &\leq \log(\chi(\mathcal{G}(\mathbb{V}, \mathbb{E}))) \leq \log(\chi(\mathcal{G}(\mathbb{V}_1, \mathbb{E})) + \chi(\mathcal{G}(\mathbb{V}_2, \mathbb{E}))) \\ &\leq \log(\chi(\mathcal{G}(\mathbb{V}_1, \mathbb{E}))) + \log(\chi(\mathcal{G}(\mathbb{V}_2, \mathbb{E}))) + 1 < k_1 + k_2 + 1 = k. \end{aligned}$$

Thus D can choose the next position of the game, $(k_i, \mathbb{V}_i, \mathbb{E})$, in such a way that $k_i \leq \log(\chi(\mathcal{G}(\mathbb{V}_i, \mathbb{E})))$. By induction hypothesis D has a winning strategy in the game $\text{FS}_{k_i}^\Phi(\mathbb{V}_i, \mathbb{E})$.

Assume then that S begins the game with a \wedge -move choosing the numbers $k_1, k_2 \in \mathbb{N}$ and the sets $\mathbb{E}_1, \mathbb{E}_2 \subseteq \mathbb{E}$. Consider now the graphs $\mathcal{G}(\mathbb{V}, \mathbb{E}) = (V, E)$ and $\mathcal{G}(\mathbb{V}, \mathbb{E}_i) = (V, E_i)$. Clearly $V_1 = V_2 = V$ and since $\mathbb{E}_1 \cup \mathbb{E}_2 = \mathbb{E}$, we have $E_1 \cup E_2 = E$. Thus by Lemma 4.8, we obtain $\chi(\mathcal{G}(\mathbb{V}, \mathbb{E})) \leq \chi(\mathcal{G}(\mathbb{V}, \mathbb{E}_1))\chi(\mathcal{G}(\mathbb{V}, \mathbb{E}_2))$. It must hold that $k_1 \leq \log(\chi(\mathcal{G}(\mathbb{V}, \mathbb{E}_1)))$ or $k_2 \leq \log(\chi(\mathcal{G}(\mathbb{V}, \mathbb{E}_2)))$, since otherwise we would have

$$\begin{aligned} k &\leq \log(\chi(\mathcal{G}(\mathbb{V}, \mathbb{E}))) \leq \log(\chi(\mathcal{G}(\mathbb{V}, \mathbb{E}_1))\chi(\mathcal{G}(\mathbb{V}, \mathbb{E}_2))) \\ &= \log(\chi(\mathcal{G}(\mathbb{V}, \mathbb{E}_1))) + \log(\chi(\mathcal{G}(\mathbb{V}, \mathbb{E}_2))) < k_1 + k_2 + 1 = k. \end{aligned}$$

Thus D can again choose the next position of the game, $(k_i, \mathbb{V}, \mathbb{E}_i)$, in such a way that $k_i \leq \log(\chi(\mathcal{G}(\mathbb{V}, \mathbb{E}_i)))$. By induction hypothesis D has a winning strategy in the game $\text{FS}_{k_i}^\Phi(\mathbb{V}, \mathbb{E}_i)$. \square

Theorem 4.10. *Let $n \in \mathbb{N}$. If a formula $\vartheta_n \in \text{ML}$ separates \mathbb{A}_n from \mathbb{B}_n , then $\text{sz}(\vartheta_n) > \text{twr}(n-1)$.*

Proof. Assume that a formula $\vartheta_n \in \text{ML}$ separates \mathbb{A}_n from \mathbb{B}_n . As observed in the end of Subsection 4.2, it holds that $\mathbb{C}_n \subseteq \mathbb{A}_n$ and $\mathbb{D}_n \subseteq \mathbb{B}_n$. Therefore ϑ_n also separates the sets \mathbb{C}_n and \mathbb{D}_n .

Assume for contradiction that $\text{sz}(\vartheta_n) \leq \text{twr}(n-1)$. By Theorem 3.2, S has a winning strategy in the game $\text{FS}_k^\Phi(\mathbb{C}_n, \mathbb{D}_n)$ for $k = \text{sz}(\vartheta_n)$.

On the other hand, by Lemma 4.6, we have $|\mathbb{C}_n| = \text{twr}(n)$ and the set \mathbb{D}_n consists of all the pointed models $\Delta\{(\mathcal{M}, w), (\mathcal{M}', w')\}$, where $\Delta\{(\mathcal{M}, w)\}, \Delta\{(\mathcal{M}', w')\} \in \mathbb{C}_n$, $(\mathcal{M}, w) \neq (\mathcal{M}', w')$. Thus the graph $\mathcal{G}(\mathbb{C}_n, \mathbb{D}_n)$ is isomorphic with the complete graph $K_{\text{twr}(n)}$. Therefore we obtain

$$\chi(\mathcal{G}(\mathbb{C}_n, \mathbb{D}_n)) = \chi(K_{\text{twr}(n)}) = \text{twr}(n).$$

By the assumption, $k \leq \text{twr}(n-1) = \log(\text{twr}(n)) = \log(\chi(\mathcal{G}(\mathbb{C}_n, \mathbb{D}_n)))$, so by Lemma 4.9, D also has a winning strategy in the game $\text{FS}_k^\Phi(\mathbb{C}_n, \mathbb{D}_n)$, which is a contradiction. \square

We now have everything we need for proving the non-elementary succinctness of FO over ML. By Lemma 4.1, for each $n \in \mathbb{N}$ there is a formula $\varphi_n(x) \in \text{FO}$ such that φ_n separates the classes \mathbb{A}_n and \mathbb{B}_n with $s(\varphi) = \mathcal{O}(n)$. On the other hand by Corollary 4.3, there is an equivalent formula $\vartheta_n \in \text{ML}$, but by Theorem 4.10 the size of ϑ_n must be at least $\text{twr}(n-1)$. So the property of all successors of a pointed model being n -bisimilar with each other can be expressed in FO with a formula of linear size, but in ML expressing it requires a formula of non-elementary size.

Corollary 4.11. *Bisimulation invariant FO is non-elementarily more succinct than ML.*

Remark 4.12. It is well known that the DAG-size of any formula φ is greater than or equal to the logarithm of the size of φ . Thus if ϑ_n is a formula as in Theorem 4.10, the DAG-size of ϑ_n must be at least $\text{twr}(n-2)$. Consequently the result of Corollary 4.11 also holds for DAG-size.

4.4 Succinctness of 2-dimensional modal logic

Our proof for the non-elementary succinctness gap between bisimulation invariant FO and ML is based on the fact that n -bisimilarity of two points $u, v \in W$ of a Kripke-model $\mathcal{M} = (W, R)$ is definable by a linear FO-formula $\psi_n(x, y)$ (see the proof of Lemma 4.1). We will now show that the property $(\mathcal{M}, u) \Leftrightarrow_n (\mathcal{M}, v)$ is succinctly expressible also in 2-dimensional modal logic.

The idea in 2-dimensional modal logic is that the truth of formulas is evaluated on pairs (u, v) of points of Kripke-models instead of single points. We refer to the book [23] of Marx and Venema and the series of papers [10], [11], [12] of Gabbay and Shehtman for a detailed exposition on 2-dimensional and multi-dimensional modal logics. For our purposes it suffices to consider the logic Gabbay and Shehtman call \mathbf{K}^2 . For consistency of notation in this paper we call the logic ML^2 and introduce it only semantically.

A Kripke-model \mathcal{T} for ML^2 consists of a set W of points, two binary accessibility relations R_1 and R_2 , and a valuation $V : \Phi \rightarrow \mathcal{P}(W)$. Correspondingly, ML^2 has two modal operators \diamond_1, \diamond_2 and their duals \square_1, \square_2 . The semantics of these operators is defined as follows:

- $(\mathcal{T}, (u, v)) \models \diamond_1 \varphi \Leftrightarrow$ there is $u' \in W$ such that uR_1u' and $(\mathcal{T}, (u', v)) \models \varphi$,
- $(\mathcal{T}, (u, v)) \models \diamond_2 \varphi \Leftrightarrow$ there is $v' \in W$ such that vR_2v' and $(\mathcal{T}, (u, v')) \models \varphi$,
- $(\mathcal{T}, (u, v)) \models \square_1 \varphi \Leftrightarrow$ for all $u' \in W$, if uR_1u' , then $(\mathcal{T}, (u', v)) \models \varphi$,
- $(\mathcal{T}, (u, v)) \models \square_2 \varphi \Leftrightarrow$ for all $v' \in W$, if vR_2v' , then $(\mathcal{T}, (u, v')) \models \varphi$.

Any pointed Kripke-model $(\mathcal{M}, w) = ((W, R, V), w)$ can be interpreted as the 2-dimensional pointed model $(\mathcal{M}_2, (w, w))$, where $\mathcal{M}_2 = (W, R, R, V)$. This gives us a meaningful way of defining properties of pointed models (\mathcal{M}, w) by formulas of ML^2 . In particular, we say that a formula $\varphi \in \text{ML}^2$ separates two classes \mathbb{A} and \mathbb{B} of pointed models if for all $(\mathcal{M}, w) \in \mathbb{A}$, $(\mathcal{M}_2, (w, w)) \models \varphi$ and for all $(\mathcal{M}, w) \in \mathbb{B}$, $(\mathcal{M}_2, (w, w)) \not\models \varphi$.

The *size* $\text{sz}(\varphi)$ of a formula $\varphi \in \text{ML}^2$ is defined in the same way as for formulas of ML; see Definition 2.3. In other words, $\text{sz}(\varphi)$ is the total number of modal operators, binary connectives and literals occurring in φ .

Observe now that two pointed models (\mathcal{M}, u) and (\mathcal{M}, v) with no propositional symbols are 1-bisimilar if and only if $(\mathcal{M}_2, (u, v)) \models \rho_1$, where $\rho_1 := \diamond_1 \top \leftrightarrow \diamond_2 \top$. Furthermore if $\rho_n \in \text{ML}_2$ defines the class of all 2-dimensional pointed models $(\mathcal{M}_2, (u, v))$ such that $(\mathcal{M}, u) \Leftrightarrow_n (\mathcal{M}, v)$, then $\rho_{n+1} := \square_1 \diamond_2 \rho_n \wedge \square_2 \diamond_1 \rho_n$ defines the class of all $(\mathcal{M}_2, (u, v))$ such that $(\mathcal{M}, u) \Leftrightarrow_{n+1} (\mathcal{M}, v)$.

Lemma 4.13. *For each $n \in \mathbb{N}$ there is a formula $\zeta_n \in \text{ML}^2$ that separates the classes \mathbb{A}_n and \mathbb{B}_n such that the size of ζ_n is exponential with respect to n , i.e., $\text{sz}(\zeta_n) = \mathcal{O}(2^n)$.*

Proof. Let ζ_n be the formula $\square_1 \square_2 \rho_n$. Then $(\mathcal{M}_2, (w, w)) \models \zeta_n$ if and only if (\mathcal{M}, u) and (\mathcal{M}, v) are n -bisimilar for all $(\mathcal{M}, u), (\mathcal{M}, v) \in \square(\mathcal{M}, w)$, whence ζ_n separates \mathbb{A}_n from its complement \mathbb{B}_n . An easy calculation shows that the size of ζ_n is $2^{n+4} - 3$. \square

By Theorem 4.3, for each $n \in \mathbb{N}$ there is a formula $\vartheta_n \in \text{ML}$ that is equivalent with ζ_n . On the other hand, by Theorem 4.10 the size of ϑ_n is at least $\text{twr}(n - 1)$. Thus, we obtain the non-elementary succinctness gap already between ML^2 and ML .

Corollary 4.14. *The 2-dimensional modal logic ML^2 is non-elementarily more succinct than ML .*

5 The formula size game for L_μ

To define a formula size game similar to the one of ML for L_μ , we will need some additional notation and concepts, since L_μ is significantly more complex than ML .

Let (V, E) be a tree and let $s, t \in V$. We say that s is *above* t if there is an E -path from s to t . We say that s is *below* t if t is above s . A triple (V, E, B) is a *tree with back edges* if (V, E) is a tree and s is below t for every $(s, t) \in B$.

We define for each formula $\varphi \in L_\mu$ its *syntax tree with back edges*, $T_\varphi = (V_\varphi, E_\varphi, B_\varphi, \text{lab}_\varphi)$ as follows. The set V_φ consists of occurrences of subformulas of φ and the relation E_φ is the subformula relation between those occurrences. Additionally lab_φ labels each vertex with its type (connective, modal operator, fixed point, literal or variable). Finally the relation B_φ contains a back edge from each vertex labelled with a variable to the successor of the fixed point binding that variable.

A *partial function* $f : M \rightarrow N$ is a function $f' : M' \rightarrow N$ for some $M' \subseteq M$. For a partial function $f : M \rightarrow N$ we denote by

$$f[b_1/a_1, \dots, b_m/a_m, -/a_{m+1}, \dots, -/a_{m+n}] := \\ (f \setminus \{(a_i, b) \mid i \in \{1, \dots, m+n\}, b \in N\}) \cup \{(a_i, b_i) \mid i \in \{1, \dots, m\}\},$$

the partial function, where values for $a_1, \dots, a_m \in M$ are set to $b_1, \dots, b_m \in N$ respectively, and the values for $a_{m+1}, \dots, a_{m+n} \in M$ are set as undefined.

We add some features to pointed Kripke-models for the game. A *clocked model* is a tuple (\mathcal{A}, w, c, a) , where (\mathcal{A}, w) is a pointed Kripke-model, $c : \text{Var} \rightarrow \kappa$ and $a \in \{\text{new}, \text{old}\}$. Here κ is a fixed cardinal larger than the size of the domain of \mathcal{A} . The partial function c associates to each fixed point a clock to show how many times the model can return to that fixed point. As clocked models traverse a graph in the game, we use the identifier *old* to keep track of where they have been previously. We suppress the age identifier a from the notation in cases where the distinction between new and old models does not matter.

For simplicity we use the symbols w and c extensively and they should be read as “the distinguished point and clocks of the model currently discussed” throughout the rest of the paper.

Let $\mathfrak{A} = (\mathcal{A}, w, c, a)$ be a clocked model and \mathbb{A} a set of clocked models. We redefine the following notations from the ML case for clocked models:

- $\Box \mathfrak{A} = \Box(\mathcal{A}, w, c, a) := \{(\mathcal{A}, w', c, a) \mid wR^{\mathcal{A}}w'\},$
- $\Box \mathbb{A} := \bigcup_{\mathfrak{A} \in \mathbb{A}} \Box \mathfrak{A}.$
- Let $f : \mathbb{A} \rightarrow \Box \mathbb{A}$ be a function such that $f(\mathfrak{A}) \in \Box \mathfrak{A}$ for every $\mathfrak{A} \in \mathbb{A}$. Then $\Diamond_f \mathbb{A} := f(\mathbb{A}).$

As for the ML -game, the \Box -notation denotes the set of all successors of a single clocked model or a set of clocked models. The clocks and age identifier are inherited. The set $\Diamond_f \mathbb{A}$ contains one successor for each clocked model in \mathbb{A} , given by the function f .

Now let \mathbb{A} be a set of clocked models, \mathbb{A}_0 a set of pointed models and $a \in \{\text{new}, \text{old}\}$. We use the following new notations:

- $\text{PM}(\mathbb{A}) := \{(\mathcal{M}, w) \mid (\mathcal{M}, w, c, a) \in \mathbb{A}\},$

- $\text{CM}(\mathbb{A}_0) := \{(\mathcal{M}, w, \emptyset, \text{new}) \mid (\mathcal{M}, w) \in \mathbb{A}_0\}$,
- $\mathbb{A}_a := \{(\mathcal{A}, w, c, b) \in \mathbb{A} \mid b = a\}$,
- $a(\mathbb{A}) := \{(\mathcal{A}, w, c, a) \mid (\mathcal{A}, w, c, b) \in \mathbb{A} \text{ for some } b \in \{\text{new}, \text{old}\}\}$.

The set $\text{PM}(\mathbb{A})$ contains the underlying pointed models of all clocked models in \mathbb{A} and the set $\text{CM}(\mathbb{A}_0)$ is the set of minimal clocked models with underlying pointed models from \mathbb{A}_0 . For an age identifier a , the set \mathbb{A}_a gives all clocked models in \mathbb{A} with that identifier and the set $a(\mathbb{A})$ gives all the models in \mathbb{A} with the age identifier changed to a .

We define for L_μ the standard approximant formulas that evaluate a fixed point only up to a bound. These approximants are formulas of *infinitary* L_μ , where infinite conjunctions and disjunctions are allowed.

Definition 5.1. Let α be an ordinal and $\psi(X)$ a formula of infinitary L_μ . Then the *approximant formulas* $\mu^\alpha X.\psi(X)$ and $\nu^\alpha X.\psi(X)$ are defined by recursion as follows:

- $\mu^0 X.\psi(X) = \perp$ and $\nu^0 X.\psi(X) = \top$,
- $\mu^{\alpha+1} X.\psi(X) = \psi(\mu^\alpha X.\psi(X))$ and $\nu^{\alpha+1} X.\psi(X) = \psi(\nu^\alpha X.\psi(X))$,
- $\mu^\lambda X.\psi(X) = \bigvee_{0 < \alpha < \lambda} \mu^\alpha X.\psi(X)$ and $\nu^\lambda X.\psi(X) = \bigwedge_{0 < \alpha < \lambda} \nu^\alpha X.\psi(X)$ for a limit ordinal λ .

This definition differs from the usual one (see e.g. [19]) in that we leave out the \perp disjunct and \top conjunct in the limit ordinal cases, and more importantly, we do not necessarily approximate all fixed points so the resulting formula is not necessarily in infinitary ML but instead in infinitary L_μ . During the game we approximate several fixed points at once, starting from a specific point in the formula. We define our own approximate formulas to reflect this.

Definition 5.2. Let $\varphi \in L_\mu$. Let $T_\varphi = (V_\varphi, E_\varphi, B_\varphi, \text{lab}_\varphi)$ be the syntax tree with back edges of φ and let $s \in V_\varphi$. Let s_1, \dots, s_n be the fixed point nodes above s in T_φ in order with s_1 being the outermost and s_n the innermost, and let $\text{lab}_\varphi(s_i) = \eta_i X_i$ for each $i \in \{1, \dots, n\}$. Let $c : \text{Var} \rightarrow \kappa$ be a partial function with $\text{dom}(c) = \{X_1, \dots, X_n\}$.

The (c, s) -*approximant* of φ , φ_s^c , is defined recursively as follows:

- if $\text{lab}(s) = l \in \text{Lit}$, then $\varphi_s^c = l$,
- if $\text{lab}(s) = \nabla \in \{\vee, \wedge\}$ and s_1, s_2 are the successors of s , then $\varphi_s^c = \varphi_{s_1}^c \nabla \varphi_{s_2}^c$,
- if $\text{lab}(s) = \Delta \in \{\diamond, \square\}$ and s_1 is the successor of s , then $\varphi_s^c = \Delta \varphi_{s_1}^c$,
- if $\text{lab}(s) = \eta X$, where $\eta \in \{\mu, \nu\}$ and $X \in \text{Var}$, and s_1 is the successor of s , then $\varphi_s^c = \eta X.\varphi_{s_1}^c$,
- if $\text{lab}(s) = X \in \text{Var} \setminus \text{dom}(c)$, then $\varphi_s^c = X$.
- Let $\text{lab}(s) = X_i$ and let u be the B_φ -successor of s .

If $c(X_i) = 0$, then if $\text{lab}(s_i) = \mu X_i$, $\varphi_s^c = \perp$ and if $\text{lab}(s_i) = \nu X_i$, $\varphi_s^c = \top$.

If $c(X_i) = \alpha + 1$ for some ordinal α , let $c_\alpha = c[\alpha/X_i, -/X_{i+1}, \dots, -/X_n]$. Now $\varphi_s^c = \varphi_u^{c_\alpha}$.

If $c(X_i)$ is a limit ordinal, let $c_\alpha = c[\alpha/X_i, -/X_{i+1}, \dots, -/X_n]$ for every $\alpha < c(X_i)$. Now

$$\varphi_s^c = \bigvee_{\alpha < c(X_i)} \varphi_u^{c_\alpha} \text{ if } \text{lab}(s_i) = \mu X_i \quad \text{and} \quad \varphi_s^c = \bigwedge_{\alpha < c(X_i)} \varphi_u^{c_\alpha} \text{ if } \text{lab}(s_i) = \nu X_i.$$

The formulas φ_s^c can contain infinite conjunctions and disjunctions but if all clocks are finite, then φ_s^c is an L_μ -formula. For instance if all models considered are finite, then finite clocks suffice. The following lemma formalizes the relationship of our approximant with the usual one.

Lemma 5.3. *Let $\varphi \in L_\mu$ and let s be a vertex in the syntax tree of φ with $\text{lab}(s) = \eta X$, where $\eta \in \{\mu, \nu\}$. Let s_1 be the successor of s and let $c : \text{Var} \rightarrow \kappa$ be a partial function with $X \notin \text{dom}(c)$. Let $c_\alpha = c[\alpha/X]$. Now*

$$\varphi_{s_1}^{c_\alpha} = \eta^{\alpha+1} X. \varphi_{s_1}^c(X).$$

Proof. Since $\eta^{\alpha+1} X. \varphi_{s_1}^c(X) = \varphi_{s_1}^c(\eta^\alpha X. \varphi_{s_1}^c(X))$, where the parentheses notation refers to substituting free occurrences of X with a formula, we may rewrite the claim in the form $\varphi_{s_1}^{c_\alpha} = \varphi_{s_1}^c(\eta^\alpha X. \varphi_{s_1}^c(X))$. We show this by transfinite induction on α .

- If $\alpha = 0$, then it is easy to see that $\varphi_{s_1}^{c_\alpha} = \varphi_{s_1}^c(\xi) = \varphi_{s_1}^c(\eta^0 X. \varphi_{s_1}^c)$, where $\xi = \perp$ if $\eta = \mu$ and $\xi = \top$ if $\eta = \nu$.
- Let $\alpha = \beta + 1$. By induction hypothesis $\varphi_{s_1}^{c_\beta} = \varphi_{s_1}^c(\eta^\beta X. \varphi_{s_1}^c(X)) = \eta^\alpha X. \varphi_{s_1}^c(X)$. We show by induction on the definition of $\varphi_t^{c_\alpha}$, where t is below s in the syntax tree of φ , that $\varphi_t^{c_\alpha} = \varphi_t^c(\eta^\alpha X. \varphi_{s_1}^c(X))$. We first note that as s_1 is the successor of the fixed point node s , the fixed point of X is the innermost one in $\text{dom}(c_\alpha)$.

- If $\text{lab}(t) = l \in \text{Lit}$, then $\varphi_t^{c_\alpha} = l = \varphi_t^c(\eta^\alpha X. \varphi_{s_1}^c(X))$.
- If $\text{lab}(t) = Y \in \text{Var} \setminus \text{dom}(c_\alpha)$, then $\varphi_t^{c_\alpha} = Y = \varphi_t^c(\eta^\alpha X. \varphi_{s_1}^c(X))$.
- Let $\text{lab}(t) = Y \in \text{dom}(c_\alpha) \setminus \{X\}$ and let u be the B -successor of t . Now for some c' , $\varphi_t^{c_\alpha} = \varphi_u^{c'}$. Since the fixed point of X is inside that of Y , $\varphi_u^{c'}$ contains no free occurrences of X . Thus $\varphi_u^{c'} = \varphi_t^c(\eta^\alpha X. \varphi_{s_1}^c(X))$.
- If $\text{lab}(t) = X$, then $\varphi_t^{c_\alpha} = \varphi_{s_1}^{c_\beta}$. Because ηX is the innermost approximated fixed point, no clocks need to be reset. Since $\varphi_t^c = X$ and by the induction hypothesis on α ,

$$\varphi_t^{c_\alpha} = \varphi_{s_1}^{c_\beta} = \eta^\alpha X. \varphi_{s_1}^c(X) = \varphi_t^c(\eta^\alpha X. \varphi_{s_1}^c(X)).$$

- If $\text{lab}(t) = \nabla \in \{\vee, \wedge\}$ and t_1 and t_2 are the successors of t , then by induction hypothesis, $\varphi_{t_1}^{c_\alpha} = \varphi_{t_1}^c(\eta^\alpha X. \varphi_{s_1}^c(X))$ and $\varphi_{t_2}^{c_\alpha} = \varphi_{t_2}^c(\eta^\alpha X. \varphi_{s_1}^c(X))$. Now

$$\begin{aligned} \varphi_t^{c_\alpha} &= \varphi_{t_1}^{c_\alpha} \nabla \varphi_{t_2}^{c_\alpha} = \varphi_{t_1}^c(\eta^\alpha X. \varphi_{s_1}^c(X)) \nabla \varphi_{t_2}^c(\eta^\alpha X. \varphi_{s_1}^c(X)) \\ &= (\varphi_{t_1}^c \nabla \varphi_{t_2}^c)(\eta^\alpha X. \varphi_{s_1}^c(X)) = \varphi_t^c(\eta^\alpha X. \varphi_{s_1}^c(X)). \end{aligned}$$

- If $\text{lab}(t) = \Delta \in \{\diamond, \square\}$ and t_1 is the successor of t , then by induction hypothesis we have $\varphi_{t_1}^{c_\alpha} = \varphi_{t_1}^c(\eta^\alpha X. \varphi_{s_1}^c(X))$. Thus

$$\varphi_t^{c_\alpha} = \Delta \varphi_{t_1}^{c_\alpha} = \Delta \varphi_{t_1}^c(\eta^\alpha X. \varphi_{s_1}^c(X)) = (\Delta \varphi_{t_1}^c)(\eta^\alpha X. \varphi_{s_1}^c(X)) = \varphi_t^c(\eta^\alpha X. \varphi_{s_1}^c(X)).$$

- If $\text{lab}(t) = \eta_1 Y$, where $\eta_1 \in \{\mu, \nu\}$ and $Y \in \text{Var}$, and t_1 is the successor of t , then by induction hypothesis $\varphi_{t_1}^{c_\alpha} = \varphi_{t_1}^c(\eta^\alpha X. \varphi_{s_1}^c(X))$. Thus

$$\varphi_t^{c_\alpha} = \eta_1 Y. \varphi_{t_1}^{c_\alpha} = \eta_1 Y. \varphi_{t_1}^c(\eta^\alpha X. \varphi_{s_1}^c(X)) = (\eta_1 Y. \varphi_{t_1}^c)(\eta^\alpha X. \varphi_{s_1}^c(X)) = \varphi_t^c(\eta^\alpha X. \varphi_{s_1}^c(X)).$$

- Now let α be a limit ordinal. By induction hypothesis $\varphi_{s_1}^{c_\beta} = \varphi_{s_1}^c(\eta^\beta X. \varphi_{s_1}^c(X))$ for all $\beta < \alpha$. The induction on $\varphi_t^{c_\alpha}$ is handled in the same way as in the previous case with the exception of the X -case.

- If $\text{lab}(t) = X$, then $\varphi_t^{c_\alpha} = \nabla_{\beta < \alpha} \varphi_{s_1}^{c_\beta}$, where $\nabla = \vee$ if $\eta = \mu$ and $\nabla = \wedge$ if $\eta = \nu$. By the induction hypothesis on α , $\varphi_{s_1}^{c_\beta} = \varphi_{s_1}^c(\eta^\beta X. \varphi_{s_1}^c(X))$ for all $\beta < \alpha$ so

$$\begin{aligned} \varphi_t^{c_\alpha} &= \nabla_{\beta < \alpha} \varphi_{s_1}^c(\eta^\beta X. \varphi_{s_1}^c(X)) = \nabla_{\beta < \alpha} \eta^{\beta+1} X. \varphi_{s_1}^c(X) = \nabla_{0 < \beta < \alpha} \eta^\beta X. \varphi_{s_1}^c(X) \\ &= \eta^\alpha X. \varphi_{s_1}^c(X) = \varphi_t^c(\eta^\alpha X. \varphi_{s_1}^c(X)). \end{aligned} \quad \square$$

The definition of the L_μ -game

Let Φ be a fixed finite set of propositional symbols. The formula size game for $L_\mu(\Phi)$, $\mu\text{-FS}_k^\Phi(\mathbb{A}_0, \mathbb{B}_0)$, has two players, S (Samson) and D (Delilah). The game has as parameters two sets of pointed Φ -models, \mathbb{A}_0 and \mathbb{B}_0 , and a natural number k . S wants to show that the sets \mathbb{A}_0 and \mathbb{B}_0 can be separated with a $L_\mu(\Phi)$ -formula of size at most k . D on the other hand wants to show this is not possible. During the game S constructs step by step the syntax tree of a formula that, he claims, separates the sets. The number k is a resource that is spent when S adds vertices to the syntax tree. If the resource k ever runs out, S loses the game. S has to simultaneously show how the models in \mathbb{A}_0 make the formula true and how the models in \mathbb{B}_0 make it false. Each model traverses the incomplete syntax tree in a fashion similar to semantic games. The role of D is to keep S honest by deciding which branch of the tree she wants to see next.

The modal μ -calculus has the special feature of fixed point formulas. In terms of models traversing the syntax tree of a formula, the truth of a least fixed point μX comes down to the model having to eventually stop returning to that fixed point. Thus, when entering such a fixed point, S must set a clock for each model, that shows how many more times he will return the model to that fixed point. On the other hand, to show a μX -formula is false in a model, the model would have to keep returning to the fixed point forever. Here it is the responsibility of D to declare how many returns is enough for her to be satisfied that the formula is indeed false. We now present the quite complex formalization of the game.

Let \mathbb{A}_0 and \mathbb{B}_0 be sets of pointed models and let $k_0 \in \mathbb{N}$. Let V^* be a predefined infinite set of vertices. The formula size game $\mu\text{-FS}_{k_0}^\Phi(\mathbb{A}_0, \mathbb{B}_0)$ for the modal μ -calculus has two players, S and D. The positions of the game are of the form $P = (V, E, B, \text{lab}, \text{res}, \text{left}, \text{right}, v)$. Here $V \subseteq V^*$ and (V, E, B) is a tree with back edges. The partial function

$$\text{lab} : V \rightarrow \{\wedge, \vee, \diamond, \square\} \cup \text{Var} \cup \{\mu X, \nu X \mid X \in \text{Var}\} \cup \text{Lit}(\Phi)$$

assigns a label to some vertices of the tree. The function $\text{res} : V \rightarrow \mathbb{N}$ assigns to each vertex the remaining resource, i.e. an upper bound for the size of the subformula starting from the vertex. The function $\text{left} : V \rightarrow \mathcal{P}(\mathbb{A}_0^*)$ assigns to each $v \in V$ its left set of clocked models $\text{left}(v)$. Here \mathbb{A}_0^* contains all the clocked models obtainable from models of \mathbb{A}_0 by altering the distinguished point, clocks and age identifier. Similarly, $\text{right} : V \rightarrow \mathcal{P}(\mathbb{B}_0^*)$ assigns the right set $\text{right}(v)$. The clock function of each model is of the form $c : \text{Var} \rightarrow \kappa$, where κ is a fixed cardinal larger than the size of the domain of any model in $\mathbb{A}_0 \cup \mathbb{B}_0$. Finally the vertex $v \in V$ is the *current vertex* of the position. We will always assume that the position P has components with these names and P' always consists of the same components with primes.

The starting position of the game is

$$(\{v_0\}, \emptyset, \emptyset, \emptyset, \{(v_0, k_0)\}, \{(v_0, \text{CM}(\mathbb{A}_0))\}, \{(v_0, \text{CM}(\mathbb{B}_0))\}, v_0).$$

The first move is always D choosing finite subsets $\mathbb{A} \subseteq \text{CM}(\mathbb{A}_0)$ and $\mathbb{B} \subseteq \text{CM}(\mathbb{B}_0)$. The following position is

$$(\{v_0\}, \emptyset, \emptyset, \emptyset, \{(v_0, k_0)\}, \{(v_0, \mathbb{A})\}, \{(v_0, \mathbb{B})\}, v_0).$$

Throughout the whole game, D wins if at any position P , $\text{res}(v) = 0$. D also wins if S is unable to make the choices required by a move. Assume the game is in position P and let $\text{left}(v) = \mathbb{A}$, $\text{right}(v) = \mathbb{B}$ and $\text{res}(v) = k > 0$. We define two cases by whether v already has a label or not. In each case we denote the following position by P' .

$v \notin \text{dom}(\text{lab})$: S has a choice of eight different moves. Note that in this case $\mathbb{A} = \mathbb{A}_{\text{new}}$ and $\mathbb{B} = \mathbb{B}_{\text{new}}$.

- **\vee -move:** S chooses sets $\mathbb{A}_1, \mathbb{A}_2 \subseteq \mathbb{A}$ s.t. $\mathbb{A}_1 \cup \mathbb{A}_2 = \mathbb{A}$, and numbers $0 < k_1, k_2 \leq k$ s.t. $k_1 + k_2 + 1 = k$. Then D chooses a number $i \in \{1, 2\}$. Let $V' = V \cup \{v_1, v_2\}$, $E' = E \cup \{(v, v_1), (v, v_2)\}$, $B' = B$, $\text{lab}' = \text{lab}[v/v]$, $\text{left}' = \text{left}[\mathbb{A}_1/v_1, \mathbb{A}_2/v_2, \text{old}(\mathbb{A})/v]$, $\text{right}' = \text{right}[\mathbb{B}/v_1, \mathbb{B}/v_2, \text{old}(\mathbb{B})/v]$, $\text{res}' = \text{res}[k_1/v_1, k_2/v_2]$ and $v' = v_i$, where v_1 and v_2 are new vertices.
- **\wedge -move:** Same as the \vee -move with the roles of \mathbb{A} and \mathbb{B} switched.
- **\diamond -move:** S chooses a function $f : \mathbb{A} \rightarrow \square\mathbb{A}$ such that $f(\mathfrak{A}) \in \square\mathfrak{A}$ for each $\mathfrak{A} \in \mathbb{A}$. Then D chooses finite subsets $\mathbb{A}' \subseteq \diamond_f\mathbb{A}$ and $\mathbb{B}' \subseteq \square\mathbb{B}$. Let $V' = V \cup \{v'\}$, $E' = E \cup \{(v, v')\}$, $B' = B$, $\text{lab}' = \text{lab}[\diamond/v]$, $\text{left}' = \text{left}[\mathbb{A}'/v', \text{old}(\mathbb{A})/v]$, $\text{right}' = \text{right}[\mathbb{B}'/v', \text{old}(\mathbb{B})/v]$ and $\text{res}' = \text{res}[k - 1/v']$, where v' is a new vertex.
- **\square -move:** Same as the \diamond -move with the roles of \mathbb{A} and \mathbb{B} switched.
- **μX -move:** S chooses a variable $X \in \text{Var}$ and for every $\mathfrak{A} = (\mathcal{A}_{\mathfrak{A}}, w_{\mathfrak{A}}, c_{\mathfrak{A}}) \in \mathbb{A}$ an ordinal $\alpha_{\mathfrak{A}}$. Then D chooses for every $\mathfrak{B} = (\mathcal{B}_{\mathfrak{B}}, w_{\mathfrak{B}}, c_{\mathfrak{B}}) \in \mathbb{B}$ an ordinal $\alpha_{\mathfrak{B}}$. Let $c'_{\mathfrak{A}} = c_{\mathfrak{A}}[\alpha_{\mathfrak{A}}/X]$ and let $\mathbb{A}' = \{(\mathcal{A}_{\mathfrak{A}}, w_{\mathfrak{A}}, c'_{\mathfrak{A}}) \mid \mathfrak{A} \in \mathbb{A}\}$. Let $c'_{\mathfrak{B}} = c_{\mathfrak{B}}[\alpha_{\mathfrak{B}}/X]$ and let $\mathbb{B}' = \{(\mathcal{B}_{\mathfrak{B}}, w_{\mathfrak{B}}, c'_{\mathfrak{B}}) \mid \mathfrak{B} \in \mathbb{B}\}$. Let $V' = V \cup \{v'\}$, $E' = E \cup \{(v, v')\}$, $B' = B$, $\text{lab}' = \text{lab}[\mu X/v]$, $\text{left}' = \text{left}[\mathbb{A}'/v', \text{old}(\mathbb{A})/v]$, $\text{right}' = \text{right}[\mathbb{B}'/v', \text{old}(\mathbb{B})/v]$ and $\text{res}' = \text{res}[k - 1/v']$.
- **νX -move:** Same as the μ -move with the roles of \mathbb{A} and \mathbb{B} switched.
- **X -move:** S chooses $X \in \text{Var}$. Let $u \in V$ be the closest vertex above v with $\text{lab}(u) \in \{\mu X, \nu X\}$. If no such vertex exists, D wins the game. Otherwise if $\mathbb{A}_{\text{new}} = \mathbb{B}_{\text{new}} = \emptyset$, S wins the game.
Let v' be the successor vertex of u . Let $V' = V$, $E' = E$, $B' = B \cup \{(v, v')\}$, $\text{lab}' = \text{lab}[X/v]$ and $\text{res}' = \text{res}$.
Assume that $\text{lab}(u) = \mu X$. If $c(X) = 0$ for some $(\mathcal{A}, w, c) \in \mathbb{A}$, D wins the game. Otherwise for every $\mathfrak{A} = (\mathcal{A}_{\mathfrak{A}}, w_{\mathfrak{A}}, c_{\mathfrak{A}}, \text{new}) \in \mathbb{A}$, S chooses $\alpha_{\mathfrak{A}} < c_{\mathfrak{A}}(X)$.
Let $\mathbb{B}_+ = \{(\mathcal{B}, w, c, \text{new}) \in \mathbb{B} \mid c(X) \neq 0\}$. For every $\mathfrak{B} = (\mathcal{B}_{\mathfrak{B}}, w_{\mathfrak{B}}, c_{\mathfrak{B}}, \text{new}) \in \mathbb{B}_+$, D chooses $\alpha_{\mathfrak{B}} < c_{\mathfrak{B}}(X)$.
Let Y_1, \dots, Y_n be the variables for which there is a node t_i on the path from u to v with $\text{lab}(t_i) \in \{\mu Y_i, \nu Y_i\}$. Let $c'_{\mathfrak{A}} = c_{\mathfrak{A}}[\alpha_{\mathfrak{A}}/X, -/Y_1, \dots, -/Y_n]$ and let $\mathbb{A}' = \{(\mathcal{A}_{\mathfrak{A}}, w_{\mathfrak{A}}, c'_{\mathfrak{A}}, \text{new}) \mid \mathfrak{A} \in \mathbb{A}\}$. Similarly let $c'_{\mathfrak{B}} = c_{\mathfrak{B}}[\alpha_{\mathfrak{B}}/X, -/Y_1, \dots, -/Y_n]$ and $\mathbb{B}' = \{(\mathcal{B}_{\mathfrak{B}}, w_{\mathfrak{B}}, c'_{\mathfrak{B}}, \text{new}) \mid \mathfrak{B} \in \mathbb{B}_+\}$. Let $\text{left}' = \text{left}[\mathbb{A}' \cup \text{left}(v')/v', \text{old}(\mathbb{A})/v]$ and $\text{right}' = \text{right}[\mathbb{B}' \cup \text{left}(v')/v', \text{old}(\mathbb{B})/v]$.
The case $\text{lab}(u) = \nu X$ is the same with the roles of \mathbb{A} and \mathbb{B} switched.
- ***Lit*-move:** S chooses a Φ -literal l . Let $\text{lab}' = \text{lab}[l/v]$ and let $\Delta' = \Delta$ for every other component Δ . In the following position P' , if l separates \mathbb{A} and \mathbb{B} , then S wins the game. Otherwise, D wins.

$v \in \text{dom}(\text{lab})$: In this case S must perform the move dictated by $\text{lab}(v)$ without creating any new vertices. These moves are essentially performed only on new models. We again denote the following position by P' and in each case we have $\nabla' = \nabla$ for $\nabla \in \{V, E, B, \text{lab}, \text{res}\}$.

- If $\text{lab}(v) = \vee$, then let v_1 and v_2 be the successors of v . S chooses sets $\mathbb{A}_1, \mathbb{A}_2 \subseteq \mathbb{A}_{\text{new}}$ s.t. $\mathbb{A}_1 \cup \mathbb{A}_2 = \mathbb{A}_{\text{new}}$. Then D chooses a number $i \in \{1, 2\}$. Let $\text{left}' = \text{left}[\mathbb{A}_1 \cup \text{left}(v_1)/v_1, \mathbb{A}_2 \cup \text{left}(v_2)/v_2, \text{old}(\mathbb{A})/v]$, $\text{right}' = \text{right}[\mathbb{B} \cup \text{right}(v_1)/v_1, \mathbb{B} \cup \text{right}(v_2)/v_2, \text{old}(\mathbb{B})/v]$ and $v' = v_i$.
- The case $\text{lab}(v) = \wedge$ is the same as \vee with the roles of \mathbb{A} and \mathbb{B} switched.

- If $\text{lab}(v) = \diamond$, then let v' be the successor of v . S chooses a function $f : \mathbb{A}_{\text{new}} \rightarrow \square \mathbb{A}_{\text{new}}$ such that $f(\mathfrak{A}) \in \square \mathfrak{A}$ for each $\mathfrak{A} \in \mathbb{A}_{\text{new}}$. Then D chooses finite $\mathbb{A}' \subseteq \diamond_f \mathbb{A}_{\text{new}}$ and $\mathbb{B}' \subseteq \square \mathbb{B}_{\text{new}}$. Let $\text{left}' = \text{left}[\mathbb{A}' \cup \text{left}(v')/v_1, \text{old}(\mathbb{A})/v]$ and $\text{right}' = \text{right}[\mathbb{B}' \cup \text{right}(v')/v_1, \text{old}(\mathbb{B})/v]$.
- The case $\text{lab}(v) = \square$ is the same as \diamond with the roles of \mathbb{A} and \mathbb{B} switched.
- If $\text{lab}(v) = \mu X$ for some $X \in \text{Var}$, then let v' be the successor of v . for every $\mathfrak{A} = (\mathcal{A}_{\mathfrak{A}}, w_{\mathfrak{A}}, c_{\mathfrak{A}}) \in \mathbb{A}_{\text{new}}$ an ordinal $\alpha_{\mathfrak{A}}$. Then D chooses for every $\mathfrak{B} = (\mathcal{B}_{\mathfrak{B}}, w_{\mathfrak{B}}, c_{\mathfrak{B}}) \in \mathbb{B}_{\text{new}}$ an ordinal $\alpha_{\mathfrak{B}}$. Let $c'_{\mathfrak{A}} = c_{\mathfrak{A}}[\alpha_{\mathfrak{A}}/X]$ and let $\mathbb{A}' = \{(\mathcal{A}_{\mathfrak{A}}, w_{\mathfrak{A}}, c'_{\mathfrak{A}}) \mid \mathfrak{A} \in \mathbb{A}_{\text{new}}\}$. Let $c'_{\mathfrak{B}} = c_{\mathfrak{B}}[\alpha_{\mathfrak{B}}/X]$ and let $\mathbb{B}' = \{(\mathcal{B}_{\mathfrak{B}}, w_{\mathfrak{B}}, c'_{\mathfrak{B}}) \mid \mathfrak{B} \in \mathbb{B}_{\text{new}}\}$. Let $\text{left}' = \text{left}[\mathbb{A}' \cup \text{left}(v')/v', \text{old}(\mathbb{A})/v]$ and $\text{right}' = \text{right}[\mathbb{B}' \cup \text{right}(v')/v', \text{old}(\mathbb{B})/v]$.
- If $\text{lab}(v) = X \in \text{Var}$, then let v' be the B -successor of v . The rest is very similar to the unlabelled X -move; the only differences are that the move is again essentially only performed on new models and the condition for an immediate win for S is $\mathbb{A}_{\text{new}} = \mathbb{B}_{\text{new}} = \emptyset$.

Note that just like in the ML-game, the \diamond -move cannot be performed if $\square \mathfrak{A} = \emptyset$ for some $\mathfrak{A} \in \mathbb{A}$, and dually for the \square -move.

Figures 4 and 5 show how the partial syntax tree (V, E, B, lab) changes when a \wedge -move or an X -move is made. The current vertex v is highlighted. The numbers near the vertices show the remaining resource given by res . In addition each vertex has a left and a right set of clocked models given by left and right . Of these only the new models affected by the move are shown in the pictures.

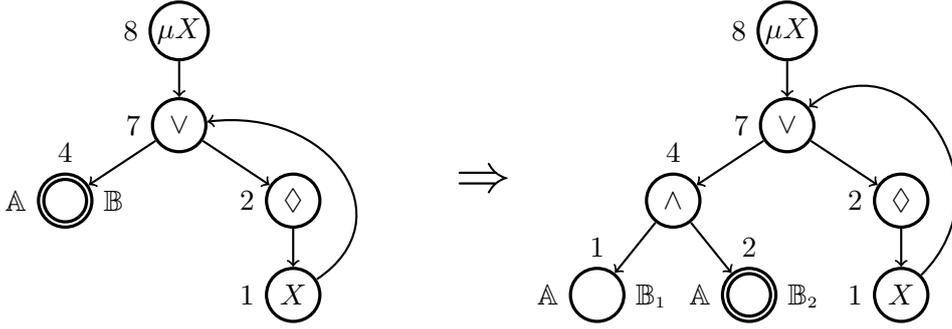


Figure 4

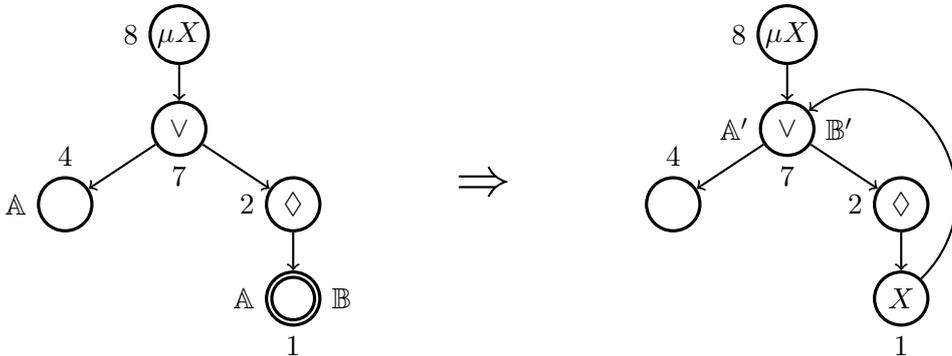


Figure 5

An important feature of our game is that, even though it contains infinite branching, every single play of the game is still finite.

Lemma 5.4. *Every play of the game $\mu\text{-FS}_k^\Phi(\mathbb{A}, \mathbb{B})$ is finite.*

Proof. A play could be infinite only if at least one variable is reached infinitely many times. Of these variables let X be the one with the outmost fixed point. Every time X is reached, if $\text{left}(v)_{\text{new}} = \text{right}(v)_{\text{new}} = \emptyset$, S wins and otherwise, the clock of at least one model is lowered. There are only finitely many models at any given position, since D always chooses finite subsets of models after modal moves. Since clocks are inherited by successor models in modal moves and ordinals are well-founded, eventually either a clock of S will reach 0 and D will win or X will be reached with empty sets of models and S wins. \square

For the essential theorem about how the game works, we assume that the strategy of S is uniform. This essentially means that S has a formula in mind and he follows the structure of that formula when constructing the syntax tree during the game.

Definition 5.5. Let $\varphi \in L_\mu$ and let $T_\varphi = (V_\varphi, E_\varphi, B_\varphi, \text{lab}_\varphi)$ be the syntax tree with back edges of φ . Let $P = (V, E, B, \text{lab}, \text{res}, \text{left}, \text{right}, v)$ be a position in a game $\mu\text{-FS}_k^\Phi(\mathbb{A}, \mathbb{B})$.

A function $g : V \rightarrow V_\varphi$ is a *position embedding* if it satisfies the following conditions:

1. $g(v_0)$ is the root of T_φ , where v_0 is the vertex of the starting position,
2. g is an embedding of (V, E) to (V_φ, E_φ) ,
3. $g \upharpoonright \text{dom}(\text{lab})$ is an embedding of (V, B, lab) to $(V_\varphi, B_\varphi, \text{lab}_\varphi)$,
4. for each $u \in V$, $\text{sz}(\varphi_{g(u)}) \leq \text{res}(u)$.

Let δ be a strategy of S. We say that δ *follows φ from position P (via the function g)* if there is a position embedding $g : V \rightarrow V_\varphi$ such that for each position P' reachable from P via the strategy δ , the function g can be extended to a position embedding $g' : V' \rightarrow V_\varphi$.

Finally δ is *uniform* if δ follows a formula $\varphi \in L_\mu$ from the starting position.

We are now ready to prove that the game indeed works as we intended. In the following if there is a vertex with label μX we shall call X a μ -*variable*, and if there is one labelled νX , we call X a ν -*variable*. Note that we assume all fixed points have separate variables.

Theorem 5.6. *Let \mathbb{A}_0 and \mathbb{B}_0 be sets of pointed models and let $k \in \mathbb{N}$. Then the following conditions are equivalent:*

1. *S has a uniform winning strategy in the game $\mu\text{-FS}_k^\Phi(\mathbb{A}_0, \mathbb{B}_0)$.*
2. *There is a sentence $\varphi \in L_\mu(\Phi)$ s.t. φ separates \mathbb{A}_0 and \mathbb{B}_0 and $\text{sz}(\varphi) \leq k$.*

Proof. (2) \Rightarrow (1). Let $\varphi \in L_\mu(\Phi)$ be a sentence such that φ separates \mathbb{A}_0 and \mathbb{B}_0 and $s(\varphi) \leq k$. Let $T_\varphi = (V_\varphi, E_\varphi, B_\varphi, \text{lab}_\varphi)$ be the syntax tree with back edges of φ . The strategy of S is to follow the structure of T_φ when forming (V, E, B, lab) , to use the resource k accordingly and to choose maximal appropriate sets of models when necessary.

If $P = (V, E, B, \text{lab}, \text{res}, \text{left}, \text{right}, v)$ is a position, let $\mathbb{A} = \text{left}(v)$, $\mathbb{B} = \text{right}(v)$ and $k = \text{res}(v)$. We define the strategy more precisely and simultaneously prove by induction that the strategy is uniform and the following condition holds for every position P of the game:

$$\begin{aligned} (\mathcal{A}, w) &\models \varphi_{g(v)}^c \text{ for every } (\mathcal{A}, w, c) \in \mathbb{A}_{\text{new}} \text{ and} \\ (\mathcal{B}, w) &\not\models \varphi_{g(v)}^c \text{ for every } (\mathcal{B}, w, c) \in \mathbb{B}_{\text{new}}, \end{aligned} \tag{*}$$

where g is a position embedding showing the uniformity of the strategy.

In the starting position, we set $g(v_0)$ as the root of T_φ . We note that $\text{sz}(\varphi_{g(v_0)}) = \text{sz}(\varphi) \leq k$ by assumption. Since there are no clocks yet, $\varphi_{g(v_0)}^c = \varphi$ for every clocked model and since the sentence φ separates the sets \mathbb{A}_0 and \mathbb{B}_0 , (*) holds no matter which subsets $\mathbb{A} \subseteq \text{CM}(\mathbb{A}_0)$ and $\mathbb{B} \subseteq \text{CM}(\mathbb{B}_0)$ D chooses.

We now divide the proof into cases based on whether v already has a label or not. We choose the move for S according to the label of $g(v)$. We only treat one of each pair of dual cases.

$v \notin \text{dom}(\text{lab}) :$

- $\text{lab}_\varphi(g(v)) = l \in \text{Lit}(\Phi)$: Then by induction hypothesis l separates the sets \mathbb{A} and \mathbb{B} so S wins by making the corresponding *Lit*-move.
- $\text{lab}_\varphi(g(v)) = \vee$: By induction hypothesis, $(*)$ holds for this position so for every $(A, w, c) \in \mathbb{A}$, $(A, w) \models \varphi_{g(v)}^c$. Now $\varphi_{g(v)}^c = \varphi_{s_1}^c \vee \varphi_{s_2}^c$, where $s_1, s_2 \in V_\varphi$ are the successors of $g(v)$, so $(A, w) \models \varphi_{s_1}^c \vee \varphi_{s_2}^c$. Let $\mathbb{A}_1 = \{(A, w, c) \in \mathbb{A} \mid (A, w) \models \varphi_{s_1}^c\}$ and $\mathbb{A}_2 = \{(A, w, c) \in \mathbb{A} \mid (A, w) \models \varphi_{s_2}^c\}$. On the other side, for every $(B, w, c) \in \mathbb{B}$, $(B, w) \not\models \varphi_{s_1}^c \vee \varphi_{s_2}^c$ so $(B, w) \not\models \varphi_{s_1}^c$ and $(B, w) \not\models \varphi_{s_2}^c$. We set $g' = g[v_1/s_1, v_2/s_2]$, where v_1, v_2 are the new vertices in V . Now $(*)$ holds in both of the possible following positions. Let $k_1 = \text{sz}(\varphi_{s_1})$ and $k_2 = k - k_1 - 1$. Since $\text{sz}(\varphi_{g(v)}) \leq k$, $\text{sz}(\varphi_{s_2}) = \text{sz}(\varphi_{g(v)}) - \text{sz}(\varphi_{s_1}) - 1 \leq k - k_1 - 1 = k_2$.
- $\text{lab}_\varphi(g(v)) = \diamond$: By induction hypothesis, for every $(A, w, c) \in \mathbb{A}$, $(A, w) \models \varphi_{g(v)}^c$. Since $\varphi_{g(v)}^c = \diamond \varphi_{s_1}^c$, where s_1 is the successor of $g(v)$, $(A, w) \models \diamond \varphi_{s_1}^c$. Thus there is $(A, w', c) \in \square \mathbb{A}$ s.t. $(A, w') \models \varphi_{s_1}^c$. Let $f : \mathbb{A} \rightarrow \square \mathbb{A}$ be a function mapping every (A, w, c) to such a (A, w', c) . Now $(A, w') \models \varphi_{s_1}^c$ for every $(A, w', c) \in \diamond_f \mathbb{A}$. On the other side, for every $(B, w, c) \in \mathbb{B}$, since $(B, w) \not\models \diamond \varphi_{s_1}^c$, for every $(B, w', c) \in \square(B, w, c)$ we get $(B, w') \not\models \varphi_{s_1}^c$. Thus $(B, w') \not\models \varphi_{s_1}^c$ for every $(B, w', c) \in \square \mathbb{B}$ so $(*)$ holds in the next position no matter which subsets $\mathbb{A}' \subseteq \diamond_f \mathbb{A}$ and $\mathbb{B}' \subseteq \square \mathbb{B}$ D chooses. For uniformity we set $g' = g[v'/s_1]$, where v' is the new vertice. Now $\text{sz}(\varphi_{s_1}) = \text{sz}(\varphi_{g(v)}) - 1 \leq k - 1$.
- Let $\text{lab}_\varphi(g(v)) = \mu X$: By induction hypothesis, for every $\mathfrak{A} = (A, w, c) \in \mathbb{A}$, $(A, w) \models \varphi_{g(v)}^c$. Since $\varphi_{g(v)}^c = \mu X.\varphi_{s_1}^c(X)$, where s_1 is the successor of $g(v)$, $(A, w) \models \mu X.\varphi_{s_1}^c(X)$. Thus there is an ordinal α such that $(A, w) \models \mu^{\alpha+1} X.\varphi_{s_1}^c(X)$. S chooses $\alpha_{\mathfrak{A}} = \alpha$ as the new clock. By Lemma 5.3, $(A, w) \models \varphi_{s_1}^{c_\alpha}$. On the other side, by induction hypothesis, for every $\mathfrak{B} = (B, w, c) \in \mathbb{B}$, $(B, w) \not\models \mu X.\varphi_{s_1}^c(X)$ so for every ordinal α , $(B, w) \not\models \mu^{\alpha+1} X.\varphi_{s_1}^c(X)$. Thus no matter which ordinal β D chooses, we get $(B, w) \not\models \mu^{\beta+1} X.\varphi_{s_1}^c(X)$ and by Lemma 5.3, $(B, w) \not\models \varphi_{s_1}^{c_\beta}$. Therefore $(*)$ holds in the following position. Uniformity is proved in the same way as in the \diamond -case.
- $\text{lab}_\varphi(g(v)) = X$, where $X \in \text{Var}$: Let u be the B -successor of $g(v)$.
Assume that X is a μ -variable. Now by induction hypothesis, for every $\mathfrak{A} = (A, w, c) \in \mathbb{A}$, $(A, w) \models \varphi_{g(v)}^c$. There are three cases according to $c(X)$.

1. If $c(X) = 0$, we get a contradiction since then $\varphi_{g(v)}^c = \perp$.
2. If $c(X) = \alpha + 1$ for some α , then S chooses α as the new clock for X in (A, w, c) . Now $\varphi_{g(v)}^c = \varphi_u^{c_\alpha}$ so $(A, w) \models \varphi_u^{c_\alpha}$.
3. If $c(X)$ is a limit ordinal, then

$$\varphi_{g(v)}^c = \bigvee_{\alpha < c(X)} \varphi_u^{c_\alpha} \quad \text{so} \quad (A, w) \models \bigvee_{\alpha < c(X)} \varphi_u^{c_\alpha}.$$

S chooses the new clock α such that $(A, w) \models \varphi_u^{c_\alpha}$ holds.

On the other side, by induction hypothesis, for every $\mathfrak{B} = (B, w, c) \in \mathbb{B}$, $(B, w) \not\models \varphi_{g(v)}^c$. We again have three cases.

1. If $c(X) = 0$, \mathfrak{B} will be removed from the game and can be disregarded.
2. Let $c(X) = \alpha + 1$. Now $\varphi_{g(v)}^c = \varphi_u^{c_\alpha}$ so $(B, w) \not\models \varphi_u^{c_\alpha}$. By Lemma 5.3, $(B, w) \not\models \mu^{\alpha+1} X.\varphi_u^{c'}(X)$, where $c' = c_\alpha[-/X]$. Let $\beta \leq \alpha$ be the choice of D for the new clock. Now by monotonicity, $(B, w) \not\models \mu^{\beta+1} X.\varphi_u^{c'}(X)$. We use Lemma 5.3 again and obtain $(B, w) \not\models \varphi_u^{c_\beta}$.

3. Finally let $c(X)$ be a limit ordinal. Now

$$\varphi_{g(v)}^c = \bigvee_{\alpha < c(X)} \varphi_u^{c_\alpha} \quad \text{so} \quad (\mathcal{B}, w) \not\models \bigvee_{\alpha < c(X)} \varphi_u^{c_\alpha}.$$

Thus $(\mathcal{B}, w) \not\models \varphi_u^{c_\alpha}$ for any $\alpha < c(X)$ D chooses.

Uniformity is trivial here since no new vertices were created.

The case where X is a ν -variable is the same with the roles of \mathbb{A} and \mathbb{B} switched.

$v \in \text{dom}(\text{lab})$: The moves are essentially the same as in the unlabelled case. The main differences are that the type of the move is already determined by $\text{lab}(v)$, and the resource splittings are already fixed. In disjunction and conjunction moves new models can be left to wait in the branch not chosen by D as the following position. We will consider only this special case of waiting new models here.

$\text{lab}(v) = \vee$: S chooses the sets \mathbb{A}_1 and \mathbb{A}_2 of new models as in the unlabelled case. There may however be some new models present in v_1 or v_2 . If so, these models are there because of previous \vee -moves, for the first of which v had no label. By induction hypothesis and the unlabelled case, (*) held for both of the possible following positions and therefore $(\mathcal{A}, w) \models \varphi_{s_i}^c$ for every model (\mathcal{A}, w, c) in the corresponding left model set. Inductively we see that $(\mathcal{A}, w) \models \varphi_{s_i}^c$ for every $(\mathcal{A}, w, c) \in \text{left}(v_i)$. The same argument shows that $(\mathcal{B}, w) \not\models \varphi_{s_i}^c$ for every $(\mathcal{B}, w, c) \in \text{right}(v_i)$. Thus (*) holds for the sets $\mathbb{A}_i \cup \text{left}(v_i)$ and $\mathbb{B} \cup \text{right}(v_i)$ in both of the possible following positions of position P .

(1) \Rightarrow (2). Let δ be a uniform winning strategy for S. Let $\varphi \in L_\mu(\Phi)$ be the formula δ follows. We denote the position embedding showing the uniformity of δ in each position by g . By Lemma 5.4 every play of the game is finite so the game tree induced by the strategy δ is well-founded. We prove by well-founded induction on the game positions reachable with δ that the same condition (*) as above holds in every position of the game.

$$\begin{aligned} (\mathcal{A}, w) \models \varphi_{g(v)}^c & \text{ for every } (\mathcal{A}, w, c) \in \mathbb{A}_{\text{new}} \text{ and} \\ (\mathcal{B}, w) \not\models \varphi_{g(v)}^c & \text{ for every } (\mathcal{B}, w, c) \in \mathbb{B}_{\text{new}}, \end{aligned} \quad (*)$$

In a position P reachable with δ , let $\text{left}(v) = \mathbb{A}$, $\text{right}(v) = \mathbb{B}$ and $\text{res}(v) = k$. We again consider the unlabelled and labelled case separately and only treat one of each pair of dual moves.

$v \notin \text{dom}(\text{lab})$:

- $\text{lab}(g(v)) \in \text{Lit}(\Phi)$: Since δ follows φ , the next move according to δ is a *Lit* move choosing that literal. Since δ is a winning strategy, that literal separates the sets \mathbb{A} and \mathbb{B} and (*) holds.
- $\text{lab}(g(v)) = \vee$: Let s_1 and s_2 be the successors of $g(v)$. Let $\mathbb{A}_1, \mathbb{A}_2, k_1$ and k_2 be the selections of S according to δ . By induction hypothesis (*) holds in both possible following positions so for every $(\mathcal{A}, w, c) \in \mathbb{A}_i$, $(\mathcal{A}, w) \models \varphi_{s_i}^c$ for $i \in \{1, 2\}$. Since $\mathbb{A} = \mathbb{A}_1 \cup \mathbb{A}_2$, for every $(\mathcal{A}, w, c) \in \mathbb{A}$, $(\mathcal{A}, w) \models \varphi_{s_1}^c \vee \varphi_{s_2}^c$. In addition $\varphi_{s_1}^c \vee \varphi_{s_2}^c = \varphi_{g(v)}^c$ so $(\mathcal{A}, w) \models \varphi_{g(v)}^c$. Let $(\mathcal{B}, w, c) \in \mathbb{B}$. By (*) in the following positions, $(\mathcal{B}, w) \not\models \varphi_{s_1}^c$ and $(\mathcal{B}, w) \not\models \varphi_{s_2}^c$ so $(\mathcal{B}, w) \not\models \varphi_{s_1}^c \vee \varphi_{s_2}^c$. Since $\varphi_{s_1}^c \vee \varphi_{s_2}^c = \varphi_{g(v)}^c$, $(\mathcal{B}, w) \not\models \varphi_{g(v)}^c$. Thus (*) holds in P .
- $\text{lab}(g(v)) = \diamond$: Let s_1 be the successor of $g(v)$. Let $f : \mathbb{A} \rightarrow \square\mathbb{A}$ be the function chosen by S according to δ . Let $(\mathcal{A}, w, c) \in \mathbb{A}$. By induction hypothesis (*) holds in the following position no matter which subsets of $\diamond_f\mathbb{A}$ and $\square\mathbb{B}$ D chooses so for $f(\mathcal{A}, w, c) = (\mathcal{A}, w', c) \in \diamond_f\mathbb{A}$, $(\mathcal{A}, w') \models \varphi_{s_1}^c$. Since w' is a successor of w , now $(\mathcal{A}, w) \models \diamond\varphi_{s_1}^c$. In addition $\diamond\varphi_{s_1}^c = \varphi_{g(v)}^c$.

so $(\mathcal{A}, w) \models \varphi_{g(v)}^c$. Let $(\mathcal{B}, w, c) \in \mathbb{B}$. By induction hypothesis $(*)$ holds for all possible following positions so for every $(\mathcal{B}, w', c) \in \square\mathbb{B}$, $(\mathcal{B}, w') \not\models \varphi_{s_1}^c$. Therefore $(\mathcal{B}, w) \not\models \diamond\varphi_{s_1}^c$ and so $(\mathcal{B}, w) \not\models \varphi_{g(v)}^c$. Thus $(*)$ holds in P .

- $\text{lab}(g(v)) = \mu X$: Let s_1 be the successor of $g(v)$. Let $\mathfrak{A} = (\mathcal{A}, w, c) \in \mathbb{A}$ and let $\alpha_{\mathfrak{A}} = \alpha$ be the choice of S according to δ . By induction hypothesis, $(*)$ holds in the following position so $(\mathcal{A}, w) \models \varphi_{s_1}^{c_\alpha}$. By Lemma 5.3, $(\mathcal{A}, w) \models \mu^{\alpha+1}X.\varphi_{s_1}^c$. Thus $(\mathcal{A}, w) \models \mu X.\varphi_{s_1}^c$. Since $\mu X.\varphi_{s_1}^c = \varphi_{g(v)}^c$, $(\mathcal{A}, w) \models \varphi_{g(v)}^c$. On the other side, by induction hypothesis, for every $\mathfrak{B} = (\mathcal{B}, w, c) \in \mathbb{B}$ and for any choice α of D for the new clock, $(\mathcal{B}, w) \not\models \varphi_{s_1}^{c_\alpha}$. Thus by Lemma 5.3, $(\mathcal{B}, w) \not\models \mu^{\alpha+1}X.\varphi_{s_1}^c(X)$ for every $\alpha < \kappa$. Since $\kappa > \text{card}(\mathcal{B})$, this means that $(\mathcal{B}, w) \not\models \mu X.\varphi_{s_1}^c$ and so $(\mathcal{B}, w) \not\models \varphi_{g(v)}^c$. Thus $(*)$ holds in P .
- $\text{lab}(g(v)) = X$: If $\mathbb{A}_{\text{new}} = \mathbb{B}_{\text{new}} = \emptyset$, S wins and $(*)$ trivially holds. Let u be the B -successor of $g(v)$. Assume that X is a μ -variable and let $(\mathcal{A}, w, c) \in \mathbb{A}$. We have three cases according to the ordinal $c(X)$.

1. If $c(X) = 0$, D wins the game, which is a contradiction, since δ is a winning strategy for S.
2. Assume that $c(X) = \alpha + 1$. Let $\beta \leq \alpha$ be the choice of S for the new clock according to δ . By induction hypothesis $(*)$ holds in the following position so $(\mathcal{A}, w) \models \varphi_u^{c_\beta}$. By Lemma 5.3, $(\mathcal{A}, w) \models \mu^{\beta+1}X.\varphi_u^{c'}$, where $c' = c_\beta[-/X]$. Thus by monotonicity, $(\mathcal{A}, w) \models \mu^{\alpha+1}X.\varphi_u^{c'}$. By Lemma 5.3 again, $(\mathcal{A}, w) \models \varphi_u^{c_\alpha}$ and so $(\mathcal{A}, w) \models \varphi_{g(v)}^c$.
3. Now assume $c(X)$ is a limit ordinal and let $\alpha < c(X)$ be the choice of S according to δ . Now by induction hypothesis $(\mathcal{A}, w) \models \varphi_u^{c_\alpha}$. Thus

$$(\mathcal{A}, w) \models \bigvee_{\alpha < c(X)} \varphi_u^{c_\alpha}$$

$$\text{so } (\mathcal{A}, w) \models \varphi_{g(v)}^c.$$

For every $(\mathcal{B}, w, c) \in \mathbb{B}$, regardless of the choice of D for the new clock, $(*)$ holds in the following position. We again have three cases.

1. If $c(X) = 0$, then since $(\mathcal{B}, w) \not\models \perp$, we get $(\mathcal{B}, w) \not\models \varphi_{g(v)}^c$.
2. If $c(X) = \alpha + 1$, then α is a choice available to D so $(\mathcal{B}, w) \not\models \varphi_u^{c_\alpha}$. Thus $(\mathcal{B}, w) \not\models \varphi_{g(v)}^c$.
3. If $c(X)$ is a limit ordinal, every $\alpha < c(X)$ is a choice available to D so $(\mathcal{B}, w) \not\models \varphi_u^{c_\alpha}$ for every $\alpha < c(X)$. Thus

$$(\mathcal{B}, w) \not\models \bigvee_{\alpha < c(X)} \varphi_u^{c_\alpha}.$$

Therefore $(\mathcal{B}, w) \not\models \varphi_{g(v)}^c$ so $(*)$ holds in P .

$v \in \text{dom}(\text{lab})$: All the moves in this case are proved the same way as in the unlabelled case. Note that since $(*)$ refers only to new models, it trivially holds for terminal positions where $\mathbb{A}_{\text{new}} = \mathbb{B}_{\text{new}} = \emptyset$ for an X -move.

For the very first move of the game, where D chooses finite subsets of the original sets of clocked models $\text{CM}(\mathbb{A}_0)$ and $\text{CM}(\mathbb{B}_0)$, by induction hypothesis $(*)$ holds in the following position no matter which subsets D chooses. Therefore all models in \mathbb{A}_0 and \mathbb{B}_0 also satisfy the condition $(*)$. Since there are no clocks in the starting position, this means that φ separates the sets \mathbb{A}_0 and \mathbb{B}_0 . By the uniformity of δ , $\text{sz}(\varphi) = \text{sz}(\varphi_{g(v_0)}) \leq \text{res}(v_0) = k$. \square

Note that condition (*) does not depend on old models and so we do not refer to them in this proof. We add old models to the game to make the proof of Lemma 6.3 in the next section easier.

Unlike other similar theorems, Theorem 5.6 has the added requirement of uniformity for the strategy of S. We conjecture that the theorem would still hold even without this condition, but proving this has turned out to be difficult. Note, however, that to prove undefinability results for $L_\mu(\Phi)$, one need only define a winning strategy for D, and so the uniformity of strategies for S need not be considered. Note further that if a property of Φ -models is not definable in $L_\mu(\Phi)$, then clearly it is not definable in L_μ .

6 Succinctness of FO over L_μ

We move on to the definitions and lemmas needed to show that FO is non-elementarily more succinct than L_μ . We need a lemma similar to Lemma 3.4 that gives D a winning strategy if bisimilar models are produced on both sides of a vertex. In the case of the L_μ -game, the clocks of the clocked models must also be taken into account. We define a sufficient condition for clocked models to be useful for D in the game and call them relevant models.

Definition 6.1. The *depth* of a pointed finite tree model (\mathcal{M}, w) , $d(\mathcal{M}, w)$, is the length of a maximal path of transitions in the model starting from w .

Definition 6.2. In a position P of a game $\mu\text{-FS}_k^\Phi(\mathbb{A}, \mathbb{B})$, let u be a vertex and let $\mathfrak{M} = (\mathcal{M}, w, c, a)$ be a clocked finite tree model in $\text{left}(u) \cup \text{right}(u)$. We say the model \mathfrak{M} is *relevant*, if for every $X \in \text{dom}(c)$, the clock of D is equal to or greater than the depth of the model, i.e.

- if $\mathfrak{M} \in \text{left}(u)$, for every ν -variable $X \in \text{dom}(c)$, $c(X) \geq d(\mathcal{M}, w)$,
- if $\mathfrak{M} \in \text{right}(u)$, for every μ -variable $X \in \text{dom}(c)$, $c(X) \geq d(\mathcal{M}, w)$.

We also say the model \mathfrak{M} is *strictly relevant* if the above condition holds for strict inequality $>$ instead of \geq .

We prove the analogue of Lemma 3.4 for relevant clocked models.

Lemma 6.3. *Let P be a position of a game $\mu\text{-FS}_k^\Phi(\mathbb{A}, \mathbb{B})$. If there are strictly relevant clocked models $(\mathcal{A}, w_{\mathcal{A}}, c_{\mathcal{A}}, \text{new}) \in \text{left}(v)$ and $(\mathcal{B}, w_{\mathcal{B}}, c_{\mathcal{B}}, \text{new}) \in \text{right}(v)$ such that $(\mathcal{A}, w_{\mathcal{A}})$ and $(\mathcal{B}, w_{\mathcal{B}})$ are bisimilar finite tree models, then D has a winning strategy from position P .*

Proof. We show that D can maintain a slightly modified condition where we only require the models to be relevant and allow one of the two models to be old. For \forall - and \wedge -moves D need only choose the side for which the two models are both present. For modal moves we see by bisimilarity that no matter which successor S chooses, a bisimilar model will be present on the opposite side. Moreover, the depth of the models is decreased and the clocks are inherited so the condition is maintained. For new fixed points D need only set her clock to be equal to the depth of the models. On a literal move, D will win since the bisimilar models cannot be separated by a literal.

For X -moves, let u be the vertex the game returns to. If this is the first time since position P the game returns to u , D will lower the clock and since the models are strictly relevant, D can now decrease the clock to the same value as the depth. If on the other hand there has been a previous return to u , then there are two cases. If there has been a modal move in between this and the previous return, then the depth has decreased and D will decrease the clock to the same value. If there have been no modal moves, the pointed models have not changed and since we allow one of the models to be old, D will now consider the old version of her model, with a clock larger by one, instead of the new one. Consider the position P' right after D switches a new model for an old one in this fashion. Assume by symmetry that this model is \mathfrak{B} on the right side.

Consider the path from the vertex just returned to, u , to the vertex s where the X -move was made. If there are no vertices with label \vee on this path, then D can just follow this path to s until the clock $c_{\mathcal{A}}(X)$ runs out and D wins. Assume there are some vertices with label \vee on the path. For each of those vertices, the child that is not on the path from u to s has a new version of \mathfrak{B} in the right set. This is because the model \mathfrak{B} has passed through the disjunction before the X -move and models on the right side are always copied on both sides of a disjunction. If S splits the left model \mathfrak{A} away from the path to s , then D will consider the new copy of \mathfrak{B} from then on. If \mathfrak{A} stays on the path from u to s indefinitely, D wins when the clock $c_{\mathcal{A}}(X)$ runs out.

If D uses this strategy, eventually S will either make a literal move and lose, or a clock of S for one of the bisimilar models will eventually run out. In either case, D wins. \square

We want to use the same graph based invariant for the proof as we did for the ML-case. The only question that remains is: which models should determine the graph of the current vertex v ? In other words, which models is S claiming to be able to separate in each position of the game? Certainly the models in $\text{left}(v)$ and $\text{right}(v)$ should be included, but they will not suffice, since other models can already be below v and thus involved with the subformula beginning from v . We define a way for D to collect all the models in the tree below a vertex s to see which models are, in a sense, "currently in s ". We define the collected sets separately for the left and the right side. Recall that $\text{PM}(\mathbb{A})$ is the set of underlying pointed models of clocked models in \mathbb{A} .

Definition 6.4. Let P be a position such that $\text{lab}(v)$ is not a literal and no vertices have label \diamond or \square . For each vertex $s \in V$, we define the *left collection of s in P* , denoted by $\mathbb{L}(s)$, recursively starting from the leaves of (V, E, B) :

- if s is an unlabelled leaf, then $\mathbb{L}(s) = \text{PM}(\text{left}(s))$,
- if $\text{lab}(s)$ is a μ -variable, then $\mathbb{L}(s) = \emptyset$,
- if $\text{lab}(s)$ is a ν -variable, then $\mathbb{L}(s) = \text{PM}(\text{left}(s))$,
- if $\text{lab}(s) = \eta X$ and the successor is s_1 , then $\mathbb{L}(s) = \text{PM}(\text{left}(s)_{\text{new}}) \cup \mathbb{L}(s_1)$,
- if $\text{lab}(s) = \vee$ and the successors are s_1, s_2 , then $\mathbb{L}(s) = \text{PM}(\text{left}(s)_{\text{new}}) \cup \mathbb{L}(s_1) \cup \mathbb{L}(s_2)$
- if $\text{lab}(s) = \wedge$ and the successors are s_1, s_2 , then $\mathbb{L}(s) = \text{PM}(\text{left}(s)_{\text{new}}) \cup (\mathbb{L}(s_1) \cap \mathbb{L}(s_2))$.

The *right collection of s in P* , $\mathbb{R}(s)$, is defined symmetrically with left and right, as well as μ and ν , switched at every point of the definition.

Note that since v does not have a literal label, a leaf of (V, E, B) can only be either an unlabelled vertex or a vertex with a variable label. This is because any vertex labelled with an operator will have at least one successor and the game ends in the next position after any literal move.

In the following, we will associate superscripted sets like $\mathbb{L}'(s)$ with the position P' with the same superscript just like we have done so far with the components of the position.

Lemma 6.5. Let P° be a position in a game $\mu\text{-FS}_k^\Phi(\mathbb{A}, \mathbb{B})$, where the current vertex $u = v^\circ$ is a successor of a fixed point vertex and no modal moves are made. Let P' be a position after P° such that no X -move has returned to a vertex above u since P° . Then $\mathbb{L}^\circ(u) \subseteq \mathbb{L}'(u)$ and $\mathbb{R}^\circ(u) \subseteq \mathbb{R}'(u)$.

Proof. The proof proceeds by induction. We assume that for position P , $\mathbb{L}^\circ(u) \subseteq \mathbb{L}(u)$ and $\mathbb{R}^\circ(u) \subseteq \mathbb{R}(u)$ and we show that the inclusion also holds for the next position P' .

If S makes a ηX -move, the new models in $\text{left}(v)$ are moved to $\text{left}'(v')$ but they still remain in $\mathbb{L}'(v)$ so $\mathbb{L}'(u) = \mathbb{L}(u)$. Note that clocks of the models do change but $\mathbb{L}'(v)$ only looks at the underlying pointed models. For the same reasons $\mathbb{R}'(u) = \mathbb{R}(u)$.

For \vee -moves the models in $\text{left}(v)$ are split among the successors v_1 and v_2 so they are still in $\mathbb{L}'(v)$ in position P' . The models in $\text{right}(v)$ are copied to both v_1 and v_2 so they are still in $\mathbb{R}'(v)$. Thus $\mathbb{L}'(u) = \mathbb{L}(u)$ and $\mathbb{R}'(u) = \mathbb{R}(u)$. The case of \wedge -moves is symmetric with the two sides switched everywhere.

If S makes an X -move, it can either be a return to u or to a vertex s below u . Assume that the return is made to u and that X is a μ -variable. Now the models in $\text{left}(v)$ are moved to $\text{left}'(u)$ so any that were already in $\mathbb{L}(u)$ stay there and more may be added so $\mathbb{L}(u) \subseteq \mathbb{L}'(u)$. The models in $\text{right}(v)$ cease to be relevant in $\text{right}'(u)$ but they remain as old models in $\text{right}'(v)$ and are still counted for $\mathbb{R}'(v)$ in P' as they were in P so $\mathbb{R}'(u) = \mathbb{R}(u)$. The case of a ν is symmetric with the two sides switched.

Finally assume that an X -move is made returning to a vertex s below u . Assume again that X is a μ -variable. The models in $\text{left}(v)$ are moved to $\text{left}'(s)$ so $\mathbb{L}(s) \subseteq \mathbb{L}'(s)$. Everything not below s remains unchanged so $\mathbb{L}(u) \subseteq \mathbb{L}'(u)$. The models in $\text{right}(v)$ cease to be relevant in $\text{right}'(s)$ but they remain as old models in $\text{right}'(v)$ and are still counted for $\mathbb{R}'(s)$ and therefore also for $\mathbb{R}'(u)$ just like in P . Thus $\mathbb{R}'(u) = \mathbb{R}(u)$. The case of a ν is again symmetric. \square

We finally have all of the required notation and lemmas to show that the non-elementary succinctness gap is present also between FO and modal μ -calculus.

Theorem 6.6. *First-order logic is non-elementarily more succinct than modal μ -calculus.*

Proof. We prove an analogous result to Lemma 4.9 for the L_μ game. We use the notation $\mathcal{G}(\mathbb{V}, \mathbb{E})$ for the same graph as in subsection 4.3. The precise statement we prove is as follows:

Let $n \in \mathbb{N}$ and let $k_0 \in \mathbb{N}$. If $k_0 < \log(\chi(\mathcal{G}(\mathbb{C}_n, \mathbb{D}_n)))$, then D has a winning strategy in the game $\mu\text{-FS}_{k_0}^\Phi(\mathbb{C}_n, \mathbb{D}_n)$.

In this proof we only consider relevant models. Many positions in the game $\mu\text{-FS}_{k_0}^\Phi(\mathbb{C}_n, \mathbb{D}_n)$ have also non-relevant models but they are not needed for the strategy of D we describe here and can safely be ignored. We will assume all models are relevant and occasionally comment on why models remain or cease to be relevant after some moves of the game.

We show by induction that D has a strategy to maintain the following condition in any position $P = (V, E, B, \text{lab}, \text{res}, \text{left}, \text{right}, v)$:

$$\text{res}(v) < \log(\chi(\mathcal{G}(\mathbb{L}(v), \mathbb{R}(v)))) \quad (*)$$

At the start of the game, $*$ holds by assumption. Since the sets \mathbb{C}_n and \mathbb{D}_n are already finite, D can keep the full sets for the first move of the game.

We first show that if S ever makes a modal move while $*$ holds, D gets a winning strategy for the game. We assume $v \notin \text{dom}(\text{lab})$ since the first modal move in a game must always be made in an unlabelled vertex. In this case there are no other vertices below v so $\mathbb{L}(v) = \text{left}(v)$ and $\mathbb{R}(v) = \text{right}(v)$. We assume $\text{res}(v) > 1$ so that S can make a modal move and not lose immediately due to the resource running out. From $*$ we obtain $\chi(\mathcal{G}(\text{left}(v), \text{right}(v))) > 2$ so there are relevant clocked models $(\Delta(\mathcal{M}_1, w_1), c_1, \text{new}), (\Delta(\mathcal{M}_2, w_2), c_2, \text{new}) \in \text{left}(v)$ and $(\Delta\{(\mathcal{M}_1, w_1), (\mathcal{M}_2, w_2)\}, c_3, \text{new}) \in \text{right}(v)$. Now if S makes a \diamond - or \square -move, then in the following position P' there is $i \in \{1, 2\}$ s.t. $(\mathcal{M}_i, w_i, c_i, \text{new}) \in \text{left}(v')$ and $(\mathcal{M}_i, w_i, c_3, \text{new}) \in \text{right}(v')$. As these two share the same underlying pointed model they are bisimilar and moreover, since the depth has decreased by at least 1 from the previous position, the models are strictly relevant. By Lemma 6.3, D now has a winning strategy from position P' .

If S makes a \vee -move, let v_1 and v_2 be the successors of v . In the following position, whichever it may be, we have $\mathbb{L}(v) = \mathbb{L}(v_1) \cup \mathbb{L}(v_2)$ and $\mathbb{R}(v) = \mathbb{R}(v_1) \cap \mathbb{R}(v_2)$. Let $\mathcal{G}_s = \mathcal{G}(\mathbb{L}(s), \mathbb{R}(s)) = (V_s, E_s)$ for $s \in \{v, v_1, v_2\}$. We obtain $V_v = V_{v_1} \cup V_{v_2}$ and $E_v \upharpoonright V_{v_i} \subseteq E_{v_i}$ for $i \in \{1, 2\}$. By Lemma 4.8,

$$\chi(\mathcal{G}_v) \leq \chi(V_{v_1}, E_v \upharpoonright V_{v_1}) + \chi(V_{v_2}, E_v \upharpoonright V_{v_2}) \leq \chi(\mathcal{G}_{v_1}) + \chi(\mathcal{G}_{v_2}).$$

Thus (just like in the proof of Theorem 4.9) we obtain $\text{res}(v_i) < \log(\chi(\mathcal{G}_{v_i}))$ for some $i \in \{1, 2\}$ so $*$ holds in the following position after D chooses this i .

If S makes a \wedge -move, let v_1 and v_2 be the successors of v . In the following position we have $\mathbb{L}(v) = \mathbb{L}(v_1) \cap \mathbb{L}(v_2)$ and $\mathbb{R}(v) = \mathbb{R}(v_1) \cup \mathbb{R}(v_2)$. We use the notation $\mathcal{G}_s = (V_s, E_s)$ from the previous case and obtain $V_v = V_{v_1} \cap V_{v_2}$ and $E_v = (E_{v_1} \upharpoonright V_v) \cup (E_{v_2} \upharpoonright V_v)$. By Lemma 4.8,

$$\chi(G_v) \leq \chi(V_v, E_{v_1} \upharpoonright V_v) \chi(V_v, E_{v_2} \upharpoonright V_v) \leq \chi(\mathcal{G}_{v_1}) \chi(\mathcal{G}_{v_2}).$$

Thus we again obtain $\text{res}(v_i) < \log(\chi(\mathcal{G}_{v_i}))$ for some $i \in \{1, 2\}$ so $*$ holds in the following position after D chooses this i .

If S makes a ηX -move, where $\eta \in \{\mu, \nu\}$, then D sets her clock for each model at the same value as the depth of the model. All relevant models remain relevant and $*$ is maintained.

If S makes an X -move, by $*$, S does not immediately win the game. Assume that u is the vertex returned to and P° is the previous position when u was the current vertex. Let P' be the position after this X -move. By Lemma 6.5 we obtain $\mathbb{L}^\circ(u) \subseteq \mathbb{L}'(u)$ and $\mathbb{R}^\circ(u) \subseteq \mathbb{R}'(u)$. By induction hypothesis, $*$ held in P° and clearly $\text{res}'(u) = \text{res}^\circ(u)$, so $*$ still holds in P' .

If S makes a *Lit*-move, by $*$, $\text{left}(v) \neq \emptyset$ and $\text{right}(v) \neq \emptyset$. Since all the models are propositionally equivalent, D wins the game.

By Theorem 5.6, we obtain that there is no sentence $\varphi \in L_\mu(\emptyset)$ that separates \mathbb{C}_n from \mathbb{D}_n with $\text{sz}(\varphi) \leq \log(\chi(\mathcal{G}(\mathbb{C}_n, \mathbb{D}_n))) = \text{twr}(n - 1)$. Thus FO is non-elementarily more succinct than L_μ . □

Remark 6.7. Just like in the case of ML, we remark that the result of Theorem 6.6 also holds for DAG-size. This is again because the difference between the size of an L_μ formula in our sense and the DAG-size of the same formula is at most exponential.

7 Conclusion

We have defined formula size games for basic modal logic and modal μ -calculus. The games utilize resource parameters to achieve a truly two-player game. In the case of modal logic the players only construct one branch of the game tree. This is in contrast with the original Adler-Immerman game, where the players form the whole tree in a single play. For modal μ -calculus the recursive nature of fixed point operators necessitates returning to previously visited nodes. However, the game still traverses only one possible path through the game tree in a single play and some branches can remain unvisited for the entire play. The μ -calculus game has infinite branching but the use of decreasing ordinal clocks, as in [16], makes each play of the game finite.

We used the games to show that the property “all successor models are n -bisimilar with each other” cannot be defined in basic modal logic or modal μ -calculus with a formula of size less than the exponential tower of height $n - 1$. On the other hand this property can be defined in FO with a formula of size linear in n . This means that FO is non-elementarily more succinct than both ML and L_μ . We also show that the same property can be defined in two-dimensional modal logic ML^2 with a formula of size exponential in n . Thus the non-elementary succinctness gap is also present between ML^2 and both ML and L_μ .

We find the ML-game to be a useful tool for proving lower bounds on the size of ML-formulas. Depending on the desired result, the game can also be modified to count a more specific parameter instead of formula size, such as the number or nesting depth of a specific operator.

The μ -calculus game is also functional for proving lower bounds, but with some caveats. The main theorem stating the usefulness of the game, Theorem 5.6, requires uniform strategies for S. This means that we assume S has a single formula in mind and always plays according to that formula. It may be that this restriction could be removed but we have been unable to prove this. However, to show succinctness results we only need one direction of the equivalence

so the issue is usually not relevant in practice. The greater concern is whether the game can be successfully used to prove succinctness results for μ -calculus in more complicated contexts. Here we only generalize a result already obtained with the ML game and we have so far failed to produce any other results with the L_μ game due to its sheer complexity. A question related to this difficulty would be whether the game could be simplified significantly while still preserving its functionality. It would be especially interesting to apply the game to open problems related to μ -calculus and succinctness, such as whether there is a polynomial transformation from L_μ to the guarded fragment or from vectorial form to regular L_μ [4].

References

- [1] M. Adler and N. Immerman. An $n!$ lower bound on formula size. *ACM Trans. Comput. Log.*, 4(3):296–314, 2003.
- [2] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, New York, NY, USA, 2001.
- [3] J. Bradfield and C. Stirling. Chapter 12 - modal mu-calculi. In P. Blackburn, J. V. Benthem, and F. Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical Reasoning*, pages 721 – 756. Elsevier, 2007.
- [4] F. Bruse, O. Friedmann, and M. Lange. On guarded transformation in the modal mu-calculus. *Logic Journal of the IGPL*, 23(2):194 – 216, 2015.
- [5] K. Etessami, M. Y. Vardi, and T. Wilke. First-order logic with two variables and unary temporal logic. *Inf. Comput.*, 179(2):279–295, 2002.
- [6] K. Etessami and T. Wilke. An until hierarchy and other applications of an ehrenfeucht-fraïssé game for temporal logic. *Inf. Comput.*, 160(1-2):88–108, 2000.
- [7] D. Fernández-Duque and P. Iliev. Succinctness in subsystems of the spatial mu-calculus. *ArXiv e-prints*, Aug. 2017.
- [8] T. French, W. van der Hoek, P. Iliev, and B. P. Kooi. Succinctness of epistemic languages. In *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 881–886, 2011.
- [9] T. French, W. van der Hoek, P. Iliev, and B. P. Kooi. On the succinctness of some modal logics. *Artif. Intell.*, 197:56–85, 2013.
- [10] D. Gabbay and V. Shehtman. Products of modal logics, part 1. *Logic Journal of the IGPL*, 6(1):73–146, 1998.
- [11] D. Gabbay and V. Shehtman. Products of modal logics. part 2: Relativised quantifiers in classical logic. *Logic Journal of the IGPL*, 8(2):165–210, 2000.
- [12] D. Gabbay and V. Shehtman. Products of modal logics. part 3: Products of modal and temporal logics. *Studia Logica*, 72(2):157–183, 2002.
- [13] G. Gogic, H. Kautz, C. Papadimitriou, and B. Selman. The comparative linguistics of knowledge representation. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence - Volume 1, IJCAI'95*, pages 862–869, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc.
- [14] M. Grohe and N. Schweikardt. The succinctness of first-order logic on linear orders. *Logical Methods in Computer Science*, 1(1), 2005.

- [15] Grohe, Martin and Schweikardt, Nicole. Comparing the succinctness of monadic query languages over finite trees. *RAIRO-Theor. Inf. Appl.*, 38(4):343–373, 2004.
- [16] L. Hella, A. Kuusisto, and R. Rönholm. Bounded game-theoretic semantics for modal mu-calculus. *ArXiv e-prints*, June 2017.
- [17] L. Hella and J. Väänänen. The size of a formula as a measure of complexity. In *Logic Without Borders - Essays on Set Theory, Model Theory, Philosophical Logic and Philosophy of Mathematics*, pages 193–214. De Gruyter, 2015.
- [18] L. Hella and M. Vilander. The succinctness of first-order logic over modal logic via a formula size game. In *Advances in Modal Logic 11, Budapest, Hungary, August 30 - September 2, 2016*, pages 401–419, 2016.
- [19] G. Lenzi. The modal mu-calculus: A survey. *Task Q*, 9:293–316, 2005.
- [20] C. Lutz. Complexity and succinctness of public announcement logic. In *5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006), Hakodate, Japan, May 8-12, 2006*, pages 137–143, 2006.
- [21] C. Lutz, U. Sattler, and F. Wolter. Modal logic and the two-variable fragment. In *Computer Science Logic, Paris, France, September 10-13, 2001, Proceedings*, pages 247–261, 2001.
- [22] N. Markey. Temporal logic with past is exponentially more succinct, concurrency column. *Bulletin of the EATCS*, 79:122–128, 2003.
- [23] M. Marx and Y. Venema. *Multi-dimensional modal logic*, volume 4 of *Applied Logic Series*. Kluwer Academic Publishers, Dordrecht, 1997.
- [24] J. C. McCabe-Dansted, T. French, S. Pinchinat, and M. Reynolds. Expressiveness and succinctness of a logic of robustness. *Journal of Applied Non-Classical Logics*, 25(3):193–228, 2015.
- [25] M. Otto. Bisimulation invariance and finite models. In Z. Chatzidakis, P. Koepke, and W. Pohlers, editors, *Logic Colloquium '02*, Lecture Notes in Logic, pages 276–298. Cambridge University Press, 2006.
- [26] A. P. Sistla and E. M. Clarke. The complexity of propositional linear temporal logics. *J. ACM*, 32(3):733–749, 1985.
- [27] L. J. Stockmeyer. *The Complexity of Decision Problems in Automata Theory and Logic*. PhD thesis, Massachusetts Institute of Technology, 1974.
- [28] W. van der Hoek and P. Iliev. On the relative succinctness of modal logics with union, intersection and quantification. In *International conference on Autonomous Agents and Multi-Agent Systems, AAMAS '14, Paris, France, May 5-9, 2014*, pages 341–348, 2014.
- [29] W. van der Hoek, P. Iliev, and B. P. Kooi. On the relative succinctness of two extensions by definitions of multimodal logic. In *How the World Computes - Turing Centenary Conference and 8th Conference on Computability in Europe, CiE 2012, Cambridge, UK, June 18-23, 2012. Proceedings*, pages 323–333, 2012.
- [30] H. van Ditmarsch, J. Fan, W. van der Hoek, and P. Iliev. Some exponential lower bounds on formula-size in modal logic. In *Advances in Modal Logic 10, held in Groningen, The Netherlands, August 5-8, 2014*, pages 139–157, 2014.
- [31] H. van Ditmarsch and P. Iliev. The succinctness of the cover modality. *Journal of Applied Non-Classical Logics*, 25(4):373–405, 2015.

- [32] T. Wilke. CTL⁺ is exponentially more succinct than CTL. In *Foundations of Software Technology and Theoretical Computer Science, 19th Conference, Chennai, India, December 13-15, 1999, Proceedings*, pages 110–121, 1999.