



HAL
open science

Frame-validity games and absolute minimality 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 absolute minimality of modal axioms. 12th International Conference on Advances in Modal Logic (AiML 2018), Aug 2018, Berne, Switzerland. pp.83-102. hal-03046246

HAL Id: hal-03046246

<https://hal.science/hal-03046246>

Submitted on 8 Dec 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 absolute minimality of modal axioms

Philippe Balbiani¹

*Institut de Recherche en Informatique de Toulouse
CNRS — Toulouse University*

David Fernández-Duque²

*Department of Mathematics
Ghent University*

Andreas Herzig³

*Institut de Recherche en Informatique de Toulouse
CNRS — Toulouse University*

Petar Iliev⁴

*Centre International de Mathématiques et d'Informatique de Toulouse
Toulouse University*

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 property. Using these games we prove that the Löb axiom $\Box(\Box p \rightarrow p) \rightarrow \Box p$ and the (m, n) -transfer axioms $\Diamond^m p \rightarrow \Diamond^n p$ are optimal among those defining their respective class of frames.

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

1 Introduction

One of the key advantages of modal logics over first-order logic is that the former are often decidable. However, decidability is not sufficient for applications: efficiency plays a huge role in determining the usefulness of a formal system.

¹ Philippe.Balbiani@irit.fr

² David.FernandezDuque@UGent.be

³ Andreas.Herzig@irit.fr

⁴ Petar.Iliev@irit.fr

Typical measures of complexity revolve around problems such as satisfiability and model-checking, but the sometimes-overlooked *succinctness* plays a crucial role as well: there is little use in a PTIME logic if properties of interest can only be defined by exponentially large formulas.

The power of first-order logic and some of its extensions to succinctly define graph properties has been investigated extensively [7], as that of the modal language and natural extensions to define properties of relational *models* [3,9]. In contrast, it seems that the only study of how succinctly *frame* properties can be expressed in modal logic is [8], where the question of how many different propositional variables are needed to modally define certain classes of Kripke frames is being considered. To increase our understanding of the succinctness of modal languages, we develop in the present paper techniques for proving lower bounds on the complexity of modal formulas defining frame properties and apply them to some well-known classes of frames.

As usual, 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. If a class of frames is definable by a modal formula, it is natural to ask how *complex* any such formula must be, where the complexity of a formula may be measured according to the total number of symbols, the modal depth, the number of occurrences of symbols of a certain type, or the number of different variables needed.

The techniques we will employ are based on *frame equivalence games*, closely related to model-equivalence games as appeared in [4,5,6]. To demonstrate the applicability of the former to both first- and second-order semantic conditions, we prove that

- (i) 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

$$\forall x \forall y (x R^m y \rightarrow x R^n y), \quad (1)$$

where R^j denotes the j -fold composition of R .

- (ii) 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.

Note that the former result applies to the well-studied axioms defining transitivity, reflexivity, and density.

2 Technical preliminaries

Our formula size games are based on formulas in negation normal form, i.e., negations appear only in front of propositional symbols. Fix a countably infinite set of *propositional variables* $P = \{p_1, p_2, \dots\}$, and let \mathcal{L}_\diamond denote the uni-modal language that has as atomic formulas 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.

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* $(W_{\mathcal{B}}, R_{\mathcal{B}})$ is a tuple $\mathcal{B} = (W_{\mathcal{B}}, R_{\mathcal{B}}, V_{\mathcal{B}})$

consisting of a frame equipped with a *valuation* $V_B: W_B \rightarrow 2^P$, and a *pointed model* is a tuple $\mathbf{c} = (\mathcal{C}, c)$ consisting of a model \mathcal{C} equipped with a designated point $c \in W_{\mathcal{C}}$; pointed models will always 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 the pointed model \mathbf{a} along the relation $R_{\mathcal{A}}$.

Given $\varphi \in \mathcal{L}_{\diamond}$ and a pointed model \mathbf{a} , we define $\mathbf{a} \models \varphi$ according to standard Kripke semantics, and 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 to denote 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 $\mathcal{X} \models \psi$.

Our goal in this paper 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: by a *complexity measure* (or just *measure*) we simply mean a function $\mu: \mathcal{L}_{\diamond} \rightarrow \mathbb{N}$. We are interested in the following measures: (i) the *length* of a formula φ , denoted $|\varphi|$ and defined as the number of nodes in its syntax tree (including leaves); (ii) the *number of occurrences* of any connective, (iii) the *modal depth*, and (iv) the *number of variables*.

Note that these are a total of nine measures, as each connective gives rise to its own measure in (ii). 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 $\Gamma \subseteq \mathcal{L}_{\diamond}$ and $\varphi \in \Gamma$, say that φ is *absolutely minimal among* Γ if for all $\psi \in \Gamma$ and any of the nine 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 developed in the setting of first-order logic in [1]. The general idea is that we have two competing 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” \mathcal{L}_{\diamond} -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 T . 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 3.1 The $(\mathcal{L}_{\diamond}, \langle \mathbf{A}, \mathbf{B} \rangle)$ *formula-complexity game on models* (denoted $(\mathcal{L}_{\diamond}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM) is played by two players, 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 $\langle L(\eta), R(\eta) \rangle$ of classes of pointed models and either a literal or a symbol from the set $\{\perp, \top, \vee, \wedge, \diamond, \Box\}$ 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 $\langle \mathbf{L}(\eta), \mathbf{R}(\eta) \rangle$. Then, he plays one of the following moves.

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

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

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

\vee -MOVE: Hercules labels η with the symbol \vee and chooses two sets $\mathbf{L}_1, \mathbf{L}_2 \subseteq \mathbf{L}$ such that $\mathbf{L}(\eta) = \mathbf{L}_1 \cup \mathbf{L}_2$. Two new heads, labelled by $\langle \mathbf{L}_1, \mathbf{R}(\eta) \rangle$ and $\langle \mathbf{L}_2, \mathbf{R}(\eta) \rangle$, 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 $\mathbf{R}_1 \cup \mathbf{R}_2 = \mathbf{R}(\eta)$.

\diamond -MOVE: Hercules labels η with the symbol \diamond and, for each pointed model $\mathbf{l} \in \mathbf{L}(\eta)$, he chooses a pointed model from $\square \mathbf{l}$; if for some $\mathbf{l} \in \mathbf{L}(\eta)$ we have $\square \mathbf{l} = \emptyset$, Hercules cannot play this move. All these new pointed models are collected in the set \mathbf{L}_1 . For each pointed model $\mathbf{r} \in \mathbf{R}(\eta)$, the Hydra replies by picking a subset of $\square \mathbf{r}$ ⁵. 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 .

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

The $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -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 (i) whenever Hercules makes a \diamond -move on a node η and adds a new node η' , then Hydra sets $\mathbf{R}(\eta') = \bigcup_{\mathbf{r} \in \mathbf{R}(\eta)} \square \mathbf{r}$, and similarly (ii) whenever Hercules makes a \square -move on a node η and adds a new node η' , then Hydra sets $\mathbf{L}(\eta') = \bigcup_{\mathbf{l} \in \mathbf{L}(\eta)} \square \mathbf{l}$.

The $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM can be used to give lower bounds on the length of formulae defining a given property [4,5,6]. Here we will generalize this to show that it can be used to give lower bounds on any complexity measure. For this, we need to view game-trees as formulae.

Definition 3.2 Given a closed $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM tree T , we define $\psi_T \in \mathcal{L}_\diamond$ 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

⁵ In particular, if $\square \mathbf{r} = \emptyset$ for some $\mathbf{r} \in \mathbf{R}(\eta)$, the Hydra does not add anything to \mathbf{R}_1 for the pointed model \mathbf{r} .

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 and \square 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 $m \in \mathbb{N}$, \mathbf{A}, \mathbf{B} are classes of models, and $\mu: \mathcal{L}_\diamond \rightarrow \mathbb{N}$ a complexity measure, we say that Hercules has a *winning strategy for the $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM with μ below m* 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 $\mu(T) < m$.

Theorem 3.3 *Let \mathbf{A}, \mathbf{B} be classes of models, μ any complexity measure, and $m \in \mathbb{N}$. Then, the following are equivalent:*

- (i) *Hercules has a winning strategy for the $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM with μ below m ;*
- (ii) *there is an \mathcal{L}_\diamond -formula φ with $\mu(\varphi) < m$ and $\mathbf{A} \models \varphi$ whereas $\mathbf{B} \models \neg\varphi$.*

We defer the proof of Theorem 3.3 to Appendix A, 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 [4,5,6]. We will also use the following easy consequence of Theorem 3.3. We assume familiarity with bisimulations [2].

Corollary 3.4 *Let \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} bisimilar to \mathbf{b} . For all complexity measures μ and for all nonnegative integers m , Hercules has no winning strategy for the (\mathbf{A}, \mathbf{B}) -FGM with μ below m .*

4 A formula-complexity game on frames

We develop 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.

Definition 4.1 Let \mathbf{A}, \mathbf{B} be classes of frames. The $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ formula-complexity game on frames (denoted $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -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 $(\mathcal{L}_\diamond, \langle \mathbf{A}^m, \mathbf{B}^m \rangle)$ -FGM.

The game tree assigned to a match of the $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF is the game tree of the subsequent $(\mathcal{L}_\diamond, \langle \mathbf{A}^m, \mathbf{B}^m \rangle)$ -FGM.

Remark 4.2 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} .

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

Theorem 4.3 *Let \mathbf{A} , \mathbf{B} be classes of frames, μ any complexity measure, and $m \in \mathbb{N}$. Then, the following are equivalent:*

- (i) *Hercules has a winning strategy for the $(\mathcal{L}_{\diamond}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF with μ below m ;*
- (ii) *there is an \mathcal{L}_{\diamond} -formula φ with $\mu(\varphi) < m$ such that $\mathbf{A} \models \varphi$ but φ is not valid in any frame of \mathbf{B} .*

Proof. (ii) IMPLIES (i). Let φ be an \mathcal{L}_{\diamond} -formula with $\mu(\varphi) < m$ that is valid on all frames in \mathbf{A} and not valid in 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 $\mathcal{B}^M \not\models \varphi$. The Hydra then responds with some set of pointed models \mathbf{A}^m ; since φ is valid on \mathbf{A} , for all $\mathcal{A} \in \mathbf{A}^m$ we have $\mathcal{A} \models \varphi$. By Theorem 3.3, it follows that Hercules has a winning strategy with μ below m for the $(\mathcal{L}_{\diamond}, \langle \mathbf{A}^m, \mathbf{B}^m \rangle)$ -FGM and thus for $(\mathcal{L}_{\diamond}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF.

(i) IMPLIES (ii). 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} ; in other words, she sets \mathbf{A}^m to be the set 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 $\mu(T) < m$; 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 m . Thus by Theorem 3.3, there is an \mathcal{L}_{\diamond} -formula φ with $\mu(\varphi) < m$ 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

5 The transfer axioms

We apply our formula-complexity games to prove the minimality of some modal 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 $\Box^m \bar{p} \vee \diamond^n p$. 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 $0 < m < n$.

In this and the following sections, we define a number of pointed models using figures for ease of understanding. We follow the convention that such pointed models consist of the relevant Kripke model and a point that is denoted




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

by the \triangleright sign next to it. Our goal is to prove the following.

Theorem 5.1 *For any $n \neq m \in \mathbb{N}$, $\square^m \bar{p} \vee \diamond^n p$ is absolutely minimal among all formulas defining (m, n) -transfer.*

The proof that for each $m, n \geq 0$, $\diamond^m p \rightarrow \diamond^n p$ is essentially the shortest formula defining (m, n) -transfer is split in four parts according to the ordering between m and n . Cases where one of the two is zero are treated in Appendix B.

5.1 Generalized density axioms

First we consider the generalized density axioms, i.e. (m, n) -transfer when $0 < m < n$. We prove that Theorem 5.1 holds in this case by considering a suitable formula-complexity game. Specifically, Hercules and the Hydra play a $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -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 1 and separated by the dotted line. \mathcal{A}_1 is constructed so that the vertical path leading from the 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. Each \mathcal{A}_i , for $2 \leq i \leq m+1$, 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 5.2 *In any winning strategy for Hercules for an $(\mathcal{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGF in which $\mathcal{A}_1 \in \mathbf{L}$ and $\mathcal{B} \in \mathbf{R}$, Hercules must pick a pointed model $(\mathcal{B}^M, \triangleright)$ based on the lowest irreflexive point in \mathcal{B} .*

Proof. It is easy to see that Hercules is not going to select a pointed model that is not based on the lowest non-reflexive point in \mathcal{B} because the Hydra can always 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 1. She has constructed 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: We 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 5.3 *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. Let T be any closed game tree on which the Hydra played greedily and η be any position of T such that $(\mathcal{B}^M, \triangleright) \in \mathbf{R}(\eta)$ while $(\mathcal{A}_i, \triangleright) \in \mathbf{L}(\eta)$ for some i with $1 \leq i \leq m+1$.*

- (i) *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 .*
- (ii) *If Hercules played a \square -move at η then he did not pick the left lowest reflexive point in \mathcal{B}^M .*

Proof. If Hercules picks the left lowest reflexive point when playing such a 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 3.4 that Hercules cannot win. The second claim is symmetric. \square

Lemma 5.4 *Suppose that \mathbf{L}, \mathbf{R} are classes of models and Hercules has a winning strategy for the $(\mathcal{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM. If T is any closed game tree in which the Hydra played greedily and η is any position of T such that $(\mathcal{B}^M, \triangleright) \in \mathbf{R}(\eta)$, then*

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

Proof. The first claim 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 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 5.3. \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 5.5 *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 and let T be a closed game tree in which the Hydra played greedily.*

- (i) *If $\{(\mathcal{A}_1, \triangleright), (\mathcal{A}_2, \triangleright)\} \subseteq \mathbf{L}$ and $(\mathcal{B}, \triangleright) \in \mathbf{R}$, then Hercules made at least one \vee -move during the game.*

- (ii) 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.
- (iii) 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.

(i) By Lemma 5.4, Hercules cannot play a modality as long as $(\mathcal{A}_1, \triangleright), (\mathcal{A}_2, \triangleright)$ are 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.

(ii) 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 5.4.i, Hercules cannot play a \square -move in any of these positions. Thus Hercules must perform 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 5.3 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.

(iii) 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, and 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 $\mathbf{w}_j \in \mathbf{L}(\nu_k)$, $\mathbf{v}_j \in \mathbf{R}(\nu_k)$ 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 that $\mathbf{w}_{j+1} \in \mathbf{L}(\eta')$ and $\mathbf{v}_{j+1} \in \mathbf{R}(\eta')$, where in the case of $j = 0$ we use Lemma 5.3 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 5.1 in the case $0 < m < n$.

Proof. If $0 < m < n$ we consider the $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF with \mathbf{A}, \mathbf{B} as depicted in Figure 1. By Lemma 5.2 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 5.5 Hercules must play at least one disjunction, one literal, n \diamond -moves, modal depth n , and at least m \square -moves. By Theorem 4.3, 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

5.2 Generalized transitivity axioms

Next we treat Theorem 5.1 in the case where $0 < n < m$. As before, we do so by considering a suitable $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -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 2. 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} .




Fig. 2. 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 5.6 *In any winning strategy for Hercules for an $(\mathcal{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGF in which $\mathcal{A}_1 \in \mathbf{L}$ and $\mathcal{B} \in \mathbf{R}$, Hercules 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 always reply with a bisimilar pointed model based on \mathcal{A}_1 . Similarly, if Hercules assigns the same valuation to the two dead-ends the Hydra 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 2.

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 2. 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 and similarly for their vertical successors. As before, if the Hydra chooses her models in this way, we say that she *mimics* Hercules’ choice.

FORMULA SIZE GAME ON MODELS: We 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 5.1.

Lemma 5.7 *Let \mathbf{L}, \mathbf{R} be classes of models so that Hercules has a winning strategy for the $(\mathcal{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -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.*

- (i) *If $(\mathcal{A}_1^M, \triangleright) \in \mathbf{L}(\eta)$ and $(\mathcal{B}^M, \triangleright) \in \mathbf{R}(\eta)$, then he picked a pointed model based on a point that lies on the hypotenuse of \mathcal{A}_1^M .*
- (ii) *If for some $i \in [3, m+1]$ we have that $(\mathcal{A}_i^M, \triangleright) \in \mathbf{L}(\eta)$ and $(\mathcal{B}^M, \triangleright) \in \mathbf{R}(\eta)$, 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 3.4 means that Hercules cannot win. \square

Lemma 5.8 *Suppose that \mathbf{L} and \mathbf{R} are classes of models and Hercules has a winning strategy for the $(\mathcal{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM. Suppose that T is a closed game tree, the Hydra played greedily, and η is a node of T .*

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

- (ii) If $(\mathcal{A}_2^M, \triangleright) \in \mathbf{L}(\eta)$ and $(\mathcal{B}^M, \triangleright) \in \mathbf{R}(\eta)$, 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 5.9 *Let \mathbf{L} and \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. Let T be a closed game tree in which the Hydra played greedily.*

- (i) *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 during the game.*
- (ii) *If $(\mathcal{A}_1^M, \triangleright) \in \mathbf{L}$, then T has at least n nested \diamond -moves and at least one literal move.*
- (iii) *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 5.5.i, except that it uses Lemma 5.8, and the proof of the third item is essentially the same as the proof of Proposition 5.5.iii. Thus we focus on the second item.

Since $(\mathcal{A}_1^M, \triangleright)$ and $(\mathcal{B}^M, \triangleright)$ satisfy the same literals, \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 perform 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 5.7, 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 decided to play a \square -move and picked a pointed model based on a point along the diagonal path in \mathcal{B}^M , the Hydra will 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 5.1 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 2 and Proposition 5.9. \square

6 The Löb axiom

Finally, we consider the Löb axiom, which defines the property of transitivity and converse-well-foundedness, i.e. that there are no infinite chains $w_0 R w_1 R \dots$. Note that this is a second-order property, and cannot be defined in first-order logic.




Fig. 3. The sets of frames $\mathbf{A} = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4\}$ and $\mathbf{B} = \{\mathcal{B}, \mathcal{B}_1\}$.

Theorem 6.1 *The formula $\Box\bar{p} \vee \Diamond(p \wedge \Box\bar{p})$ is absolutely minimal among all formulas defining the class of transitive and converse well-founded frames.*

We have already shown that $\Box\Box\bar{p} \vee \Diamond p$ is absolutely minimal among those formulas defining transitivity, so our strategy will be to expand on the frames and pointed models in Figure 2 to additionally force Hercules to play a conjunction. Since these models were already well-founded we can use previous results.

Let us consider an $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF played by Hercules and the Hydra with the frames shown in Figure 3. Obviously, $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$, and \mathcal{B} are obtained from the frames in Figure 2 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 or 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. Intuitively, 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 . It is obvious that 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 6.2 *In any winning strategy for Hercules in the $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -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.

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

Lemma 6.3 *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 on which the Hydra played greedily, then for any game position η and any non-leaf point w of \mathcal{A}_4 ,*

if $(\mathcal{A}_4^M, w) \in \mathbf{L}(\eta)$, $(\mathcal{B}_1^M, \triangleright) \in \mathbf{R}(\eta)$, 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 6.4 *Suppose that \mathbf{L} , \mathbf{R} are classes of models for which Hercules has a winning strategy for the $(\mathcal{L}_\diamond, (\mathbf{L}, \mathbf{R}))$ -FGM and let T be a closed game tree on which the Hydra played greedily.*

- (i) *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 $\mathbf{L}(\eta)$ contains a pointed model based on \mathcal{A}_4^M whereas $(\mathcal{B}_1^M, \triangleright) \in \mathbf{R}(\eta)$.*
- (ii) *If Hercules plays a \diamond -move in a position η in which $\mathbf{L}(\eta)$ contains a pointed model based on \mathcal{A}_4^M while $(\mathcal{B}_1^M, \triangleright)$ is on the right, he must play at least one subsequent \wedge -move.*

Proof.

(i) 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 6.3, he selected $(\mathcal{B}_1^M, \triangleright)$ again.

When Hercules plays such a move, the Hydra would reply 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.

(ii) 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 6.1.

Proof. Consider a $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -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 Figures 2 and 3. Hercules must choose his pointed models according to Lemmas 5.6 and 6.2, and Hydra replies by mimicking Hercules. Using Proposition 5.9 we see that if the Hydra plays greedily then any closed game tree must have modal depth at least two, contain two instances of \Box , one instance of each \diamond and \vee , and one variable. By Proposition 6.4, it also contains one conjunction, as required. \square

7 Conclusion

The present work was motivated to a large degree by ideas and results from [8], where the notion of minimal modal equivalent of a first-order conditions was introduced. Note however that the term minimal is used in [8] only with respect to the number of different variables needed to modally define a first-order condition which 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. With this we have shown that several familiar modal axioms are minimal with respect to all measures considered, including the Löb axiom, which is 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.

Question 1 *Is there a 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 exponentially in $\mu(\varphi_n)$?*

References

- [1] Adler, M. and N. Immerman, *An $n!$ lower bound on formula size*, ACM Transactions on Computational Logic **4** (2003), pp. 296–314.
- [2] Chagrov, A. and M. Zakharyashev, “Modal Logic,” Oxford logic guides **35**, Oxford University Press, 1997.
- [3] Fernández-Duque, D. and P. Iliev, *Succinctness in subsystems of the spatial mu-calculus*, Journal of Applied Logics (2018).
- [4] French, T., W. van der Hoek, P. Iliev and B. Kooi, *Succinctness of epistemic languages*, in: T. Walsh, editor, *Proceedings of IJCAI*, 2011, pp. 881–886.
- [5] French, T., W. van der Hoek, P. Iliev and B. Kooi, *On the succinctness of some modal logics*, Artificial Intelligence **197** (2013), pp. 56–85.
- [6] Hella, L. and M. Vilander, *The succinctness of first-order logic over modal logic via a formula size game*, Advances in Modal Logic **11**, 2016, pp. 401–419.
- [7] Immerman, N., “Descriptive complexity,” Graduate texts in computer science, Springer, 1999.
- [8] Vakarelov, D., *Modal definability in languages with a finite number of propositional variables and a new extension of the Sahlqvist’s class*, Advances in Modal Logic **4** (2003), pp. 499–518.
- [9] van Ditmarsch, H. and P. Iliev, *The succinctness of the cover modality*, Journal of Applied Non-Classical Logics **25** (2015), pp. 373–405.

Appendix

A Properties of the formula-complexity game on models

We have seen that a closed game tree T induces a formula ψ_T . Under certain conditions, we can also turn formulae into game trees.

Lemma A.1 *Let \mathbf{A}, \mathbf{B} be classes of models and $\varphi \in \mathcal{L}_\diamond$ be so that $\mathbf{A} \models \varphi$ and $\mathbf{B} \models \neg\varphi$. Then, Hercules has a strategy for the $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM so that any match 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 the game tree T is closed with $\psi_T = \perp$, as required.

$\varphi = \varphi_1 \vee \varphi_2$. 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 $(\mathcal{L}_\diamond, \langle \mathbf{A}_i, \mathbf{B}_i \rangle)$ -FGM with resulting closed game trees 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.

$\varphi = \diamond\theta$. 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 $\psi_T = \diamond\theta$.

OTHER CASES. Each of the remaining cases is dual to one discussed above. \square

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

Lemma A.2 *Let \mathbf{A}, \mathbf{B} be classes of models and let T be a closed game tree for the $(\mathcal{L}_\diamond, \langle \mathbf{A}_i, \mathbf{B}_i \rangle)$ -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 η starting from the leaves we show that $\mathbf{L}(\eta) \models \psi_\eta$ and $\mathbf{R}(\eta) \models \neg\psi_\eta$. The base case is immediate since Hercules can only play a literal when it is true on the left but false on the right, and the inductive steps for \perp, \top, \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. If Hercules played a \diamond -move on η with daughter η' then for each $\mathbf{l} \in \mathbf{L}(\eta)$ he chose $\mathbf{l}' \in \Box\mathbf{L}(\eta)$ and placed $\mathbf{l}' \in \mathbf{L}(\eta')$; by the induction hypothesis $\mathbf{l}' \models \psi_{\eta'}$, so that $\mathbf{l} \models \diamond\psi_{\eta'} = \psi_\eta$ by the semantics of \diamond . Meanwhile for $\mathbf{r} \in \mathbf{R}(\eta)$, if $\mathbf{r}' \in \Box\mathbf{r}$ then since the Hydra




Fig. B.1. The frames \mathcal{A}_1 , \mathcal{A}_2 and \mathcal{B} and the pointed models based on them.

played greedily $\mathbf{r}' \in \mathbf{R}(\eta')$, and since \mathbf{r}' was arbitrary we see that $\mathbf{r} \models \neg \diamond \psi_{\eta'}$. The case for a \square -move is symmetric. \square

With this we prove Theorem 3.3

Proof. Let \mathbf{A} , \mathbf{B} be classes of models, μ any complexity measure, and $m \in \mathbb{N}$. Recall that Theorem 3.3 states that the following are equivalent:

- (i) Hercules has a winning strategy for the $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM with μ below m , and
- (ii) there is an \mathcal{L}_\diamond -formula φ with $\mu(\varphi) < m$ and $\mathbf{A} \models \varphi$ whereas $\mathbf{B} \models \neg \varphi$.

First assume that (i) holds, and let Hydra play the $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGM greedily. By using his winning strategy, Hercules can ensure that the game terminates on some closed tree T with $\mu(T) < m$. But by definition this means that $\mu(\psi_T) < m$, and by Lemma A.2, $\mathbf{A} \models \psi_T$ while $\mathbf{B} \models \neg \psi_T$.

Conversely, if (ii) holds, by Lemma A.1 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 $\mu(T) < m$. \square

B Generalized reflexivity and recurrence

In this appendix we prove Theorem 5.1 in cases where one of the parameters is zero.

B.1 The generalized reflexivity axioms

Recall that we write *n-reflexivity* instead of $(0, n)$ -transfer. In order to prove that Theorem 5.1 holds in this case, we consider a $(\mathcal{L}_\diamond, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF where $\mathbf{A} = \{\mathcal{A}_1, \mathcal{A}_2\}$ and \mathbf{B} contains a single element \mathcal{B} . These frames are shown in the left rectangle in Figure B.1 and separated by the dotted line. The “highest” point in \mathcal{A}_2 can be reached 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 immediate that $p \rightarrow \diamond^n p$ is valid on both \mathcal{A}_1 and \mathcal{A}_2 and not valid on \mathcal{B} .

Next we study Hercules’ possible strategies. We begin with his 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 B.1 *In any winning strategy for Hercules for an $(\mathcal{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGF in which $\mathcal{A}_1 \in \mathbf{L}$ and $\mathcal{B} \in \mathbf{R}$,*

- (i) *Hercules chooses the valuation on \mathcal{B} so that at least one literal is true in one point but not on the other, and*
- (ii) *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 B.1. 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 can reply with the pointed models shown on the left of the dotted line in the right half in Figure B.1. She selects these pointed models so that 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: Let us consider now 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 that the modal moves that Hercules may make have some restrictions. The following can be seen by observing that playing otherwise would produce bisimilar pointed models on each side.

Lemma B.2 *Let \mathbf{L}, \mathbf{R} be classes of models so that Hercules has a winning strategy for the $(\mathcal{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -FGM and T a closed game tree in which the Hydra played greedily.*

- (i) *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 .*
- (ii) *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 B.3 *Suppose that \mathbf{L}, \mathbf{R} are classes of models and that Hercules has a winning strategy for the $(\mathcal{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -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 B.2 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.




Fig. B.2. The frames $\mathcal{A}_1, \dots, \mathcal{A}_{m+1}$ and \mathcal{B} and the pointed models based on them.

Since Hercules cannot play \perp or \top , he must use at least one variable. \square

With this we are ready to prove Theorem 5.1 in the case where $m = 0$.

Proof. Let \mathbf{A} and \mathbf{B} be as depicted in the left rectangle in Figure B.1; since the frames of \mathbf{A} are n -reflexive but the ones in \mathbf{B} are not, by Theorem 4.3 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 B.2 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 B.1. Lemma B.2 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 B.3 he can win the former only by playing a literal-move whereas the latter can be won only by playing a sequence of n \diamond -moves that must be made in order to perform a cycle leading back to the black point in \mathcal{A}_2 , giving us at least n occurrences of \diamond and modal depth at least n . We can then use Theorem 4.3 to conclude that $\bar{p} \vee \diamond^n p$ is absolutely minimal. \square

B.2 The generalized recurrence axioms

Now we treat the m -recurrence axioms, where $n = 0$. This time Hercules and the Hydra play a $(\mathcal{L}_{\diamond}, \langle \mathbf{A}, \mathbf{B} \rangle)$ -FGF where $\mathbf{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_{m+1}\}$ while \mathbf{B} contains a single element \mathcal{B} , as depicted in the left rectangle in Figure B.2. 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 all the frames in \mathbf{A} and it is not valid in the frame \mathcal{B} .

SELECTION OF THE MODELS ON THE RIGHT: It follows from Lemma B.1 that Hercules must pick the pointed model $(\mathcal{B}^M, \triangleright)$ shown in the right half of Fig-

ure B.2. 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 with the pointed models shown on the left of the dotted line in the right half in Figure B.2. Again, she picks these pointed models so that points that satisfy the same set of literals have the same colour.

FORMULA SIZE GAME ON MODELS: Let us 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 B.4 *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. Using Lemma B.2, we see that in order to win a FGM with a starting position η in which $(\mathcal{A}_1^M, \triangleright)$ is on the left and $(\mathcal{B}^M, \triangleright)$ is on the right, Hercules must not play either a \diamond - or a \square -move at η . On the other hand, for every game position ν in which there is some $(\mathcal{A}_i^M, \triangleright)$ on the left for $2 \leq i \leq m+1$ and $(\mathcal{B}^M, \triangleright)$ on the right, Hercules must play 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 every $(\mathcal{A}_i^M, \triangleright)$ for $2 \leq i \leq m+1$ from $(\mathcal{A}_1^M, \triangleright)$. \square

Lemma B.5 *Let \mathbf{L}, \mathbf{R} be classes of models so that Hercules has a winning strategy for the $(\mathcal{L}_\diamond, \langle \mathbf{L}, \mathbf{R} \rangle)$ -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 5.5.iii. With this we are ready to prove Theorem 5.1 for the case where $n = 0$.

Proof. Consider the (\mathbf{A}, \mathbf{B}) -FGF where \mathbf{A}, \mathbf{B} are as depicted in Figure B.2 on the left: by Lemma B.1 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 B.3 Hercules must play at least one variable, by Lemma B.4 he must play at least one \vee -move, by Lemma B.5 he must play at least m \square -moves and modal depth at least m on the resulting FGM, and we can apply Theorem 4.3. \square