Exponential Lower Bounds on Definable Fixed Points

Konstantinos Papafilippou
Department of Mathematics, University of Ghent, Belgium

David Fernández-Duque
Department of Philosophy, University of Barcelona, Spain

Abstract

It is known that the μ-calculus is no more expressive than basic modal logic over the class of finite partial orders, as well as over the class of finite, strict partial orders. Nevertheless, we show that the μ-calculus is exponentially more succinct, even when a reflexive modality is added as primitive. As corollaries, we obtain a lower bound for the fixed-point theorem for Gödel-Löb logic and a variant for Grzegorczyk logic, as well as lower bounds on interpolants for the interpolation theorem of Gödel-Löb logic.

2012 ACM Subject Classification Theory of computation → Modal and temporal logics; Theory of computation → Complexity theory and logic

Keywords and phrases Modal logic, Provability Logic, Fixed-Point Theorem, mu-calculus, Interpolation, Succinctness

Digital Object Identifier 10.4230/LIPIcs.CSL.2025.25

Funding This work has been partially supported by FWO-FWF grant G030620N/I4513N and SNSF–FWO Lead Agency Grant 200021L_196176/G0E2121N.

1 Introduction

Expressivity and complexity are two crucial criteria in the design of formal languages for logical reasoning. However, these two properties alone only paint part of the picture, as a more *succinct* language may have advantages over an equally (or even more) expressive language if formula size is reduced sufficiently to considerably save on storage space and improve processing time. Two formal languages $\mathcal{L}_1$ and $\mathcal{L}_2$ may be equally expressive, yet certain properties may be expressed in $\mathcal{L}_1$ by much shorter expressions than in $\mathcal{L}_2$; when the size difference is e.g. exponential, it may dwarf any potential advantage offered by $\mathcal{L}_2$ on account of purely complexity-theoretic considerations.

Modal logics are an appealing framework for computational logic precisely due to the balance between expressivity and complexity, making them more adaptable than propositional logic but more tractable than first or higher order logic. But a well-informed choice of the "right" modal logic for a given task should also involve an understanding of how it fares in terms of succinctness.

In particular, Gödel-Löb logic (**GL**) provides a textbook example of a success story in modal logic: it is the logic of finite (or, more generally, converse well-founded) strict partial orders, hence it governs the behaviour of computational processes that terminate in finite time. It is obtained from the basic modal logic **K** by adding *Löb's axiom,* $\Box(\Box\varphi \to \varphi) \to \Box\varphi$. It is also the logic of provability in Peano arithmetic and related theories, as well as the logic of scattered topological spaces, granting it applications in foundations of mathematics and spatial reasoning. For the first setting, one interprets modal formulas as arithmetical statements, with variables representing arbitrary statements in the language of **PA** and $\Box\varphi$ being interpreted as Gödel's $\mathsf{Bew}(\ulcorner\varphi\urcorner)$, which formalises the statement that φ is provable in **PA**. Solovay [23] showed that the set of valid formulas (i.e., those that correspond to

33rd EACSL Annual Conference on Computer Science Logic (CSL 2025).
Editors: Jörg Endrullis and Sylvain Schmitz; Article No. 25; pp. 25:1–25:19
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

theorems of **PA**) in this setting are precisely those provable in **GL**. For the second, one interprets $\Diamond$ as a topological Cantor derivative operator, e.g. $\Diamond\varphi$ is the set of limit points of those points satisfying φ. **GL** once again captures the set of validities in this context (see e.g. [6]).

Moreover, **GL** is remarkably well-behaved, being decidable, finitely axiomatizable, and enjoying Craig interpolation [8] and definable fixed points [21]. The latter in particular means that the μ-calculus adds no expressive power to **GL** [2]. This is also true for the class **Grz** of Grzegorczyk frames, based on *Noetherian* posets; essentially, the reflexive closures of **GL** frames [9]. **Grz** is also the logic of "provably true" over Peano arithmetic and is the greatest modal companion of intuitionistic propositional logic [26, 8] and is characterised by the axiom $\Box(\Box(\varphi \to \Box\varphi) \to \varphi) \to \Box\varphi$.

One could thus jump to the conclusion that the μ-calculus over finite posets (either reflexive or irreflexive) is not worth considering. However, such disinterest would be misguided, as it does not take questions of succinctness into account. The fixed-point theorem for **GL** states that for any formula $\varphi(x)$ where x occurs only in the scope of $\Box$ (or $\Diamond$), there is a formula ψ such that $\psi \leftrightarrow \varphi(\psi)$ is derivable. All proofs [7, 16, 20, 21, 22] yield some ψ that is at least exponentially larger than φ. This raises the question of whether this bound is optimal, to which we provide a positive answer. In contrast, the μ-calculus formula $\mu x.\varphi(x)$ yields a fixed point of φ and is only slightly larger than φ itself.[1] Thus we conclude that, despite fixed points already being definable in the basic modal language, there is much to be gained by passing to a language with explicit fixed point constructors.

Research in succinctness involves delicate techniques and it has been an active area in the last decades; see e.g. [17, 10, 1]. Closest to the present work, [11] show that over **GL** frames, a language with the reflexive modality $\diamonddot$ is exponentially more succinct than a language with $\Diamond$. As a corollary, exponential succinctness of the μ-calculus is obtained for a language with $\Diamond$ alone, given that $\diamonddot$ can be defined succinctly in the μ-calculus. However, this result has two shortcomings with regards to our current goal. First, it does not clarify if the μ-calculus is more succinct than a language with the reflexive modality $\diamonddot$, so that the results cannot be applied to the logic **Grz** which enjoys a restricted version of the fixed-point theorem. Second, succinctness is obtained via nested fixed point operators, and the lower bound for the fixed-point theorem would require a *single* application of μ. We thus aim for a sharper result for the μ-calculus, for which we extend known techniques and provide new constructions not contingent on the distinction between $\Diamond$ and $\diamonddot$.

Intuitively, proving that one language $\mathcal{L}_1$ is more succinct than another language $\mathcal{L}_2$ ultimately boils down to proving a sufficiently big lower bound on the size of $\mathcal{L}_2$-formulas expressing some semantic property. If we want to show that $\mathcal{L}_1$ is exponentially more succinct than $\mathcal{L}_2$, we must find an infinite sequence of semantic properties (i.e., classes of models) $\mathbf{P}_1, \mathbf{P}_2, \ldots$ definable in both $\mathcal{L}_1$ and $\mathcal{L}_2$, show that there are $\mathcal{L}_1$-formulas $\varphi_0, \varphi_1, \ldots$ defining $\mathbf{P}_1, \mathbf{P}_2, \ldots$ and prove that, for every n, every $\mathcal{L}_2$-formula ψ_n defining $\mathbf{P}_n$ has size exponential in the size of φ_n. There are various techniques used for achieving such results; here we use formula-size games developed in the setting of Boolean function complexity by [19] and in the setting of first-order logic and some temporal logics by [1]. By now, the formula-size games have been adapted to a host of modal logics (see for example [14], [18], [13], [25]) and used to obtain lower bounds on modal formulas expressing properties of Kripke models.

[1] Normally the μ-calculus requires that x appear only positively in $\varphi(x)$, but over **GL** this condition can be weakened to allow for modalized formulas.

2 Modal logic

In this section, we present the modal μ-calculus and formalize its Kripke semantics. Let us begin by defining the base modal language we will work with. We will consider logics over variants of the language $\mathcal{L}_\Diamond$ given by the following grammar (in Backus-Naur form). Fix a set $\mathbb{P}$ of *propositional variables* (also called *atoms*), and define:

$$\varphi, \psi := \quad \top \quad | \quad \bot \quad | \quad p \quad | \quad \overline{p} \quad | \quad \varphi \vee \psi \quad | \quad \varphi \wedge \psi \quad | \quad \Diamond\varphi \quad | \quad \Box\varphi$$

Here, $p \in \mathbb{P}$ and $\overline{p}$ denotes the negation of p. For the game-theoretic techniques we will use, it is convenient to allow negations only at the atomic level, and thus we include all duals as primitives, but not negation or implication; however, we may use the latter as shorthands, defined via De Morgan's laws. Formulas of the forms $p, \overline{p}$ are *literals.* The *size* of a formula φ is denoted $|\varphi|$ and is defined as follows.

▶ **Definition 1.** *We define a function* $|\cdot| : \mathcal{L}_\Diamond \to \mathbb{N}$ *recursively by*

- $|p| = |\overline{p}| = 1$
- $|\varphi \wedge \psi| = |\varphi \vee \psi| = |\varphi| + |\psi| + 1$
- $|\Diamond\varphi| = |\Box\varphi| = |\varphi| + 1.$

Next we review semantics for modal logic in general, and for **GL** in particular.

▶ **Definition 2.** *A* Kripke frame *is a structure* $\mathcal{A} = (|\mathcal{A}|, R_\mathcal{A})$ *where* $R_\mathcal{A}$ *is a binary relation on* $|\mathcal{A}|$*. If* $\mathcal{A}$ *is a Kripke frame, a* valuation *on* $\mathcal{A}$ *is a function* $V : |\mathcal{A}| \to 2^\mathbb{P}$ *(recall that* $\mathbb{P}$ *is the set of atoms). A frame* $\mathcal{A}$ *equipped with a valuation* V *(often denoted* $V_\mathcal{A}$*) is a* Kripke model.

By abuse on notation we will write $x \in \mathcal{A}$ instead of $x \in |\mathcal{A}|$. The valuation V can be extended recursively to define truth of all formulas of the modal language.

▶ **Definition 3.** *Let* $\mathcal{A} = (\mathcal{A}, R_\mathcal{A})$ *be any Kripke frame and* V *a valuation. We define the truth set*

$$\|\varphi\|_\mathcal{A} =: \{w \in \mathcal{A} : (\mathcal{A}, w) \Vdash \varphi\}$$

by structural induction on φ*:*

$$\begin{array}{lcl} w \in \|p\|_\mathcal{A} & \Leftrightarrow & p \in V(w) \\ w \in \|\overline{p}\|_\mathcal{A} & \Leftrightarrow & p \notin V(w) \\ w \in \|\varphi \wedge \psi\|_\mathcal{A} & \Leftrightarrow & w \in \|\varphi\|_\mathcal{A} \cap \|\psi\|_\mathcal{A} \\ w \in \|\varphi \vee \psi\|_\mathcal{A} & \Leftrightarrow & w \in \|\varphi\|_\mathcal{A} \cup \|\psi\|_\mathcal{A} \\ w \in \|\Diamond\varphi\|_\mathcal{A} & \Leftrightarrow & \exists v(wR_\mathcal{A}v \,\&\, w \in \|\varphi\|_\mathcal{A}) \\ w \in \|\Box\varphi\|_\mathcal{A} & \Leftrightarrow & \forall v(wR_\mathcal{A}v \Rightarrow w \in \|\varphi\|_\mathcal{A}) \end{array}$$

Given a model $\mathcal{A}$ *and formulas* φ, ψ*, we say that* φ *is* equivalent to ψ *on* $\mathcal{A}$ *if* $\|\varphi\|_\mathcal{A} = \|\psi\|_\mathcal{A}$*. If* $\mathbf{A}$ *is a class of models, we say that* φ, ψ *are equivalent over* $\mathbf{A}$ *if they are equivalent on any element of* $\mathbf{A}$*. We may also say that* $\varphi \equiv \psi$ *over* $\mathbf{A}$ *and omit mention of* $\mathbf{A}$ *if it is the class of all Kripke models.*

We will focus our attention mostly on the logics **GL** and **Grz**, which as we will see can be regarded as a fragment. **GL** may be interpreted over structures with a converse well-founded relation and for our purposes we may restrict our attention to models based on trees, presented as strict partial orders.

▶ **Definition 4.** *A* tree *is a pair $(T, \prec)$, where T is a set and $\prec$ is a strict partial order such that, if $\eta \in T$ then $\{\zeta \in T : \zeta \prec \eta\}$ is finite and linearly ordered, and T has a minimum element called its* root*. We will sometimes notationally identify $(T, \prec)$ as T, and write $\preceq$ for the reflexive closure of $\prec$.*

Maximal elements of T are leaves. For $\eta, \zeta \in T$, we say that ζ is a child of η if ζ is the least element ξ (if it exists) such that $\eta \prec \xi$. A path (of length m) on T is a sequence $\vec{\eta} = (\eta_i)_{i \leq m}$ such that η_{i+1} is the child of η_i.

For our purposes, a **GL** *model* is a model $\mathcal{A}$ where $(\mathcal{A}, R_{\mathcal{A}})$ is a finite tree, in which case we write $\sqsubset_{\mathcal{A}}$ instead of $R_{\mathcal{A}}$. As we will be working exclusively with **GL** frames and models, in the sequel we write simply *frame* or *model* instead of **GL** frame or **GL** model.

▶ Remark 5. It should be stressed that working in a more restrictive class of models yields *stronger* results as far as succinctness is concerned: for example, if no small modal formula ψ is equivalent to some μ-calculus expression φ over the class of **GL** models as we have defined them, then certainly no small ψ' is equivalent to φ over the class of *all* Kripke models, as in particular ψ' would still have to be equivalent to φ over the smaller class of **GL** models.

3 Extensions and fixed points

The modal language, as we have presented it, may be naturally extended to include other operations. Even when these operations do not add expressive power to our language, they can yield considerable gains in terms of succinctness, as we will see later in the text. We begin by discussing the reflexive modality.

3.1 The reflexive modality

We may define a modality based on $\sqsubseteq$ rather than $\sqsubset$. This may be defined in $\mathcal{L}_{\Diamond}$ by letting $\boxdot\varphi$ be a shorthand for $\varphi \wedge \Box\varphi$. Dually, $\diamonddot\varphi$ is defined as a shorthand for $\varphi \vee \Diamond\varphi$. Let $\mathcal{L}_{\Diamond\diamonddot}$ be the extension of $\mathcal{L}_{\Diamond}$ that includes $\diamonddot, \boxdot$ as primitives. Semantics for $\mathcal{L}_{\Diamond\diamonddot}$ are defined by setting $\|\varphi\|_{\mathcal{A}} = \|\varphi'\|_{\mathcal{A}}$, where φ' is obtained by replacing instances of $\diamonddot, \boxdot$ by their definitions; note that in general, φ' tends to be exponentially larger than φ [11]. We extend Definition 1 to $\mathcal{L}_{\Diamond\diamonddot}$ in the obvious way, by

$$|\diamonddot\varphi| = |\boxdot\varphi| = |\varphi| + 1.$$

Closely related to **GL** is the logic **Grz** of Noetherian (reflexive) partial orders, but it is easy to see that **Grz** is also the logic of **GL** frames, although based on $\mathcal{L}_{\diamonddot}$ rather than $\mathcal{L}_{\Diamond}$. By working over the combined language $\mathcal{L}_{\Diamond\diamonddot}$, our succinctness results apply to both **GL** and **Grz**, as well as many weaker logics.

3.2 Fixed Point Theorems

The celebrated De Jongh-Sambin theorem states that fixed points for modalized formulas are definable in **GL**, where x is *modalized* in φ if it only appears in the scope of $\Diamond$ or $\Box$.[2] For example, the formula $\neg\Box p$ has a fixed point ψ such that $\psi \equiv \neg\Box\psi$ over **GL**; in this case, we can take $\psi = \Diamond\top$. An upper bound on the size of ψ can be obtained by analyzing existing proofs.

[2] In other words, φ is of the form $\psi(\Box(\chi_1(x)), \dots, \Box(\chi_n(x)))$ with x not occuring in $\psi(p_1, \dots, p_n)$.

▶ **Theorem 6** (De Jongh ~1975, Sambin [21]). *Given a formula $\varphi(x)$ in which x is modalized, there is a formula ψ such that $\varphi(\psi) \equiv \psi$ over the class of* **GL** *models. The formula ψ is unique up to equivalence, and is of size $2^{O(|\varphi|\log(|\varphi|))}$.*

Proof. We follow the construction of the fixed point formula in [21]. Since x is modalized in φ, we have that $\varphi = \psi\big(\Box\chi_1(x), \ldots, \Box\chi_n(x)\big)$. Let σ_i^n be the fixed point of $\psi[\Box\chi_i(x)/\top]$, then the fixed point σ^{n+1} of φ is $\psi\big(\Box\chi_1(\sigma_1^n), \ldots, \Box\chi_n(\sigma_n^n)\big)$. Let $m = |\varphi|$ and $k \le |\psi|$. Then by induction on n, one readily shows that $|\sigma^1| \le m + k$ and then $|\sigma^n + 1| \le m + n \cdot |\sigma^{n-1}| \le m \cdot 2^{(n+1)\log(n+1)}$. ◀

The logic **Grz** also enjoys a fixed point property, but in this case for formulas $\varphi(x)$ where x is *positive,* i.e. with the restriction that $\overline{x}$ may not appear in φ.

Observe how in the case of **GL**, we ask for x to be modalized in $\varphi(x)$ while in **Grz** and in the μ-calculus we want it to be positive. Positivity is typically required in order to avoid pathological cases. For example if $\Box\neg p$ were to have a definable fixed point ψ in **Grz**, then over **Grz** the following would hold $\psi \equiv \Box\neg\psi \equiv \neg\psi \wedge \Box\neg\psi$, a contradiction. Naturally, there is always a fixed point for any $\varphi(x)$ over **GL** as well when x is positive in $\varphi(x)$, however it is not necessarily unique up to equivalence.

▶ **Theorem 7.** *Given a formula $\varphi(x)$ in which x is positive, there is a formula ψ such that $\varphi(\psi) \equiv \psi$ over the class of* **Grz** *models. The size of ψ may be bounded by a $2^{O(|\varphi|^3)}$ function.*

This does not seem to have been stated in this form in the literature, but it is a consequence of Theorem 8 below, which states that the μ-calculus is no more expressive than modal logic over the classes of **GL** or **Grz** models. In order to make this precise, let us review the μ-calculus.

3.3 The μ-calculus

The μ-calculus is obtained from $\mathcal{L}_{\Diamond\diamonddot}$ by adding formula constructors $\mu x.\varphi$ and $\nu x.\varphi$, where x is positive in φ. We denote the resulting language by $\mathcal{L}^{\mu}_{\Diamond\diamonddot}$. For a model $\mathcal{A}$, a variable x and $X \subseteq |\mathcal{A}|$, let $\mathcal{A}_{[x/X]}$ be a model which is the same as $\mathcal{A}$ except that $V_{\mathcal{A}_{[x/X]}}(x) = X$. Then, $\|\mu x.\varphi\|_{\mathcal{A}}$ is the least fixed point of the map $X \mapsto \|\varphi\|_{\mathcal{A}_{[x/X]}}$; in other words, $\|\mu x.\varphi\|_{\mathcal{A}} = \|\varphi\|_{\mathcal{A}_{[x/\|\mu x.\varphi\|_{\mathcal{A}}]}}$, and every other set with this property contains $\|\mu x.\varphi\|_{\mathcal{A}}$. Syntactically, we obtain $\mu x.\varphi(x) \equiv \varphi(\mu x.\varphi(x))$. Similarly, $\|\nu x.\varphi\|_{\mathcal{A}}$ is the *greatest* fixed point of the map $X \mapsto \|\varphi\|_{\mathcal{A}_{[x/X]}}$. This definition is known to be sound due to the Knaster-Tarski theorem [24], which entails that monotone operators on a powerset always have least and greatest fixed points.

Sub-languages of the μ-calculus are denoted by indicating the modalities allowed, e.g. $\mathcal{L}^{\mu}_{\Diamond}$ allows the modalities $\Diamond, \Box$ but not $\diamonddot, \boxdot$. Formula complexity is extended by setting

$$|\mu x.\varphi| = |\nu x.\varphi| = |\varphi| + 1.$$

Normally the μ-calculus provides a far-reaching extension of modal logic, but surprisingly this is no longer the case over **GL** [2] and **Grz** [9]. The bounds given are obtained from a separate construction by [12].

▶ **Theorem 8.** *Given a formula $\varphi \in \mathcal{L}^{\mu}_{\Diamond\diamonddot}$ of the μ-calculus, there are formulas $\varphi_{\mathbf{GL}}$ and $\varphi_{\mathbf{Grz}}$ of $\mathcal{L}_{\Diamond}$ such that $\varphi_{\mathbf{GL}} \equiv \varphi$ over* **GL** *and $\varphi_{\mathbf{Grz}} \equiv \varphi$ over* **Grz**. *The sizes of $\varphi_{\mathbf{GL}}$ and $\varphi_{\mathbf{Grz}}$ are of size $2^{O(|\varphi|^3)}$.*

Proof. In [12], an explicit translation from the μ-calculus into an extension of modal logic is given, with an operator $\blacklozenge^{\infty}(\varphi_0, \varphi_1, \dots, \varphi_{n-1})$. Over both **GL** and **Grz**, this operator is equivalent to $\Diamonddot \bigwedge \varphi_i$, and thus one obtains a translation of the μ-calculus into $\mathcal{L}_{\Diamond}$. By examining the formulas involved, and essentially repeating the same calculation as in [12], we may compute an upper bound of $2^{O(|\varphi|^3)}$. ◀

Note that in general $\varphi_{\mathbf{GL}}$ and $\varphi_{\mathbf{Grz}}$ may be distinct. In view of these results, it may seem that over **GL** and **Grz**, the μ-calculus is merely a cosmetic extension of $\mathcal{L}_{\Diamond}$. However, note that the fixed-point formulas provided by Theorems 6 and 7 are quite large. As we will see, this is unavoidable and thus the μ-calculus offers a substantial advantage when succinctness is taken into account. In order to simultaneously provide lower bounds for Theorems 6–8, in our succinctness results, we will work with formulas that are both positive and modalized.

4 Model equivalence games

In this section, we set up the model equivalence games tailored for a language with $\Diamond$ and $\Diamonddot$ interpreted over **GL** models. The game is based on sets of rooted **GL** models; that is, **GL** models $\mathcal{A}$ which are *generated* from some $a \in \mathcal{A}$ in the sense that $\mathcal{A} = a{\uparrow}_{\mathcal{A}} := \{b \in \mathcal{A} : a \sqsubseteq b\}$. Henceforth we simply call these *rooted models* and we will usually designate the root by writing the model as the pair $(\mathcal{A}, a)$. The following operations will be useful in describing the game.

▶ **Definition 9.** *Given a set of rooted models* $\mathbf{A}$*, we define:*

- $\Box\mathbf{A} := \{(b{\uparrow}_{\mathcal{A}}, b) : a \sqsubset_{\mathcal{A}} b$ *for some* $(\mathcal{A}, a) \in \mathbf{A}\}$*;*
- $\boxdot\mathbf{A} := \{(b{\uparrow}_{\mathcal{A}}, b) : a \sqsubseteq_{\mathcal{A}} b$ *for some* $(\mathcal{A}, a) \in \mathbf{A}\}$*;*
- *given* $f : \mathbf{A} \to \bigcup_{\mathcal{A} \in \mathbf{A}} \mathcal{A}$ *where* $f(\mathcal{A}) \in \Box\mathcal{A}$*, then* $\Diamond_f \mathbf{A} := rng(f)$*;*
- *given* $f : \mathbf{A} \to \bigcup_{\mathcal{A} \in \mathbf{A}} \mathcal{A}$ *where* $f(\mathcal{A}) \in \boxdot\mathcal{A}$*, then* $\Diamonddot_f \mathbf{A} := rng(f)$*.*

We write $\Box\mathcal{A}$ *for* $\Box\{\mathcal{A}\}$ *and* $\boxdot\mathcal{A}$ *for* $\boxdot\{\mathcal{A}\}$ *respectively. We also write* $\mathbf{A} \Vdash \varphi$ *to mean that* φ *holds in the root of every model* $\mathcal{A} \in \mathbf{A}$*.*

▶ **Definition 10.** *Let* $\mathbf{M}$ *be a class of rooted models and* φ *be a formula. The* $(\varphi, \mathbf{M})$ *model equivalence game* ($(\varphi, \mathbf{M})$-MEG) *is played by two players,* Hercules *and the* Hydra*, according to the following rules.*

SETTING UP THE PLAYING FIELD.

The Hydra's only move is in the initiation of the game by choosing two sets of models $\mathbf{A}, \mathbf{B} \subseteq \mathbf{M}$ *such that* $\mathbf{A} \Vdash \varphi$ *and* $\mathbf{B} \Vdash \neg\varphi$*.*

After that, Hercules constructs a finite game-tree T*. Each node* $\eta \in T$ *will be labelled with a pair* $(\mathbf{L}(\eta), \mathbf{R}(\eta))$ *of sets of rooted models, and a symbol that is either a literal or one of* $\{\wedge, \vee, \Diamond, \Box, \Diamonddot, \boxdot\}$*. We will usually write* $\mathbf{A}(\eta) \circ \mathbf{B}(\eta)$ *instead of* $(\mathbf{A}, \mathbf{B})$ *for pairs of sets of rooted models.*

At each step of the construction, a leaf η *can be either declared a* head *or a* stub *in accordance to the rules of the game. Once it has been declared a stub, no further moves can be played on it. The root* λ *of the tree is labelled as* $\mathbf{L}(\lambda) \circ \mathbf{R}(\lambda) = \mathbf{A} \circ \mathbf{B}$ *and declared a head.*

Afterwards, the game continues so long as there is at least one head. In each turn, Hercules chooses a head η *labelled by* $\mathbf{L} \circ \mathbf{R}$ *and plays one of the following moves.*

literal-move.
Hercules chooses a literal ι such that $\mathbf{L} \Vdash \iota$ and $\mathbf{R} \nVdash \iota$. The node η is declared a stub and labelled with the symbol ι.

$\vee$-move.
Hercules labels η with the symbol $\vee$ and chooses two sets $\mathbf{L}_1, \mathbf{L}_2 \subseteq \mathbf{L}$ such that $\mathbf{L} = \mathbf{L}_1 \cup \mathbf{L}_2$. Two new heads, labelled by $\mathbf{L}_1 \circ \mathbf{R}$ and $\mathbf{L}_2 \circ \mathbf{R}$, are added to the tree as children of η.

$\wedge$-move.
Analogous to a $\vee$-MOVE, except that Hercules instead chooses $\mathbf{R}_1, \mathbf{R}_2 \subseteq \mathbf{R}$.

$\Diamond$-move.
Hercules labels η with the symbol $\Diamond$ and chooses a function f, as in Definition 9, for which $\Diamond_f \mathbf{L}$ exists (if it does not exist i.e. for some $\mathcal{A} \in \mathbf{L}$ we have $\Box \mathcal{A} = \emptyset$, Hercules cannot play this move). We let $\mathbf{L}_1$ be $\Diamond_f \mathbf{L}$ and $\mathbf{R}_1$ to be $\Box \mathbf{R}$.[3] *A new head labelled by $\mathbf{L}_1 \circ \mathbf{R}_1$ is added as a child to η.*

$\Box$-move.
Analogous to a $\Diamond$-MOVE, except that Hercules instead chooses a function f for which $\Diamond_f \mathbf{R}$ exists and the new head is labelled by $\Box \mathbf{L} \circ \Diamond_f \mathbf{R}$.

$\diamonddot$-move and $\boxdot$-move.
Analogous to $\Diamond$- and $\Box$-MOVES, but with $\boxdot$ and $\diamonddot$ in place of $\Box$ and $\Diamond$ respectively.

The $(\varphi, \mathbf{M})$-MEG game concludes when there are no heads. If the game-tree is finite (in size) and it has no heads, we call it closed *and Hercules has won. We say that* Hercules has a winning strategy in n moves *in the $(\varphi, \mathbf{M})$-MEG if no matter how the Hydra sets up the playing field, the resulting game tree has at most n nodes and is closed.*

If there is an $\mathcal{L}_{\Diamond\diamonddot}$ formula ψ equivalent to φ on $\mathbf{M}$, Hercules can read a winning strategy off of ψ for the $(\varphi, \mathbf{M})$-MEG. Conversely, if Hercules has a winning strategy then such a ψ can be read off of the game tree when Hydra plays optimally, i.e. always choosing as many rooted models as allowed. We thus obtain the following.

▶ **Theorem 11.** *Hercules has a winning strategy in n moves in the $(\varphi, \mathbf{M})$-MEG iff there is a $\mathcal{L}_{\Diamond\diamonddot}$ formula ψ equivalent to φ on $\mathbf{M}$ such that $|\psi| \leq n$. (See e.g. [18])*

It moreover should be clear that Hercules cannot win if there are isomorphic models on the left and right, since there will be no formula distinguishing them.

▶ **Proposition 12.** *No closed game tree contains a node η such that there are $\mathcal{A} \in \mathbf{L}(\eta)$, $\mathcal{B} \in \mathbf{R}(\eta)$ that are isomorphic. (See e.g. [18])*

5 The playing field

We will use the model equivalence games to show that the μ-calculus with a *single* application of the least fixed point operator is exponentially more succinct than modal logic over the class of **GL** frames, even when equipped with *both* $\Diamond$ and $\diamonddot$. Our proof will be based on an infinite sequence of formulas that convey the existence of a certain binary tree: the root is labelled by q^0_n, with two children labelled by q^0_{n-1} and q^1_{n-1}, respectively, and so on until we reach leaves labelled by q^0_0 and q^1_0. To maintain control over the tree structure, we use auxiliary variables p^k which "remember" the parent's label.

[3] In particular, if $\Box \mathcal{B} = \emptyset$ for some $\mathcal{B} \in \mathbf{R}$, then nothing is added to $\mathbf{R}_1$ for the rooted model $\mathcal{B}$.

▶ **Definition 13.** *For every $n \geq 0$, let the open formulas $\varphi_n^*(x)$ be defined as follows. First, for $n \in \mathbb{N}$ set $\theta_n(x)$ to be the formula*

$$\bigwedge_{j\leq 2}\left(q^j_{n+1} \to \bigwedge_{k\leq 2} \Diamond(q^k_n \wedge p^j \wedge x \wedge \neg q^j_{n+1})\right)$$

and define $\varphi_n^*(x) = \bigwedge_{i\leq n}\theta_i(x)$. Then, for $n \geq 0$, let φ_n and $\overline{\varphi_n}$ be defined as:

$$\varphi_n := q^0_{n+1} \wedge \mu x.\varphi^*_n(x);$$
$$\overline{\varphi_n} := q^1_{n+1} \wedge \mu x.\varphi^*_n(x).$$

By the fixed-point theorem for **GL**, we know there are $\mathcal{L}_{\Diamond\diamond}$ formulae $\psi_n \equiv \varphi_n$ over **GL** of size at most $2^{O(n\log(n))}$.[4] Due to the occurrence of $\neg q^j_{n+1}$, the formulae θ_n are equivalent to the formulae θ'_n obtained by substituting $\diamond$ for $\Diamond$. As such, any lower bound results for the size of $\mathcal{L}_{\Diamond\diamond}$ formulae equivalent to the formulae φ_n in **GL** will also produce succinctness results for $\mathcal{L}^{\mu}_{\diamond}$, hence also for **Grz**.

Observe that the formulas $\varphi^*_n(x)$ are all positive over x and modalized. We therefore know by the fixed-point theorem for **GL** that their fixed-point is unique and since they are positive for x, their fixed point will also be equivalent to their greatest and least fixed point.

As promised, the above formulas will define a tree embedding property as this is a sufficient condition a rooted model should satisfy in order for some φ_n to hold in its root. If we are to be more precise, consider the following model $\mathcal{T}_n = \langle \mathcal{T}_n, \prec, V_{\mathcal{T}_n}\rangle$, where $\mathcal{T}_n$ is the set of binary sequences of length $\leq n+1$, rooted at the empty sequence $\langle\rangle$ and:

1. $V_{\mathcal{T}_n}(q^0_{n+1}) := \{\langle\rangle\}$;
2. $V_{\mathcal{T}_n}(q^j_i) = \{s \in \mathcal{T}_n : |s| := n+1-i \wedge s(|s|-1) = j\}$ for $i \leq n$;
3. $V_{\mathcal{T}_n}(p^j) = \{s \in \mathcal{T}_n : \exists\, k\ s := r \frown \langle k\rangle \wedge r \in V_{\mathcal{T}_n}(q^j_{n-|r|})\}$;

where $s \frown r$ denotes the concatenation of the sequences s and r. Intuitively, q^0_i is true on paths of length $n+1-i$ that go "left" on the last step, and q^1_i holds instead when they go "right" on the last step. The truth value of p^j on the world mimics that of q^j_{i+1} on its parent. Note that the lower index of the variables goes "backward" from the root to the leaves.

A *model embedding* is a function $f : \mathcal{M} \to \mathcal{N}$ such that:

1. For every $a, b \in \mathcal{M}$, if $a \sqsubset_M b$ then $f(a) \sqsubset_N f(b)$;
2. For every $p \in \mathbb{P}(\mathcal{M})$ and $a \in \mathcal{M}$, $a \in V_{\mathcal{M}}(p)$ iff $f(a) \in V_{\mathcal{N}}(p)$;

where $\mathbb{P}(\mathcal{M})$ denotes the set of propositional variables occurring in the valuation of some world in $\mathcal{M}$; i.e. $\mathbb{P}(\mathcal{M}) := \{p \in \mathbb{P} : \exists v \in \mathcal{M} v \Vdash p\}$.

By construction, for every n, $\mathcal{T}_n \Vdash \varphi_n$. While we will mostly focus on φ_n in our inductive arguments, down the line we will need models satisfying $\overline{\varphi_n}$ at the root. Models of φ_n and $\overline{\varphi_n}$ vary on whether q^0_n or q^1_n holds at the root, i.e. in φ_n the tree starts on the "left" and in $\overline{\varphi_n}$ on the "right". Hence we use the following notational convention:

Given a rooted model $(\mathcal{A}, a)$, let $N^+(a)$ be the set of children of a, i.e. the set of direct successors of a. Define $\overline{\mathcal{A}} := (\mathcal{A}, \sqsubset_{\mathcal{A}}, V_{\overline{\mathcal{A}}})$ where for $\overline{j} = 1-j$:

- The valuations of q^j_i and $q^{\overline{j}}_i$ are swapped at the root;
- The valuations of p^j and $p^{\overline{j}}$ are swapped at the children of a.

All unmentioned propositional variables will be evaluated the same as before.
Observe that in the case of the tree models we have defined, the symmetric counterpart $\overline{\mathcal{T}_n}$ satisfies $\overline{\varphi_n}$ at its root.

[4] Since the size of φ_n is linear in n, we substitute n for $|\varphi_n|$ in the above expression.

▶ **Lemma 14.** *Let $n \in \mathbb{N}$ and $(\mathcal{A}, a)$ be a rooted model such that for all $b \in \mathcal{A}$, there is at most one pair $\langle i, j\rangle$ such that $b \in V_{\mathcal{A}}(q_i^j)$, and at most one j such that $b \in V_{\mathcal{A}}(p^j)$.*

1. *The model $(\mathcal{A}, a)$ satisfies φ_n iff there is a model embedding $f : \mathcal{T}_n \to \mathcal{A}$ such that $f(\langle\rangle) = a$.*
2. *The model $(\mathcal{A}, a)$ satisfies $\overline{\varphi_n}$ iff there is a model embedding $f : \overline{\mathcal{T}_n} \to \mathcal{A}$ such that $f(\langle\rangle) = a$.*

Proof. We prove both items simultaneously by induction on n.

Left-to-Right. For $n = 0$ assume first that $(\mathcal{A}, a)$ is a model satisfying the assumptions of the Lemma and additionally $\mathcal{A}, a \Vdash \varphi_0$. Thus $\mathcal{A}, a \Vdash q_1^0$ and since $\mathcal{A}, a \Vdash \mu x.\varphi_0^*(x)$, then $\mathcal{A}, a \Vdash \varphi_0^*(\mu x.\varphi_0^*(x))$ and hence for $k \in 2$ there are $b_k \sqsupset a$ such that $\mathcal{A}, b_k \Vdash q_0^k \wedge p^0$. This naturally gives us the embedding from $\mathcal{T}_1$ into $\mathcal{A}$.
Assume that the induction step holds for n and we will show the equivalence for $n+1$. Let $\mathcal{A}, a \Vdash \varphi_{n+1}$, then $\mathcal{A}, a \Vdash q_{n+1}^0$ and $A, a \Vdash \varphi_{n+1}^*(\mu x.\varphi_{n+1}^*(x))$, so there are $b_k \sqsupset a$ for $k \in 2$ such that $\mathcal{A}, b_k \Vdash q_n^k \wedge p^j \wedge \mu x.\varphi_{n+1}^*(x)$ while also $\mathcal{A}, b_k \nVdash q_{n+1}^0$. Since $\mathcal{A}, b_k \nVdash q_{n+1}^0$, it follows that $\mathcal{A}, b_k \Vdash \varphi_{n+1}^*(x)$ iff $\mathcal{A}, b_k \Vdash \varphi_n^*(x)$. Therefore:

$$
\begin{aligned}
&\mathcal{A}, b_k \Vdash \mu x.\varphi_{n+1}^*(x) \Leftrightarrow\\
&\mathcal{A}, b_k \Vdash \varphi_{n+1}^*(\mu x.\varphi_{n+1}^*(x)) \Leftrightarrow\\
&\mathcal{A}, b_k \Vdash \varphi_n^*(\mu x.\varphi_{n+1}^*(x)) \Rightarrow\\
&\mathcal{A}, b_k \Vdash \varphi_n^*(\mu x.\varphi_n^*(x)).
\end{aligned} \tag{1}
$$

The last implication holds by monotonicity. Thus, we can use the induction hypothesis for $b_0{\uparrow}_{\mathcal{A}}$ and $b_1{\uparrow}_{\mathcal{A}}$ to obtain the desired embedding. The case for $\overline{\varphi_n}$ is symmetrical.

Right-to-Left. Let $f : \mathcal{T}_0 \to \mathcal{A}$ be a model embedding; we just need to show that $\mathcal{A}, a \Vdash \varphi_0^*(\mu x.\varphi_0^*(x))$. By our assumption there will be elements $b_k \sqsupset a$ for $k \in 2$ such that $\mathcal{A}, b_k \Vdash q_0^k \wedge p^0$ and $\mathcal{A}, b_k \nVdash q_1^j$ for any $j \in 2$. Therefore $\mathcal{A}, b_k \Vdash \varphi_0^*(\bot)$ and since x occurs positively in φ_0^*, it follows that $\mathcal{A}, b_k \Vdash \varphi_0^*(\mu x.\varphi_0^*(x))$. Hence $\mathcal{A}, a \Vdash \varphi_0^*(\varphi_0^*(\mu x.\varphi_0^*(x)))$, and so $\mathcal{A}, a \Vdash \mu x.\varphi_0^*(x)$.
Assume now that the induction step holds for n and let $f : \mathcal{T}_{n+1} \to \mathcal{A}$ be a model embedding. Thus $\mathcal{A}, a \Vdash q_{n+1}^0$ and there are $b_k \sqsupset a$ for $k \leq 2$ and model embeddings $g_0 : \mathcal{T}_n \to b_0{\uparrow}_{\mathcal{A}}$, $g_1 : \overline{\mathcal{T}_n} \to b_1{\uparrow}_{\mathcal{A}}$. By the induction hypothesis, $\mathcal{A}, b_0 \Vdash \varphi_n$ and $\mathcal{A}, b_1 \Vdash \overline{\varphi_n}$. Item 2 is proved similarly. ◀

Our models will be designed so that they have a *critical branch* on which Hercules will be forced to play. This critical branch is described using a "successor" function which goes up the tree along said branch. This and other technical notions needed to describe Hercules' strategy are given by the following definition.

▶ **Definition 15.** *Given a rooted model $(\mathcal{A}, a)$ and a propositional variable p, we will denote by $\mathcal{A}(p)$ the model $\langle \mathcal{A}, \sqsubset_A, V'\rangle$ where V' is the same as $V_{\mathcal{A}}$ except that p will also hold in the root of the model. In the interest of distinguishing the root of $\mathcal{A}(p)$ from $\mathcal{A}(q)$ we will write them as $a(p)$ and $a(q)$, respectively (even though it's technically the same element).*

A model with successors *is a model $\mathcal{A}$ equipped with a partial function $S_{\mathcal{A}} : \mathcal{A} \to \mathcal{A}$ such that $S_{\mathcal{A}}(a)$ is always a child of a.*

If $\mathcal{A}$ is a rooted model with successors, the critical branch *of $\mathcal{A}$ is the maximal path $\vec{w} = (w_i)_{i \leq m}$ such that w_0 is the root of $\mathcal{A}$ and $w_{i+1} = S_{\mathcal{A}}(w_i)$ for all $i < m$; we say that m is the* critical height *of $\mathcal{A}$.*

We denote by $S[\mathcal{A}]$ *the generated submodel of* $S_{\mathcal{A}}(w_0)$ *with* $S_{\mathcal{A}}(w_0)$ *as its root and its induced successor function being* $S_{S[\mathcal{A}]}(w_0)$. *For a natural number* r, *we define the* r*th iteration of* $S_{\mathcal{A}}$ *by induction so that* $S^{(0)}[\mathcal{A}] := \mathcal{A}$, $S^{(0)}_{\mathcal{A}}(w) := w$ *and on the inductive step* $S^{(r+1)}[\mathcal{A}] := S[S^{(r)}[\mathcal{A}]]$, $S^{(r+1)}_{\mathcal{A}}(w) := S_{\mathcal{A}}(S^{(r)}_{\mathcal{A}}(w))$.

The partial function $S_{\mathcal{A}}$ will not be used in the semantics, but it will help us to describe Hercules' strategy. We will begin by defining sets of rooted models $\mathbf{A}^n$ and $\mathbf{B}^n$ recursively on n containing 2^{n+1} models each. Of those, the former 2^n will be used to prove our succinctness lower bound while the latter 2^n are auxiliary[5] and used solely in our recursive construction. The following definition is illustrated in Figures 1 and 2.

▶ **Definition 16.** *First, for* $n = 0$ *and for* $i < 2$ *we define* $\mathcal{A}^0_i$ *with domains the sequences* s *of length at most* 1 *on the natural numbers* $\{0, 1, 2\}$. *Then* $\sqsubset_{\mathcal{A}^0_i}$ *is the prefix relation and valuations* $V_{\mathcal{A}^0_i}(p^i) = \{\langle k\rangle : k \le 2\}$, $V_{\mathcal{A}^0_i}(q^i_1) = \{\langle\rangle\}$ *and* $V_{\mathcal{A}^0_i}(q^j_0) = \{\langle j\rangle\}$ *for* $j < 2$. *Set* $S_{\mathcal{A}^0_i}(\langle\rangle) = \langle 1 - i\rangle$ *and* $S_{\mathcal{A}^0_i}$ *is undefined otherwise. The* $\mathcal{B}^0_i$ *have as domain the sequences* s *of length at most* 1 *in the natural numbers* ≤ 1. *Their relations are the prefix relations and the valuations are the following:* $V_{\mathcal{B}^0_i}(p^i) = \{\langle k\rangle : k \le 1\}$, $V_{\mathcal{B}^0_i}(q^i_1) = \{\langle\rangle\}$ *and* $V_{\mathcal{B}^0_i}(q^i_0) = \{\langle 0\rangle\}$. *Set* $S_{\mathcal{B}^0_i} = \{(\langle\rangle, \langle 1\rangle)\}$.
For $2 \le i < 4$, *we let* $\mathcal{A}^0_i := \overline{\mathcal{A}^0_{i-2}}$ *and* $\mathcal{B}^0_i := \overline{\mathcal{B}^0_{i-2}}$. *Their successor functions remain the same. Now, given* n, *we will define the models* $\mathcal{A}^{n+1}_i$ *and* $\mathcal{B}^{n+1}_i$ *with a case distinction in* i. *In this paper, given models* $\mathcal{A}$ *and* $\mathcal{B}$, *we let* $\mathcal{A} \amalg \mathcal{B}$ *be the model with domain the disjoint union of* $\mathcal{A}$ *and* $\mathcal{B}$,[6] *accessibility relation* $\sqsubset_{\mathcal{A}\amalg\mathcal{B}}$ *being the disjoint union of the accessibility relations* $\sqsubset_{\mathcal{A}}$ *and* $\sqsubset_{\mathcal{B}}$, *and similarly the valuation. This is the set of all the elements in* $\mathcal{A}$ *and* $\mathcal{B}$ *with each element labelled by the set to which it belongs.*

Case $i < 2^{n+1}$. *First, set* $\mathcal{X}$ *to be*

$$\Big(\coprod_{k=0}^{2^{n+1}-1} \mathcal{A}^n_k(p^1)\Big) \amalg \mathcal{A}^n_i(p^0) \amalg \overline{\mathcal{A}^n_i}(p^0) \amalg \mathcal{B}^n_i(p^0),$$

and construct $\mathcal{A}^{n+1}_i$ *by adding a (fresh) irreflexive root* a^{n+1}_i *which is below all elements of* $\mathcal{X}$ *and satisfies* q^0_{n+2}. *We set*

$$S_{\mathcal{A}^{n+1}_i} := S_{\mathcal{A}^n_i} \cup \{(a^{n+1}_i, a^n_i(p^0))\}.$$

The models $\mathcal{B}^{n+1}_i$ *are defined similarly by setting*

$$\mathcal{Y} = \Big(\coprod_{k=0}^{2^{n+1}-1} \mathcal{A}^n_k(p^1)\Big) \amalg \overline{\mathcal{A}^n_i}(p^0) \amalg \mathcal{B}^n_i(p^0),$$

and, as before, adding an irreflexive root b^{n+1}_i *that satisfies* q^0_{n+2}. *Set*

$$S_{\mathcal{B}^{n+1}_i} := S_{\mathcal{B}^n_i} \cup \{(b^{n+1}_i, b^n_i(p^0))\}.$$

Case $2^{n+1} \le i < 2^{n+2}$. *Here we construct our auxiliary models where* $\mathcal{A}^{n+1}_i = \overline{A^{n+1}_{i-2^{n+1}}}$ *and* $\mathcal{B}^{n+1}_i = \overline{B^{n+1}_{i-2^{n+1}}}$.

We remark that the auxiliary models were only used to help us inductively construct the $\mathcal{A}^n_i$ and $\mathcal{B}^n_i$ models. Letting $\mathbf{A}^n = \{\mathcal{A}^n_i : i < 2^n\}$, $\mathbf{B}^n = \{\mathcal{B}^n_i : i < 2^n\}$ and $\mathbf{M}_n = \mathbf{A}^n \cup \mathbf{B}^n$, we will study the $(\varphi_n, \mathbf{M}_n)$-MEG where φ_n are the formulae in Definition 13.

[5] They are the “right” versions of the former models.
[6] Hence elements in the intersection $\mathcal{A} \cap \mathcal{B}$ will appear twice; once labelled by $\mathcal{A}$ and once labelled by $\mathcal{B}$.

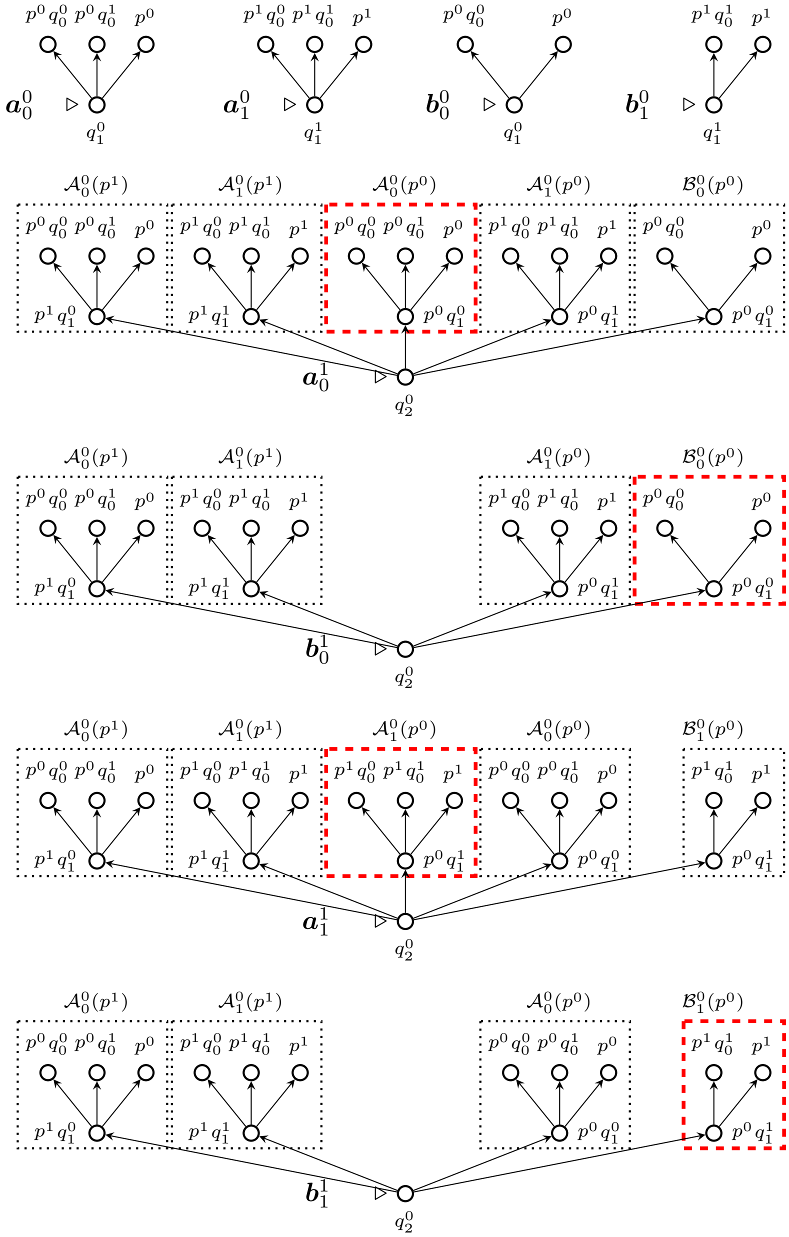

Figure 1 The figure illustrates models in $\mathbf{M}_0$ and $\mathbf{M}_1$. At the very top, we see the rooted models of $\mathbf{A}^0, \mathbf{B}^0$ as well as their auxiliary models. These are then used in the construction of the models of $\mathbf{A}^1, \mathbf{B}^1$ that we can see in the rest of the graphic.

Each copy of the smaller models being used in the construction is indicated by a box and a label. It is easy to see that $\mathcal{A}_0^0$ embeds $\mathcal{T}_0$ (with the image being all but the rightmost leaf) and $\mathcal{A}_1^0$ embeds $\overline{\mathcal{T}}_0$, but $\mathcal{B}_0^0$ and $\mathcal{B}_1^0$ do not, hence $\mathcal{A}_0^0$ satisfies φ_0 while $\mathcal{B}_0^0$ does not, and similarly $\mathcal{A}_1^0$ satisfies $\overline{\varphi}_0$ while $\mathcal{B}_1^0$ does not. A similar analysis shows that e.g. $\mathcal{A}_0^1$ satisfies φ_1 but $\mathcal{B}_0^1$ does not.

The successor function will point from the root towards the sub-model in the red, dashed boxes; for example, $S[\mathcal{A}_0^1] = \mathcal{A}_0^0(p^0)$. The models $S[\mathcal{A}_0^1]$ and $S[\mathcal{B}_0^1]$ are similar enough that Hercules is not able to tell them apart before reaching their topmost worlds. The remaining branches are there to make this task as difficult as possible.

6 The lower bound

In this section, we show that if Hydra sets up the playing field with $\mathbf{M}^n$, then Hercules cannot win the game in fewer than 2^n moves. This is our main technical result and will require several preparatory lemmas.

As mentioned, each model has a critical branch. More specifically, each $\mathcal{A}_i^{n+1}$ is very similar to the respective $\mathcal{B}_i^{n+1}$, and can only be distinguished by Hercules if he plays along the critical branch. The number i can be seen as coding a binary string simply by writing $i = e_{n+1}2^{n+1} + \ldots + e_0 2^0$ in binary. Then each e_{n+1-r} indicates whether the critical branch goes left or right at step i; note that this includes step 0, corresponding to the label of the root, i.e. the critical branch of $\mathcal{A}_i^{n+1}$ goes left first while that of $\overline{\mathcal{A}_i^{n+1}}$ goes right.

▶ **Lemma 17.** *For all n, $i < 2^{n+2}$ and $r < n+2$, the following hold:*
1. *$S^{(r)}[\mathcal{A}_i^{n+1}]$ is isomorphic to $\mathcal{A}_k^{n+1-r}(p^e)$;*
2. *$S^{(r)}[\mathcal{B}_i^{n+1}]$ is isomorphic to $\mathcal{B}_i^{n+1-r}(p^e)$;*

where $k \equiv i \mod 2^{n+2-r}$ with $k < 2^{n+2-r}$ and e is the digit $e_{n+2-r+1}$ in the binary expansion of i.

Proof. We will only prove Item 1 as the case for the $\mathcal{B}_i^{n+1}$ model is identical. The case for $r = 1$ is immediate from the definitions of the models and the S function. So assume that the lemma holds for $r < n+1$. By the induction hypothesis, $S^{(r+1)}[\mathcal{A}_i^{n+1}]$ is isomorphic to $S[\mathcal{A}_k^{n+1-r}(p^{e_{n+1-r}})]$, which by the induction hypothesis for $r = 1$, this is in turn isomorphic to $\mathcal{A}_k^{n+1-r}(p^{e_{n-r}})$. ◀

By construction, $\mathcal{T}_n$ embeds into $\mathcal{A}_i^n$ but not into $\mathcal{B}_i^n$, yielding the following.

▶ **Lemma 18.** *For all n and all $i < 2^n$, $\mathcal{A}_i^n \Vdash \varphi_n$ and $\mathcal{B}_i^n \Vdash \neg\varphi_n$.*

Proof. For this proof, we will extend our definition of $\bar{\cdot}$ into embeddings as follows: Given an embedding $f : \mathcal{A} \to \mathcal{B}$, we define $\overline{f} : \overline{\mathcal{A}} \to \overline{\mathcal{B}}$ to be such that $\overline{f}(a) = f(a)$ for all $a \in A$. Notice that $\overline{f}$ will still be an embedding if it preserves the root r and maps children of the root of A into children of the root of B (i.e. $f[N^+(r)] \subseteq N^+(f(r))$).

We will make use of Lemma 14. First, we show $\mathcal{A}_i^n \Vdash \varphi_n$ by proving the existence of embeddings $f_i^n : \mathcal{T}_n \to \mathcal{A}_i^n$ for $i < 2^n$ by induction on n. Notice that since $\mathcal{A}_i^n$ and $\mathcal{T}_n$ have the same depth, we will also obtain $f_i^n[N^+(\langle\rangle)] \subseteq N^+(f_i^n(\langle\rangle))$; thus, $\overline{f_i^n}$ will also be an embedding from $\overline{\mathcal{T}_n}$ to $\mathcal{A}_{i+2^n}^n = \overline{\mathcal{A}_i^n}$.
For $n = 0$, one needs only look at the definition of $\mathcal{A}_0^0$. Now assume that the statement holds for n, let $i < 2^{n+1}$ be arbitrary, $j \equiv i \mod 2^n$ and $f_j^n : \mathcal{T}_n \to \mathcal{A}_j^n$ be an embedding given from our induction hypothesis. We can then define $f_i^{n+1} : \mathcal{T}_{n+1} \to \mathcal{A}_i^{n+1}$ by mapping:
- $\langle\rangle$ to a_i^{n+1};
- $\langle 0 \rangle \frown s$ to the point corresponding to $f_j^n(s)$ on the copy of $\mathcal{A}_j^n(p^0)$ in $\mathcal{A}_i^{n+1}$;
- $\langle 1 \rangle \frown s$ to the point corresponding to $\overline{f_j^n}(s)$ on the copy of $\overline{\mathcal{A}_j^n}(p^0)$ in $\mathcal{A}_i^{n+1}$.

$\mathcal{A}_i^{n+1}$ satisfies the base conditions of Lemma 14 and thus our induction step holds for all $i < 2^{n+1}$ and so $\mathcal{A}_i^n \Vdash \varphi_n$.

We now show $\mathcal{B}_i^n \Vdash \neg\varphi_n$ by proving that there are no embeddings as in Lemma 14 $f : \mathcal{T}_n \to \mathcal{B}_i^n$ for all $i < 2^n$ by induction on n.
For $n = 0$, this is clear. Assume that the induction hypothesis holds for n, and suppose, for a contradiction, that there is an embedding $f : \mathcal{T}_n \to \mathcal{B}_i^{n+1}$ with $f(\langle\rangle) = b_i^{n+1}$ for some $i < 2^n$. As f is an embedding, $\{q_{n+1}^0, p^0\} \subseteq V_{\mathcal{B}_i^{n+1}}^{-1}(f(\langle 0 \rangle))$ and the only world that satisfies this condition in $\mathcal{B}_i^{n+1}$ is the root b_i^n of the $\mathcal{B}_i^n(p^0)$ part of the model. This implies the existence of a root-preserving embedding $f' : \mathcal{T}_n \to \mathcal{B}_i^n$ which contradicts our induction hypothesis. The case for $2^n \leq i < 2^{n+1}$ is done in a similar way. ◀

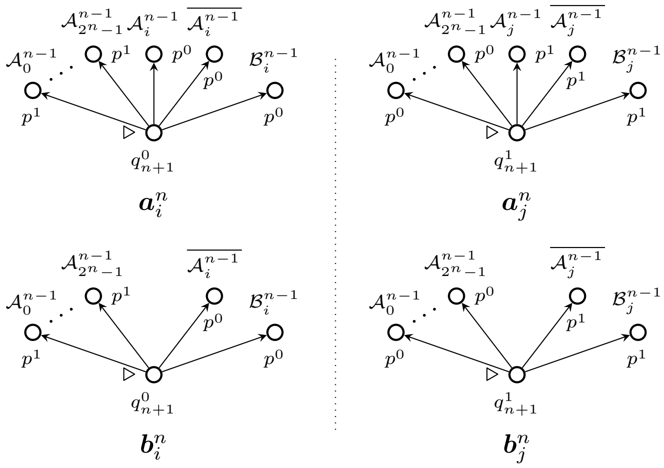

Figure 2 The inductive structure of the models for $n > 0$ with the models of $\mathcal{A}^n$ and $\mathcal{B}^n$ to the left and the auxiliary models to the right ($0 \leq i < 2^n \leq j < 2^{n+1}$).

Hence, Hydra can set up the playing field by placing the models $\mathcal{A}^n_i$ on the left and the models $\mathcal{B}^n_i$ on the right. In this case, Hercules requires exponentially many moves to win the game.

▶ **Definition 19.** *Suppose that $\mathcal{M}$, $\mathcal{N}$ are two finite rooted models with successors. We say that $r \in \mathbb{N}$* distinguishes *$\mathcal{M}$ and $\mathcal{N}$ if $S^{(r)}[\mathcal{M}]$ and $S^{(r)}[\mathcal{N}]$ differ on the truth of a propositional variable at their roots, but whenever $i < r$, then $S^{(i)}[\mathcal{M}]$ and $S^{(i)}[\mathcal{N}]$ agree on the truth of all propositional variables at their respective roots. We call r the* distinguishing value *of $\mathcal{M}$ and $\mathcal{N}$.*

Note that the distinguishing value of two models $\mathcal{M}, \mathcal{N}$ need not be defined, but when it is, it is unique. Moreover, the distinguishing values of the models we have constructed always exist.

▶ **Lemma 20.** *Fix $n \geq 1$ and $0 \leq i < j < 2^{n+1}$. Then, $\mathcal{A}^n_i$ and $\mathcal{A}^n_j$ are distinguished at some $r < n$, satisfying the following properties:*

(a) *If $i < 2^n$ and $2^n \leq j$, then $\mathcal{A}^n_i$ and $\mathcal{A}^n_j$ have distinguishing value 0.*

(b) *If $\mathcal{A}^n_i$ and $\mathcal{A}^n_j$ have distinguishing value r, then $\mathcal{A}^{n+1}_i$, $\mathcal{A}^{n+1}_j$ have distinguishing value $r+1$. The same holds for $\mathcal{A}^{n+1}_{2^{n+1}+i}$, $\mathcal{A}^{n+1}_{2^{n+1}+j}$.*

Proof. Item a is immediate since the roots of $\mathcal{A}^n_i$ and $\mathcal{A}^n_j$ are evaluated differently. For Item b, observe that $i, j < 2^{n+1}$ implies that $\mathcal{A}^{n+1}_i$ and $\mathcal{A}^{n+1}_j$ have roots with the same valuation. Thus, they are distinguished at $r+1$. Since $\mathcal{A}^{n+1}_{2^{n+1}+i} = \overline{\mathcal{A}^{n+1}_i}$ and $\mathcal{A}^{n+1}_{2^{n+1}+j} = \overline{\mathcal{A}^{n+1}_j}$, these are also distinguished at $r+1$. ◀

▶ **Lemma 21.** *Fix n and $i < 2^{n+1}$. Then, $\mathcal{A}^n_i$ and $\mathcal{B}^n_i$ are distinguished at $n+1$.*

Proof. By an easy induction on n. ◀

By *twins of height k* we mean a pair of the form $(S^{(k)}[\mathcal{A}^n_i], S^{(k)}[\mathcal{B}^n_i])$, where $i < 2^n$ and $k \leq n+1$. If $\mathbf{L}$ is a set of rooted models from $\mathbf{A}^n$ and $\mathbf{R}$ a set of rooted models from $\mathbf{B}^n$, we say that there are twins of height k in $\mathbf{L} \circ \mathbf{R}$ if there are twins $(S^{(k)}[\mathcal{A}^n_i], S^{(k)}[\mathcal{B}^n_i])$ such that $S^{(k)}[\mathcal{A}^n_i] \in \mathbf{L}$ and $S^{(k)}[\mathcal{B}^n_i] \in \mathbf{R}$.

We will study how pairs of the form $(S^{(k)}[\mathcal{A}_i^n], S^{(r)}[\mathcal{B}_i^n])$ in $\mathbf{L}(\eta) \circ \mathbf{R}(\eta)$ affect the viability of the various modal moves for Hercules. This will place restrictions on the relationship between k and r. For example, we see that Item b below states that if $k < r$ then Hercules couldn't play any $\boxdot$ moves as Hydra can get from $S^{(k)}[\mathcal{A}_i^n]$ into a model isomorphic to any choice of Hercules in $\boxdot S^{(r)}[\mathcal{B}_i^n]$. Clearly, that also excludes any $\Box$ moves for Hercules, as two isomorphic models are of the same height and hence any isomorphic model the Hydra would produce for the corresponding $\boxdot$ move, would also show up in a $\Box$ move. In Lemma 22, all of the restrictions applying to a reflexive modality will not only just apply to the irreflexive modality, but also, they will apply even if we substitute $\leq$ for $<$.

At this point, let us fix $n \geq 0$ and assume that Hydra labels the root with $\mathbf{A}^n \circ \mathbf{B}^n$.

▶ **Lemma 22.** *For any node η in a closed game tree $(T, \prec)$ for the $(\varphi_n, \mathbf{GL})$-MEG*

(a) *If there are twins $S^{(k)}[\mathcal{A}_i^n], S^{(k)}[\mathcal{B}_i^n]$ in $\mathbf{L}(\eta) \circ \mathbf{R}(\eta)$ with $k < n+1$ then no literal move was played in the node η.*

(b) *If there are $S^{(k)}[\mathcal{A}_i^n], S^{(r)}[\mathcal{B}_i^n]$ in $\mathbf{L}(\eta) \circ \mathbf{R}(\eta)$ and $k < r$ then $\boxdot$ was not played in the node η. If $k \leq r$ then no $\Box$ move was played in η either.*

(c) *If there are twins $S^{(k)}[\mathcal{A}_i^n], S^{(k)}[\mathcal{B}_i^n]$ in $\mathbf{L}(\eta) \circ \mathbf{R}(\eta)$ and a $\boxdot$ move was played in the node η, then $S^{(k)}[\mathcal{B}_i^n]$ was chosen.*

(d) *If there are twins $S^{(k)}[\mathcal{A}_i^n], S^{(k)}[\mathcal{B}_i^n]$ in $\mathbf{L}(\eta) \circ \mathbf{R}(\eta)$ and a $\diamonddot$ move was played in the node η, then either $S^{(k)}[\mathcal{A}_i^n]$ or $S^{(k+1)}[\mathcal{A}_i^n]$ was chosen. Hence, if a $\Diamond$ move was played in η, then $S^{(k+1)}[\mathcal{A}_i^n]$ was chosen.*

(e) *If there are two distinct twins $S^{(k)}[\mathcal{A}_i^n], S^{(k)}[\mathcal{B}_i^n]$ and $S^{(r)}[\mathcal{A}_j^n], S^{(r)}[\mathcal{B}_j^n]$ in $\mathbf{L}(\eta) \circ \mathbf{R}(\eta)$, then*

  (i) *if $r+1 < k$, then no $\diamonddot$ or $\boxdot$ move was played in η. Similarly, if $r < k$, then no $\Diamond$ or $\Box$ move was played in η.*

  (ii) *If $\mathcal{A}_i^n, \mathcal{B}_j^n$ are distinguished at r and $r = k$ then no $\Diamond$ move was played in η. If instead $k = r+1$, then no $\diamonddot$-move was played in η either.*

Proof. Item a is immediate by Lemma 21 as no literal move can be played if there are models $(\mathcal{A}, \mathcal{B}) \in \mathbf{L}(\eta) \circ \mathbf{R}(\eta)$ with distinguishing value $r > 0$.

For Item b, if $r = n+1$, then the statement is clear by the definition of the $\mathcal{B}_j^0$. Thus, assume $r < n+1$. Then observe that $S^{(r)}[\mathcal{B}_i^n] \subseteq \boxdot S^{(k)}[\mathcal{A}_i^n]$ and hence $\boxdot S^{(r)}[\mathcal{B}_i^n] \subseteq \boxdot S^{(k)}[\mathcal{A}_i^n]$. Hence, no matter where $S^{(r)}[\mathcal{B}_i^n]$ is mapped in $\boxdot S^{(r)}[\mathcal{B}_i^n]$ by Hercules, Hydra will include an isomorphic model in its response. Observe that if $k = r$ then we still have that $\Box S^{(k)}[\mathcal{B}_i^n] \subseteq \Box S^{(k)}[\mathcal{A}_i^n]$.

In Item c, since $\boxdot \mathcal{A} = \Box \mathcal{A} \cup \{\mathcal{A}\}$ for any model $\mathcal{A}$, we get that $\Box S^{(k)}[\mathcal{B}_i^n] \subseteq \boxdot S^{(k)}[\mathcal{A}_i^n]$ and thus only $S^{(k)}[\mathcal{B}_i^n]$ can be used by Hercules.

Moving into Item d, the cases for $k = n+1$ and for $n = 0$ are trivial; therefore, let $k < n+1$ and $0 < n$. By Lemma 17 it is sufficient to prove the statement for $k = 0$.[7] Now assume, aiming towards a contradiction, that neither of $\mathcal{A}_i^n$ and $S[\mathcal{A}_i^n]$ were chosen by Hercules. We will show that all of the remaining alternatives are isomorphic to some model in $\boxdot \mathcal{B}_i^n$. Let us assume that $i < 2^n$, then Hercules could not choose any model in $\boxdot\{\mathcal{A}_j^{n-1}(p^1), \overline{\mathcal{A}_i^{n-1}}(p^0), \mathcal{B}_i^{n-1}(p^0)\}$ as all those models belong by definition in $\mathcal{B}_i^n$. Finally, Hercules could not have chosen a model in $\Box \mathcal{A}_i^{n-1}(p^0)$ since $\Box \mathcal{A}_i^{n-1}(p^0) = \Box \mathcal{A}_i^{n-1}(p^1) \subseteq \boxdot \mathcal{B}_i^n$.

[7] Formally we should prove it for $\mathcal{A}_i^n(p^e)$, $\mathcal{B}_i^n(p^e)$, but the proof is otherwise identical.

In Item e-i, we can assume by Lemma 17 that $r = 0$ and $k \geq 2$, and we will show that $S^{(2)}[\mathcal{A}_i^n] \in \boxdot\mathcal{B}_j^n$. This will imply that $\boxdot S^{(2)}[\mathcal{A}_i^n] \subseteq \boxdot\mathcal{B}_j^n$, thus giving us $S^{(k)}[\mathcal{A}_i^n] \in \boxdot\mathcal{B}_j^n$. Since $i \neq j$, $\mathcal{A}_i^{n-1}(p^e)$ is one of the model branches of $\mathcal{B}_j^n$ by definition for some e. Then $S^{(2)}[\mathcal{A}_i^n] = S[\mathcal{A}_i^{n-1}(p^e)] \in \boxdot\mathcal{B}_j^n$.

Finally, we prove Item e-ii. The case for $k = n+1$ is trivial by the definition of the game. Since $r = k < n+1$ is the distinguishing value, by Lemma 17 it follows that $S^{(k)}[\mathcal{A}_i^n]$ is isomorphic to $\mathcal{A}_{i'}^{n-k}(p^e)$ and $S^{(k)}[\mathcal{A}_j^n]$ is isomorphic to $\mathcal{A}_{j'}^{n-k}(p^e)$ for some i, j, e. Thus we can without loss of generality assume that $r = k = 0$. We can assume without loss of generality that $i < 2^n \leq j$, then by definition $\mathcal{A}_i^{n-1} \in \boxdot\mathcal{B}_i^n$. ◀

▶ **Definition 23.** *In the closed game tree $(T, \prec)$, let $\Lambda(i)$ be the set of leaves η such that*

1. *for every $\eta' \preceq \eta$ there is $r \geq 0$ such that the twins $(S^{(r)}[\mathcal{A}_i^n], S^{(r)}[\mathcal{B}_i^n])$ appear in $\mathbf{L}(\eta') \circ \mathbf{R}(\eta')$ and,*
2. *for every $\zeta \prec \eta$, every child σ of ζ with $\zeta \prec \sigma \preceq \eta$ and every other child σ' of ζ, if $S^{(r)}[\mathcal{A}_i^n], S^{(r)}[\mathcal{B}_i^n]$ are in $\mathbf{L}(\sigma') \circ \mathbf{R}(\sigma')$ then $S^{(k)}[\mathcal{A}_i^n], S^{(k)}[\mathcal{B}_i^n]$ are in $\mathbf{L}(\sigma) \circ \mathbf{R}(\sigma)$ for some $k < r$.*

More informally, if $\eta \in \Lambda(i)$ then the path to η from the root is exactly the path that by *Condition* 2 "locally" minimises the height r of the $(S^{(r)}[\mathcal{A}_i^n], S^{(r)}[\mathcal{B}_i^n])$ twins is chosen.

▶ **Lemma 24.** *For a closed $(\varphi_n, \mathbf{GL})$-MEG game tree $(T, \prec)$ in which the Hydra plays optimally, the following hold:*

(a) $\forall\, i < 2^n$, $\Lambda(i) \neq \emptyset$;

(b) $\forall\, i < 2^n\ \forall\, \eta \in \Lambda(i)\ \forall\, k \leq n+1$ *there is a $\zeta \preceq \eta$ such that k is least with the property that $S^{(k)}[\mathcal{A}_i^n], S^{(k)}[\mathcal{B}_i^n]$ are in $\mathbf{L}(\zeta) \circ \mathbf{R}(\zeta)$;*

(c) *if $0 \leq i < j < 2^n$, then $\Lambda(i) \cap \Lambda(j) = \emptyset$.*

Proof. We will prove each Item by contradiction, starting with Item a. Assume otherwise and let ζ be a maximal node of $(T, \prec)$ for which Conditions 1 and 2 of Definition 23 hold. Let k be the least natural number such that $(S^{(k)}[\mathcal{A}_i^n], S^{(k)}[\mathcal{B}_i^n]) \in \mathbf{L}(\zeta) \circ \mathbf{R}(\zeta)$. As ζ is not a leaf, Hercules has not played a literal move. If Hercules plays a $\vee$-move, then at least one of the two children of ζ, call it ζ', will have $S^{(k)}[\mathcal{A}_i^n] \in \mathbf{L}(\zeta')$. But since $\mathbf{R}(\zeta) = \mathbf{R}(\zeta')$, we get that $(S^{(k)}[\mathcal{A}_i^n], S^{(k)}[\mathcal{B}_i^n]) \in \mathbf{L}(\zeta') \circ \mathbf{R}(\zeta')$ and this is the least k with such property since $\mathbf{L}(\zeta') \subseteq \mathbf{L}(\zeta)$. The case for the $\wedge$-move comes contrary to our maximality assumption for ζ in the same way as the $\vee$-move case. If Hercules plays a $\Diamond$ or a $\diamonddot$-move, then by Lemma 22 $(S^{(k)}[\mathcal{A}_i^n], S^{(k)}[\mathcal{B}_i^n])$ or $(S^{(k+1)}[\mathcal{A}_i^n], S^{(k+1)}[\mathcal{B}_i^n])$ will belong to $\mathbf{L}(\zeta') \circ \mathbf{R}(\zeta')$ with ζ' being the only child of ζ. This satisfies Condition 1, while Condition 2 holds trivially; this contradicts our minimality assumption for ζ. By Lemma 22, Hercules has not played a $\Box$-move, and this leaves only the $\boxdot$-move. In this case, Lemma 22 once more dictates that $(S^{(k)}[\mathcal{A}_i^n], S^{(k)}[\mathcal{B}_i^n])$ will belong in the only child of ζ. As with the $\diamonddot$-case, this comes contrary to the minimality condition for ζ. As such, we have reached a contradiction and so $\forall i < 2^n$, $\Lambda(i) \neq \emptyset$.

Now for Item b, assume otherwise for some η and k and let $\zeta \preceq \eta$ be greatest with an $r < k$ where $(S^{(r)}[\mathcal{A}_i^n], S^{(r)}[\mathcal{B}_i^n]) \in \mathbf{L}(\zeta) \circ \mathbf{R}(\zeta)$. Clearly $\zeta \prec \eta$ as otherwise a literal move could be played, something only possible if $n+1 = r < k \leq n+1$. Thus there is a child $\sigma \preceq \eta$ of ζ. By Lemma 22 and by the definition of $\Lambda(i)$, no matter what move Hercules performs, $(S^{(s)}[\mathcal{A}_i^n], S^{(s)}[\mathcal{B}_i^n]) \in \mathbf{L}(\sigma) \circ \mathbf{R}(\sigma)$ for some $s \leq k$, contradicting either the maximality of ζ or our original assumption.

Finally we prove Item c. Assume the statement doesn't hold, then there are $i < j < 2^n$ such that $\eta \in \Lambda(i) \cap \Lambda(j)$ for some leaf $\eta \in T$. By Lemma 20 $\mathcal{A}_i^n$ and $\mathcal{A}_j^n$ have some distinguishing value $m \leq n$. Let $\zeta \prec \eta$ be maximal such that the least r and k with

$$(S^{(r)}[\mathcal{A}_i^n], S^{(r)}[\mathcal{B}_i^n]), (S^{(k)}[\mathcal{A}_j^n], S^{(k)}[\mathcal{B}_j^n]) \in \mathbf{L}(\zeta) \circ \mathbf{R}(\zeta) \tag{2}$$

are such that $r \leq m$ or $k \leq m$. Assume that $r = m \leq k$; it will always be the case that one of them will be m from the proof of Item b. Clearly no $\vee$ or $\wedge$-move could have been played at ζ, as then the child $\sigma \prec \eta$ of ζ would either invalidate the choice of ζ, or otherwise, σ would then be a witness of a failure of Condition 2 for η, thus invalidating the assumption $\eta \in \Lambda(i) \cap \Lambda(j)$. Similarly, no $\boxdot$-move could have been played either as it would contradict our choice of the node ζ. If a $\diamonddot$-move was played at ζ then by Lemma 22 it has to be that $r = k$. Hence, the choice of ζ will be violated as (2) will also hold for its child $\sigma \preceq \eta$. Finally, the $\Diamond$ and $\Box$-moves are simpler cases of the $\diamonddot$ and $\boxdot$-moves. ◀

It follows that any closed game tree has at least 2^n leaves, yielding our main technical result.

▶ **Proposition 25.** *For every $n \geq 1$, Hercules has no winning strategy of less than $2^n + 2$ moves on the $(\varphi_n, \mathbf{M}_n)$-MEG.*

Proof. Let $(T, \prec)$ be a closed game tree of $(\varphi_n, \mathbf{M}_n)$-MEG. By Lemma 24, the sets of leaves $\{\Lambda(i) : i < 2^n\}$ are non-empty and disjoint. Since each leaf represents a literal move being played, Hercules must have played at least 2^n literal moves. As there are at least 2^n leaves, at least one branching move must have been played. Furthermore, by the definition of $\Lambda(i)$, at least one modal rule must have been played. As such, Hercules must have played at least $2^n + 2$ moves. ◀

7 Succinctness

This section contains the main results of our study, first showing an exponential succinctness result in a wide range of Kripke frames. An example of such an application can be found in [12]. We will then examine some additional benefits we can obtain from the connection of the interpolation and the fixed point theorems of **GL**.

7.1 Succinctness of definable fixed points

Proposition 25 is a powerful tool for proving succinctness results. In general, succinctness results for a class of models apply to any larger class. Thus, our results apply not only to **GL** models, but also to a wide range of classes of Kripke models.

▶ **Theorem 26.** *Let $\mathbf{C}$ be any class of Kripke models containing all finite **GL** models or all finite **Grz** models. Then, there is a sequence of formulas $(\varphi_0(x), \varphi_1(x), \dots)$ which are both modalized and positive on x with $|\varphi_i(x)| = O(i)$ such that for any $i \in \mathbb{N}$ and any $\psi \in \mathcal{L}_{\Diamond\diamonddot}$, if $\varphi_i(\psi) \equiv \psi$ over $\mathbf{C}$, then $|\psi| \geq 2^i$.*

Proof. The sequence consists of the formulae φ_i^* in Definition 13. Counting the symbols present in the formulas, we get, by a simple induction, $|\varphi_n^*| \leq 41 \cdot n$. Assume, towards a contradiction, that there is some $\psi_n \in \mathcal{L}_\Diamond$ such that $\psi \equiv \mu x.\varphi_n^*(x)$ over $\mathbf{C}$ and $|\psi_n| < 2^n$. But $\psi_n' := q_n^0 \wedge \psi_n \equiv \varphi_n$ by Definition 13, making it an $\mathcal{L}_\Diamond$ equivalent to φ_n of size $< 2^n + 2$. However, by Theorem 11, this contradicts Proposition 25. ◀

▶ **Corollary 27.** *Let* $\mathbf{C}$ *be any class of Kripke models containing all finite* $\mathbf{GL}$ *models or all finite* $\mathbf{Grz}$ *models. Then, the languages* $\mathcal{L}^{\mu}_{\Diamond}$ *and* $\mathcal{L}^{\mu}_{\diamondsuit}$ *are exponentially more succinct than* $\mathcal{L}_{\Diamond\diamondsuit}$ *over* $\mathbf{C}$*. To be precise, for* $\mathcal{L}^{\mu} \in \{\mathcal{L}^{\mu}_{\Diamond}, \mathcal{L}^{\mu}_{\diamondsuit}\}$*, there is a sequence of* $\mathcal{L}^{\mu}$ *formulas* $(\varphi_0, \varphi_1, \dots)$ *with* $|\varphi_i| = O(i)$ *such that for any* $i \in \mathbb{N}$ *and any* $\psi \in \mathcal{L}_{\Diamond\diamondsuit}$*, if* $\varphi_i \equiv \psi$ *over* $\mathbf{C}$*, then* $|\psi| \geq 2^i$*.*

Proof. Observe that if $\varphi \equiv \psi$ over $\mathbf{C}$, then it is also the case over $\mathbf{GL}$ or $\mathbf{Grz}$. As such, by Theorem 26, the sequence of $\mathcal{L}^{\mu}_{\Diamond}$ formulae φ_i of Definition 13 is exponentially more succinct than their $\mathcal{L}_{\Diamond\diamondsuit}$ counterparts. Finally, if we consider the sequence of $\mathcal{L}^{\mu}_{\diamondsuit}$ formulae ψ_i that are the same as φ_i if we were to substitute $\Diamond$ by $\diamondsuit$, we know that $\psi_i \equiv \varphi_i$. ◀

7.2 Size of interpolants

A somewhat surprising link to this study is that with the interpolation theorem. The interpolation has been studied in many logics and via both model theoretic and proof theoretic means [8, 15]. However, while proof theoretic proofs of the interpolation theorem give us bounds for the proof size, no good bounds on the interpolants can be immediately derived. The link in this case is primarily tied to one of the proofs of the fixed-point theorem which we will briefly present here. Let us first recall the interpolation theorem.

▶ **Theorem 28** (Craig interpolation for $\mathbf{GL}$)**.** *Let* φ *and* ψ *be such that* $\mathbf{GL} \vdash \varphi \to \psi$*. There exists some formula* σ *containing only variables occurring in both* φ *and* ψ *such that*

$$\mathbf{GL} \vdash \varphi \to \sigma \text{ and } \mathbf{GL} \vdash \sigma \to \psi.$$

Proof. See e.g. [8]. ◀

Interpolation then easily implies the definability theorem of Beth.

▶ **Theorem 29** (Beth's definability theorem for $\mathbf{GL}$ [5])**.** *For any* $\varphi(x)$ *and* y *different from* x*, if* $\mathbf{GL} \vdash \varphi(x) \wedge \varphi(y) \to (x \leftrightarrow y)$ *then there is some formula* ψ *containing only variables in* $\varphi(x)$ *excluding* x *such that* $\mathbf{GL} \vdash \varphi(x) \to (\psi \leftrightarrow x)$*.*

Proof. By the assumptions $\mathbf{GL} \vdash \varphi(x) \wedge x \to (\varphi(y) \to y)$, then ψ is the formula given by Craig's interpolation theorem. For more details, see e.g. [7]. ◀

The proof proceeds by one proving uniqueness of fixed points in the following sense by Bernardi.

▶ **Theorem 30** (Bernardi [3])**.** *Let* $\varphi(x)$ *be modalized in* x*. Then*

$$\mathbf{GL} \vdash \boxdot(\varphi(x) \leftrightarrow x) \wedge \boxdot(\varphi(y) \leftrightarrow y) \to (x \leftrightarrow y).$$

Proof. See [3, 4, 7]. ◀

Then by the Beth definability theorem we can find some appropriate ψ such that $\mathbf{GL} \vdash \boxdot(\varphi(x) \leftrightarrow x) \to (\psi \leftrightarrow x)$. As a result, succinctness results for the fixed-point theorem of $\mathbf{GL}$ can be directly applied to provide succinctness results on the size of the interpolants.

▶ **Corollary 31.** *There exist sequences of formulae* $(\varphi_0, \varphi_1, \dots)$ *and* $(\psi_0, \psi_1, \dots)$ *both of size* $|\varphi_i|, |\psi_i| \leq O(i)$ *and such that* $\mathbf{GL} \vdash \varphi_i \to \psi_i$ *for every* i*, while every interpolant* σ_i *of* φ_i *and* ψ_i *is of size* $|\sigma_i| = 2^{\Omega(i)}$*.*

Proof. The sequences consist of the formulae $\varphi_i(x) := \boxdot(\varphi_i^*(x) \leftrightarrow x) \wedge x$ and $\psi_i(y) := \boxdot(\varphi_i^*(y) \leftrightarrow y) \to y$. Then any interpolant σ_i of φ_i and ψ_i is necessarily a fixed point of $\varphi_i^*(x)$ which by Theorem 26 must have size at least exponential in i. ◀

8 Conclusion

We have shown that the μ-calculus with only one occurrence of a fixed point operator is exponentially more succinct than basic modal logic, even when equipped with a reflexive modality and even in the setting of **GL**, where fixed points are already definable. This yields a lower bound on the fixed point formulas provided by Theorems 6 and 7, hence providing the first lower bounds for these celebrated results. This places the μ-calculus over **GL** or **Grz** as a powerful formal system in terms of expressivity, despite the theoretical definability of fixed points.

There is a small gap between the upper and lower bound for the fixed-point theorem for **GL**, with the lower bound being $2^{\Omega(n)}$ and the upper $2^{O(n\log(n))}$; it is unclear which of the two is tighter. In contrast, for **Grz** we obtain a larger gap of $2^{\Omega(n)}$ vs. $2^{O(n^3)}$; in this case we believe that the upper bound can be significantly improved, which we plan to address in future work.

Our proof of succinctness of the interpolants is a rather Post Hoc expansion of our succinctness for the fixed-point results coming directly from the bibliography. As such it is restricted to interpolants over **GL**. As far as we know, a result of this form is new and hence, a lucrative open problem would be expanding the methods of model equivalence games to get succinctness lower bounds for interpolants over **S**4 or **K**4 frames as an example.

Finally, we note that our techniques provide lower bound on formula length but *not* on the number of subformulas, or equivalently, the size of dag-like representations of formulas. This is particularly relevant since issues such as complexity can be bounded with respect to the latter measure, which may in fact be much smaller. Finding lower bounds on the number of subformulas would require a non-trivial modification of the model equivalence game; we leave the development of such games and the question of whether the μ-calculus remains exponentially succinct over dag-like formulas as challenging avenues for future research.

References

1. M. Adler and N. Immerman. An $n!$ lower bound on formula size. *ACM Transactions on Computational Logic*, 4(3):296–314, 2003. doi:10.1145/772062.772064.
2. Gaëlle Fontaine Balder ten Cate and Tadeusz Litak. Some modal aspects of xpath. *Journal of Applied Non-Classical Logics*, 20(3):139–171, 2010. doi:10.3166/JANCL.20.139-171.
3. Claudio Bernardi. The fixed-point theorem for diagonalizable algebras. *Studia Logica*, 34(3):239–251, 1975. doi:10.1007/bf02125226.
4. Claudio Bernardi. The uniqueness of the fixed-point in every diagonalizable algebra. *Studia Logica*, 35(4):335–343, 1976. doi:10.1007/bf02123401.
5. E.W. Beth. On padoa's method in the theory of definition. *Indagationes Mathematicae (Proceedings)*, 56:330–339, 1953. doi:10.1016/S1385-7258(53)50042-3.
6. A. Blass. Infinitary combinatorics and modal logic. *Journal of Symbolic Logic*, 55(2):761–778, 1990. doi:10.2307/2274663.
7. G. S. Boolos. *The Logic of Provability*. Cambridge University Press, 1993.
8. A. Chagrov and M. Zakharyaschev. *Modal Logic*, volume 35 of *Oxford logic guides*. Oxford University Press, 1997.
9. A. Dawar and M. Otto. Modal characterisation theorems over special classes of frames. *Annals of Pure and Applied Logic*, 161:1–42, 2009. doi:10.1016/J.APAL.2009.04.002.
10. K. Eickmeyer, M. Elberfeld, and F. Harwath. Expressivity and succinctness of order-invariant logics on depth-bounded structures. In *Proceedings of MFCS 2014*, pages 256–266, 2014.
11. David Fernández-Duque and Petar Iliev. Succinctness in subsystems of the spatial μ-calculus. *FLAP*, 5(4):827–874, 2018. URL: https://www.collegepublications.co.uk/downloads/ifcolog00024.pdf.

12 David Fernández-Duque and Konstnatinos Papafilippou. The universal tangle for spatial reasoning. In Sarah Alice Gaggl, Maria Vanina Martinez, and Magdalena Ortiz, editors, *Logics in Artificial Intelligence - 18th European Conference, JELIA 2023, Dresden, Germany, September 20-22, 2023, Proceedings*, volume 14281 of *Lecture Notes in Computer Science*, pages 814–827. Springer, 2023. `doi:10.1007/978-3-031-43619-2_55`.

13 S. Figueira and D. Gorín. On the size of shortest modal descriptions. *Advances in Modal Logic*, 8:114–132, 2010.

14 T. French, W. van der Hoek, P. Iliev, and B. Kooi. Succinctness of epistemic languages. In T. Walsh, editor, *Proceedings of IJCAI*, pages 881–886, 2011.

15 Dov M. Gabbay and Larisa Maksimova. *Interpolation and Definability: Modal and Intuitionistic Logics*. Oxford University Press, May 2005. `doi:10.1093/acprof:oso/9780198511748.001.0001`.

16 Zachary Gleit and Warren Goldfarb. Characters and fixed-points in provability logic. *Notre Dame Journal of Formal Logic*, 31(1):26–36, 1989. `doi:10.1305/ndjfl/1093635330`.

17 M. Grohe and N. Schweikardt. The succinctness of first-order logic on linear orders. *Logical Methods in Computer Science*, 1:1–25, 2005.

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*, volume 11, pages 401–419, 2016.

19 A. Razborov. Applications of matrix methods to the theory of lower bounds in computational complexity. *Combinatorica*, 10(1):81–93, 1990. `doi:10.1007/BF02122698`.

20 Lisa Reidhaar-Olson. A new proof of the fixed-point theorem of probability logic. *Notre Dame J. Formal Log.*, 31:37–43, 1989. URL: `https://api.semanticscholar.org/CorpusID:7685381`, `doi:10.1305/NDJFL/1093635331`.

21 Giovanni Sambin. An effective fixed-point theorem in intuitionistic diagonalizable algebras. *Studia Logica*, 35(4):345–361, 1976. `doi:10.1007/bf02123402`.

22 Craig Smoryński. Modal logic and self-reference. In D. M. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, pages 1–53. Springer Netherlands, Dordrecht, 2004. `doi:10.1007/978-94-017-0466-3_1`.

23 R. M. Solovay. Provability interpretations of modal logic. *Israel Journal of Mathematics*, 25:287–304, 1976.

24 A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific J. Math.*, 5(2):285–309, 1955. URL: `https://projecteuclid.org:443/euclid.pjm/1103044538`.

25 W. van der Hoek, P. Iliev, and B. Kooi. On the relative succinctness of modal logics with union, intersection and quantification. In *Proceedings of AAMAS*, pages 341–348, 2014.

26 VII Soviet Symposium on Logic. *On modal "companions" of superintuitionistic logics*, Kiev, 1976. Russian.