



HAL
open science

Frame-validity games and lower bounds on the complexity of modal axioms

Philippe Balbiani, David Fernández Duque, Andreas Herzig, Petar Iliev

► **To cite this version:**

Philippe Balbiani, David Fernández Duque, Andreas Herzig, Petar Iliev. Frame-validity games and lower bounds on the complexity of modal axioms. *Logic Journal of the IGPL*, 2020, 30 (1), 10.1093/jigpal/jzaa068 . hal-02936458

HAL Id: hal-02936458

<https://hal.science/hal-02936458v1>

Submitted on 11 Sep 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Frame-validity games and lower bounds on the complexity of modal axioms

Philippe Balbiani

IRIT, CNRS

Toulouse University

Philippe.Balbani@irit.fr

David Fernández-Duque

Department of Mathematics

Ghent University

David.FernandezDuque@UGent.be

Andreas Herzig

IRIT, CNRS

Toulouse University

Andreas.Herzig@irit.fr

Petar Iliev

Institute of Philosophy and Sociology

Bulgarian Academy of Sciences

petar.iliev@gmail.com

Abstract

We introduce frame-equivalence games tailored for reasoning about the size, modal depth, number of occurrences of symbols, and number of different propositional variables of modal formulae defining a given frame-property. Using these games, we prove lower bounds on the above measures for a number of well-known modal axioms; what is more, for some of the axioms, we show that they are optimal among the formulae defining the respective class of frames.

Keywords: modal logic, correspondence theory, formula-size games, lower bounds on formula-size.

1 Introduction

The study of the descriptive complexity of a class of structures \mathbf{S} relative to a class of formulae Φ from a logic L is described in [7], [18], and [26] as revolving around the question: what can we say about the definability of \mathbf{S} with formulae from Φ ? For example, we might want to know whether there is a formula $\varphi \in \Phi$ defining \mathbf{S} or, if not, whether there is a family $\varphi_1, \varphi_2, \dots$ of Φ -formulae such that each φ_i defines a subset \mathbf{S}_i of \mathbf{S} and the union of all \mathbf{S}_i is \mathbf{S} . If we have the former situation, it is natural to ask about the minimal number of variables in any $\psi \in \Phi$ defining \mathbf{S} (where different occurrences of the same variable are either counted or not), or its minimal length (i.e., the total number of symbols in ψ), or the minimal number of operators like quantifiers, disjunctions and conjunctions, and the minimal depth of their nesting. In the latter situation, we might want to know how (some of) these measures scale with the index of individual formulae.

The study of the descriptive complexity of the class of finite graphs relative to the formulae of existential second-order logic is motivated by long-standing open problems in logic and computational complexity theory (see section 7 in [7]), while results in the descriptive complexity of the same class relative to first-order formulae in the language of graph theory are driven among others by the study of algorithms for graph-isomorphism testing (see for example [24]).

Of course, the importance of the correspondence between descriptive complexity and computational complexity cannot be overstated and several monographs are devoted to the subject [10], [19], [22]. However, even if we disregard both the computational complexity and the algorithmic angles, the purely model-theoretic side of descriptive complexity questions is interesting in its own right, as argued in [7] (pp. 25–26).

In the present paper, we take some initial steps in the study of the descriptive complexity of classes of Kripke frames (directed graphs) relative to the standard unimodal language. In contrast to the area dealing with the descriptive complexity of classes of Kripke models (edge- and vertex-coloured directed graphs) that has been very active in recent years (e.g., [5], [8], [9], [11], [15], [16]), as far as we are aware, it seems that the only investigation of how succinctly *frame* properties can be expressed with modal formulae is [27], where the question of how many different propositional variables are needed to modally define certain classes of frames is being considered.

The main technical tool we develop in order to obtain our results is called *frame equivalence game*. We define this game by adding valuation-picking moves to the model equivalence game from [9, 11, 12, 14, 15]. This is analogous to the addition of subset selecting moves to the Ehrenfeucht-Fraïssé game in order to define its (weak) monadic second-order version as explained already in the classic [6]. The model equivalence game can be seen as a modal version of the generalisation of Ehrenfeucht-Fraïssé games given in [1]. This generalisation of Ehrenfeucht-Fraïssé games in turn, can be viewed in a certain sense as "lifting" to first-order logic the Boolean game FORMULA that appeared in [25]. Finally, the game from [25] is related to the communication complexity game from [21].

To demonstrate the applicability of the frame equivalence game to both first- and second-order semantic conditions, we prove the following.

- (a) Any modal formula in a language with the universal box- and diamond-modalities that defines the class of graphs that are not n -colourable contains at least $\log_2 n$ different propositional variables and has size that is at least linear in n . As a counterpart to this lower bound, we provide a formula of a quasilinear length that contains $\log_2 n$ different propositional variables. As far as we are aware, previously known modal formulae defining this second-order property contain at least n different propositional variables (e.g., [3, 13]).
- (b) For each $m, n \geq 0$, the (m, n) -transfer axiom $\diamond^m p \rightarrow \diamond^n p$ is essentially the shortest modal formula defining the first-order condition

$$\forall x \forall y (x R^m y \rightarrow x R^n y), \tag{1}$$

where R^j denotes the j -fold composition of R . Note that this result applies to the well-studied axioms defining transitivity, reflexivity, and density.

- (c) The Löb axiom $\Box(\Box p \rightarrow p) \rightarrow \Box p$ is essentially the shortest modal formula defining transitivity plus the second-order property of converse well-foundedness.
- (d) The **S4** axiom $(p \vee \diamond \diamond p) \rightarrow \diamond p$ is the shortest among those defining reflexivity plus transitivity.
- (e) The axiom $p \rightarrow \Box \diamond p$ is the shortest modal formula that defines symmetry.

Since descriptive complexity questions in first- and second-order logic are driven by the study of finite structures, we would like to point out that, with the exception of the lower bound in (c), all our results are established with the help of finite frames. We leave it as an open problem whether the Löb axiom is the shortest modal formula defining transitivity plus converse well-foundedness on the class of finite frames.

Our investigation was motivated to a large degree by ideas and results from [27], where the notion of minimal modal equivalent of a first-order condition was introduced. Note however that

the term minimal is used in [27] only with respect to the number of different variables needed to modally define a first-order condition: this does not tell us much about the length, modal depth, or the number of Boolean connectives required and that is why we have extended the notion of minimality to cover these as well.

Although this work is concerned with the model theory of modal logic, it is worth stressing that results about the descriptive complexity of classes of frames or models have clear complexity-theoretic relevance. In particular, there are well-known examples of modal logics L_1 and L_2 with the same expressive power and the same complexity of the satisfiability problem that have exponential gap in succinctness [23]. This means that we have an infinite sequence of classes of models (semantic properties) C_1, C_2, \dots with descriptive complexities that grow linearly in the lengths of defining L_1 -formulae but exponentially in the lengths of defining L_2 -formulae. To be precise, there is a sequence $\varphi_1, \varphi_2, \dots$ of L_1 -formulae with lengths growing linearly in their indices such that the lengths of the L_2 -formulae in any sequence of equivalent formulae ψ_1, ψ_2, \dots grow exponentially. Even if we had an algorithm for the satisfiability problems for both L_1 and L_2 working in linear time, this algorithm would finish work on say φ_{500} in approximately 500 steps, whereas, when applied to any equivalent L_2 -formula, the number of steps required would be roughly 2^{500} . In short, there is little use in a PTIME logic if properties of interest can only be defined by exponentially large formulae.

The rest of the paper is organised as follows. The next section recalls some standard definitions. Section 3 describes formula-bound games on models and Section 4 turns them into games on frames. Section 5 provides shortest axioms for the n non-colourability property. Section 6 provides a general result for transfer axioms. Section 7 is about the **S4** axiom and Section 8 is about the Löb axiom. Section 9 concludes. The appendix contains the proofs of sections 3 and an analysis of the symmetry axiom.

The present paper extends the paper [2].

2 Technical preliminaries

Fix a countably infinite set of *propositional variables* $P = \{p_1, p_2, \dots\}$, and let L_\diamond denote the uni-modal language that has as atomic formulae the *literals* p, \bar{p} for each $p \in P$ as well as \perp, \top and as primitive connectives \vee, \wedge, \diamond , and \Box . The expressions $\neg\varphi$ and $\varphi \rightarrow \psi$ will be regarded as abbreviations defined using De Morgan's rules. We will also be interested in the language L_\diamond^\forall that extends L_\diamond with the universal modalities \exists and \forall .

As usual, a frame is a pair $\mathcal{A} = (W_{\mathcal{A}}, R_{\mathcal{A}})$ where $W_{\mathcal{A}}$ is a nonempty set and $R_{\mathcal{A}} \subseteq W_{\mathcal{A}} \times W_{\mathcal{A}}$; a *model based on* $(W_{\mathcal{A}}, R_{\mathcal{A}})$ is a tuple $\mathcal{A} = (W_{\mathcal{A}}, R_{\mathcal{A}}, V_{\mathcal{A}})$ consisting of a frame equipped with a *valuation* $V_{\mathcal{A}}: W_{\mathcal{A}} \rightarrow 2^P$; and a *pointed model* $\mathbf{a} = (\mathcal{A}, a)$ consists of a model \mathcal{A} with a designated point $a \in W_{\mathcal{A}}$. Pointed models will be denoted by $\mathbf{a}, \mathbf{b}, \dots$ and frames or models by $\mathcal{A}, \mathcal{B}, \dots$. For a pointed model $\mathbf{a} = (\mathcal{A}, a)$, we denote by $\Box\mathbf{a}$ the set $\{(\mathcal{A}, b) : a R_{\mathcal{A}} b\}$, i.e., the set of all pointed models that are *successors* of \mathbf{a} along the relation $R_{\mathcal{A}}$. Analogously, we use $\forall\mathbf{a}$ to denote the set $\{(\mathcal{A}, b) : b \in W_{\mathcal{A}}\}$.

Given $\varphi \in L_\diamond^\forall$ and a pointed model \mathbf{a} , we define $\mathbf{a} \models \varphi$ according to standard Kripke semantics. As usual, if \mathcal{A} is a model, we write $\mathcal{A} \models \varphi$ if $(\mathcal{A}, a) \models \varphi$ for all $a \in W_{\mathcal{A}}$, and if \mathcal{A} is a frame, $\mathcal{A} \models \varphi$ if $(\mathcal{A}, V) \models \varphi$ for every valuation V . We use *structure* as an umbrella term for either a model, a frame, or a pointed model. For a class of structures \mathbf{A} and a formula φ , we write $\mathbf{A} \models \varphi$ when $\mathcal{X} \models \varphi$ for all $\mathcal{X} \in \mathbf{A}$, and say that the formulae φ and ψ are *equivalent* on \mathbf{A} when for all $\mathcal{X} \in \mathbf{A}$, $\mathcal{X} \models \varphi$ if and only if $\mathcal{X} \models \psi$. We say that a modal formula φ defines a class \mathbf{F} of frames if \mathbf{F} exactly consists of the frames on which φ is valid.

Our goal is to develop techniques to establish when a formula φ is of minimal complexity

among those defining some class of frames. Here *complexity* could mean many things: a *complexity measure* (or just *measure*) is a function $\mu: \mathbf{L} \rightarrow \mathbb{N}$, where \mathbf{L} is either \mathbf{L}_\diamond or \mathbf{L}_\forall , and such that, for any formulae φ and ψ , if ψ is a subformula of φ , then $\mu(\psi) \leq \mu(\varphi)$. We are interested in the following measures which we call *natural measures*: (1) the *length* $|\varphi|$ of a formula φ defined as the number of nodes in its syntax tree (including leaves); (2) the *number of occurrences* of any connective, (3) the *modal depth*, and (4) the *number of variables* (which, depending on the context, we can use in the sense of the number of different variables or in the sense of the total number of not necessarily different variables occurring in a formula).

Note that each connective gives rise to its own measure in (2). We will show that several modal axioms of interest are minimal with respect to all of these measures simultaneously. To this end, given a set $\Gamma \subseteq \mathbf{L}_\forall$ and $\varphi \in \Gamma$, we say that φ is *absolutely minimal among* Γ if for all $\psi \in \Gamma$ and any of the respective measures μ described above, $\mu(\varphi) \leq \mu(\psi)$.

3 A formula-bound game on models

The game described below is the modal analogue of the formula-size game for first-order logic developed in [1]. The idea is that we have two players, *Hercules* and the *Hydra*. Given two classes of pointed models \mathbf{A} and \mathbf{B} , Hercules is trying to show that there is a “small” \mathbf{L}_\forall -formula φ such that $\mathbf{A} \models \varphi$ but $\mathbf{B} \models \neg\varphi$ whereas the Hydra is trying to show that any such φ is “big”. The players move by adding and labelling nodes on a game-tree $\langle T, \preceq \rangle$. For our purposes a *tree* is a finite set partially ordered by some order \preceq such that if $\eta \in T$ then $\downarrow\eta = \{\nu : \nu \preceq \eta\}$ is linearly ordered; any set of the form $\downarrow\eta$ is a *branch* of T .

Definition 1. *The $(\mathbf{L}_\forall, \langle \mathbf{A}, \mathbf{B} \rangle)$ formula-complexity game on models (denoted $(\mathbf{L}_\forall, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM) is played by Hercules and the Hydra who construct a game-tree T in such a way that each node $\eta \in T$ is labelled with a pair $\langle \mathfrak{L}(\eta), \mathfrak{R}(\eta) \rangle$ of classes of pointed models and either a literal or a symbol from $\{\perp, \top, \vee, \wedge, \diamond, \square, \exists, \forall\}$ according to the rules below.*

Any leaf η can be declared either a head or a stub. Once η has been declared a stub, no further moves can be played on it. The construction of T begins with a root labelled by $\langle \mathbf{A}, \mathbf{B} \rangle$ that is declared a head. Afterwards, the game continues as long as there is at least one head. In each turn, Hercules goes first by choosing a head η labelled by $\langle \mathfrak{L}(\eta), \mathfrak{R}(\eta) \rangle$. Hercules then plays one of the following moves to which the Hydra possibly replies.

LITERAL-MOVE: *Hercules chooses a literal ι such that $\mathfrak{L}(\eta) \models \iota$ and $\mathfrak{R}(\eta) \models \neg\iota$. The node η is declared a stub and labelled with the symbol ι .*

\perp -MOVE: *Hercules can play this move only if $\mathfrak{L}(\eta) = \emptyset$. The node η is declared a stub and labelled with the symbol \perp .*

\top -MOVE: *Hercules can play this move only if $\mathfrak{R}(\eta) = \emptyset$. The node η is declared a stub and labelled with the symbol \top .*

\vee -MOVE: *Hercules labels η with \vee and chooses two subclasses $\mathbf{L}_1, \mathbf{L}_2 \subseteq \mathfrak{L}(\eta)$ with $\mathfrak{L}(\eta) = \mathbf{L}_1 \cup \mathbf{L}_2$. Two new heads, labelled by $\langle \mathbf{L}_1, \mathfrak{R}(\eta) \rangle$ and $\langle \mathbf{L}_2, \mathfrak{R}(\eta) \rangle$, are added to T as daughters of η .*

\wedge -MOVE: *Dual to the \vee -move, except that in this case Hercules chooses $\mathbf{R}_1, \mathbf{R}_2$ such that $\mathbf{R}_1 \cup \mathbf{R}_2 = \mathfrak{R}(\eta)$.*

\diamond -MOVE: *Hercules labels η with \diamond and, for each pointed model $\mathbf{l} \in \mathfrak{L}(\eta)$, chooses a pointed model from $\square\mathbf{l}$; if for some $\mathbf{l} \in \mathfrak{L}(\eta)$ we have $\square\mathbf{l} = \emptyset$, Hercules cannot play this move. All these new pointed models are collected in the class \mathbf{L}_1 . For each $\mathbf{r} \in \mathfrak{R}(\eta)$, the Hydra replies by picking a*

subset of $\square\mathbf{r}$. All the pointed models chosen by the Hydra are collected in the class \mathbf{R}_1 .¹ A new head labelled by $\langle \mathbf{L}_1, \mathbf{R}_1 \rangle$ is added as a daughter to η .

\square -MOVE: Dual to the \diamond -move, except that Hercules first chooses a successor for each $\mathbf{r} \in \mathfrak{R}(\eta)$ and Hydra chooses her successors for frames in $\mathfrak{L}(\eta)$.

\exists -MOVE: Hercules labels η with \exists and, for each $\mathbf{l} \in \mathfrak{L}(\eta)$, he chooses a pointed model from $\forall\mathbf{l}$. All these new pointed models are collected in the class \mathbf{L}_1 . For each $\mathbf{r} \in \mathfrak{R}(\eta)$, the Hydra replies by picking a subset of $\forall\mathbf{r}$. All the pointed models chosen by the Hydra are collected in the class \mathbf{R}_1 . A new head labelled by $\langle \mathbf{L}_1, \mathbf{R}_1 \rangle$ is added as a daughter to η .

\forall -MOVE: Dual to the \exists -move, except that Hercules first chooses a successor for each $\mathbf{r} \in \mathfrak{R}(\eta)$ and Hydra chooses her successors for frames in $\mathfrak{L}(\eta)$.

The $(\mathbf{L}_{\diamond}^{\forall}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM concludes when there are no heads and we say in this case that T is a closed game tree.

Note that the Hydra has no restrictions on the number of pointed models she chooses on modal moves; in fact, she can choose all of them, and it is often convenient to assume that she always does so. To be precise, say that the Hydra *plays greedily* if (1) whenever Hercules makes a \diamond -move on a node η and a new node η' is added then $\mathfrak{R}(\eta') = \bigcup_{\mathbf{r} \in \mathfrak{R}(\eta)} \square\mathbf{r}$, and similarly (2) whenever Hercules makes a \square -move on a node η and a new node η' is added then $\mathfrak{L}(\eta') = \bigcup_{\mathbf{l} \in \mathfrak{L}(\eta)} \square\mathbf{l}$. Analogously for \exists - and \forall -moves.

The $(\mathbf{L}_{\diamond}^{\forall}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM can be used to give lower bounds on the length of $\mathbf{L}_{\diamond}^{\forall}$ -formulae defining a given property of Kripke models; if we are interested in the length of formulae in the sub-language \mathbf{L}_{\diamond} of $\mathbf{L}_{\diamond}^{\forall}$ that does not have \exists and \forall operators, we simply do not allow the corresponding \exists and \forall -moves and this new game, which is the one from [12, 11, 14, 15], is denoted $(\mathbf{L}_{\diamond}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM. Here we will show how to generalize these games so that they can be used to give lower bounds on any complexity measure. For this, we need to view game-trees as formulae.

Definition 2. Given a closed $(\mathbf{L}_{\diamond}^{\forall}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM tree T , we define $\psi_T \in \mathbf{L}_{\diamond}^{\forall}$ to be the unique formula whose syntax tree is given by T .

Formally speaking, ψ_T is defined by recursion on T starting from leaves: if T is a single leaf then it must be labelled by a literal ι , or by \perp , or by \top , so we respectively set $\psi_T = \iota$, or $\psi_T = \perp$, or $\psi_T = \top$; if T has a root η labelled by \vee , then η has two daughters η_1, η_2 . Letting T_1, T_2 be the respective generated subtrees, we define $\psi_T = \psi_{T_1} \vee \psi_{T_2}$. The cases for $\wedge, \diamond, \square, \exists$, and \forall are all analogous. Then, given a complexity measure μ , we extend the domain of μ to include the set of closed game trees by defining $\mu(T) = \mu(\psi_T)$.

If $\mathbf{L} \in \{\mathbf{L}_{\diamond}, \mathbf{L}_{\diamond}^{\forall}\}$, $k \in \mathbb{N}$, \mathbf{A}, \mathbf{B} are classes of models, and $\mu: \mathbf{L} \rightarrow \mathbb{N}$ a complexity measure (including but not restricted to the four measures that we have defined in Section 2), we say that Hercules has a *winning strategy for the $(\mathbf{L}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM with μ below k* if Hercules has a strategy so that no matter how Hydra plays, the game terminates in finite time with a closed tree T so that $\mu(T) < k$.

Theorem 1. Let $\mathbf{L} \in \{\mathbf{L}_{\diamond}, \mathbf{L}_{\diamond}^{\forall}\}$, \mathbf{A}, \mathbf{B} be classes of pointed models, $\mu: \mathbf{L} \rightarrow \mathbb{N}$ any complexity measure, and $k \in \mathbb{N}$. Then the following are equivalent:

1. Hercules has a winning strategy for the $(\mathbf{L}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM with μ below k ;
2. there is an \mathbf{L} -formula φ with $\mu(\varphi) < k$ and $\mathbf{A} \models \varphi$ whereas $\mathbf{B} \models \neg\varphi$.

¹In particular, if $\square\mathbf{r} = \emptyset$ for some $\mathbf{r} \in \mathfrak{R}(\eta)$ then $\mathbf{R}_1 = \emptyset$, i.e., the Hydra does not add anything to \mathbf{R}_1 . For example, when $\mathfrak{L}(\eta)$ is the set of all serial models and $\mathfrak{R}(\eta)$ contains a model built on the irreflexive singleton frame $\langle \{w\}, \emptyset \rangle$ then a \diamond -move on η results in a new head labelled $\langle \mathfrak{L}(\eta), \emptyset \rangle$.

We defer the proof of Theorem 1 to Appendix 10, where we also establish some useful properties of the formula-complexity game. However, we remark that the proof is essentially the same as that of the special case where $\mu(\varphi) = |\varphi|$, which can be found in any of [12, 11, 14, 15]. We will also use the following easy consequence of Theorem 1. We assume familiarity with bisimulations ([4] p. 54).

Corollary 1. *Let $L \in \{L_\diamond, L_\diamond^\forall\}$, \mathbf{A} and \mathbf{B} be classes of pointed models such that there are $\mathbf{a} \in \mathbf{A}$ and $\mathbf{b} \in \mathbf{B}$ with \mathbf{a} L -bisimilar to \mathbf{b} . For all complexity measures μ and all non-negative integers k , Hercules has no winning strategy for the $(L, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM with μ below k .*

4 A formula-complexity game on frames

We develop now an analogous game to the one above that is played on frames instead of models in order to reason about the “resources” needed to modally define properties of frames with L_\diamond - or L_\diamond^\forall -formulae.

Definition 3. *Let \mathbf{A}, \mathbf{B} be classes of frames. The $(L_\diamond^\forall, \langle \mathbf{A}, \mathbf{B} \rangle)$ formula-complexity game on frames (denoted $(L_\diamond^\forall, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF) is played by Hercules and the Hydra as follows.*

HERCULES SELECTS MODELS: *For each $\mathcal{B} \in \mathbf{B}$ Hercules chooses a model \mathcal{B}^M based on \mathcal{B} and a point $\triangleright_{\mathcal{B}} \in W_{\mathcal{B}}$ and then sets $\mathbf{B}^m = \{(\mathcal{B}^M, \triangleright_{\mathcal{B}}) : \mathcal{B} \in \mathbf{B}\}$.*

THE HYDRA SELECTS MODELS: *The Hydra replies by choosing a class of pointed models \mathbf{A}^m of the form (\mathcal{A}, V, a) with $\mathcal{A} \in \mathbf{A}$.*

FORMULA GAME ON MODELS: *Hercules and the Hydra play the $(L_\diamond^\forall, \langle \mathbf{A}^m, \mathbf{B}^m \rangle)$ -FGM.*

The game tree assigned to a match of the $(L_\diamond^\forall, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF is the game tree of the subsequent $(L_\diamond^\forall, \langle \mathbf{A}^m, \mathbf{B}^m \rangle)$ -FGM. As before, if we are interested in the length of L_\diamond -formulae, we do not allow \exists - and \forall -moves, and the relevant game is denoted $(L_\diamond, \langle \mathbf{A}^m, \mathbf{B}^m \rangle)$ -FGF.

Remark 1. *The Hydra is free to assign as many models as she wants to each $\mathcal{A} \in \mathbf{A}$, even no model at all. We say that the Hydra plays functionally if she chooses \mathbf{A}^m so that for each $\mathcal{A} \in \mathbf{A}$ there is exactly one pointed model $(\mathcal{A}^M, \triangleright_{\mathcal{A}}) \in \mathbf{A}^m$ with \mathcal{A}^M based on \mathcal{A} . In this text the Hydra will often play functionally.*

As was the case for the FGM, for $L \in \{L_\diamond, L_\diamond^\forall\}$, $k \in \mathbb{N}$, classes of frames \mathbf{A}, \mathbf{B} , and $\mu: L_\diamond^\forall \rightarrow \mathbb{N}$ a complexity measure, Hercules has a winning strategy for the $(L_\diamond^\forall, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF with μ below k if Hercules has a strategy such that, no matter how the Hydra plays, the game terminates in finite time with a closed tree T so that $\mu(T) < k$.

Theorem 2. *Let $L \in \{L_\diamond, L_\diamond^\forall\}$, \mathbf{A}, \mathbf{B} be classes of frames, μ any complexity measure, and $k \in \mathbb{N}$. Then, the following are equivalent:*

1. *Hercules has a winning strategy for the $(L, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF with μ below k ;*
2. *there is an L -formula φ with $\mu(\varphi) < k$ that is valid on every frame of \mathbf{A} and non-valid on every frame of \mathbf{B} .*

Proof. 2 IMPLIES 1. Let φ be an L -formula with $\mu(\varphi) < k$ that is valid on all frames in \mathbf{A} and not valid on any frame in \mathbf{B} . For each $\mathcal{B} \in \mathbf{B}$, Hercules can choose a pointed model $\mathcal{B}^M = (\mathcal{B}, V, b)$ based on \mathcal{B} so that $\mathcal{B}^M \not\models \varphi$. Let \mathbf{B}^m be the class of models chosen by Hercules. The Hydra then responds with some class of pointed models \mathbf{A}^m based on the frames in \mathbf{A} ; since φ is valid on all frames in \mathbf{A} , we have $\mathbf{A}^m \models \varphi$. By Theorem 1, it follows that Hercules has a winning strategy with μ below k for the $(L, \langle \mathbf{A}^m, \mathbf{B}^m \rangle)$ -FGM and thus for $(L, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF.

1 IMPLIES 2. Now assume that Hercules has such a strategy, and that he chooses \mathbf{B}^m according to this strategy. Then Hydra *opens greedily* by choosing every pointed model based on a frame in \mathbf{A} ; i.e., she sets \mathbf{A}^m to be the class of all (\mathcal{A}, V, a) with $\mathcal{A} \in \mathbf{A}$, V a valuation on \mathcal{A} and $a \in W_{\mathcal{A}}$.

By playing according to his strategy, Hercules can win the $(\mathbf{A}^m, \mathbf{B}^m)$ -FGM with a closed game tree T such that $\mu(T) < k$; but this is only possible if his sub-strategy for the $(\mathbf{A}^m, \mathbf{B}^m)$ -FGM is a winning strategy with μ below k . Thus by Theorem 1, there is an L_{\diamond} -formula φ with $\mu(\varphi) < k$ such that $\mathbf{A}^m \models \varphi$ and $\mathbf{B}^m \models \neg\varphi$. Since Hercules chose one pointed model for each $\mathcal{B} \in \mathbf{B}$ it follows that φ is not valid in any frame in \mathbf{B} , while since Hydra chose all possible pointed models it follows that $\mathcal{A} \models \varphi$. \square

Next, we apply our formula-complexity games to prove lower bounds on the complexity of some modal axioms. For ease of understanding, we define the pointed models employed in our proofs using figures. We follow the convention that such pointed models consist of the relevant Kripke model and a point denoted by the \triangleright sign next to it.

5 The non-colourability property

For a natural number $n \geq 1$, let us consider the property of a graph being not n -colourable, i.e, the set of its vertices cannot be partitioned in at most n equivalence classes so that no two vertices sharing an edge are in the same class. This property is modally definable with a formula from L_{\diamond}^{\forall} . A natural way of finding such a formula is to reason as follows. To encode the n colours, we use the propositional symbols p_1, \dots, p_n , respectively. Using \exists and \forall , we say “if every node of the graph is coloured with exactly one colour, then there are two edge-related nodes that have the same colour”. Formally,

$$\forall \left((p_1 \vee \dots \vee p_n) \wedge \left(\bigwedge_{1 \leq i < j \leq n} \neg(p_i \wedge p_j) \right) \right) \rightarrow \exists \left(\bigvee_{1 \leq i \leq n} (p_i \wedge \diamond p_i) \right).$$

A version of the above formula can be found in [3]. The subformula $\bigwedge_{1 \leq i < j \leq n} \neg(p_i \wedge p_j)$ is the reason why the length of the whole formula is quadratic in n . One of the referees alerted us to the fact that the family of formulae $\chi_n = \exists((p_1 \rightarrow \diamond p_1) \wedge \dots \wedge (p_n \rightarrow \diamond p_n))$, which have linear length, also define non- n -colourability. They can be found in [13] where it is stated that χ_n is a version of the axiom MT_n of [17]. We show below that we can find a formula of a quasilinear length and exponentially smaller number of variables that expresses the non-colourability property.

Recall that P denotes the set of propositional variables. For $k \geq 1$, let $P_k \subset P$ be the subset of P containing only the first k variables in P .

Definition 4. We define a sequence of formulae $(\varphi_n)_{n=1}^{\infty}$ as follows.

We set $\varphi_1 = \exists \diamond \top$. If $n \geq 2$, let $k = \lceil \log_2 n \rceil$ (so that $2^{k-1} < n \leq 2^k$). Fix an enumeration $\{S_1, \dots, S_{2^k}\}$ of 2^{P_k} and associate an elementary conjunction \hat{E} to every $E \subseteq P^k$ defined by $\hat{E} = \bigwedge_{p \in E} p \wedge \bigwedge_{p \in P_k \setminus E} \bar{p}$. Then let φ_n be the formula

$$\exists \left(\bigvee_{1 \leq i \leq n} (\hat{S}_i \wedge \diamond \hat{S}_i) \vee \bigvee_{n+1 \leq j \leq 2^k} \hat{S}_j \right).$$

For example, for $n = 4$ we have $k = 2$, so φ_4 is

$$\exists \left(\begin{array}{c} (\overline{p_1} \wedge \overline{p_2} \wedge \diamond(\overline{p_1} \wedge \overline{p_2})) \vee (\overline{p_1} \wedge p_2 \wedge \diamond(\overline{p_1} \wedge p_2)) \\ \vee \\ (p_1 \wedge \overline{p_2} \wedge \diamond(p_1 \wedge \overline{p_2})) \vee (p_1 \wedge p_2 \wedge \diamond(p_1 \wedge p_2)) \end{array} \right).$$

Since $2^k < 2n$ and each S_i contains less than $\log_2 n + 1$ propositional variables, it is easily seen there are less than $2 \times 2n(\log_2 n + 1)$ occurrences of propositional variables in φ_n , and similarly, that the lengths of φ_n are bounded from above by a function in $O(n \log_2 n)$. Moreover, the formulae φ_n characterise non- n -colourability. Below, note that directed graphs are just Kripke frames, hence we can speak of validity of formulae on directed graphs. We will moreover regard non-directed graphs as directed graphs with a symmetric edge relation.

Proposition 1. *For any graph \mathcal{G} , φ_n is valid in \mathcal{G} iff \mathcal{G} is not n -colourable.*

Proof. We begin by showing that if \mathcal{G} is n -colourable then φ_n is not valid in \mathcal{G} . Suppose that $W_{\mathcal{G}}$ can be partitioned in m equivalence classes C_1, \dots, C_n so that no two vertices sharing the same edge belong to the same class. Recall that $\{S_1, \dots, S_{2^k}\}$ is an enumeration of all subsets of P_k . Define a valuation on \mathcal{G} by setting $p \in V(w)$ if and only if, for the unique i such that $w \in C_i$, we have $p \in S_i$. It is immediate that the negation of φ_n ,

$$\forall \left(\bigwedge_{1 \leq i \leq n} (\hat{S}_i \rightarrow \neg \diamond \hat{S}_i) \wedge \bigwedge_{n+1 \leq j \leq 2^k} \neg \hat{S}_j \right),$$

is true in the model (\mathcal{G}, V) .

Conversely, assume that φ_n is not valid in \mathcal{G} . Therefore, there is a valuation V such that $\neg \varphi_n$ holds. i.e.,

$$\forall \left(\bigwedge_{1 \leq i \leq n} (\hat{S}_i \rightarrow \neg \diamond \hat{S}_i) \wedge \bigwedge_{n+1 \leq j \leq 2^k} \neg \hat{S}_j \right)$$

is true in the resulting Kripke model (\mathcal{G}, V) . It is easily seen that this implies that \mathcal{G} can be n -coloured by defining for $i \in [1, n]$, C_i to be the set of all $w \in W_{\mathcal{G}}$ such that, for all $p \in P_k$, $p \in V(w)$ if and only if $p \in S_i$. \square

In the rest of this section, we establish a linear lower bound on the size of L_{\diamond}^{\forall} -formulae that define non- n -colourability.

Theorem 3. *For any natural numbers $n \geq 2$, any L_{\diamond}^{\forall} -formula φ that defines the property of a graph being non- n -colourable contains at least $\lceil \log_2 n \rceil$ different propositional symbols, at least one occurrence of the \exists operator, and has size at least n .*

We begin the proof of Theorem 3 by proving the bound on the number of variables.

Recall that the complete graph on n nodes, usually denoted \mathcal{K}_n , is an undirected, irreflexive graph with n vertices in which every pair of distinct vertices is connected by a unique edge. Since Kripke semantics are based on directed graphs, we will regard $\mathcal{K}_n = (W_n, R_n)$ as a directed graph, albeit with a symmetric relation, so that wR_nv if and only if $w \neq v$. Clearly, every \mathcal{K}_n is n -colourable.

For $n \geq 1$, let $\tilde{\mathcal{K}}_n$ be a graph that consists of two disjoint copies of \mathcal{K}_n so that only one of the copies of \mathcal{K}_n contains exactly one reflexive node. The definition is as follows.

Definition 5. *Let $n \geq 1$ and fix $s \in W_n$. We define $\tilde{\mathcal{K}}_n = (\tilde{W}_n, \tilde{R}_n)$, where $\tilde{W}_n = W \times \{i, r\}$ and $(w, x)\tilde{R}_n(v, y)$ if and only if either $w \neq v$ and $x = y$ or $w = v = s$ and $x = y = r$. We call $W_n \times \{i\}$ the irreflexive component of $\tilde{\mathcal{K}}_n$ and $W_n \times \{r\}$ the reflexive component of $\tilde{\mathcal{K}}_n$.*




Figure 1: The graph $\tilde{\mathcal{K}}_3$.

Example 1. The graph $\tilde{\mathcal{K}}_3$ is shown in Figure 1.

Obviously, because of the reflexive point, any $\tilde{\mathcal{K}}_n$ is a non- n -colourable graph.

Lemma 1. For every valuation V on \mathcal{K}_n , if there are two vertices in \mathcal{K}_n that satisfy the same propositional variables, then there is a valuation \tilde{V} on $\tilde{\mathcal{K}}_n$, such that the model $\mathcal{K}_n^M = (\mathcal{K}_n, V)$ is L_{\diamond}^{\forall} -bisimilar to the model $\tilde{\mathcal{K}}_n^M = (\tilde{\mathcal{K}}_n, \tilde{V})$.

Proof. Let us fix a pair of nodes u, s in \mathcal{K}_n that satisfy the same propositional variables. The model $\tilde{\mathcal{K}}_n^M$ consists of two disjoint copies of the model \mathcal{K}_n^M but in one of the copies one of the points u or s is reflexive. It is easy to see that \mathcal{K}_n^M is L_{\diamond}^{\forall} -bisimilar to $\tilde{\mathcal{K}}_n^M$. \square

Example 2. The bisimilar $\tilde{\mathcal{K}}_3^M$ and \mathcal{K}_3^M are shown in Figure 2 on the left and right of the dotted line, respectively. Nodes with the same colour satisfy the same variables.




Figure 2: $\tilde{\mathcal{K}}_3^M$ and \mathcal{K}_3^M .

Proposition 2. For any $n \geq 2$, any L_{\diamond}^{\forall} -formula φ that defines the property of a graph being non- n -colourable contains at least $\lceil \log_2 n \rceil$ different propositional symbols.

Proof. Suppose that $l < \lceil \log_2 n \rceil$ and let ψ be a L_{\diamond}^{\forall} -formula containing only l different propositional variables, say p_1, \dots, p_l . We are going to show that this formula is either valid on \mathcal{K}_n or not valid on $\tilde{\mathcal{K}}_n^M$, and hence ψ does not define the property of not being n -colourable. Assume that ψ is not valid on \mathcal{K}_n and let V, w be a valuation and a point, respectively such that $(\mathcal{K}_n, V, w) \not\models \psi$. It follows from $l < \lceil \log_2 n \rceil$ that $n > 2^l$ and, therefore, there are at least two different nodes u and v in \mathcal{K}_n that satisfy the same subset of $\{p_1, \dots, p_l\}$. Applying Lemma 1, we obtain a valuation \tilde{V} on $\tilde{\mathcal{K}}_n$ and a point \tilde{w} in $\tilde{\mathcal{K}}_n$ such that $(\tilde{\mathcal{K}}_n, \tilde{V}, \tilde{w})$ is L_{\diamond}^{\forall} -bisimilar to (\mathcal{K}_n, V, w) . Hence, $(\tilde{\mathcal{K}}_n, \tilde{V}, \tilde{w}) \not\models \psi$. \square

This establishes the lower bound on the number of variables of Theorem 3. For the rest of the properties, fix an $n \geq 1$ and consider a $(L_{\diamond}^{\forall}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF with $\mathbf{A} = \{\tilde{\mathcal{K}}_n\}$ and $\mathbf{B} = \{\mathcal{K}_n\}$. Clearly, the formula φ_n of Definition 4 is valid on the frame in \mathbf{A} and not valid on the frame in \mathbf{B} . Below we give the strategy Hercules must follow if he wishes to win the game when the Hydra plays greedily. We begin with his selection of models.

SELECTION OF THE MODELS ON THE RIGHT: It follows from Lemma 1 that if Hercules wants to win the subsequent FGM, he must choose his model $\mathcal{B} = (\mathcal{K}_n, V)$ so that any two different vertices of \mathcal{B} satisfy different sets of literals. Let the singleton set \mathbf{B}_n^m contain the pointed model (\mathcal{B}, w) chosen by Hercules, where $w \in W_n$ is arbitrary.

SELECTION OF THE MODELS ON THE LEFT: The Hydra constructs a set \mathbf{A}_n^m of n different pointed models based on $\tilde{\mathcal{K}}_n$ as follows. Intuitively, for each $w \in W_n$, she forms a model \mathcal{A}_w consisting of two copies of \mathcal{B} , where in the second copy the reflexive point satisfies the same propositional variables as w . Formally, let (s, r) be the unique reflexive point of $\tilde{\mathcal{K}}_n$. For each $w \in W_n$, let π_w be a permutation of W_n such that $\pi_w(s) = w$. We define $V_w(u, x) = V(\pi_w(u))$ and $\mathcal{A}_w = (\tilde{\mathcal{K}}_n, V_w)$. Finally, we set $\mathbf{A}_n^m = \{\mathcal{A}_w \mid w \in W_n\}$.

Convention: We will henceforth identify a vertex $w \in W_n$ with the set of propositional variables $V(w)$; Note that, since Hercules assigns different valuations to different points, a set of variables E can name at most one vertex. Similarly, we will denote a vertex (v, x) of \mathcal{A}_w by E^x if $E = V_w(v, x)$.

Example 3. The sets of pointed models \mathbf{A}_3^m and \mathbf{B}_3^m are shown in Figure 3. Nodes that satisfy the same literals are given identical colours. Let B , G , and W denote the set of literals true on the black, grey, and white point, respectively. Suppose that Hercules has chosen the pointed model (\mathcal{B}, B) shown on the right of the dotted line. The Hydra responds with the pointed models (\mathcal{A}_B, B^i) , $((\mathcal{A}_W, B^i))$, and (\mathcal{A}_G, B^i) shown on the left.




Figure 3: The sets \mathbf{A}^3 and \mathbf{B}^3 .

FORMULA SIZE GAME ON MODELS: We consider $(L_{\diamond}^{\forall}, \langle \mathbf{A}_n^m, \mathbf{B}_n^m \rangle)$ -FGM.

Definition 6. A special pair of pointed models is a pair $\langle (\mathcal{A}_S, E), (\mathcal{B}, E') \rangle$, where $E = E'$.

Proposition 3. For any game tree T for a FGM and any node η of T , if there is a special pair $\langle (\mathcal{A}_S, E^x), (\mathcal{B}, E) \rangle$ with $(\mathcal{A}_S, E^x) \in \mathcal{L}(\eta)$ and $(\mathcal{B}, E) \in \mathfrak{R}(\eta)$, then

1. Hercules did not play a literal move at η ;
2. if $x = i$ and Hercules did not play an \exists -move at η , then, for at least one daughter η_1 of η , there is a special pair $\langle (\mathcal{A}_S, U^1), (\mathcal{B}, U) \rangle$ such that $(\mathcal{A}_S, U^1) \in \mathcal{L}(\eta_1)$ and $(\mathcal{B}, U) \in \mathfrak{R}(\eta_1)$.

Proof. The first item is obvious. For the second item we have to consider \vee -, \wedge -, \diamond -, \square -, and \forall -moves. If Hercules played either an \vee - or an \wedge -move at η , the statement is clearly true. If Hercules played a \diamond -move, since E^1 is a point in the irreflexive component of $\tilde{\mathcal{K}}_n$, he must have picked a successor (\mathcal{A}_S, U^i) of (\mathcal{A}_S, E^i) with $U \neq E$. Since the Hydra plays greedily, she is going to pick, among others, the pointed model $(\mathcal{B}, U) \in \square(\mathcal{B}, E)$ and the statement follows. The cases for \square - and \forall -moves are treated similarly. \square

Lemma 2. For any classes of pointed models \mathbf{L} and \mathbf{R} such that $\mathbf{A}_n^m \subseteq \mathbf{L}$, $\mathbf{B}_n^m \subseteq \mathbf{R}$, and Hercules has a winning strategy in the $(L_{\diamond}^{\forall}, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM, if T is a closed game tree for this game and the Hydra played greedily, then at least one node of T is an \exists -move.

Proof. Let ρ denote the root of T and let $\langle(\mathcal{A}_S, E^i), (\mathcal{B}, E)\rangle$ be a special pair with $(\mathcal{A}_S, E^i) \in \mathfrak{L}(\rho)$ and $(\mathcal{B}, E) \in \mathfrak{R}(\rho)$. If Hercules did not play an \exists -move during the game, we see, using Proposition 3, that T is not a closed game tree which is a contradiction. \square

Lemma 3. *For any classes of pointed models \mathbf{L} and \mathbf{R} such that $\mathbf{A}_n^m \subseteq \mathbf{L}$, $\mathbf{B}_n^m \subseteq \mathbf{R}$, and Hercules has a winning strategy in the $(\mathbf{L}_{\diamond}^{\forall}, \langle\mathbf{L}, \mathbf{R}\rangle)$ -FGM, if T is a closed game tree for this game and the Hydra played greedily, then T has at least n nodes.*

The proof of the lemma revolves around the notion of *weight function* or *complexity functional* – a popular tool in Boolean function complexity [20] – that allows us to formulate proofs by induction on a notion of “progress during a FGM”.

Definition 7. *For any finite binary tree T , a weight function f for T is a function that assigns to any node η of T a non-negative real number such that*

1. *if η is a leaf, then $f(\eta) \leq 1$;*
2. *if η has two immediate successors η_1 and η_2 , then $f(\eta) \leq f(\eta_1) + f(\eta_2) + 1$;*
3. *if η has one immediate successor η_1 , then $f(\eta) \leq f(\eta_1) + 1$.*

Lemma 4. *For any finite binary tree T and any weight function f for T , if ρ is the root of T , then T has at least $f(\rho)$ nodes.*

Proof. An easy induction on the number of nodes in T . \square

In order to prove Lemma 3 we define a suitable weight function on the nodes of T as follows. For a node η of T , we let $f(\eta)$ be the number of $S \in W_n$ such that there is at least one special pair $\langle(\mathcal{A}_S, E^x), (\mathcal{B}, E)\rangle$ with $(\mathcal{A}_S, E^x) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}, E) \in \mathfrak{R}(\eta)$.

Obviously, if ρ denotes the root of T , then $f(\rho) = n$. It remains to check the following.

Lemma 5. *The function f defined above is a weight function.*

Proof. We need to show that f satisfies the three items from Definition 7.

1. If η is a leaf, then it follows from item 1 of Proposition 3 that there is no special pair $\langle(\mathcal{A}_S, E^x), (\mathcal{B}, E)\rangle$ with $(\mathcal{A}_S, E^x) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}, E) \in \mathfrak{R}(\eta)$. Hence, $f(\eta) \leq 1$.
2. If η has two successors η_1 and η_2 , then η is either a \vee - or a \wedge -move. It is immediate that $f(\eta_1) + f(\eta_2) \geq f(\eta)$. Thus, the second condition of Definition 7 is fulfilled.
3. If η has one immediate successor η' , then η is a \forall -, \exists -, \square -, or \diamond -move. Let Γ be the set of all $S \in W_n$ such that there is a special pair $\langle(\mathcal{A}_S, E^x), (\mathcal{B}, E)\rangle$ with $(\mathcal{A}_S, E^x) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}, E) \in \mathfrak{R}(\eta)$ and define Γ' with η' in place of η . In each case we claim that there is $\Delta \subseteq \Gamma$ with $|\Delta| \leq 1$ such that $\Gamma \setminus \Delta \subseteq \Gamma'$, from which we obtain $f(\eta') \geq f(\eta) + 1$. We consider the following cases.

η IS A \forall -MOVE. Let $S \in \Gamma$ so that, for some E , we have that $(\mathcal{A}_S, E^x) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}, E) \in \mathfrak{R}(\eta)$. If Hercules picks (\mathcal{B}, U) as a successor of (\mathcal{B}, E) , then, since the Hydra is playing greedily, the pointed model (\mathcal{A}_S, U^x) , where $x = i$, is going to be in $\mathfrak{L}(\eta')$. Therefore, U witnesses that $S \in \Gamma'$ and $\Gamma \subseteq \Gamma'$.

η IS AN \exists -MOVE. Let us consider a special pair $\langle(\mathcal{A}_S, E^x), (\mathcal{B}, E)\rangle$ with $(\mathcal{A}_S, E^x) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}, E) \in \mathfrak{R}(\eta)$. Let $y \in \{i, r\}$ and suppose that Hercules places (\mathcal{A}_S, U^y) in $\mathfrak{L}(\eta')$ as the successor of (\mathcal{A}_S, E^x) , then the Hydra’s greedy strategy guarantees that $(\mathcal{B}, U) \in \mathfrak{R}(\eta')$ and $\Gamma \subseteq \Gamma'$.

η IS A \square -MOVE. Let $S \in \Gamma$, so that $(\mathcal{A}_S, E^x) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}, E) \in \mathfrak{R}(\eta)$ for some E . If Hercules picks (\mathcal{B}, U) as a successor of (\mathcal{B}, E) , then, the Hydra, by playing greedily, puts a pointed model (\mathcal{A}_S, U^x) in $\mathfrak{L}(\eta')$ because there are U -points in both the reflexive and the irreflexive component of \mathcal{K}_S . Therefore, $\Gamma \subseteq \Gamma'$.

η IS A \diamond -MOVE. Partition Γ into two subsets Σ and $\Delta = \Gamma \setminus \Sigma$, where $S \in \Sigma$ if there are $E \neq F$ and x, y such that $(\mathcal{A}_S, E^x) \in \mathfrak{L}(\eta)$, $(\mathcal{B}, E) \in \mathfrak{R}(\eta)$, and $(\mathcal{A}_S, F^y) \in \mathfrak{L}(\eta')$. For such an S , the Hydra's greedy strategy implies that (\mathcal{B}, F) is also among the successors of (\mathcal{B}, E) picked by her and thus $S \in \Gamma'$. Since $S \in \Sigma$ was arbitrary, $\Sigma \subseteq \Gamma'$.

If $|\Delta| \leq 1$, we are done, since then $\Gamma \setminus \Delta \subseteq \Gamma'$. So, assume otherwise and let $S \in \Delta$. Since $S \in \Gamma$, there is a set of variables E and $x \in \{\text{i,r}\}$ such that $(\mathcal{A}_S, E^x) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}, E) \in \mathfrak{R}(\eta)$. Let $(\mathcal{A}_S, F^x) \in \mathfrak{L}(\eta')$ be the successor chosen by Hercules; since $S \notin \Sigma$, we have $F = E$, hence E^x is the unique reflexive point of \mathcal{A}_S so that $E^x = S^x$. Using the assumption that $|\Delta| > 1$, let $U \neq S$ be another element of Δ . As above, we have that $(\mathcal{B}, U) \in \mathfrak{R}(\eta)$, hence the Hydra's greedy strategy implies that $(\mathcal{B}, S) \in \mathfrak{R}(\eta')$. Since also $(\mathcal{A}_S, S^x) \in \mathfrak{L}(\eta')$, we have that $S \in \Gamma'$, as needed. \square

Lemma 3 is an immediate consequence of Lemmas 4 and 5. With this, we have established all claims of Theorem 3.

6 The transfer axioms

We begin with what we call the *transfer axioms*, defined as $\text{TA}(m, n) = \diamond^m p \rightarrow \diamond^n p$, where $m \neq n \in \mathbb{N}$; since we treat $\varphi \rightarrow \psi$ as an abbreviation, we can rewrite these axioms as $\square^m \bar{p} \vee \diamond^n p$. It is well-known that $\text{TA}(m, n)$ defines the first-order property of (m, n) -transfer (1) from the introduction. As special cases we have that $(2, 1)$ -transfer is just transitivity and $(0, 1)$ -transfer is reflexivity. Instead of (m, n) -transfer we write n -reflexivity when $m = 0$, m -recurrence when $n = 0$, (m, n) -transitivity when $m > n > 0$ and (m, n) -density when $0 < m < n$.

Our goal is to prove the following.

Theorem 4. *For any $n \neq m \in \mathbb{N}$, $\square^m \bar{p} \vee \diamond^n p$ is absolutely minimal among all formulae defining (m, n) -transfer.*

The proof that for $m, n \geq 0$, $\diamond^m p \rightarrow \diamond^n p$ is essentially the shortest formula defining (m, n) -transfer is split in four parts according to the ordering between m and n .

6.1 Generalized density axioms

First, we study the (m, n) -density axioms, i.e., (m, n) -transfer with $0 < m < n$. We prove that Theorem 4 holds in this case by considering a $(\mathbb{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF where $\mathbf{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_{m+1}\}$ and \mathbf{B} contains a single element \mathcal{B} . These frames are shown in the left rectangle in Figure 4 and separated by the dotted line. \mathcal{A}_1 is constructed so that the vertical path leading from the




Figure 4: The frames $\mathcal{A}_1, \dots, \mathcal{A}_{m+1}$ and \mathcal{B} and the pointed models based on them.

lowest non-reflexive point to the uppermost non-reflexive one consists of m steps whereas the rightmost path that starts and ends respectively with these two points consists of n steps (not counting the reflexive steps) and every point on this rightmost path is reflexive. The frame \mathcal{B} is obtained from \mathcal{A}_1 by simply erasing the latter path. For $2 \leq i \leq m+1$, each \mathcal{A}_i contains a vertical path of $i-2$ steps. Obviously, $\diamond^m p \rightarrow \diamond^n p$ is valid in all frames in \mathbf{A} and not valid on \mathcal{B} .

SELECTION OF THE MODELS ON THE RIGHT: If Hercules wishes to win the game, he must choose his pointed models with some care.

Lemma 6. *In any winning strategy for Hercules for an $(\mathbf{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGF in which $\mathcal{A}_1 \in \mathbf{L}$ and $\mathcal{B} \in \mathbf{R}$, he picks a pointed model $(\mathcal{B}^M, \triangleright)$ based on the lowest irreflexive point in \mathcal{B} .*

Proof. It is easily seen that Hercules is not going to select a pointed model that is not based on the lowest irreflexive point in \mathcal{B} because the Hydra can reply with a bisimilar pointed model based on \mathcal{A}_1 . \square

SELECTION OF MODELS ON THE LEFT: The Hydra replies with the pointed models shown on the left of the dotted line in the right rectangle in Figure 4. She constructs them as follows. Using the fact that \mathcal{B} is a sub-structure of \mathcal{A}_1 , the Hydra makes sure that the same points in \mathcal{A}_1^M and \mathcal{B}^M satisfy the same literals; moreover, the black points in both models satisfy the same literals, too. The models \mathcal{A}_i^M for $2 \leq i \leq m+1$ receive valuations that make them initial segments of the vertical path in \mathcal{B}^M , i.e., the lowest non-reflexive point in any \mathcal{A}_i^M and the lowest non-reflexive point in \mathcal{B}^M satisfy the same literals and similarly for their vertical successors. When the Hydra chooses her pointed models in this way, we say she *mimics* Hercules' choice.

FORMULA SIZE GAME ON MODELS: Consider the FGM starting with $(\mathcal{A}_1^M, \triangleright), \dots, (\mathcal{A}_{m+1}^M, \triangleright)$ on the left and $(\mathcal{B}^M, \triangleright)$ on the right. First, we show that there are some constraints on the moves that Hercules may make.

Lemma 7. *Let \mathbf{L}, \mathbf{R} be classes of pointed models and let Hercules has a winning strategy for the $(\mathbf{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM. Let T be any closed game tree on which the Hydra played greedily and η be any position of T such that $(\mathcal{B}^M, \triangleright) \in \mathfrak{R}(\eta)$ while $(\mathcal{A}_i, \triangleright) \in \mathfrak{L}(\eta)$ for some i with $1 \leq i \leq m+1$.*

1. *If Hercules played a \diamond -move at η , then he did not pick the left lowest reflexive point in \mathcal{A}_i^M , and if $i = 1$, then he picked the bottom-right reflexive point on \mathcal{A}_1^M .*
2. *If Hercules played a \square -move at η , he did not pick the left lowest reflexive point in \mathcal{B}^M .*

Proof. If Hercules picks the left lowest reflexive point when playing a \diamond -move, the Hydra is going to reply with the same point in \mathcal{B}_1^M and obtain bisimilar pointed models on each side. If $i = 1$ and Hercules picks the unique irreflexive successor on \mathcal{A}_1^M , then Hydra can reply with the irreflexive successor on \mathcal{B}^M , which means by Corollary 1 that Hercules cannot win. The second claim is symmetric. \square

Lemma 8. *Suppose that \mathbf{L}, \mathbf{R} are classes of pointed models and Hercules has a winning strategy for the $(\mathbf{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM. If T is any closed game tree in which the Hydra played greedily and η is any position of T such that $(\mathcal{B}^M, \triangleright) \in \mathfrak{R}(\eta)$, then*

1. *if $(\mathcal{A}_1^M, \triangleright) \in \mathfrak{L}(\eta)$, then Hercules did not play a \square -move on η ;*
2. *if $(\mathcal{A}_2^M, \triangleright) \in \mathfrak{L}(\eta)$, then Hercules did not play a \diamond -move on η .*

Proof. The first claim is immediate since if Hercules played a \square -move, the Hydra could reply with the same point in \mathcal{A}_1^M and obtain bisimilar pointed models on each side. For the second, Hercules is forced to pick the reflexive point in \mathcal{A}_2^M when playing a \diamond -move which contradicts Lemma 7. \square

With this we can establish lower bounds on the number of moves of each type that Hercules must make, as established by the proposition below.

Proposition 4. *Let \mathbf{L}, \mathbf{R} be classes of pointed models such that Hercules has a winning strategy for the $(\mathbf{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM and let T be a closed game tree in which the Hydra played greedily.*

1. *If $\{(\mathcal{A}_1^M, \triangleright), (\mathcal{A}_2^M, \triangleright)\} \subseteq \mathbf{L}$ and $(\mathcal{B}, \triangleright) \in \mathbf{R}$, then Hercules made at least one \vee -move.*
2. *If $(\mathcal{A}_1^M, \triangleright) \in \mathbf{L}$, and $(\mathcal{B}^M, \triangleright) \in \mathbf{R}$, then T has modal depth at least n , at least n \diamond -moves and one literal.*
3. *If $\{(\mathcal{A}_2^M, \triangleright), \dots, (\mathcal{A}_{m+1}^M, \triangleright)\} \subseteq \mathbf{L}$ and $(\mathcal{B}^M, \triangleright) \in \mathbf{R}$, then Hercules made at least m \square -moves during the game.*

Proof. (1) By Lemma 8, Hercules cannot play a modality as long as $(\mathcal{A}_1^M, \triangleright), (\mathcal{A}_2^M, \triangleright)$ are both on the left and $(\mathcal{B}, \triangleright)$ on the right, and the three satisfy the same literals, so that he cannot play a literal either. Playing a \wedge -move would lead to at least one new game position that is the same as the previous one. Hence, every winning strategy for Hercules must ‘separate’ $(\mathcal{A}_1, \triangleright)$, from $(\mathcal{A}_2, \triangleright)$ with an \vee -move.

(2) Note that $(\mathcal{A}_1^M, \triangleright)$ and $(\mathcal{B}^M, \triangleright)$ satisfy the same literals and \vee - and \wedge -moves lead to at least one new game-position in which $(\mathcal{A}_1^M, \triangleright)$ is on the left and $(\mathcal{B}^M, \triangleright)$ is on the right. By Lemma 8.1, Hercules cannot play a \square -move in any of these positions. Thus Hercules must make a \diamond -move in a position in which $(\mathcal{A}_1^M, \triangleright)$ is on the left and $(\mathcal{B}^M, \triangleright)$ is on the right. By Lemma 7 he is going to pick the first reflexive point on the rightmost path in \mathcal{A}_1^M . The Hydra replies with, among others, the left lowest reflexive point in \mathcal{B}^M . Since this point satisfies the same literals as the reflexive points lying on the rightmost path in \mathcal{A}_1^M , Hercules cannot play a literal-move; moreover, \vee -, \wedge - and \square -moves lead to at least one new game position that is essentially the same as the previous one. In the case of \square -moves this is true because, when playing such a move, Hercules must stay in the lowest reflexive point in \mathcal{B}^M while the Hydra can stay in the current reflexive point on the rightmost path in \mathcal{A}_1^M . Hence, he must make at least $n - 1$ subsequent \diamond -moves to reach a point in \mathcal{A}_1^M that differs on a literal from the lowest reflexive point in \mathcal{B}^M . Finally he must play a literal, as no other move can close the tree.

(3) Fix $i \in [2, m + 1]$. Let w_1, \dots, w_{i-1} enumerate the vertical path of \mathcal{A}_i starting at the root; similarly, let v_1, \dots, v_m enumerate the vertical path of \mathcal{B} . Let $\mathbf{w}_j = (\mathcal{A}_i^M, w_j)$ and $\mathbf{v}_j = (\mathcal{B}^M, v_j)$.

Say that a branch $\vec{\nu} = (\nu_0, \dots, \nu_k)$ on T is *i-critical* if there exists $j \in [1, i]$ with $\mathbf{w}_j \in \mathfrak{L}(\nu_k)$, $\mathbf{v}_j \in \mathfrak{R}(\nu_k)$ and Hercules has played exactly $j - 1$ modal moves on ν_1, \dots, ν_{k-1} . Since T is finite and the singleton branch consisting of the root is *i-critical*, we can pick a maximal *i-critical* branch $\vec{\eta} = (\eta_0, \dots, \eta_\ell)$ for some value of j .

We claim that $j = i - 1$ and Hercules plays a \square -move on η_ℓ . Since T is closed, η_ℓ cannot be a head, but \mathbf{w}_j and \mathbf{v}_j share the same valuation so it cannot be a stub either, thus η_ℓ is not a leaf. If Hercules played a \wedge - or a \vee -move, then η_ℓ would have a daughter giving us a longer *i-critical* branch. Thus Hercules played a modality on η_ℓ . If $j < i - 1$, then, for the unique daughter η' of η_ℓ , we have $\mathbf{w}_{j+1} \in \mathfrak{L}(\eta')$ and $\mathbf{v}_{j+1} \in \mathfrak{R}(\eta')$, where in the case of $j = 0$ we use Lemma 7 and otherwise there simply are no other options for Hercules; but this once again gives us a longer *i-critical* branch. Thus $j = i - 1$; but then Hercules is not allowed to play \diamond , as there is a pointed model on the left without successors, so he played a \square -move on η_ℓ .

We conclude that for each $i \in [2, m + 1]$ there is an instance of \square of modal depth exactly $i - 1$, which implies that each instance is distinct. \square

With this we prove Theorem 4 in the case $0 < m < n$.

Proof. If $0 < m < n$ we consider the $(\mathbf{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF with \mathbf{A}, \mathbf{B} as shown in Figure 4. By Lemma 6 Hercules chooses some pointed model \mathcal{B}^M based on the irreflexive point at the bottom

of \mathcal{B} , and Hydra replies by mimicking Hercules' pointed models. Then by Proposition 4 Hercules must play at least one \vee , one literal, n \diamond -moves, the modal depth of the tree is n , and at least m \square -moves. By Theorem 2, any formula valid on every frame of \mathbf{A} and no frame of \mathbf{B} must satisfy these bounds; but the frames in \mathbf{A} satisfy the (m, n) -transfer property while those in \mathbf{B} do not. \square

6.2 Generalized transitivity axioms

Next we treat Theorem 4 in the case where $0 < n < m$. As before, we do so by considering a suitable $(L_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF where $\mathbf{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_{m+1}\}$ and \mathbf{B} contains a single element \mathcal{B} , but now using the frames shown in Figure 5. The frame \mathcal{A}_1 is based on a right-angled triangle in which the sum of the relation steps in the legs is m whereas the number of relation steps in the hypotenuse is n ; moreover, each path on the left of the hypotenuse that shares nodes with it consist of n relation steps, too. The frame \mathcal{B} is obtained from \mathcal{A}_1 by "separating" the hypotenuse from the horizontal leg and erasing the points that do not lie either on the hypotenuse or on the legs of \mathcal{A}_1 . Each \mathcal{A}_i , for $2 \leq i \leq m+1$, contains a vertical path of $i-2$ relation steps and a diagonal one of n relation steps. Obviously, $\diamond^m p \rightarrow \diamond^n p$ is valid in all frames in \mathbf{A} and not valid on \mathcal{B} .

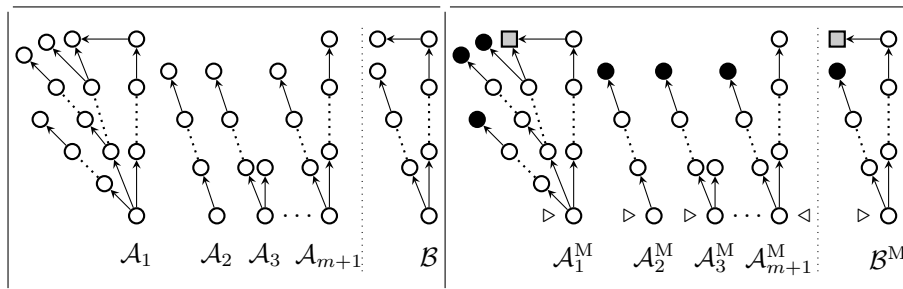


Figure 5: The frames $\mathcal{A}_1, \dots, \mathcal{A}_{m+1}$ and \mathcal{B} and the pointed models based on them.

SELECTION OF THE MODELS ON THE RIGHT: In this case, Hercules must choose his models according to the following.

Lemma 9. *In any winning strategy for Hercules for an $(L_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGF in which $\mathcal{A}_1 \in \mathbf{L}$ and $\mathcal{B} \in \mathbf{R}$ he picks a pointed model $(\mathcal{B}^M, \triangleright)$ based on the lowest point in \mathcal{B} , and assigns different valuations to the two dead-end points of \mathcal{B} .*

Proof. Hercules is not going to select a pointed model that is not based on the lowest point in \mathcal{B} because the Hydra can reply with a bisimilar pointed model based on \mathcal{A}_1 ; similarly, if he assigns the same valuation to the two dead-ends, she can choose a bisimilar model based on \mathcal{A}_1 by copying the valuations from the hypotenuse onto all paths of length n , and copying the valuations from the legs onto the path of length m ; since the valuations coincide on the end-points, there is no clash at the top left of the triangle. \square

To indicate that the two end-points of \mathcal{B} receive different valuations, we have drawn one of them black while the other is shaped as a rectangle. The literals true in the rest of the points are immaterial. Thus, Hercules constructs the pointed model $(\mathcal{B}^M, \triangleright)$ shown in the right rectangle in Figure 5.

SELECTION OF MODELS ON THE LEFT: The Hydra replies with the pointed models shown on the left of the dotted line in the right rectangle in Figure 5. The pointed model based on \mathcal{A}_1 is defined so that the set of literals true in the points on a diagonal path that shares points with the hypotenuse but do not coincide with it copy the respective sets of literals true in the points of the diagonal path in \mathcal{B} .

The models \mathcal{A}_i for $2 \leq i \leq m+1$ receive valuations so that their diagonal paths coincide with the diagonal path in the model \mathcal{B} whereas their vertical paths are ‘initial segments’ of the vertical path in \mathcal{B} , i.e., the lowest point in any \mathcal{A}_i for $2 \leq i \leq m+1$ and the lowest point in \mathcal{B} satisfy the same literals; similarly for their vertical successors. As before, if the Hydra picks her models in this way, we say that she *mimics* Hercules’ choice.

FORMULA SIZE GAME ON MODELS: Consider the FGM starting with $(\mathcal{A}_1^M, \triangleright), \dots, (\mathcal{A}_{m+1}^M, \triangleright)$ on the left and $(\mathcal{B}^M, \triangleright)$ on the right. These lemmas are analogous to those in Section 6.1.

Lemma 10. *Let \mathbf{L}, \mathbf{R} be classes of models so that Hercules has a winning strategy for the $(\mathbf{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM. Let T be any closed game in which the Hydra played greedily and η be a node on which Hercules played a \diamond -move.*

1. *If $(\mathcal{A}_1^M, \triangleright) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}^M, \triangleright) \in \mathfrak{R}(\eta)$, then he picked a pointed model based on a point that lies on the hypotenuse of \mathcal{A}_1^M .*
2. *If for some $i \in [3, m+1]$ we have that $(\mathcal{A}_i^M, \triangleright) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}^M, \triangleright) \in \mathfrak{R}(\eta)$, then he picked the rightmost daughter as a successor of $(\mathcal{A}_i^M, \triangleright)$.*

Proof. Both items hold because if Hercules picked a different point, the Hydra replied with the same point in \mathcal{B}^M . In either case we obtain bisimilar models on each side, which by Corollary 1 means that Hercules cannot win. \square

Lemma 11. *Suppose that \mathbf{L} and \mathbf{R} are classes of models and Hercules has a winning strategy for the $(\mathbf{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM. Suppose that T is a closed game tree, the Hydra played greedily, and η is a node of T .*

1. *If $(\mathcal{A}_1^M, \triangleright) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}^M, \triangleright) \in \mathfrak{R}(\eta)$, then Hercules did not play a \square -move at η .*
2. *If $(\mathcal{A}_2^M, \triangleright) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}^M, \triangleright) \in \mathfrak{R}(\eta)$, then Hercules did not play a \diamond -move at η .*

Proof. The first item is immediate from the fact that if Hercules played a \square -move, the Hydra can reply with the same point in \mathcal{A}_1^M , and similarly in the second case the Hydra would reply with the same pointed model based on \mathcal{B}^M . \square

As was the case for the generalized density axioms, Hercules must play at least one \vee -move to separate \mathcal{A}_1^M from \mathcal{A}_2^M .

Proposition 5. *Let \mathbf{L} and \mathbf{R} be classes of models such that Hercules has a winning strategy for the $(\mathbf{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM. Let T be a closed game tree in which the Hydra played greedily.*

1. *If $(\mathcal{A}_1, \triangleright), (\mathcal{A}_2, \triangleright) \in \mathbf{L}$ and $(\mathcal{B}, \triangleright) \in \mathbf{R}$, then Hercules made at least one \vee -move.*
2. *If $(\mathcal{A}_1^M, \triangleright) \in \mathbf{L}$, then T has at least n nested \diamond -moves and at least one literal move.*
3. *If $\{(\mathcal{A}_2^M, \triangleright), \dots, (\mathcal{A}_{m+1}^M, \triangleright)\} \subseteq \mathbf{L}$, then T has at least m \square -moves.*

Proof. The proof of the first item is analogous to that of Proposition 4.1, except that it uses Lemma 18, and the proof of the third item is essentially the same as the proof of Proposition 4.3. Thus we focus on the second item.

Since $(\mathcal{A}_1^M, \triangleright)$ and $(\mathcal{B}^M, \triangleright)$ satisfy the same literals and since \vee - and \wedge -moves lead to at least one new game-position in which $(\mathcal{A}_1^M, \triangleright)$ is on the left and $(\mathcal{B}^M, \triangleright)$ is on the right, Hercules must make a \diamond -move in a position in which $(\mathcal{A}_1^M, \triangleright)$ is on the left and $(\mathcal{B}^M, \triangleright)$ is on the right. It follows from Lemma 10, that he is going to pick the immediate successor along the hypotenuse of

\mathcal{A}_1^M . The Hydra replies, with among others, the immediate successor along the diagonal path in \mathcal{B}^M . Since the new pointed models satisfy the same literals, Hercules cannot play a literal-move; moreover, \vee - and \wedge -moves lead to at least one new game position that is essentially the same as the previous one. If he played a \square -move and picked a pointed model based on a point along the diagonal path in \mathcal{B}^M , the Hydra would reply with the same point along a path that is different from the hypotenuse because such paths are always available. Hence, he must make at least $n - 1$ subsequent \diamond -moves to reach the point in which the hypotenuse of \mathcal{A}_1^M and its horizontal leg meet. Finally, at this point Hercules must play a literal, as this is the only move that will lead to a closed game-tree. \square

With this we conclude the proof of Theorem 4 in the case $0 < n < m$.

Proof. Similar to the proof for the case $0 < m < n$, except that we use the classes \mathbf{A} , \mathbf{B} of Figure 5 and Proposition 6. \square

We proceed to proving Theorem 4 in the cases where one of the parameters is zero.

6.3 The generalized reflexivity axioms

Recall that we write *n-reflexivity* instead of $(0, n)$ -transfer. In order to prove that Theorem 4 holds in this case, we consider a $(L_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF where $\mathbf{A} = \{\mathcal{A}_1, \mathcal{A}_2\}$ and $\mathbf{B} = \{\mathcal{B}\}$. These frames are shown in the left rectangle in Figure 6 and separated by the dotted line. The “highest” point in \mathcal{A}_2 is reachable in $n - 1$ relation steps from the lowest one and then we can return back to the latter in one additional relation step, i.e, the points in \mathcal{A}_2 that are different from the reflexive one form a cycle of length n . It is obvious that $p \rightarrow \diamond^n p$ is valid on both \mathcal{A}_1 and \mathcal{A}_2 and not valid on \mathcal{B} .




Figure 6: The frames \mathcal{A}_1 , \mathcal{A}_2 and \mathcal{B} and the pointed models based on them.

Consider Hercules’ choice of models on the right.

SELECTION OF THE POINTED MODELS ON THE RIGHT: If Hercules is to win the formula-complexity game, he must choose his models in a specific way.

Lemma 12. *In any winning strategy for Hercules for an $(L_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGF in which $\mathcal{A}_1 \in \mathbf{L}$ and $\mathcal{B} \in \mathbf{R}$,*

1. *he chooses the valuation on \mathcal{B} so that at least one literal is true in one point but not on the other, and*
2. *he picks the pointed model based on the irreflexive point in \mathcal{B} .*

The pointed model based on \mathcal{B} and its irreflexive point chosen by Hercules is shown in the right half of Figure 6. We indicate that the two points in \mathcal{B} satisfy different sets of literals by making one of them black and the other white.

SELECTION OF THE POINTED MODELS ON THE LEFT: The Hydra replies with the pointed models shown on the left of the dotted line in the right half in Figure 6. Two points in any two models satisfy the same set of literals iff they have the same colour. As usual, we say that she *mimics* Hercules if she chooses her pointed models in this way.

FORMULA SIZE GAME ON MODELS: Consider the FGM starting with $(\mathcal{A}_1^M, \triangleright), (\mathcal{A}_2^M, \triangleright)$ on the left and $(\mathcal{B}^M, \triangleright)$ on the right. We first note some restrictions on Hercules's modal moves. The following can be seen by observing that playing otherwise would produce bisimilar pointed models on each side.

Lemma 13. *Let \mathbf{L}, \mathbf{R} be classes of models so that Hercules has a winning strategy for the $(\mathbf{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM and T a closed game tree in which the Hydra played greedily.*

1. *If there is a game position η in which any pointed model based on either \mathcal{A}_1^M or \mathcal{A}_2^M is on the left and any pointed model based on \mathcal{B}^M is on the right, then Hercules did not play a \square -move at η .*
2. *If there is a game position η in which $(\mathcal{A}_1^M, \triangleright)$ is on the left and a pointed model based on \mathcal{B}^M is on the right, then Hercules did not play a \diamond -move at η .*

From this it is easy to see that Hercules must play at least one variable.

Lemma 14. *Suppose that \mathbf{L}, \mathbf{R} are classes of models and that Hercules has a winning strategy for the $(\mathbf{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM. Let T be a closed game tree in which the Hydra played greedily and such that there is a position η in which $(\mathcal{A}_1^M, \triangleright)$ is on the left and $(\mathcal{B}^M, \triangleright)$ is on the right. Then, the number of literal moves in T is at least one.*

Proof. By Lemma 13 Hercules cannot play any \diamond - or \square -moves, and \wedge - or \vee -moves result in at least one new position with both of these pointed models. Since Hercules cannot play \perp or \top , he must use at least one variable. \square

With this we are ready to prove Theorem 4 in the case where $m = 0$.

Proof. Let \mathbf{A} and \mathbf{B} be as depicted in the left rectangle in Figure 6; since the frames of \mathbf{A} are n -reflexive but the ones in \mathbf{B} are not, by Theorem 2 it suffices to show that the Hydra can play so that any closed game tree has at least one \vee -move, one literal move, and modal depth at least n .

Let $\mathbf{B}^m = \{(\mathcal{B}^M, \triangleright_{\mathcal{B}})\}$ be the singleton set of pointed models chosen by Hercules, which by Lemma 12 must be so that the top and bottom points have different valuations, and let Hydra choose \mathbf{A}^m as depicted in the right-hand side of Figure 6. Lemma 13 implies that Hercules cannot begin the FGM starting with $(\mathcal{A}_1^M, \triangleright), (\mathcal{A}_2^M, \triangleright)$ on the left and $(\mathcal{B}^M, \triangleright)$ on the right by playing either a \diamond - or a \square -move. Playing an \wedge -move will result in at least one new position that is the same as the previous one. Therefore, Hercules must play an \vee -move and he and the Hydra will have to compete in two new sub-games: the first one starting with $(\mathcal{A}_1^M, \triangleright)$ on the left and $(\mathcal{B}^M, \triangleright)$ on the right while the second starts with $(\mathcal{A}_2^M, \triangleright)$ on the left and $(\mathcal{B}^M, \triangleright)$ on the right.

By Lemma 14 he can win the former only by playing a literal-move; the latter can be won only by playing a sequence of n \diamond -moves that must be made to perform a cycle leading back to the black point in \mathcal{A}_2 , giving at least n occurrences of \diamond and modal depth at least n . We use Theorem 2 to conclude that $\bar{p} \vee \diamond^n p$ is absolutely minimal. \square




Figure 7: The frames $\mathcal{A}_1, \dots, \mathcal{A}_{m+1}$ and \mathcal{B} and the pointed models based on them.

6.4 The generalized recurrence axioms

We treat the m -recurrence axioms when $n = 0$. This time Hercules and the Hydra play a $(\mathbb{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF where $\mathbf{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_{m+1}\}$ while \mathbf{B} has a single element \mathcal{B} as depicted in the left rectangle in Figure 7. For $2 \leq i \leq m+1$, each \mathcal{A}_i is a path of $i-2$ relation steps. Clearly, $\diamond^m p \rightarrow p$ is valid in the frames in \mathbf{A} and it is not valid in the frame \mathcal{B} .

SELECTION OF THE MODELS ON THE RIGHT: Lemma 12 implies that Hercules must pick the pointed model $(\mathcal{B}^M, \triangleright)$ shown in the right half of Figure 7. We colour one of the points of black and the other white to indicate that they satisfy different sets of literals.

SELECTION OF THE POINTED MODELS ON THE LEFT: The Hydra replies with the pointed models shown on the left of the dotted line in the right half of Figure 7. She picks these pointed models so that points that satisfy the same set of literals have the same colour.

FORMULA SIZE GAME ON MODELS: Consider the FGM starting with $\mathbf{A}^m = \{(\mathcal{A}_1^M, \triangleright), \dots, (\mathcal{A}_{m+1}^M, \triangleright)\}$ on the left and $\mathbf{B}^m = \{(\mathcal{B}^M, \triangleright)\}$ on the right.

Lemma 15. *In any closed game tree T for the $(\mathbf{A}^m, \mathbf{B}^m)$ -FGM in which the Hydra played greedily, Hercules played at least one \vee -move.*

Proof. Lemma 13 implies that in order to win a FGM with a starting position η where $(\mathcal{A}_1^M, \triangleright) \in \mathcal{L}(\eta)$ and $(\mathcal{B}^M, \triangleright) \in \mathcal{R}(\eta)$, Hercules must not play either a \diamond - or a \square -move at η . On the other hand, for every game in which there is some $(\mathcal{A}_i^M, \triangleright)$ for $2 \leq i \leq m+1$ among the pointed models chosen by the Hydra and $(\mathcal{B}^M, \triangleright)$ is among the models chosen by Hercules, if he wants to win the game, then there is at least one game position ν such that $(\mathcal{A}_i^M, \triangleright)$ is on the left and $(\mathcal{B}^M, \triangleright)$ is on the right and Hercules played at least one \diamond - or \square -move at ν . This implies that in any FGM with a starting position in which the pointed models selected by the Hydra are on the left and $(\mathcal{B}^M, \triangleright)$ is on the right, Hercules must play at least one \vee to separate the set of $(\mathcal{A}_i^M, \triangleright)$, for $2 \leq i$, from $(\mathcal{A}_1^M, \triangleright)$. \square

Lemma 16. *Let \mathbf{L}, \mathbf{R} be classes of models so that Hercules has a winning strategy for the $(\mathbb{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM. Let T be a closed game tree in which the Hydra played greedily. If all $(\mathcal{A}_i^M, \triangleright)$ for $2 \leq i \leq m+1$ are in \mathbf{L} and $(\mathcal{B}^M, \triangleright) \in \mathbf{R}$, Hercules must have played at least m \square -moves and the modal depth of T must be at least m .*

We omit the proof, which is similar to that of Proposition 4.3. With this we are ready to prove Theorem 4 for the case where $n = 0$.

Proof. Consider the (\mathbf{A}, \mathbf{B}) -FGF where \mathbf{A}, \mathbf{B} are as depicted in Figure 7 on the left: by Lemma 12, Hercules must choose different valuations for the points of \mathbf{B} and choose the bottom point. Let Hydra reply as depicted on the right-hand side of the figure.

By Lemma 14, Hercules must play at least one variable, by Lemma 15 he must play at least one \vee -move, by Lemma 19 he must play at least m \square -moves and modal depth at least m on the resulting FGM, and we can apply Theorem 2. \square

7 The S4 axiom

Theorem 5. *The formula $(\bar{p} \wedge \square \square \bar{p}) \vee \diamond p$ is absolutely minimal among the set of L_\diamond -formulas defining the class of reflexive and transitive frames.*

It is tempting to try to prove Theorem 5 by modifying slightly the frames in Figure 5 for $m = 2$ and $n = 1$ so as to take care of reflexivity and applying more or less the reasoning from Subsection 6.2. However, a closer look reveals that we have to ensure that Hercules is forced to make at least one \wedge -move. Hence, we must have at least two models on the right for the model equivalence game; moreover, there are no guarantees that Lemma 18 will remain true if we make the relations in the frames of \mathbf{A} reflexive (in fact it does not). Nevertheless, we make this strategy work by taking some extra care.

Let us consider a $(L_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF where $\mathbf{A} = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3\}$ and $\mathbf{B} = \{\mathcal{B}_1, \mathcal{B}_2\}$ as shown in the left rectangle in Figure 8. Obviously, $(\bar{p} \wedge \square \square \bar{p}) \vee \diamond p$ is valid on all the frames in \mathbf{A} and not




Figure 8: The frames $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ and $\mathcal{B}_1, \mathcal{B}_2$, and the pointed models based on them.

valid on any frame in \mathbf{B} .

SELECTION OF THE MODELS ON THE RIGHT: Hercules must choose his models as follows.

Lemma 17. *In any winning strategy for Hercules for an $(L_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGF,*

- if $\mathcal{A}_1 \in \mathbf{L}$ and $\mathcal{B}_1 \in \mathbf{R}$, Hercules picks a pointed model $(\mathcal{B}_1^M, \triangleright)$ based on the lowest point in \mathcal{B}_1 and assigns different valuations to the two leftmost points;
- if $\mathcal{A}_2 \in \mathbf{L}$ and $\mathcal{B}_2 \in \mathbf{R}$, Hercules picks a pointed model $(\mathcal{B}_2^M, \triangleright)$ based on the lowest point in \mathcal{B}_2 and assigns different valuations to the points of \mathcal{B}_2 .

Proof. The proof of the first item is the same as the proof of Lemma 9. For the second item, if Hercules picked a pointed model based on the reflexive point in \mathcal{B}_2 or assigned the same valuations to the two points in \mathcal{B}_2 , the Hydra would reply with a bisimilar pointed model based on \mathcal{A}_2 by making both points in \mathcal{A}_2 satisfy the same literals as the ones satisfied by the reflexive point in \mathcal{B}_2 . \square

As before, to indicate that the two leftmost points of \mathcal{B}_1 receive different valuations, we have drawn one of them black while the other is shaped as a grey rectangle. The literals true in the

rest of the points in \mathcal{B}_1 and in the points in \mathcal{B}_2 are not important for our discussion (provided, of course, that the two points in \mathcal{B}_2 satisfy different sets of literals which we have depicted by drawing them with different shapes and colours) but we have depicted them in different shapes and colours because this is a convenient way of visualising the pointed models chosen by the Hydra. Thus, Hercules constructs the pointed model $(\mathcal{B}_1^M, \triangleright)$ and $(\mathcal{B}_2^M, \triangleright)$ shown in the right rectangle in Figure 8.

SELECTION OF MODELS ON THE LEFT: The Hydra replies by mimicking Hercules' choice as shown on the left of the dotted line in the right rectangle in Figure 8. Points with the same shape and colour satisfy the same literals. The only possible exception to this rule is the grey point in \mathcal{B}_2^M which must differ on some literal with the black rectangle point but apart from that it may satisfy the same literals as some point in \mathcal{B}_1^M .

FORMULA SIZE GAME ON MODELS: We consider the FGM starting with $(\mathcal{A}_1^M, \triangleright), \dots, (\mathcal{A}_5^M, \triangleright)$ on the left and $(\mathcal{B}_1^M, \triangleright), (\mathcal{B}_2^M, \triangleright)$ on the right.

Lemma 18. *Suppose that \mathbf{L} and \mathbf{R} are classes of models and Hercules has a winning strategy for the $(\mathbf{L}_\diamond, (\mathbf{L}, \mathbf{R}))$ -FGM. Suppose that T is a closed game tree and the Hydra played greedily.*

1. *If $(\mathcal{A}_1^M, \triangleright) \in \mathbf{L}$ and $(\mathcal{B}_1^M, \triangleright) \in \mathbf{R}$, then there is a node η in T with $(\mathcal{A}_1^M, \triangleright) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\eta)$, and Hercules played a \diamond -move at η .*
2. *If η is a node of T with $(\mathcal{A}_5^M, \triangleright) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\eta)$, then Hercules did not play a \diamond -move at η .*
3. *If η is a node of T with $(\mathcal{A}_4^M, \triangleright) \in \mathfrak{L}(\eta)$ and there is a pointed model based on \mathcal{B}_2^M in $\mathfrak{R}(\eta)$, then Hercules did not play a \square -move at η .*
4. *If η is a node of T with $(\mathcal{A}_4^M, \triangleright) \in \mathfrak{L}(\eta)$ and there is a pointed model based on \mathcal{B}_2^M in $\mathfrak{R}(\eta)$, then if Hercules played a \diamond -move at η , he selected again $(\mathcal{A}_4^M, \triangleright)$.*
5. *Let η be a node on which Hercules played a \square -move. If, for some $i \in \{2, 3\}$, $(\mathcal{A}_i^M, \triangleright) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\eta)$, then he did not pick a pointed model based on the black point as a successor of $(\mathcal{B}_1^M, \triangleright)$.*
6. *Let η be a node on which Hercules played a \diamond -move. If, for some $i \in \{2, 3\}$, $(\mathcal{A}_i^M, \triangleright) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\eta)$, then he did not pick a pointed model based on the black point as a successor of $(\mathcal{A}_i^M, \triangleright)$.*

Proof. 1. Since $(\mathcal{A}_1^M, \triangleright)$ and $(\mathcal{B}_1^M, \triangleright)$ satisfy the same literals, Hercules cannot play a literal move at a node χ with $(\mathcal{A}_1^M, \triangleright) \in \mathfrak{L}(\chi)$ and $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\chi)$. Playing a \vee - or a \wedge -move at such a node χ will result in at least one new game position κ such that $(\mathcal{A}_1^M, \triangleright) \in \mathfrak{L}(\kappa)$ and $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\kappa)$. If Hercules played a \square -move at χ , he must have picked again $(\mathcal{B}_1^M, \triangleright)$ as a successor of $(\mathcal{B}_1^M, \triangleright)$ because if he selected a pointed model based on either the black or the white point, the Hydra would reply with the same point in \mathcal{A}_1^M which contradicts the fact that T is closed. Therefore, playing a \square -move at χ would result in a game position γ with $(\mathcal{A}_1^M, \triangleright) \in \mathfrak{L}(\gamma)$ and $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\gamma)$. Thus, T must contain a node η with $(\mathcal{A}_1^M, \triangleright) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\eta)$ at which Hercules played a \diamond -move.

2. If Hercules played such a move he must select again $(\mathcal{A}_5^M, \triangleright)$ as a successor of $(\mathcal{A}_5^M, \triangleright)$ and the Hydra would reply with, among others, a bisimilar pointed model based on the black point in \mathcal{B}_1^M .
3. If Hercules played a \square -move he must pick a pointed model based on the black rectangle point in \mathcal{B}_1^M to which the Hydra would reply with the same point in \mathcal{A}_4^M ($\mathcal{A}_5^M, \triangleright$) as a successor of $(\mathcal{A}_5^M, \triangleright)$.
4. If Hercules picked a pointed model based on the black rectangle point in \mathcal{A}_4^M , the Hydra would reply with the same point in \mathcal{B}_2^M .

5. If Hercules picked a pointed model based on the black point the Hydra would reply with the same point in \mathcal{B}_1^M . In either case we obtain bisimilar models on each side, which by Corollary 1 means that Hercules cannot win. □

Lemma 19. *Let \mathbf{L} and \mathbf{R} be classes of pointed models for which Hercules has a winning strategy in the $(\mathbf{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM. Let T be a closed game tree for this game and let us suppose that the Hydra played greedily. Let (\mathcal{A}_3^M, w) and (\mathcal{B}_1^M, w) denote the pointed models based on the respective model and the white circular point in it.*

1. *If $(\mathcal{A}_2^M, \triangleright) \in \mathbf{L}$ and $(\mathcal{B}_1^M, \triangleright) \in \mathbf{R}$, then there is a node η in T with $(\mathcal{A}_2^M, \triangleright) \in \mathfrak{L}(\eta)$, $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\eta)$, and Hercules played a \square -move at η so that he selected (\mathcal{B}_1^M, w) ;*
2. *If $(\mathcal{A}_3^M, \triangleright) \in \mathbf{L}$ and $(\mathcal{B}_1^M, \triangleright) \in \mathbf{R}$, then there is a node η in T such that $(\mathcal{A}_3^M, w) \in \mathfrak{L}(\eta)$, $(\mathcal{B}_1^M, w) \in \mathfrak{R}(\eta)$;*
3. *If $(\mathcal{A}_3^M, w) \in \mathbf{L}$ and $(\mathcal{B}_1^M, w) \in \mathbf{R}$, then there is a node η in T such that $(\mathcal{A}_3^M, w) \in \mathfrak{L}(\eta)$, $(\mathcal{B}_1^M, w) \in \mathfrak{R}(\eta)$ and Hercules played a \square -move at η .*

Proof. 1. Since $(\mathcal{A}_2^M, \triangleright)$ and $(\mathcal{B}_1^M, \triangleright)$ satisfy the same literals, Hercules cannot play a literal move at a node η with $(\mathcal{A}_2^M, \triangleright) \in \mathfrak{L}(\eta)$ and $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\eta)$. Playing a \vee - or a \wedge -move would result in at least one new tree-node χ with $(\mathcal{A}_2^M, \triangleright) \in \mathfrak{L}(\chi)$ and $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\chi)$. Using the last item of Lemma 18, we see that, if Hercules plays a \diamond -move at such a node, he is going to pick $(\mathcal{A}_2^M, \triangleright)$ again to which the Hydra is going to reply with, among others, $(\mathcal{B}_1^M, \triangleright)$ and thus we are back in essentially the same game position. The same is true if Hercules plays a \square -move and selects $(\mathcal{B}_1^M, \triangleright)$. Therefore, using the fifth item of Lemma 18, we conclude that there must be a node η in T such that $(\mathcal{A}_2^M, \triangleright) \in \mathfrak{L}(\eta)$, $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\eta)$, and Hercules played a \square -move at η so that he selected (\mathcal{B}_1^M, w) .

2. The proof of this item is immediate with the help of the last two items of Lemma 18.

3. Since (\mathcal{A}_3^M, w) and (\mathcal{B}_1^M, w) satisfy the same literals, Hercules cannot play a literal move at a node χ with $(\mathcal{A}_3^M, w) \in \mathfrak{L}(\chi)$ and $(\mathcal{B}_1^M, w) \in \mathfrak{R}(\chi)$. Playing a \vee - or a \wedge -move would result in at least one new tree-node κ with $(\mathcal{A}_3^M, w) \in \mathfrak{L}(\kappa)$ and $(\mathcal{B}_1^M, w) \in \mathfrak{R}(\kappa)$. Obviously, if Hercules plays a \diamond -move at such a node, he is going to pick (\mathcal{A}_3^M, w) again to which the Hydra is going to reply with, among others, (\mathcal{B}_1^M, w) and thus we are back in essentially the same game position. Hence, there must be a node η in T such that $(\mathcal{A}_3^M, w) \in \mathfrak{L}(\eta)$, $(\mathcal{B}_1^M, w) \in \mathfrak{R}(\eta)$, and Hercules played a \square -move at η . □

Proposition 6. *Let \mathbf{L} and \mathbf{R} be classes of models such that Hercules has a winning strategy for the $(\mathbf{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM. Let T be a closed game tree in which the Hydra played greedily.*

1. *If $(\mathcal{A}_1^M, \triangleright) \in \mathbf{L}$ and $(\mathcal{B}_1^M, \triangleright) \in \mathbf{R}$, then Hercules has made at least one \diamond -move.*
2. *If $(\mathcal{A}_1^M, \triangleright), (\mathcal{A}_5^M, \triangleright) \in \mathbf{L}$ and $(\mathcal{B}_1^M, \triangleright) \in \mathbf{R}$, then Hercules made at least one \vee -move.*
3. *If $(\mathcal{A}_4^M, \triangleright) \in \mathbf{L}$ and $(\mathcal{B}_1^M, \triangleright), (\mathcal{B}_2^M, \triangleright) \in \mathbf{R}$, then Hercules made at least one \wedge -move.*
4. *If $\{(\mathcal{A}_2^M, \triangleright), (\mathcal{A}_3^M, \triangleright)\} \subseteq \mathbf{L}$ and $(\mathcal{B}_1^M, \triangleright) \in \mathbf{R}$, then Hercules played at least two \square -moves.*

Proof. The first and the second item follow from the first two items of Lemma 18. For the third item, let us suppose that Hercules did not play a \wedge -move. Since, $(\mathcal{A}_4^M, \triangleright)$ and $(\mathcal{B}_1^M, \triangleright)$ satisfy the same literals, Hercules cannot play a literal move at a node χ with $(\mathcal{A}_4^M, \triangleright) \in \mathfrak{L}(\chi)$ while $(\mathcal{B}_1^M, \triangleright)$ and a pointed model \mathcal{M} based on \mathcal{B}_2^M are in $\mathfrak{R}(\chi)$. Playing a \vee -move at such a node χ will result in at least one new game position κ such that $(\mathcal{A}_4^M, \triangleright) \in \mathfrak{L}(\kappa)$ and $(\mathcal{B}_1^M, \triangleright)$ and \mathcal{M} are in $\mathfrak{R}(\kappa)$. According to the third item of Lemma 18, Hercules is not going to play a \square -move at such a node whereas according to the fourth item of the same Lemma, if Hercules plays a \diamond -move, he must select $(\mathcal{A}_4^M, \triangleright)$, to which the Hydra is going to reply with among others $(\mathcal{B}_1^M, \triangleright)$ and a pointed

model based on \mathcal{B}_2^M and we are back in the previous situation. Thus in the absence of a \wedge -move we see that Hercules has no winning strategy which contradicts our assumption.

For the last item, it follows from Lemma 19 that

- there is a node η in T such that $(\mathcal{A}_2^M, \triangleright) \in \mathfrak{L}\eta$, $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\eta)$, and Hercules played a \square -move at η so that he selected (\mathcal{B}_1^M, w) ;
- there is a node χ in T such that $(\mathcal{A}_3^M, w) \in \mathfrak{L}(\chi)$, $(\mathcal{B}_1^M, w) \in \mathfrak{R}(\chi)$ and Hercules played a \square -move at η .

If η and χ do not coincide, then, obviously, Hercules played at least two \square -moves. Suppose now that η and χ coincide and let κ be its daughter. According to the first item, $(\mathcal{B}_1^M, w) \in \mathfrak{R}(\kappa)$. Since the Hydra plays greedily and using the second item, we see that $(\mathcal{A}_3^M, w) \in \mathfrak{L}(\kappa)$. It is immediate from the third item of Lemma 19 that in the sub-game starting at the node κ , Hercules played at least one additional \square -move. \square

With this we conclude the proof of Theorem 5.

8 The Löb axiom

The Löb axiom, which defines the property of transitivity and converse-well-foundedness, (i.e., there are no infinite chains $w_0 R w_1 R \dots$). This is a conjunction of two properties and the resulting *Gödel-Löb* logic is often presented with the additional axiom $\square p \rightarrow \square \square p$ but it is a non-trivial exercise to show that this is already a consequence of the Löb axiom $\square(\square p \rightarrow p) \rightarrow \square p$. Note that the well-foundedness is a second-order property and cannot be defined in first-order logic.

Theorem 6. *The formula $\square \bar{p} \vee \diamond(p \wedge \square \bar{p})$ is absolutely minimal among all formulas defining the class of transitive and converse well-founded frames.*

We have already shown that $\square \square \bar{p} \vee \diamond p$ is absolutely minimal among those formulas defining transitivity, so our strategy will be to expand on the frames and pointed models in Figure 5 to additionally force Hercules to play a conjunction. Since these models were already well-founded we can use previous results.

Consider an $(L_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF played by Hercules and the Hydra with the frames shown in Figure 9. Obviously, $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$, and \mathcal{B} are obtained from the frames in Figure 5 for $m = 2$ and $n = 1$. Additionally, \mathbf{A} contains the frame \mathcal{A}_4 that is a transitive tree with infinitely many branches such that, for every natural number $n > 0$, there is a branch for which the maximum number of relation steps from the root to its leaf is n . Similarly, \mathbf{B} contains the frame \mathcal{B}_1 shown on the right of the dotted line in the same figure. We are going to use \mathcal{A}_4 and \mathcal{B}_1 in order to force Hercules to play an \wedge -move.

SELECTION OF THE MODELS ON THE RIGHT: We only consider the choice of pointed model for the frame \mathcal{B}_1 . Obviously, Hercules is not going to base a pointed model on the dead-end point in \mathcal{B}_1 because the Hydra would reply with a bisimilar pointed model based on one of the leaves of \mathcal{A}_4 .

Lemma 20. *In any winning strategy for Hercules in the $(L_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF, Hercules will choose a pointed model based on the reflexive point on \mathcal{B}_1 .*

SELECTION OF MODELS ON THE LEFT: Hydra will choose her pointed models based on $\mathcal{A}_1, \mathcal{A}_2$ and \mathcal{A}_3 as before. For her pointed model based on \mathcal{A}_4 , she picks a pointed model based on the root of the tree in which all leaves of \mathcal{A}_4 satisfy the same literals as the ones satisfied by the dead-end point in \mathcal{B}_1 whereas the rest of the points satisfy the same literals as the ones satisfied

by the reflexive point in \mathcal{B}_1 . Once again, if Hydra plays in this way we say that she *mimics* Hercules' selection.




Figure 9: The sets of frames $\mathbf{A} = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4\}$ and $\mathbf{B} = \{\mathcal{B}, \mathcal{B}_1\}$.

FORMULA SIZE GAME ON MODELS: The next lemmas will be used to prove that Hercules must play an \wedge -move.

Lemma 21. *Let \mathbf{L}, \mathbf{R} be classes of models such that Hercules has a winning strategy for the $(\mathcal{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM. If T is a closed game tree on which the Hydra played greedily, then for any game position η and any non-leaf point w of \mathcal{A}_4 , if $(\mathcal{A}_4^M, w) \in \mathfrak{L}(\eta)$, $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\eta)$, and Hercules played a \square -move at η , then he selected $(\mathcal{B}_1^M, \triangleright)$ again.*

Proof. If Hercules picked the dead-end point in \mathcal{B}_1^M , the Hydra, using the transitivity of the relation, would reply with a bisimilar pointed model based on a leaf in \mathcal{A}_4^M . \square

Proposition 7. *Suppose that \mathbf{L}, \mathbf{R} are classes of models for which Hercules has a winning strategy for the $(\mathcal{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM and let T be a closed game tree on which the Hydra played greedily.*

1. *If $(\mathcal{A}_4^M, \triangleright) \in \mathbf{L}$ and $(\mathcal{B}_1^M, \triangleright) \in \mathbf{R}$, Hercules played at least one \diamond -move on a node η such that $\mathfrak{L}(\eta)$ contains a pointed model based on \mathcal{A}_4^M whereas $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\eta)$.*
2. *If Hercules plays a \diamond -move in a position η in which $\mathfrak{L}(\eta)$ contains a pointed model based on \mathcal{A}_4^M while $(\mathcal{B}_1^M, \triangleright) \in \mathfrak{R}(\eta)$, he must play at least one subsequent \wedge -move.*

Proof. (1) Let us suppose that Hercules plays without \diamond -moves. Since $(\mathcal{A}_4^M, \triangleright)$ and $(\mathcal{B}_1^M, \triangleright)$ satisfy the same literals, no literal move is possible in a game position η in which $(\mathcal{A}_4^M, \triangleright)$ is on the left and $(\mathcal{B}_1^M, \triangleright)$ on the right. Playing a \wedge - or a \vee -move results in at least one new position in which $(\mathcal{A}_4^M, \triangleright)$ is on the left and $(\mathcal{B}_1^M, \triangleright)$ is on the right. Hence a \square -move is inevitable and by Lemma 21, he selected $(\mathcal{B}_1^M, \triangleright)$ again.

When Hercules plays such a move, the Hydra replies with all infinitely many pointed models based on \mathcal{A}_4^M and an immediate successor of the root of the tree. From this new position on any finite number of \vee -, \wedge - and \square -moves are going to result in at least one new position that contains $(\mathcal{B}_1^M, \triangleright)$ on the right whereas on the left we have infinitely many pointed models based on \mathcal{A}_4^M and a non-leaf point. Obviously, none of the \top -, \perp -, and literal-moves are possible in such a position. Hence, Hercules has no winning strategy without \diamond -moves.

(2) Let us suppose that Hercules plays a \diamond -move in such a position. The Hydra is going to respond with both $(\mathcal{B}_1^M, \triangleright)$ and a pointed model based on the dead-end point in \mathcal{B}_1^M . Let us suppose now that Hercules is not going to play any subsequent \wedge -move. Obviously, \perp , \top , and literal moves are impossible; moreover, the presence of a dead-end pointed model on the right prevents \square -moves. Clearly, playing an \vee -move would result in at least one new game position which is the same as the previous one. Therefore, Hercules can only play \diamond -moves until he reaches a pointed model (\mathcal{A}_4, v) such that the only successor of v is a leaf. Playing a \diamond -move in

such a position would lead to a loss in the next step because of the presence of bisimilar pointed models on the left and right. Since (\mathcal{A}_4^M, v) and $(\mathcal{B}_1^M, \triangleright)$ satisfy the same literals no literal moves are possible either. Therefore, Hercules has no winning strategy without playing at least one \wedge -move. \square

With this we can prove Theorem 6.

Proof. Consider a $(L_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF where $\mathbf{A} = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4\}$ and $\mathbf{B} = \{\mathcal{B}, \mathcal{B}_1\}$ as given in Figure 9. Hercules must choose his pointed models according to Lemmas 9 and 20, and Hydra replies by mimicking Hercules. Using Proposition 6 we see that if the Hydra plays greedily then any closed game tree must have modal depth at least two, contain two instances of \square , one instance of each \diamond and \vee , and one variable. By Proposition 7, it also contains one conjunction, as required. \square

9 Conclusion

We have shown that several familiar modal axioms are minimal with respect to all measures considered, including the non-colourability and the Löb axiom, which are not first-order definable. It is obvious that once we have shown that a given frame property is modally definable, we can study its minimal modal complexity with respect to different complexity measures and therefore there are many natural open problems related to the present work. We would like to mention one in particular.

The importance of the Sahlqvist formulae cannot be overstated and they have been studied extensively over the years. However, it seems that a very basic question about them has not received the attention it deserves. Namely, since these formulae have a specific “syntactic shape”, it is natural to ask whether this syntactic restriction leads to an increase of their complexity. For example, Vakarelov conjectured in [27] that there is no Sahlqvist formula in a language with two propositional symbols that defines the first-order condition

$$\forall y_1 \forall y_2 \forall y_3 \forall y_4 ((xRy_1 \wedge xRy_2 \wedge xRy_3 \wedge xRy_4) \rightarrow \exists z (y_1Rz \wedge y_2Rz \wedge y_3Rz \wedge y_4Rz)).$$

Note, however, that it can be modally defined with the help of the non-Sahlqvist formula

$$(\diamond\square(p_1 \vee p_2) \wedge \diamond\square(p_1 \vee \overline{p_2}) \wedge \diamond\square(\overline{p_1} \vee p_2)) \rightarrow \square\diamond(p_1 \wedge p_2)$$

that contains only two different propositional symbols. Vakarelov’s conjecture is an instance of the following general problem.

Question 1. *Is there a natural complexity measure² μ with respect to which Sahlqvist formulae are asymptotically more complex than non-Sahlqvist ones and by how much? In particular, can this complexity gap be “big”, i.e., is there a natural complexity measure μ and an infinite sequence of formulae $\varphi_1, \varphi_2, \dots$ such that if ψ_1, ψ_2, \dots is a sequence of equivalent Sahlqvist formulae then $\mu(\psi_n)$ grows super-polynomially or even exponentially in $\mu(\varphi_n)$?*

The above question seems very difficult but the next one might be more approachable.

Question 2. *Can the proofs we employed in the case of the (m, n) -transfer axioms be extended to show that the Lemmon-Scott’s axioms, $\diamond^m \square^i p \rightarrow \square^j \diamond^n p$, are absolutely minimal among those defining the first-order condition $xR^m y \wedge xR^j z \rightarrow \exists t (yR^i t \wedge zR^n t)$?*

²Recall that we call natural measures the length of a formula, the number of occurrences of any connective, the modal depth, and the number of variables.

An (admittedly weak) indication that the answer to the second question might be “yes” is the fact that a slight modification of some of our frames and models can be used to establish Theorem 7 below whose proof is presented in Appendix 11

Theorem 7. *The formula $p \rightarrow \Box\Diamond p$ is absolutely minimal among the L_\Diamond -formulae that define symmetry.*

Acknowledgements

We are grateful to the anonymous referees for the many valuable remarks and suggestions.

References

- [1] M. Adler and N. Immerman. An $n!$ lower bound on formula size. *ACM Transactions on Computational Logic*, 4(3):296–314, 2003.
- [2] P. Balbiani, D. Fernández-Duque, A. Herzig, and P. Iliev. Frame validity games and absolute minimality of modal axioms. In G. Bezhanishvili, G. D’Agostino, G. Metcalfe, and T. Studer, editors, *Proceedings of Advances in Modal Logic*, volume 12, pages 83–102. College Publications, 2018.
- [3] P. Balbiani, T. Tinchev, and D. Vakarelov. Modal logics for region-based theories of space. *Fundamenta Informaticae*, 81(1-3):29–82, 2007.
- [4] A. Chagrov and M. Zakharyashev. *Modal Logic*, volume 35 of *Oxford logic guides*. Oxford University Press, 1997.
- [5] H. van Ditmarsch and P. Iliev. The succinctness of the cover modality. *Journal of Applied Non-Classical Logics*, 25(4):373–405, 2015.
- [6] A. Ehrenfeucht. An application of games to the completeness problem for formalized theories. *Fundamenta Mathematicae*, 49:129–141, 1961.
- [7] R. Fagin. Finite model theory - a personal perspective. *Theoretical Computer Science*, 116:3–31, 1993.
- [8] D. Fernández-Duque and P. Iliev. Succinctness in subsystems of the spatial μ -calculus. *Journal of Applied Logics - IfCoLoG Journal of Logics and their Applications*, 5(4):827–874, 2018.
- [9] S. Figueira and D. Gorín. On the size of shortest modal descriptions. In L. Beklemishev, V. Goranko, and V. Shehtman, editors, *Proceedings of Advances in Modal Logic*, volume 8, pages 114–132, 2010.
- [10] H.-D. Ebbinghaus; J. Flum. *Finite Model Theory*. Springer, 1995.
- [11] T. French, W. van der Hoek, P. Iliev, and B. Kooi. On the succinctness of some modal logics. *Artificial Intelligence*, 197:56–85, 2013.
- [12] T. French, W. van der Hoek, P., and B. Kooi. Succinctness of epistemic languages. In T. Walsh, editor, *Proceedings of IJCAI*, pages 881–886, 2011.
- [13] R. Goldblatt, I. Hodkinson, and Y. Venema. On canonical modal logics that are not elementarily determined. *Logique et Analyse Nouvelle Série*, 46(181):77–101, 2003.

- [14] L. Hella and M. Vilander. The succinctness of first-order logic over modal logic via a formula size game. In L. Beklemishev, S. Demri, and A. Máté, editors, *Proceedings of Advances in Modal Logic*, volume 11, pages 401–419. College Publications, 2016.
- [15] L. Hella and M. Vilander. Formula size games for modal logic and μ -calculus. *Journal of Logic and Computation*, 29(8):1311–1344, 2019.
- [16] W. van der Hoek and P. Iliev. On the relative succinctness of modal logics with union, intersection and quantification. In A. Lomuscio, P. Scerri, A. Bazzan, and M. Huhns, editors, *Proceedings of AAMAS*, pages 341–348, 2014.
- [17] G. E. Hughes. Every world can see a reflexive world. *Studia Logica*, 49:175–181, 1990.
- [18] N. Immerman. Descriptive and computational complexity. In J. Hartmanis, editor, *Computational Complexity Theory, Proc. Symp. in Applied Math.*, volume 38, pages 75–91, 1989.
- [19] N. Immerman. *Descriptive complexity*. Springer-Verlag, 1999.
- [20] S. Jukna. *Boolean Function Complexity: Advances and Frontiers*. Springer, 2012.
- [21] M. Karchmer and A. Wigderson. Monotone circuits for connectivity require super-logarithmic depth. In *Proceedings of the 20th ACM STOC*, pages 539–550, 1988.
- [22] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
- [23] C. Lutz. Complexity and succinctness of public announcement logic. In P. Stone and G. Weiss, editors, *Proceedings of AAMAS*, pages 137–144, 2006.
- [24] O. Pikhurko and O. Verbitsky. Logical complexity of graphs: a survey. In *Model Theoretic Methods in Finite Combinatorics, Contemporary Mathematics*, volume 558, pages 129–179. American Mathematical Society, 2011.
- [25] A. Razborov. Applications of matrix methods to the theory of lower bounds in computational complexity. *Combinatorica*, 10(1):81–93, 1990.
- [26] G. Turán. On the definability of properties of finite graphs. *Discrete Mathematics*, 49:291–302, 1984.
- [27] D. Vakarelov. Modal definability in languages with a finite number of propositional variables and a new extension of the Sahlqvist’s class. In P. Balbiani, N.-Y. Suzuki, F. Wolter, and M. Zakharyashev, editors, *Proceedings of Advances in Modal Logic*, volume 4, pages 499–518. College Publications, 2003.

APPENDIX

10 Properties of the formula-complexity game on models

We have seen that a closed game tree T induces a formula ψ_T . The other way round, we can also turn formulae into game trees.

Lemma 22. *Let \mathbf{A}, \mathbf{B} be classes of models and $\varphi \in \mathcal{L}_{\diamond}^{\forall}$ be so that $\mathbf{A} \models \varphi$ and $\mathbf{B} \models \neg\varphi$. Then Hercules has a winning strategy for the $(\mathcal{L}_{\diamond}^{\forall}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM so that any game terminates on a closed game tree T with $\psi_T = \varphi$.*

Proof. We proceed by induction on the structure of φ .

φ IS A LITERAL. If φ is a literal ι , then Hercules plays the literal-move by choosing ι and the game tree T is closed with $\psi_T = \iota$, as required.

φ IS \perp . If φ is \perp , then Hercules plays the \perp -move and (as \mathbf{B} must be empty) the game tree T is closed with $\psi_T = \perp$, as required.

φ IS $\varphi_1 \vee \varphi_2$. Hercules can play the \vee -move and add two nodes η_1, η_2 labelled by $\langle \mathbf{A}_1, \mathbf{B} \rangle$ and $\langle \mathbf{A}_2, \mathbf{B} \rangle$, respectively, where $\mathbf{A} = \mathbf{A}_1 \cup \mathbf{A}_2$, $\mathbf{A}_1 \models \varphi_1$ and $\mathbf{A}_2 \models \varphi_2$. Applying the induction hypothesis to each sub-game, for $i \in \{1, 2\}$ Hercules has a strategy for the $(\mathcal{L}_{\diamond}, \langle \mathbf{A}_i, \mathbf{B} \rangle)$ -FGM with resulting closed game tree T_i so that $\psi_{T_i} = \varphi_i$. This yields a closed game tree T for the original game with $\psi_T = \varphi$, as desired.

φ IS $\diamond\theta$. For each $\mathbf{a} \in \mathbf{A}$, Hercules chooses a pointed model from $\Box\mathbf{a}$ that satisfies θ and collects all these pointed models in the class \mathbf{A}_1 . Hydra replies by choosing a subset of $\Box\mathbf{b}$ for each $\mathbf{b} \in \mathbf{B}$ and collects these pointed models in \mathbf{B}_1 . A new node η labelled with $\langle \mathbf{A}_1, \mathbf{B}_1 \rangle$ is added to the game tree as a successor to the one labelled with $\langle \mathbf{A}, \mathbf{B} \rangle$. Obviously, $\mathbf{A}_1 \models \theta$ and $\mathbf{B}_1 \models \neg\theta$. Applying the induction hypothesis, we conclude that Hercules has a strategy for the sub-game starting at η so that the resulting game tree S is closed with $\psi_S = \theta$. This yields a closed tree T for the original game with $\psi_T = \diamond\theta$.

φ IS $\exists\theta$. The proof of this case follows the lines of that of $\diamond\theta$.

OTHER CASES: Each of the remaining cases is dual to one discussed above. □

Next we show that if the Hydra plays greedily, then any closed game tree T for the $(\mathcal{L}_{\diamond}^{\forall}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM is such that $\mathbf{A} \models \psi_T$ and $\mathbf{B} \models \neg\psi_T$.

Lemma 23. *Let \mathbf{A}, \mathbf{B} be classes of models and let T be a closed game tree for the $(\mathcal{L}_{\diamond}^{\forall}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM on which the Hydra played greedily. Then, $\mathbf{A} \models \psi_T$ and $\mathbf{B} \models \neg\psi_T$.*

Proof. For a node η of T , let T_{η} be the subtree with root η , and let $\psi_{\eta} = \psi_{T_{\eta}}$. By induction on the size of T_{η} starting from the leaves we show that $\mathcal{L}(\eta) \models \psi_{\eta}$ and $\mathfrak{R}(\eta) \models \neg\psi_{\eta}$. The base case is immediate since Hercules can only play a literal or \perp or \top when it is true on the left but false on the right, and the inductive steps for \vee and \wedge are straightforward. The critical case is when Hercules plays a modality on η , which is when we use that the Hydra plays greedily. For a \diamond -move on η with daughter η' , for each $\mathbf{l} \in \mathcal{L}(\eta)$, he chose $\mathbf{l}' \in \Box\mathcal{L}(\eta)$ and placed $\mathbf{l}' \in \mathcal{L}(\eta')$; by the induction hypothesis $\mathbf{l}' \models \psi_{\eta'}$, so that by the semantics of \diamond , $\mathbf{l} \models \diamond\psi_{\eta'} = \psi_{\eta}$. Meanwhile for $\mathbf{r} \in \mathfrak{R}(\eta)$, if $\mathbf{r}' \in \Box\mathbf{r}$ then since the Hydra played greedily $\mathbf{r}' \in \mathfrak{R}(\eta')$, and since \mathbf{r}' was arbitrary we see that $\mathbf{r} \models \neg\diamond\psi_{\eta'}$. The case for a \Box -move is symmetric and the cases of the \exists - and \forall -moves are analogous. □

With this we prove Theorem 1. Recall that Theorem 1 states that the following are equivalent:

1. Hercules has a winning strategy for the $(L_{\diamond}^{\forall}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM with μ below k , and
2. there is an L_{\diamond}^{\forall} -formula φ with $\mu(\varphi) < k$ and $\mathbf{A} \models \varphi$ whereas $\mathbf{B} \models \neg\varphi$.

Proof. Let \mathbf{A}, \mathbf{B} be classes of models, μ any complexity measure, and $k \in \mathbb{N}$.

First assume that (1) holds, and let Hydra play the $(L_{\diamond}^{\forall}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM greedily. By using his winning strategy, Hercules can ensure that the game terminates on some closed tree T with $\mu(T) < k$. But by definition this means that $\mu(\psi_T) < k$, and by Lemma 23, $\mathbf{A} \models \psi_T$ while $\mathbf{B} \models \neg\psi_T$.

Conversely, if (2) holds, by Lemma 22 Hercules has a strategy so that no matter how the Hydra plays, any match ends with a closed tree T with $\psi_T = \varphi$, so that in particular $\mu(T) < k$. \square

11 The symmetry axiom

This appendix contains the proof of Theorem 7, that is, we show that the formula $\bar{p} \vee \Box \diamond p$ is absolutely minimal among the L_{\diamond} -formulas defining symmetry.

Let us consider a $(L_{\diamond}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF where $\mathbf{A} = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3\}$ while \mathbf{B} contains a single element \mathcal{B} , as depicted in the left rectangle in Figure 10. Note that all frames in \mathbf{A} are symmetric whereas this is not true about the frame \mathcal{B} . Hence the formula $\bar{p} \vee \Box \diamond p$ is valid on the frames in \mathbf{A} and not valid on \mathcal{B} .




Figure 10: The sets of frames $\mathbf{A} = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3\}$ and $\mathbf{B} = \{\mathcal{B}\}$ and the respective models based on them.

SELECTION OF THE MODELS ON THE RIGHT: Using Lemma 12, we see that Hercules must pick the pointed model $(\mathcal{B}^M, \triangleright)$ shown in the right half of Figure 10. Again, to indicate that the two points of \mathcal{B}^M satisfy different sets of literals, we colour one of them black and the other white.

SELECTION OF THE POINTED MODELS ON THE LEFT: The Hydra replies as shown on the left of the dotted line in the right half in Figure 10. Recall, that points satisfying the same set of literals have the same colour.

FORMULA SIZE GAME ON MODELS: We are going to consider the FGM starting with $\mathbf{A}^m = \{(\mathcal{A}_1^M, \triangleright), (\mathcal{A}_2^M, \triangleright), (\mathcal{A}_3^M, \triangleright)\}$ on the left and $\mathbf{B}^m = \{(\mathcal{B}^M, \triangleright)\}$ on the right.

Lemma 24. *In any closed game tree T for the $(\mathbf{A}^m, \mathbf{B}^m)$ -FGM, Hercules played at least one \vee -move and at least one \Box -move.*

Proof. The proof is almost identical to the proof of Lemma 15. Indeed, it is immediate from Lemma 13, that if Hercules wants to win the game, he must not play either a \diamond - or a \Box -move at a position η in which $(\mathcal{A}_1^M, \triangleright)$ is on the left and $(\mathcal{B}^M, \triangleright)$ is on the right. On the other hand, for every game in which $(\mathcal{A}_2^M, \triangleright)$ is among the pointed models chosen by the Hydra and $(\mathcal{B}^M, \triangleright)$ is among the models chosen by Hercules, if he wants to win the game, then there is at least one

game position ν such that $(\mathcal{A}_2^M, \triangleright)$ is on the left and $(\mathcal{B}^M, \triangleright)$ is on the right and Hercules played a \square -move at ν . This implies that in any FGM with a starting position in which the pointed models selected by the Hydra are on the left and $(\mathcal{B}^M, \triangleright)$ is on the right, Hercules must play at least one \vee to separate $(\mathcal{A}_1^M, \triangleright)$ from $(\mathcal{A}_2^M, \triangleright)$ and one subsequent \square -move. \square

Lemma 25. *In any closed game tree T for the $(\mathbf{A}^m, \mathbf{B}^m)$ -FGM where the Hydra played greedily, Hercules played at least one \diamond -move.*

Proof. Let us consider a game position η with $(\mathcal{A}_3^M, \triangleright)$ on the left and $(\mathcal{B}^M, \triangleright)$ on the right and let us suppose that Hercules attempts to win the FGM with η as a starting position without playing a \diamond -move. Clearly, a literal move is impossible at η . By playing a \vee - or a \wedge -move, he will arrive to at least one new position that is essentially the same as η . If he plays a \square -move he must select the successor of $(\mathcal{B}^M, \triangleright)$ based on the reflexive white point in \mathcal{B}^M . The Hydra is going to reply with the successor of $(\mathcal{A}_3^M, \triangleright)$ based on the reflexive white point in \mathcal{A}_3^M . It is immediate that in this new game position a literal move is impossible; moreover, no amount of \square -moves are going to help Hercules win the game. Hence, Hercules must make at least one \diamond -move. \square

Thus, Theorem 7 is immediate from Lemma 24 and Lemma 25.