Modal Logic Is More Succinct Iff Bi-Implication Is Available in Some Form

Christoph Berkholz
Technische Universität Ilmenau, Germany

Dietrich Kuske
Technische Universität Ilmenau, Germany

Christian Schwarz
Technische Universität Ilmenau, Germany

Abstract

Is it possible to write significantly smaller formulae, when using more Boolean operators in addition to the De Morgan basis (and, or, not)? For propositional logic a negative answer was given by Pratt: every formula with additional operators can be translated to the De Morgan basis with only polynomial increase in size.

Surprisingly, for modal logic the picture is different: we show that adding bi-implication allows to write exponentially smaller formulae. Moreover, we provide a complete classification of finite sets of Boolean operators showing they are either of no help (allow polynomial translations to the De Morgan basis) or can express properties as succinct as modal logic with additional bi-implication. More precisely, these results are shown for the modal logic T (and therefore for K). We complement this result showing that the modal logic S5 behaves as propositional logic: no additional Boolean operators make it possible to write significantly smaller formulae.

2012 ACM Subject Classification Theory of computation → Logic

Keywords and phrases succinctness, modal logic

Digital Object Identifier 10.4230/LIPIcs.STACS.2024.12

Funding *Christoph Berkholz*: Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - project number 414325841.

1 Introduction

Many classical logics such as propositional logic, first-order and second-order logic, temporal and modal logics incorporate a complete set of Boolean operators in their definitions – mostly the De Morgan basis ($\wedge$, $\vee$, $\neg$). While for the expressiveness it is clearly irrelevant which complete operator set is used, this choice may have an impact on the succinctness of formulae. The main aim of this paper is to understand which additional operators allow to express properties with strictly more succinct formulae.

A first simple observation is that if an additional operator defines a *read-once function*, i.e., it can be expressed in the De Morgan basis in such a way that every variable occurs at most once, then it can easily be eliminated without blowing-up the formula too much. Thus, read-once operators such as $x \to y \equiv \neg x \vee y$ are really just syntactic sugar. For operators that are not read-once, such as bi-implication $x \leftrightarrow y$ or the ternary majority operator $\mathrm{maj}(x, y, z)$, the situation is less clear, because mindlessly replacing them with any equivalent De Morgan formula may lead to an exponential explosion of the formula size. So can it be that such additional operators actually allow to write exponentially more succinct formulae? For propositional logic, a negative answer was given by Pratt [11]: first balance the formula so that it has logarithmic depth and then replace the additional operators by any De Morgan translation. This clearly leads to a linear increase in formula depth and therefore only to a polynomial increase in formula size.

41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024).
Editors: Olaf Beyersdorff, Mamadou Moustapha Kanté, Orna Kupferman, and Daniel Lokshtanov;
Article No. 12; pp. 12:1–12:17
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Balancing a formula is, however, not possible for logics that contain quantifiers. For such logics it is still possible to efficiently remove certain operators that are not read-once. We show that if an operator $\mathrm{op}(x_1, \ldots, x_k)$ is "locally read-once", that is, has for every $i \in [k]$ an equivalent De Morgan formula in which x_i appears only once, then it can be efficiently eliminated. While we prove this result explicitly for modal logic in Section 2, it actually holds for other logics with quantifiers as well. An example of an operator that is locally read-once, but not read-once, is $\mathrm{maj}(x, y, z)$ (see Example 3).

One of our main results is that this characterisation is tight for modal logic: for operators that are not locally read-once, there is no way of removing them without increasing the formula size exponentially in the worst-case. Thus, adding any operator that is not locally read-once to the De Morgan basis allows to write exponentially more succinct formulae. One example of such a useful operator is bi-implication $x \leftrightarrow y$, for which it is not hard to show that any equivalent expression over $(\wedge, \vee, \neg)$ contains x *and* y twice and hence it is not locally read-once. Furthermore, we analyse the succinctness classes of modal logic wrt. polynomial translations. Can it be that adding even more Boolean operators allows to express properties even more succinct? Here we show that *any* extension of modal logic by a set of operators containing at least one that is not locally read-once has the same succinctness. Thus there are exactly two succinctness classes that are exponentially separated: one containing standard modal logic and the other containing its extension by bi-implication.

Since this dichotomy is in contrast with propositional logic, where only one succinctness class exists, we also investigate what happens for fragments of modal logic defined by restrictions on the Kripke structures. Here we obtain the same dichotomy for structures with a reflexive accessibility relation. But upon considering equivalence relations only, we can show that the two succinctness classes collapse, as they do in propositional logic.

Related work. It seems that this paper is the first to consider the influence of Boolean operators to the succinctness of modal logics. Other aspects have been studied in detail.

Pratt [11] studied the effect of complete bases of binary operators on the size of propositional formulae and proved in particular that there are always polynomial translations. Wilke [15] proved a succinctness gap between two branching time temporal logics, Adler and Immerman [1] developed a game-theoretic method and used it to improve Wilke's result and to show other succinctness gaps. The succinctness of further temporal logics was considered, e.g., in [2, 10].

Lutz et al. [9, 8] study the succinctness and complexity of several modal logics. French et al. [3] consider multi-modal logic with an abbreviation that allows to express "for all $i \in \Gamma$ and all i-successors, φ holds" where Γ is some set of modalities. Using Adler-Immerman-games, they prove (among other results in similar spirit) that this abbreviation allows exponentially more succinct formulae than plain multi-modal logic.

Grohe and Schweikardt [5] study the succinctness of first-order logic with a bounded number of variables and, for that purpose, develop extended syntax trees as an alternative view on Adler-Immerman-games. These extended syntax trees were used by van Ditmarsch et al. [14] to prove an exponential succinctness gap between a logic of contingency (public announcement logic, resp.) and modal logic.

Hella and Vilander [6] define a formula size game (modifying the Adler-Immerman-game) and use it to show that bisimulation invariant first-order logic is non-elementarily more succinct than modal logic.

2 Where "all" logics coincide

Suppose $\mathcal{L}$ is some classical logic like propositional logic, predicate logic, second-order predicate logic, temporal logic (linear or branching), or modal logic. All these logics use Boolean connectives, usually $\{\vee, \neg\}$, $\{\vee, \wedge, \neg\}$, $\{\rightarrow\}$, or $\{\vee, \wedge, \neg, \rightarrow, \leftrightarrow\}$. One could also allow connectives like majority ("at least two of the statements s_1, s_2, and s_3 are true") or divisibility by three ("all or none of the statements s_1, s_2, and s_3 are true") without changing the expressive power. But what about the succinctness? More precisely: if, in addition to the De Morgan basis $\{\vee, \wedge, \neg\}$, we allow Boolean connectives from the sets F and G, respectively, is the resulting logic $\mathcal{L}[F]$ more succinct than $\mathcal{L}[G]$? The main result of this section demonstrates that there are at most two "succinctness classes" (up to a polynomial), namely the one containing plain logic $\mathcal{L}$ and the other one containing the extension of $\mathcal{L}$ with bi-implication $\leftrightarrow$.

Formulae. Let P be a countably infinite set of propositional variables and let $\mathbb{B} = \{\top, \bot\}$ be the Boolean domain where we assume $\top > \bot$. For a set F of Boolean functions, we let $\mathrm{ML}[F]$ be the set of all formulae in modal logic that may use operators from F in addition to the constants $\top$ and $\bot$ as well as the Boolean operators $\neg, \wedge$, and $\vee$. Formally, $\mathrm{ML}[F]$ is defined by the syntax

$$\varphi ::= \ \bot \mid \top \mid p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid f(\varphi, \dots, \varphi) \mid \Diamond\varphi,$$

for propositional variables $p \in P$ and operators $f \in F$[1]. We write ML for $\mathrm{ML}[\emptyset]$, the set of formulae in standard modal logic, and $\mathrm{ML}[f]$ for $\mathrm{ML}[\{f\}]$. Furthermore, $\mathrm{PL}[F] \subseteq \mathrm{ML}[F]$ denotes the set of propositional formulae that may use functions from F as well; more precisely, $\mathrm{PL}[F]$ is the set of formulae from $\mathrm{ML}[F]$ that do not use the operator $\Diamond$. Since we always include the null-ary functions $\top$ and $\bot$ and the unary functions p and $\neg p$, we can assume that all functions in F are of arity at least two.

The *size* $|\varphi|$ of a formula from $\mathrm{ML}[F]$ is the number of nodes in its syntax tree.

Semantics. Formulae are interpreted over *pointed Kripke structures*, i.e., over tuples $S = (W, R, V, \iota)$, consisting of a set W of possible worlds, a binary accessibility relation $R \subseteq W \times W$, a valuation $V : W \to \mathcal{P}(P)$, assigning to every world in W the set of propositional variables that are declared to be true at this world, and an initial world $\iota \in W$.

The satisfaction relation $\models$ between a world w of S and an $\mathrm{ML}[F]$-formula is defined inductively, where

- $S, w \models p$ if $p \in V(w)$,
- $S, w \models \Diamond\varphi$ if $S, w' \models \varphi$ for some $w' \in W$ with $(w, w') \in R$, and
- $S, w \models f(\alpha_1, \dots, \alpha_k)$ if $f(b_1, \dots, b_k) = \top$ where, for all $i \in [k]$, $b_i = \top$ iff $S, w \models \alpha_i$

(the definitions of $S, w \models \varphi$ for $\varphi \in \{\top, \bot, \neg\alpha, \alpha \vee \beta, \alpha \wedge \beta\}$ are as expected). A pointed Kripke structure S is a *model* of φ ($S \models \varphi$) if φ holds in its initial world, i.e., $S, \iota \models \varphi$.

Now let $\mathcal{C}$ be some class of pointed Kripke structures. A formula φ is *satisfiable in* $\mathcal{C}$ if it has a model in $\mathcal{C}$ and φ *holds in* $\mathcal{C}$ if every structure from $\mathcal{C}$ is a model of φ. The formula φ *entails* the formula ψ *in* $\mathcal{C}$ (written $\varphi \models_{\mathcal{C}} \psi$) if any model of φ from $\mathcal{C}$ is also a model of ψ; φ and ψ are *equivalent over* $\mathcal{C}$ (denoted $\varphi \equiv_{\mathcal{C}} \psi$) if $\varphi \models_{\mathcal{C}} \psi$ and $\psi \models_{\mathcal{C}} \varphi$.

[1] Depending on the context, we consider an element $f \in F$ as a Boolean function or as a symbol in a formula.

Classes of Kripke structures. For different application areas (i.e., interpretations of the operator $\Diamond$), the following classes of Kripke structures have attracted particular interest. For convenience, we define them as classes of *pointed* Kripke structures.

- The class $\mathcal{S}_{\mathrm{K}}$ of all pointed Kripke structures.
- The class $\mathcal{S}_{\mathrm{T}}$ of all pointed Kripke structures with reflexive accessibility relation.
- The class $\mathcal{S}_{\mathrm{S5}}$ of pointed Kripke structures where the accessibility relation is an equivalence relation.

Suppose φ is a propositional formula. Then $S, \iota \models \varphi$ only depends on the set $V(\iota)$ of variables that hold in the world ι. Thus, instead of evaluating propositional formulae (as special modal formulae) in Kripke structures, it suffices to evaluate them (as is usually done) in sets of propositional variables or, equivalently, mappings from the set of variables into the Boolean domain $\mathbb{B}$.

▶ **Definition 1** (Translations). *Let F and G be sets of Boolean functions, $\mathcal{C}$ a class of pointed Kripke structures, and $\kappa \colon \mathbb{N} \to \mathbb{N}$ some function. Then* $\mathrm{ML}[F]$ *has* κ-translations wrt. $\mathcal{C}$ *in* $\mathrm{ML}[G]$ *if, for every formula $\varphi \in \mathrm{ML}[F]$, there exists a formula $\psi \in \mathrm{ML}[G]$ with $\varphi \equiv_{\mathcal{C}} \psi$ and $|\psi| \leq \kappa(|\varphi|)$.*

The logic $\mathrm{ML}[F]$ *has* polynomial translations wrt. $\mathcal{C}$ *in* $\mathrm{ML}[G]$ *if it has κ-translations wrt. $\mathcal{C}$ for some polynomial function κ;* sub-exponential *and* exponential *translations are defined similarly.*

In this section, we aim at sufficient conditions for the existence of polynomial translations wrt. $\mathcal{S}_{\mathrm{K}}$ of $\mathrm{ML}[F]$ in $\mathrm{ML}[G]$, i.e., we will always consider the class of *all* pointed Kripke structures. For notational convenience, we will regularly omit the explicit reference to the class $\mathcal{S}_{\mathrm{K}}$, e.g., "equivalent" means "equivalent over $\mathcal{S}_{\mathrm{K}}$", $\varphi \models \psi$ means $\varphi \models_{\mathcal{S}_{\mathrm{K}}} \psi$, and "$\kappa$-translations" means "$\kappa$-translations wrt. $\mathcal{S}_{\mathrm{K}}$".

In this paper, $[n] = \{1, 2, \ldots, n\}$ for all $n \in \mathbb{N}$.

2.1 Polynomial translations

Suppose F and G are sets of Boolean functions. Recall that formulae from $\mathrm{ML}[F]$ can use operators from F as well as $\neg, \wedge$, and $\vee$ (and similarly for $\mathrm{ML}[G]$). Since the De Morgan basis is complete, for every function $f \in F$, there is some formula $\omega \in \mathrm{PL} \subseteq \mathrm{ML}[G]$ such that the formulae $f(p_1, \ldots, p_k)$ and $\omega(p_1, \ldots, p_k)$ are equivalent. Consequently, to translate a formula $\varphi \in \mathrm{ML}[F]$ into a formula $\psi \in \mathrm{ML}[G]$, we only need to replace every sub-formula $f(\alpha_1, \ldots, \alpha_k)$ in φ by $\omega(\alpha_1, \ldots, \alpha_k)$. In general, this translation leads to an exponential size increase. But if, in the formula ω, every variable p_i appears only once, we obtain a linear translation. In this section, we provide polynomial (and in general non-linear) translations of $\mathrm{ML}[F]$ in $\mathrm{ML}[G]$ under the following weaker assumption.

▶ **Definition 2** (Representations). *Let G be a set of Boolean functions, f a Boolean operator of arity k, and $i \in [k]$.*

A $\mathrm{PL}[G]$-representation *of (f, i) is a $\mathrm{PL}[G]$-formula $\omega_i(p_1, \ldots, p_k)$ that is equivalent to the $\mathrm{PL}[f]$-formula $f(p_1, \ldots, p_k)$ and uses the variable p_i at most once.*

A set F of Boolean functions has $\mathrm{PL}[G]$*-representations if there are* $\mathrm{PL}[G]$*-representations for all $f \in F$ and $i \in [\mathrm{ar}(f)]$.*

▶ **Example 3.** Consider the majority function $\mathrm{maj}(p, q, r)$ that is true iff at least two arguments are true. Then $(\mathrm{maj}, 1)$ has the PL-representation $\big(p \wedge (q \vee r)\big) \vee (q \wedge r)$. Using the symmetry of maj, it follows that $\{\mathrm{maj}\}$ has PL-representations.

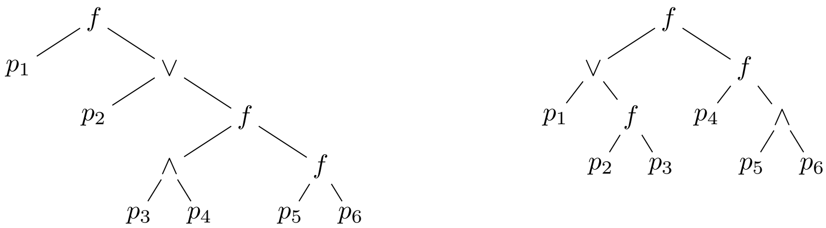

Figure 1 Syntax trees of $\varphi = f(p_1, p_2 \vee f(p_3 \wedge p_4, f(p_5, p_6))$ and $\psi = f(p_1 \vee f(p_2, p_3), f(p_4, p_5 \wedge p_6))$.

Next, consider bi-implication $\leftrightarrow$, where the following observation seems folklore: if the PL-formula $\psi(p, q)$ is equivalent to $p \leftrightarrow q$ and mentions p only once (say, under an even number of negations), then $\top \equiv (\bot \leftrightarrow \bot) \equiv \psi(\bot, \bot) \leq \psi(\top, \bot) \equiv \bot$, a contradiction.

Assuming that F has PL[G]-representations, we will construct, from a formula in ML[F], an equivalent formula in ML[G] of polynomial size. Since this will be done inductively, we will have to deal with formulae from ML[$F \cup G$] and the task then is better described as elimination of functions $f \in F$ from formulae in ML[$F \cup G$].

Before we present the details of our construction, we briefly demonstrate the main idea behind the proof (for $F = \{f\}$ and $G = \emptyset$). The main results (Lemma 7 and Proposition 8) will appear at the end of the section.

Assume that f is of arity 2 and consider the two formulae

$$\varphi = f\big(\ p_1,\ p_2 \vee f\big(\ p_3 \wedge p_4,\ f(p_5, p_6)\ \big)\ \big) \text{ and } \psi = f\big(\ p_1 \vee f(p_2, p_3),\ f(p_4, p_5 \wedge p_6)\ \big),$$

whose syntax trees are depicted in Fig. 1. A distinguishing property of the left tree is the existence of a branch (a path from the root to a leaf) that contains all F-vertices. Assuming f to have PL[G]-representations (with $G = \emptyset$), there exist Boolean combinations $\omega_1(x, y) \equiv \omega_2(x, y) \equiv f(x, y)$ of the variables x and y, such that x occurs only once in $\omega_1(x, y)$ and y only once in $\omega_2(x, y)$. Proceeding bottom-up, we now replace each F-vertex $f(\alpha, \beta)$ in the syntax tree of φ by either $\omega_1(\alpha, \beta)$ or $\omega_2(\alpha, \beta)$, depending on whether we have previously modified the left or the right sub-tree (regarding $f(p_5, p_6)$, we are free to choose between ω_1 and ω_2). Note that, although ω_1 and ω_2 may duplicate some parts of φ, our choice ensures that we never duplicate such parts whose size has already changed. Consequently, this procedure results in a linear increase $\ell \cdot |\varphi|$ in the size of φ, where the coefficient ℓ essentially depends on how often y occurs in $\omega_1(x, y)$ and how often x occurs in $\omega_2(x, y)$. The resulting formula φ' belongs to ML[G] and is equivalent to φ. Hence ML[$F \cup G$]-formulae for which all F-vertices lie on some common branch have ML[G]-translations of linear size.

The formula ψ on the other hand does not have the property that all F-vertices lie on some common branch, but the two sub-formulae α and β rooted at the children of the root do. Hence we can apply the above transformation to them separately, yielding equivalent ML[G]-formulae α' and β' whose size increases at most by a factor of ℓ. Then $\psi' = f(\alpha', \beta') \equiv \psi$ is of size $|\psi'| \leq \ell \cdot |\psi|$. Note that this step reduces the total number of F vertices – in particular, ψ' now contains only a single operator from F. Applying the step again yields an equivalent ML[G]-formula ψ'' of size $|\psi''| \leq \ell \cdot |\psi'| \leq \ell^2 \cdot |\psi|$. For arbitrary ML[$F \cup G$]-formulae, the number of steps relates to the "nesting depth" D of those F-vertices that have at least two arguments in ML[$F \cup G$] $\setminus$ ML[G], thus resulting in a formula of size $\ell^D \cdot |\psi|$. In this section, we will show that D is at most logarithmic in the size of ψ, thus giving a polynomial bound and establishing the main part of the succinctness result.

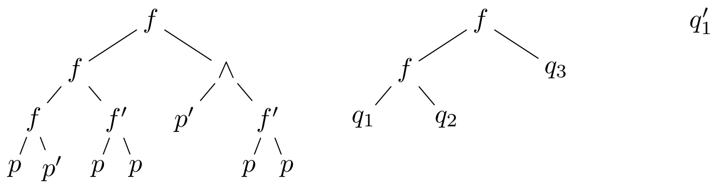

Figure 2 Syntax trees of $\varphi = f\big(f\big(f(p,p'), f'(p,p)\big),\ p' \wedge f'(p,p)\big)$, $d_F(\varphi) = f\big(f(q_1,q_2), q_3\big)$, and $d_F^2(\varphi) = q'_1$.

Let F and G be disjoint sets of Boolean functions and let $N_{F,G} \subseteq \mathrm{ML}[F \cup G]$ be described by the syntax

$$\varphi ::= \ \psi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid \neg\varphi \mid f(\psi, \ldots, \psi, \varphi, \psi, \ldots, \psi) \mid g(\varphi, \ldots \varphi) \mid \Diamond\varphi,$$

where $\psi \in \mathrm{ML}[G]$, $f \in F$, and $g \in G$. The formulae from $N_{F,G}$ have a slightly more general property than ψ in the example above. More precisely, a formula φ belongs to $N_{F,G}$ if, and only if, for any sub-formula $f(\alpha_1, \ldots, \alpha_k)$ of φ with $f \in F$, at most one α_i contains some operator from F. Conjunction, disjunction, and operators from G on the other hand may have occurrences of operators from F in any number of arguments.

▶ **Definition 4** (Derivative). *Let $\varphi = \varphi(\overline{p}) \in \mathrm{ML}[F \cup G]$ with $\overline{p} = (p_1, \ldots, p_m)$. The* F-derivative $d_F(\varphi)$ of φ *is the smallest* $\mathrm{ML}[F \cup G]$*-formula $\gamma(\overline{p}, q_1, \ldots, q_n)$ (up to renaming of the variables $q_1, \ldots, q_n$), such that*

- *q_i occurs exactly once in $\gamma(\overline{p}, q_1, \ldots q_n)$ for all $i \in [n]$ and*
- *there exist $\alpha_1, \ldots, \alpha_n \in N_{F,G} \setminus \mathrm{ML}[G]$ such that φ and $\gamma(\overline{p}, \alpha_1, \ldots, \alpha_n)$ are identical.*

Intuitively, $\gamma(\overline{p}, q_1, \ldots, q_n)$ is obtained from $\varphi(\overline{p})$ by simultaneously replacing all "maximal $(N_{F,G} \setminus \mathrm{ML}[G])$-formulae" by distinct fresh variables $q_1, \ldots, q_n$ (where multiple occurrences of the same formula are replaced by different variables). An example is depicted in Fig. 2 (with $F = \{f, f'\}$ and $G = \emptyset$).

Let $\varphi \in \mathrm{ML}[F \cup G] \setminus \mathrm{ML}[G]$. Then $d_F(\varphi)$ contains fewer occurrences of operators from F than φ. Hence there exists a smallest integer $r \geq 0$ for which the r-th derivative $d_F^r(\varphi) = d_F(d_F(\cdots d_F(\varphi)\cdots))$ is an $\mathrm{ML}[G]$-formula, where $d_F^0(\varphi) = \varphi$.

▶ **Definition 5** (Rank). *Let $\varphi \in \mathrm{ML}[F \cup G]$. The* F-rank $\mathrm{rank}_F(\varphi)$ of φ *is the smallest integer $r \geq 0$ for which $d_F^r(\varphi) \in \mathrm{ML}[G]$.*

We first show that a formula with high F-rank must also be large.

▶ **Lemma 6.** *Let F and G be disjoint sets of Boolean functions and $\varphi \in \mathrm{ML}[F \cup G]$. Then $|\varphi| \geq 2^{\mathrm{rank}_F(\varphi)}$.*

Proof. Recall that we can assume that all functions in F have arity at least 2. Since we only refer to the F-rank of a formula, we will simply speak of the rank of a formula.

We prove the stronger claim that the syntax tree of a formula φ of positive rank has at least $2^{\mathrm{rank}_F(\varphi)-1}$ sub-trees of the form $f(\beta_1, \ldots, \beta_k)$ with $f \in F$ and $\beta_1, \ldots, \beta_k \in \mathrm{ML}[G]$ (in the following, we call such a sub-tree an *F-leaf*).

Since counting the F-leaves in all derivatives of φ results in a lower bound on the total number of operators from F in φ, it follows that φ contains at least $1+2+\ldots+2^{\mathrm{rank}_F(\varphi)-1} = 2^{\mathrm{rank}_F(\varphi)} - 1$ operators from F. Hence $|\varphi| \geq 2^{\mathrm{rank}_F(\varphi)}$ since none of the functions in F has arity zero.

It now remains to prove the bound on the number of F-leaves. Recall that we consider formulae of rank at least one. If $\mathrm{rank}_F(\varphi) = 1$, φ contains at least one operator from F and hence has at least $2^{1-1} = 1$ F-leaf. Now, assume that $\varphi = \varphi(\overline{p})$ is of rank at least two. Let $\gamma(\overline{p}, q_1, \dots, q_n)$ be the derivative of φ, and let $\alpha_1, \dots, \alpha_n \in N_{F,G}$ such that $\varphi = \gamma(\overline{p}, \alpha_1, \dots, \alpha_n)$. Then, for every F-leaf $f(\beta_1, \dots, \beta_k)$ of $\gamma(\overline{p}, q_1, \dots, q_n)$, there exist indices $i \neq j$ such that $\beta_i, \beta_j \in \{q_1, \dots, q_n\}$. By induction hypothesis, $\gamma(\overline{p}, q_1, \dots, q_n)$ has at least $2^{\mathrm{rank}_F(\varphi)-2}$ F-leaves, each of which contains at least two of the fresh variables $\{q_1, \dots, q_n\}$. Since each α_i contains at least one operator from F, $\varphi = \gamma(\overline{p}, \alpha_1, \dots, \alpha_n)$ has at least twice the number of F-leaves compared to $\gamma(\overline{p}, q_1, \dots, q_n)$, i.e., $2^{\mathrm{rank}_F(\varphi)-1}$ F-leaves. ◀

We can now turn to the main ingredient for our succinctness result.

▶ **Lemma 7.** *Let F and G be disjoint sets of Boolean functions and, for $f \in F$ and $i \in [\mathrm{ar}(f)]$, let $\omega_{f,i} \in \mathrm{PL}[G]$ be a $\mathrm{PL}[G]$-representation of (f, i). Let, furthermore, $\kappa\colon \mathbb{N} \to \mathbb{N}$ be a monotone function such that $1 \leq \kappa(1)$ and $|\omega_{f,i}| \leq \kappa(\mathrm{ar}(f))$ for any $f \in F$ and $i \in [\mathrm{ar}(f)]$. Finally, let $\kappa'\colon \mathbb{N} \to \mathbb{N}\colon n \mapsto \kappa(n)^{\log_2 n} \cdot n$.*

Then $\mathrm{ML}[F \cup G]$ has κ'-translations in $\mathrm{ML}[G]$.

Proof. For $\varphi \in \mathrm{ML}[F \cup G]$, let K_φ denote the maximal arity of any $f \in F$ occurring in φ (or 1 if $\varphi \in \mathrm{ML}[G]$). Note that $\kappa(K_\varphi) \geq 1$ for all $\varphi \in \mathrm{ML}[F \cup G]$.

We prove the following claim: every $\varphi \in \mathrm{ML}[F \cup G]$ is equivalent to an $\mathrm{ML}[G]$-formula of size at most $\big(\kappa(K_\varphi)\big)^{\mathrm{rank}_F(\varphi)} \cdot |\varphi|$. Since $K_\varphi \leq |\varphi|$ and $\mathrm{rank}_F(\varphi) \leq \log_2 |\varphi|$ by Lemma 6, this claim ensures that every $\mathrm{ML}[F \cup G]$-formula φ of size n has an equivalent $\mathrm{ML}[G]$-formula of size at most $\kappa(n)^{\log_2 n} \cdot n = \kappa'(n)$.

The proof of the claim proceeds by induction on the F-rank of φ. Since F remains fixed, we will refer to the F-rank simply as rank. As before, we assume that all operators in F are at least binary.

Let $\varphi \in \mathrm{ML}[F \cup G]$ be of rank at most one. We show by induction on the structure of φ that there exists an equivalent $\mathrm{ML}[G]$-formula φ' of size $|\varphi'| \leq \kappa(K_\varphi) \cdot |\varphi|$. If φ is a propositional variable or one of the constants $\top$ or $\bot$, $|\varphi| = 1 \leq \kappa(1) \cdot |\varphi|$ since $\kappa(1) \geq 1$.

Now, assume that $\varphi = \alpha_1 \wedge \alpha_2$. By induction hypothesis, there exist $\mathrm{ML}[G]$-formulae β_1, β_2 with $|\beta_j| \leq \kappa(K_{\alpha_j}) \cdot |\alpha_j|$ and $\beta_j \equiv \alpha_j$ for $j \in [2]$. Since κ is monotone and $K_{\alpha_j} \leq K_\varphi$, $|\beta_j| \leq \kappa(K_{\alpha_j}) \cdot |\alpha_j| \leq \kappa(K_\varphi) \cdot |\alpha_j|$ for $j \in [2]$. Set $\varphi' = \beta_1 \wedge \beta_2$. Then φ' is an $\mathrm{ML}[G]$-formula, equivalent to φ, and of size

$$|\varphi'| = |\beta_1| + |\beta_2| + 1 \leq \kappa(K_\varphi) \cdot (|\alpha_1| + |\alpha_2|) + 1 \leq \kappa(K_\varphi) \cdot |\varphi|\,, \qquad \text{since } \kappa(K_\varphi) \geq 1.$$

A similar argument establishes the cases $\alpha_1 \vee \alpha_2$, $\neg\alpha$, $\Diamond\alpha$, and $g(\alpha_1, \dots, \alpha_n)$ for $g \in G$.

Finally, assume that φ is of the form $f(\alpha_1, \dots, \alpha_k)$ with $f \in F$. Since φ is of rank one and therefore a formula in $N_{F,G}$, there exists an index $i \in [k]$ such that no argument other than α_i contains an operator from f, i.e., with $\alpha_j \in \mathrm{ML}[G]$ for all $j \neq i$ (and $\alpha_i \in \mathrm{ML}[F \cup G]$). By induction hypothesis, α_i is equivalent to an $\mathrm{ML}[G]$-formula β_i of size $|\beta_i| \leq \kappa(K_{\alpha_i}) \cdot |\alpha_i| \leq \kappa(K_\varphi) \cdot |\alpha_i|$. Set $\beta_j = \alpha_j$ for all $j \neq i$. Then $\beta_1, \dots, \beta_k \in \mathrm{ML}[G]$. Recall that $\omega_{f,i}(p_1, \dots, p_k)$ is a $\mathrm{PL}[G]$-representation of (f, i) that uses the variable p_i at most once and has size $\leq \kappa(k) \leq \kappa(K_\varphi)$. Set $\varphi' = \omega_{f,i}(\beta_1, \dots, \beta_k)$. Then φ' is equivalent to $f(\alpha_1, \dots, \alpha_k) = \varphi$ and belongs to $\mathrm{ML}[G]$. Furthermore, φ' is obtained from $\omega_{f,i}(p_1, \dots, p_k)$ by replacing one variable with β_i and all others with formulae of size at most $\sum_{j \neq i} |\beta_j|$. Since all functions in F are at least binary, it follows that $\sum_{j \neq i} |\beta_j| \geq 1$. Hence

$$\begin{aligned}|\varphi'| &= |\omega_{f,i}(\beta_1,\dots,\beta_k)| \\ &\le |\beta_i| + \kappa(k)\cdot \sum_{j\in[k]\setminus\{i\}} |\beta_j| && \text{, since } \sum_{j\in[k]\setminus\{i\}} |\beta_j| \ge 1 \\ &\le \kappa(K_\varphi)\cdot|\alpha_i| + \kappa(K_\varphi)\cdot \sum_{j\in[k]\setminus\{i\}} |\alpha_j| \\ &\le \kappa(K_\varphi)\cdot|f(\alpha_1,\dots,\alpha_k)| = \kappa(K_\varphi)\cdot|\varphi|.\end{aligned}$$

This shows the claim for formulae of rank at most one.

We proceed by induction on the rank of φ. Let $\varphi = \varphi(\overline{p}) \in \mathrm{ML}[F\cup G]$ be of rank $r \ge 2$. Let $\gamma(\overline{p}, q_1,\dots,q_n)$ be the derivative of $\varphi(\overline{p})$ and $\alpha_1,\dots,\alpha_n$ be $N_{F,G}$-formulae, such that $\varphi = \gamma(\overline{p},\alpha_1,\dots,\alpha_n)$. Since the α_i are of rank one, there exist $\beta_1,\dots,\beta_n \in \mathrm{ML}[G]$ with $\alpha_i \equiv \beta_i$ and $|\beta_i| \le \kappa(K_{\alpha_i})\cdot|\alpha_i| \le \kappa(K_\varphi)\cdot|\alpha_i|$ for $i\in[k]$. Set $\psi = \psi(\overline{p}) = \gamma(\overline{p},\beta_1,\dots,\beta_k)$. Then ψ is equivalent to φ and of size $|\psi| \le \kappa(K_\varphi)\cdot|\varphi|$. Intuitively, ψ is obtained from φ by replacing the "maximal" $\mathrm{ML}[F\cup G]$ sub-formulae of rank 1 by equivalent $\mathrm{ML}[G]$-formulae. Since ψ is of rank $r-1$, it follows by induction hypothesis that ψ is equivalent to a formula $\varphi' \in \mathrm{ML}[G]$ of size $|\varphi'| \le \kappa(K_\psi)^{r-1}\cdot|\psi| \le \kappa(K_\varphi)^r\cdot|\varphi|$. Since $\varphi\equiv\psi\equiv\varphi'$, this finishes the verification of the claim from the beginning of this proof. ◀

From Lemma 7, we can get the main result of this section, stating that $\mathrm{ML}[F\cup G]$ is not more succinct than $\mathrm{ML}[G]$, provided F is a finite set of Boolean functions with $\mathrm{PL}[G]$-representations.

▶ **Proposition 8.** *Let F and G be disjoint finite sets of Boolean functions such that F has $\mathrm{PL}[G]$-representations. Then $\mathrm{ML}[F\cup G]$ has polynomial translations in $\mathrm{ML}[G]$.*

Since $\mathrm{ML}[F] \subseteq \mathrm{ML}[F\cup G]$, this implies in particular that $\mathrm{ML}[F]$ has polynomial translations in $\mathrm{ML}[G]$. In view of Example 3, it ensures specifically that ML[maj] has polynomial translations in plain ML.

Proof. Since F is finite, there is some constant c such that any $f\in F$ and $i\in[\mathrm{ar}(f)]$ have a $\mathrm{PL}[G]$-representation of size at most c. By the previous lemma, any $\varphi\in\mathrm{ML}[F\cup G]$ is thus equivalent to an $\mathrm{ML}[G]$-formula of size at most $c^{\log_2|\varphi|}\cdot|\varphi| = |\varphi|^{1+\log_2 c} \le |\varphi|^d$ for some constant d. ◀

2.2 A decidable characterisation of representations

Proposition 8 gives a condition ("has $\mathrm{PL}[G]$-representations") for the existence of polynomial translations. In this section, we demonstrate that this condition is decidable.

▶ **Definition 9** (Local monotonicity). *A Boolean function $f\colon \mathbb{B}^k \to \mathbb{B}$ is* monotone in the i-th argument *if either*

(M1) *for all $\overline{a}\in\mathbb{B}^{i-1}$, $\overline{b}\in\mathbb{B}^{k-i}$, $f(\overline{a},\bot,\overline{b}) \le f(\overline{a},\top,\overline{b})$ or*

(M2) *for all $\overline{a}\in\mathbb{B}^{i-1}$, $\overline{b}\in\mathbb{B}^{k-i}$, $f(\overline{a},\bot,\overline{b}) \ge f(\overline{a},\top,\overline{b})$.*

That is, when changing the i-th argument from $\bot$ to $\top$, while keeping the remaining ones fixed, the truth value of f uniformly increases or decreases (where, in both cases, the value may also remain unchanged). A function is called *locally monotone* if it is monotone in every argument and *non-locally-monotone* otherwise. Then conjunction, disjunction, negation, implication, as well as majority are locally monotone functions, while bi-implication is not.

▶ **Proposition 10.** *Let F and G be disjoint sets of Boolean functions. Then F has* PL$[G]$*-representations if, and only if, all functions in F are locally monotone or some function from G is non-locally-monotone.*

Proof. First, suppose that all functions in F are locally monotone. We prove that, under this assumption, F even has PL-representations, which ensures the claim since $\mathrm{PL} \subseteq \mathrm{PL}[G]$. So, let $f \in F$ be of arity k (as before, we can assume $k \geq 2$). To simplify notation, we will only construct a PL-representation of (f, k). In addition, we assume that f is increasing in the k-th argument, i.e., $f(\overline{a}, \bot) \leq f(\overline{a}, \top)$ for all $\overline{a} \in \mathbb{B}^{k-1}$. There exists a PL-formula $\omega(x_1, \ldots, x_k)$ that is equivalent to $f(x_1, \ldots, x_k)$. Since $f(\overline{a}, \bot) \leq f(\overline{a}, \top)$ for any $\overline{a} \in \mathbb{B}^{k-1}$, it follows that

$$f(x_1, \ldots, x_k) \equiv \Big(\omega(x_1, \ldots, x_{k-1}, \top) \wedge x_k\Big) \vee \omega(x_1, \ldots, x_{k-1}, \bot)\,.$$

In particular, the formula on the right uses the variable x_k only once and therefore forms a PL-representation of (f, k).

Next, suppose there is some $g \in G$ that is non-locally-monotone. We have to provide PL$[G]$-representations for all functions $f \in F$. So let $f \in F$ be arbitrary and of arity k; for notational simplicity, we prove that there is some PL$[G]$-representation of (f, k).

Let A denote the set of tuples $\overline{a} \in \mathbb{B}^{k-1}$ with $f(\overline{a}, \top) = f(\overline{a}, \bot) = \top$ and let B denote the set of tuples $\overline{b} \in \mathbb{B}^{k-1}$ with $f(\overline{b}, \top) > f(\overline{b}, \bot)$. Let $\overline{x} \in A$ abbreviate the PL-formula

$$\bigvee_{\overline{a} \in A} \left(\bigwedge_{i \in [k-1], a_i = \top} x_i \wedge \bigwedge_{i \in [k-1], a_i = \bot} \neg x_i \right)$$

and let $\overline{x} \in B$ be defined likewise. Then the formula $f(\overline{x}, x_k)$ is equivalent to the formula

$$\overline{x} \in A \vee \big(\overline{x} \notin A \wedge (\overline{x} \in B \leftrightarrow x_k)\big)\,, \tag{1}$$

which belongs to PL$[\leftrightarrow]$ (and mentions the variable x_k only once).

Let ℓ be the arity of the function $g \in G$ and suppose, for notational simplicity, that it is not monotone in its last argument. Hence there are $\overline{a}, \overline{b} \in \mathbb{B}^{\ell-1}$ such that $g(\overline{a}, \bot) < g(\overline{a}, \top)$ and $g(\overline{b}, \bot) > g(\overline{b}, \top)$. For $i \in [\ell - 1]$, set

$$\theta_i = \begin{cases} a_i & \text{if } a_i = b_i \\ \overline{x} \in B & \text{if } a_i > b_i \\ \overline{x} \notin B & \text{if } a_i < b_i \end{cases}$$

and write $\overline{\theta}$ for the tuple $(\theta_i)_{i \in [\ell-1]}$. Note that, for all $i \in [\ell-1]$, θ_i is equivalent to a_i if $\overline{x} \in B$, and to b_i if $\overline{x} \notin B$. By choice of $\overline{a}$ and $\overline{b}$ it follows that the formulae $(\overline{x} \in B \leftrightarrow x_k) \in \mathrm{PL}[\leftrightarrow]$ and $g(\overline{\theta}, x_k) \in \mathrm{PL}[G]$ are equivalent. We finally replace $(\overline{x} \in B \leftrightarrow x_k)$ in the formula (1) by $g(\overline{\theta}, x_k)$ which yields the PL$[G]$-representation $\overline{x} \in A \vee (\overline{x} \notin A \wedge g(\overline{\theta}, x_k))$ of (f, k).

It remains to be shown that the existence of PL$[G]$-representations implies that (i) or (ii) holds. So assume that F has PL$[G]$-representations and that (ii) does not hold, i.e., that all functions from G are locally monotone. We prove by induction on the size of a formula $\varphi(p_1, \ldots, p_k) \in \mathrm{PL}[G]$ the following: if the variable p_k appears only once in φ, then the function represented by φ is monotone in its k-th argument. The claim is trivial for formulae of the form $\top$, $\bot$, and p_i.

For the induction step, let $\varphi = g(\alpha_1, \ldots, \alpha_\ell)$ with $g \in G$. Since the formula φ contains the variable p_k only once, it appears in at most one of the arguments α_i; for notational simplicity, we assume it appears in α_ℓ. By the induction hypothesis, there is $\otimes \in \{\leq, \geq\}$ such that

$$\forall \overline{a} \in \mathbb{B}^{\ell-1} \colon \alpha_\ell(\overline{a}, \bot) \otimes \alpha_\ell(\overline{a}, \top)\,. \tag{2}$$

Since we assumed all functions from G to be locally monotone, there also is $\odot \in \{\leq, \geq\}$ such that

$$\forall \overline{a} \in \mathbb{B}^{\ell-1} \colon g(\alpha_1(\overline{a}), \ldots, \alpha_{\ell-1}(\overline{a}), \bot) \odot g(\alpha_1(\overline{a}), \ldots, \alpha_{\ell-1}(\overline{a}), \top)\,. \tag{3}$$

Putting (2) and (3) together, we obtain

$$\begin{aligned} &\forall \overline{a} \in \mathbb{B}^{\ell-1} \colon \varphi(\overline{a}, \bot) \leq \varphi(\overline{a}, \top) \quad \text{if } \otimes = \odot\text{, and} \\ &\forall \overline{a} \in \mathbb{B}^{\ell-1} \colon \varphi(\overline{a}, \bot) \geq \varphi(\overline{a}, \top) \quad \text{if } \otimes \neq \odot\,. \end{aligned}$$

Intuitively, $\odot$ indicates whether φ increases (when going from $\bot$ to $\top$ in the last argument of g, i.e., in α_ℓ), while $\otimes$ may flip the direction if the truth-value of α_ℓ decreases, when increasing p_k. Hence, the formula φ represents a function that is locally monotone in its last argument provided φ is of the form $g(\alpha_1, \ldots, \alpha_\ell)$ for some $g \in G$. The arguments are similar if φ is of the form $\alpha_1 \wedge \alpha_2$, $\alpha_1 \vee \alpha_2$, or $\neg\alpha_1$.

This finishes the inductive proof.

Recall that F has PL$[G]$-representations. Since each PL$[G]$-representation of (f, i) describes a function (namely f) that is monotone in the i-th argument, we obtain that all functions from F are locally monotone, which completes the proof. ◀

Now Propositions 8 and 10 yield the following central result.

▶ **Theorem 11.** *Let $\mathcal{C}$ be some class of pointed Kripke structures. Let F and G be disjoint finite sets of Boolean functions such that all functions from F are locally monotone or some function from G is non-locally-monotone. Then* ML$[F]$ *has polynomial translations wrt. $\mathcal{C}$ in* ML$[G]$.

Proof. By the assumptions on F and G, F has PL$[G]$-representations by Proposition 10. Hence, by Proposition 8, ML$[F]$ has polynomial translations wrt. $\mathcal{S}_\mathrm{K}$ in ML$[G]$, i.e., for any formula $\varphi \in \mathrm{ML}[F]$, there exists a formula $\psi \in \mathrm{ML}[G]$ of polynomial size with $\varphi \equiv_{\mathcal{S}_\mathrm{K}} \psi$. Since $\mathcal{C} \subseteq \mathcal{S}_\mathrm{K}$, this implies $\varphi \equiv_{\mathcal{C}} \psi$. Hence, indeed, ML$[F]$ has polynomial translations wrt. $\mathcal{C}$ in ML$[G]$. ◀

As we have mentioned at the beginning of this section, all the results also hold for other logics $\mathcal{L}$ such as predicate logic, second-order predicate logic, temporal logic (linear or branching), and probably many more. Suppose that some function from G is non-locally-monotone. Since $\leftrightarrow$ is non-locally-monotone, it follows that $\mathcal{L}[G]$ and $\mathcal{L}[\leftrightarrow]$ have polynomial translations in each other, i.e., they are equally succinct (up to a polynomial). Similarly, if all functions from G are locally monotone, then $\mathcal{L}[G]$ and $\mathcal{L} = \mathcal{L}[\emptyset]$ have polynomial translations in each other since all functions from $\emptyset$ are locally monotone. In other words, for any set of Boolean functions G, $\mathcal{L}[G]$ is as succinct as $\mathcal{L}[\leftrightarrow]$ or as $\mathcal{L}$.

Thus, the situation looks similar for many logics: there are at most two "succinctness classes". For propositional logic, it was shown by Pratt [11] that also PL$[\leftrightarrow]$ has polynomial translations in PL, i.e., that there is just one such class. In the following section, we will show that, whether or not modal logic has two "succinctness classes", depends on the class of Kripke structures.

3 Where modal logics diverge

So far, we saw that for many logics $\mathcal{L}$, there are at most two "succinctness classes" (up to a polynomial), namely those of $\mathcal{L}$ and $\mathcal{L}[\leftrightarrow]$, respectively. In this section, we will show that these classes differ for the modal logic T (and hence also for K) but coincide for the modal logic S5. We will therefore, from now on, be precise and return to the original notation, i.e., write $\equiv_{\mathcal{S}_\mathrm{K}}$ instead of $\equiv$ etc.

3.1 Where the "succinctness classes" differ

Our aim in this section is to show that $\mathrm{ML}[\leftrightarrow]$ is exponentially more succinct than ML. Formally, we prove $\mathrm{ML}[\leftrightarrow]$ does not have sub-exponential translations wrt. $\mathcal{S}_\mathrm{T}$ in ML (recall that $\mathcal{S}_\mathrm{T}$ is the class of pointed Kripke structures with reflexive accessibility relation).

▶ **Lemma 12.** *The logic* $\mathrm{ML}[\leftrightarrow]$ *does not have sub-exponential translations wrt.* $\mathcal{S}_\mathrm{T}$ *in* ML.

Proof. Let $\varphi_0 = p_0$ and $\varphi_{n+1} = p_{n\oplus 1} \wedge (p \leftrightarrow \Diamond\varphi_n)$ for $n \geq 0$, where $m \oplus n := (m+n) \bmod 2$. We will prove that $|\psi| \geq 2^n$ for any $n \geq 0$ and any $\psi \in \mathrm{ML}$ with $\psi \equiv_{\mathcal{S}_\mathrm{T}} \varphi_n$. Since $|\varphi_n|$ is linear in n, this ensures the lemma's claim.

In this proof, we use the following notation. Let $\psi \in \mathrm{ML}$ be any ML-formula. Then ψ is a Boolean combination of formulae of the form $\top$, $\bot$, $p \in P$, and $\Diamond\lambda$ with $\lambda \in \mathrm{ML}$. By E_ψ, we denote the set of all formulae $\lambda \in \mathrm{ML}$ such that $\Diamond\lambda$ appears in this Boolean combination with even negation depth. Similarly, O_ψ denotes the set of all formulae $\lambda \in \mathrm{ML}$ such that $\Diamond\lambda$ appears in this Boolean combination with odd negation depth.

A more formal definition proceeds as follows by induction:

$$E_\psi = \begin{cases} \emptyset & \text{if } \psi \in \{\top, \bot\} \cup P \\ \{\lambda\} & \text{if } \psi = \Diamond\lambda, \lambda \in \mathrm{ML} \\ O_\alpha & \text{if } \psi = \neg\alpha \\ E_\alpha \cup E_\beta & \text{if } \psi \in \{\alpha \wedge \beta, \alpha \vee \beta\} \end{cases} \qquad O_\psi = \begin{cases} \emptyset & \text{if } \psi \in \{\top, \bot\} \cup P \\ \emptyset & \text{if } \psi = \Diamond\lambda, \lambda \in \mathrm{ML} \\ E_\alpha & \text{if } \psi = \neg\alpha \\ O_\alpha \cup O_\beta & \text{if } \psi \in \{\alpha \wedge \beta, \alpha \vee \beta\} \end{cases}$$

We now show, by induction on n, that any ML-formula ψ with $\psi \equiv_{\mathcal{S}_\mathrm{T}} \varphi_n$ satisfies $|\psi| \geq 2^n$. Since φ_n uses (at most) the propositional variables p, p_0, and p_1, we can assume the same about ψ.

The case $n = 0$ is easy to see since any formula has size at least $1 = 2^0$. Let now $n \geq 0$ and consider the formula $\varphi_{n+1} = p_{n\oplus 1} \wedge (p \leftrightarrow \Diamond\varphi_n)$ and assume there were an ML-formula $\psi \equiv_{\mathcal{S}_\mathrm{T}} \varphi_{n+1}$ that satisfies $|\psi| < 2^{n+1}$.

Note that, although the sets E_ψ and O_ψ may have non-empty intersection,

$$\left|\bigvee E_\psi\right| + \left|\bigvee O_\psi\right| \leq |\psi| < 2^{n+1},$$

hence at least one of these disjunctions must be of size $< 2^n$.

Case 1: $|\bigvee O_\psi| < 2^n$. Then, by the induction hypothesis, $\varphi_n \not\equiv_{\mathcal{S}_\mathrm{T}} \bigvee O_\psi$. Let m denote the size of the formula $\varphi_n \wedge \neg\psi$.

Let α be a formula from ML that is satisfiable in $\mathcal{S}_\mathrm{T}$, of size at most m, and that uses no propositional variables other than p, p_0, p_1. Then α has a model $A_\alpha \in \mathcal{S}_\mathrm{T}$ that uses, at most, the propositional variables from α. Since there are only finitely many such formulae α, there exists a finite set $\mathcal{C}_m \subseteq \mathcal{S}_\mathrm{T}$ of pointed Kripke structures such that every ML-formula has a model in $\mathcal{C}_m$, provided it is satisfiable in $\mathcal{S}_\mathrm{T}$, of size at most m, and uses at most the propositional variables from $\{p, p_0, p_1\}$.

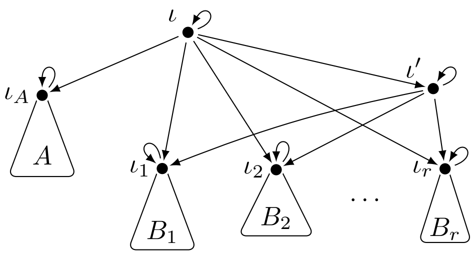

Figure 3 Schematic representation of Kripke structure S with $V(\iota) = V(\iota') = \{p_{n\oplus 1}\}$.

Let O^+_ψ be the set of formulae λ from O_ψ with $\lambda \models_{\mathcal{S}_{\mathrm{T}}} \varphi_n$. Since $\left|\bigvee O^+_\psi\right| \le |\bigvee O_\psi| < 2^n$, the induction hypothesis ensures $\bigvee O^+_\psi \not\equiv_{\mathcal{S}_{\mathrm{T}}} \varphi_n$. On the other hand, $\bigvee O^+_\psi \models_{\mathcal{S}_{\mathrm{T}}} \varphi_n$, hence the formula $\varphi_n \wedge \neg \bigvee O^+_\psi$ is satisfiable in $\mathcal{S}_{\mathrm{T}}$ and of size at most m. Consequently, there exists a structure $A = (W_A, V_A, R_A, \iota_A) \in \mathcal{C}_m$ with $A \models \varphi_n$ but $A \not\models \bigvee O^+_\psi$.

Let $B_1, \dots, B_r \in \mathcal{C}_m$ with $B_i = (W_i, R_i, V_i, \iota_i)$ be the models of $\neg\varphi_n$ from $\mathcal{C}_m$. We assume that the sets W_i for $i \in [r]$ and W_A are mutually disjoint.

We now define a Kripke structure $S = (W, R, V) \in \mathcal{S}_{\mathrm{T}}$ as follows (cf. Fig. 3):

$$W = \{\iota, \iota'\} \uplus \bigcup_{i\in[r]} W_i \cup W_A$$

$$R = \big\{(\iota,\iota), (\iota,\iota'), (\iota',\iota'), (\iota,\iota_A)\big\} \cup \big(\{\iota,\iota'\} \times \{\iota_i \mid i \in [r]\}\big) \cup \bigcup_{i\in[r]} R_i \cup R_A$$

$$V(w) = \begin{cases} \{p_{n\oplus 1}\} & \text{if } w \in \{\iota, \iota'\} \\ V_i(w) & \text{if } w \in W_i \text{ for some } i \in [r] \\ V_A(w) & \text{if } w \in W_A \end{cases}$$

From S, we obtain the pointed Kripke structures (S, ι) and (S, ι') by choosing the initial world as ι and ι' respectively. Note that both structures belong to $\mathcal{S}_{\mathrm{T}}$ since the accessibility relation R is reflexive. Our aim is to prove $(S, \iota) \models \neg\varphi_{n+1} \wedge \psi$, which contradicts the equivalence of φ_{n+1} and ψ.

First, we show $(S, \iota) \models \neg\varphi_{n+1}$. Recall that $A \models \varphi_n$ and therefore $(S, \iota_A) \models \varphi_n$. From $(\iota, \iota_A) \in R$, we obtain $(S, \iota) \models \neg p \wedge \Diamond\varphi_n$, implying $(S, \iota) \models \neg\varphi_{n+1}$.

To prove $(S, \iota) \models \psi$, we first show $(S, \iota') \models \varphi_{n+1}$, implying $(S, \iota') \models \psi$ since we assumed $\varphi_{n+1} \equiv_{\mathcal{S}_{\mathrm{T}}} \psi$. From this, we will infer that also $(S, \iota) \models \psi$.

Recall that $B_i \models \neg\varphi_n$ and therefore $(S, \iota_i) \models \neg\varphi_n$ for all $i \in [r]$. Furthermore, $(S, \iota') \models \neg p_{n \bmod 2}$ implies $(S, \iota') \models \neg\varphi_n$. Since $\{\iota', \iota_i \mid i \in [r]\}$ is the set of worlds accessible from ι', this implies $(S, \iota') \models \neg\Diamond\varphi_n$. Since, in addition, $(S, \iota') \models p_{n\oplus 1} \wedge \neg p$, we obtain $(S, \iota') \models \varphi_{n+1}$. Now $\varphi_{n+1} \equiv_{\mathcal{S}_{\mathrm{T}}} \psi$ and $(S, \iota') \in \mathcal{S}_{\mathrm{T}}$ imply $(S, \iota') \models \psi$.

The final step in our proof is the verification of $(S, \iota) \models \psi$. Recall that ψ is a Boolean combination of atomic formulae and of formulae $\Diamond\lambda$ with $\lambda \in O_\psi \cup E_\psi$. Note that (S, ι) and (S, ι') agree in the atomic formulae holding there. Since O_ψ and E_ψ are the formulae $\Diamond\lambda$ appearing in the Boolean combination with odd and even, respectively, negation depth, it suffices to show the following:

1. If $\lambda \in E_\psi$ with $(S, \iota') \models \Diamond\lambda$, then $(S, \iota) \models \Diamond\lambda$.
2. If $\lambda \in O_\psi$ with $(S, \iota') \not\models \Diamond\lambda$, then $(S, \iota) \not\models \Diamond\lambda$.

To demonstrate the former, suppose $\lambda \in E_\psi$ with $(S, \iota') \models \Diamond\lambda$. Then $(S, \iota') \models \lambda$ or there is $i \in [r]$ with $(S, \iota_i) \models \lambda$. Since $(\iota, \iota') \in R$ and $(\iota, \iota_i) \in R$, we obtain $(S, \iota) \models \Diamond\lambda$ in either case.

To demonstrate the second claim regarding $\lambda \in O_\psi$, we proceed by contraposition: so let $\lambda \in O_\psi$ with $(S,\iota) \models \Diamond\lambda$. Here, we distinguish two cases.

- Suppose $\lambda \models_{\mathcal{S}_\mathrm{T}} \varphi_n$, i.e., $\lambda \in O_\psi^+$. From $(S,\iota) \models \Diamond\lambda$, we obtain that the formula λ holds in one of the worlds ι, ι', ι_A, or ι_i for some $i \in [r]$. But $(S,\iota_A) \models \lambda$ implies $A \models \lambda$, which is impossible since $A \not\models \bigvee O_\psi^+$ and $\lambda \in O_\psi^+$. But also $(S,\iota_i) \models \lambda$ and therefore $B_i \models \lambda$ is impossible since $\lambda \models_{\mathcal{S}_\mathrm{T}} \varphi_n$, $B_i \in \mathcal{S}_\mathrm{T}$, and $B_i \models \neg\varphi_n$. Consequently, the formula λ holds in one of the worlds ι and ι'. Again using $\lambda \models_{\mathcal{S}_\mathrm{T}} \varphi_n$, we obtain that also φ_n holds in ι or in ι'. But this cannot be the case since $p_{n \bmod 2}$ does not hold in either of the two worlds – a contradiction. This finishes the verification of the second claim above in case $\lambda \in O_\psi^+ \subseteq O_\psi$.
- Finally, the case $\lambda \not\models_{\mathcal{S}_\mathrm{T}} \varphi_n$ remains to be considered. But then the formula $\lambda \wedge \neg\varphi_n$ is satisfiable in $\mathcal{S}_\mathrm{T}$. Since the size of this formula is bounded by m, there is some pointed Kripke structure $B \in \mathcal{C}_m$ with $B \models \lambda \wedge \neg\varphi_n$. The choice of the pointed Kripke structures $B_1, \dots B_r$ implies $B = B_i$ for some $i \in [r]$. Hence $(S,\iota_i) \models \lambda$. From $(\iota',\iota_i) \in R$, we obtain $(S,\iota') \models \Diamond\lambda$.

This finishes the proof of the two numbered claims above. As explained there, they imply $(S,\iota) \models \psi$.

So, we proved $(S,\iota) \models \neg\varphi_{n+1} \wedge \psi$ in case $\left|\bigvee O_\psi\right| < 2^n$ contradicting the equivalence of φ_{n+1} and ψ.

Case 2: $|\bigvee E_\psi| < 2^n$. We consider the formula $\neg p_{n\oplus 1} \vee \big((\neg p) \leftrightarrow \Diamond\varphi_n\big) \equiv_{\mathcal{S}_\mathrm{K}} \neg\varphi_{n+1} \equiv_{\mathcal{S}_\mathrm{T}} \neg\psi$. Observe that $O_{\neg\psi} = E_\psi$, such that $|\bigvee O_{\neg\psi}| < 2^n$. Hence we can use the same argument as before for $\neg\psi$ to obtain a contradiction: simply label the worlds ι and ι' with $\{p, p_{n\oplus 1}\}$ instead of $\{p_{n\oplus 1}\}$. ◀

We now come to the classification of modal logics ML$[F]$ that have polynomial translations in ML.

▶ **Theorem 13.** *Let $\mathcal{C}$ be either of the classes $\mathcal{S}_\mathrm{K}$ or $\mathcal{S}_\mathrm{T}$. Let F be a finite set of Boolean functions. Then the following are equivalent:*

(1) *All functions from F are locally monotone.*

(2) *The set F has* PL*-representations.*

(3) ML$[F]$ *has polynomial translations wrt. $\mathcal{C}$ in* ML*.*

(4) ML$[\leftrightarrow]$ *does not have sub-exponential translations wrt. $\mathcal{C}$ in* ML$[F]$*.*

(5) ML$[\leftrightarrow]$ *does not have polynomial translations wrt. $\mathcal{C}$ in* ML$[F]$*.*

Proof. The implication (1)⇒(2) follows from Proposition 10 (with $G = \emptyset$), the implication (2)⇒(3) is Proposition 8 (again, with $G = \emptyset$). Now suppose (3) and assume, towards a contradiction, that (4) does not hold. Then ML$[\leftrightarrow]$ has sub-exponential translations in ML$[F]$ and ML$[F]$ has polynomial translations in ML. Since $f(n)^k$ is sub-exponential for any sub-exponential function f and any constant k, we get that ML$[\leftrightarrow]$ has sub-exponential translations in ML, contradicting Lemma 12. Thus, the implication (3)⇒(4) holds. The implication (4)⇒(5) is trivial. Finally, suppose (5) holds. Then, by Proposition 8 again, $\{\leftrightarrow\}$ does not have PL$[F]$-representations. Hence, by Proposition 10 (with $F = \{\leftrightarrow\}$ and $G = F$), all functions in F are locally monotone. ◀

3.2 Where the "succinctness classes" collapse

In this section we show that, for every finite set of operators F, ML[F] has polynomial translations wrt. $\mathcal{S}_{\mathrm{S5}}$ in ML. Recall that $\mathcal{S}_{\mathrm{S5}}$ is the set of pointed Kripke structures whose accessibility relation is an equivalence relation. Note that two formulae φ and ψ are equivalent wrt. $\mathcal{S}_{\mathrm{S5}}$ if, and only if, they are equivalent wrt. the class of pointed Kripke structures S from $\mathcal{S}_{\mathrm{S5}}$ whose accessibility relation is total. These pointed Kripke structures have the pleasant property that a $\Diamond$-quantified formula either holds in every world or in none, i.e.,

$$S, w \models \Diamond\varphi \quad \text{iff} \quad S, w' \models \Diamond\varphi\,. \tag{4}$$

We use the above equivalence to prove that ML[$\leftrightarrow$]-formulae can be "balanced" when considering structures from $\mathcal{S}_{\mathrm{S5}}$ only. That ML[$\leftrightarrow$] has polynomial translations wrt. $\mathcal{S}_{\mathrm{S5}}$ in ML then forms an easy corollary. The general result for any set F then follows from Theorem 11.

For a formula $\varphi \in \mathrm{ML}[F]$, let $\|\varphi\|$ denote the number of leaves of the syntax tree of φ, i.e., the total number of occurrences of propositional variables and constants $\top$ and $\bot$. Furthermore, let $d(\varphi)$ be the depth of the syntax tree of φ, that is, the length of a longest path from the root to a leaf. In particular, $d(\varphi) = 0$ if, and only if, $\varphi \in P \cup \{\top, \bot\}$. If all operators in φ have arity at most r, the number of leaves, the size, and the depth of φ satisfy $\|\varphi\| \leq |\varphi| \leq r^{d(\varphi)+1}$.

▶ **Lemma 14.** *For every $\varphi \in \mathrm{ML}[\leftrightarrow]$ there exists a formula $\varphi' \in \mathrm{ML}[\leftrightarrow]$ with $\varphi' \equiv_{\mathcal{S}_{\mathrm{S5}}} \varphi$ and $d(\varphi') \leq 8 \cdot \big(1 + \log_2 \|\varphi\|\big)$.*

▶ Remark 15. Since $\mathrm{PL}[\leftrightarrow] \subseteq \mathrm{ML}[\leftrightarrow]$, the above lemma implies that each formula $\varphi \in \mathrm{PL}[\leftrightarrow]$ is equivalent to some PL[$\leftrightarrow$]-formula of depth logarithmic in $\|\varphi\|$. According to Gashkov and Sergeev [4], a more general form of this result for propositional logic was known to Khrapchenko in 1967, namely that it holds for any complete basis of Boolean functions (e.g., for $\{\wedge, \vee, \neg, \leftrightarrow\}$ as here). They also express their regret that the only source for this is a single paragraph in a survey article by Yablonskii and Kozyrev [16], see also [7]. Often, it is referred to as Spira's theorem who published it in 1971 [13], assuming that all at most binary Boolean functions are allowed in propositional formulae. Khrapchenko's general form was then published by Savage [12].

Proof. Throughout the proof, we consider the logic ML[$\leftrightarrow$] only and therefore simply speak of formulae when referring to ML[$\leftrightarrow$]-formulae.

The proof proceeds by induction on $\|\varphi\|$. First assume that $\|\varphi\| = 1$. Then $\varphi = \mathrm{op}_1\mathrm{op}_2 \cdots \mathrm{op}_r\lambda$ with $r \geq 0, \mathrm{op}_i \in \{\neg, \Diamond\}$ for all $i \in [r]$, and $\lambda \in P \cup \{\top, \bot\}$. Using the equivalences $\neg\neg\alpha \equiv_{\mathcal{S}_{\mathrm{S5}}} \alpha$, $\Diamond\Diamond\alpha \equiv_{\mathcal{S}_{\mathrm{S5}}} \Diamond\alpha$, and $\Diamond\neg\Diamond\alpha \equiv_{\mathcal{S}_{\mathrm{S5}}} \neg\Diamond\alpha$ (the latter two following from (4)), the formula φ is equivalent over $\mathcal{S}_{\mathrm{S5}}$ to a formula ψ of depth at most $3 \leq 8 \cdot \big(1 + \log_2 \|\varphi\|\big)$. This establishes the case $\|\varphi\| = 1$.

Otherwise, $\|\varphi\| \geq 2$ and φ contains some binary operator. Let $m = \|\varphi\|$. Intuitively, we split the formula φ into two parts, each containing about half the leaves from φ. Formally, there are formulae $\alpha(x)$ with only one occurrence of x and β such that $\varphi = \alpha(\beta)$,

- $\|\beta\| > \frac{m}{2}$, and
- $\beta = \mathrm{op}(\beta_1, \beta_2)$ with $\mathrm{op} \in \{\wedge, \vee, \leftrightarrow\}$ and $\|\beta_1\|, \|\beta_2\| \leq \frac{m}{2}$.

It is not difficult to find such formulae: simply start at the root of the syntax tree of φ and proceed towards the leaves in the direction of the child that contains more than half the leaves of φ (while there is one). The vertex, in which the procedure stops, corresponds to the operator op in $\beta = \mathrm{op}(\beta_1, \beta_2)$ above. Note that $\|\alpha(x)\| = m - \|\beta\| + 1 < m - \frac{m}{2} + 1$, i.e., $\|\alpha(x)\| \leq \frac{m}{2}$.

First assume that x does not occur under a $\Diamond$-operator in $\alpha(x)$. Then

$$\varphi = \alpha(\beta) \equiv_{\mathcal{S}_{\mathrm{S5}}} \big(\alpha(\bot) \wedge \neg\beta\big) \vee \big(\alpha(\top) \wedge \beta\big) \tag{5}$$

since, in both formulae, β is interpreted in the initial-world (hence the formulae are even equivalent over $\mathcal{S}_{\mathrm{K}}$). By induction hypothesis, there exist formulae $\alpha'(x), \beta_1'$, and β_2' with

- $\alpha'(x) \equiv_{\mathcal{S}_{\mathrm{S5}}} \alpha(x)$ and $d(\alpha'(x)) \leq 8 \cdot (1 + \log \frac{m}{2}) = 8 \cdot \log_2 m$ as well as
- $\beta_i' \equiv_{\mathcal{S}_{\mathrm{S5}}} \beta_i$ and $d(\beta_i') \leq 8 \cdot (1 + \log \frac{m}{2}) = 8 \cdot \log_2 m$ for $i \in \{1, 2\}$.

Set $\beta' = \mathrm{op}(\beta_1', \beta_2')$. Then β' is equivalent to β over $\mathcal{S}_{\mathrm{S5}}$ and of depth at most $1 + 8 \cdot \log_2 m$. From (5), it follows that $\varphi \equiv_{\mathcal{S}_{\mathrm{S5}}} \big(\alpha'(\bot) \wedge \neg\beta'\big) \vee \big(\alpha'(\top) \wedge \beta'\big) =: \varphi'$. Furthermore, the depth of φ' satisfies

$$\begin{aligned} d(\varphi') &= \max\big\{2 + d\big(\alpha'(x)\big),\ 3 + d\big(\beta'\big)\big\} \\ &= \max\{2 + 8 \cdot \log_2 m,\ 3 + 1 + 8 \cdot \log_2 m\} \leq 4 + 8 \cdot \log_2 m \leq 8 \cdot (1 + \log_2 m), \end{aligned}$$

which completes the first case, where x does not occur under a $\Diamond$ in $\alpha(x)$.

Otherwise, the variable x occurs in the scope of a $\Diamond$-operator. Now, we split $\alpha(x)$ at the last such $\Diamond$, i.e., there exist formulae $\alpha_1(y)$ with only one occurrence of y and $\alpha_2(x)$ with only one occurrence of x that, furthermore, does not lie under a $\Diamond$-operator, such that $\alpha(x) = \alpha_1(\Diamond\alpha_2(x))$. In particular, $\|\alpha_1(y)\|, \|\alpha_2(x)\| \leq \|\alpha(x)\| \leq \frac{m}{2}$. By induction hypothesis, there exist $\alpha_1'(y), \alpha_2'(x), \beta_1'$, and β_2' with

- $\alpha_i'(z) \equiv_{\mathcal{S}_{\mathrm{S5}}} \alpha_i(z)$ and $d(\alpha_i'(z)) \leq 8 \cdot \log_2 m$ for $i \in \{1, 2\}$ as well as
- $\beta_i' \equiv_{\mathcal{S}_{\mathrm{S5}}} \beta_i$ and $d(\beta_i') \leq 8 \cdot \log_2 m$ for $i \in \{1, 2\}$.

As before, let $\beta' = \mathrm{op}(\beta_1', \beta_2')$ with $\beta' \equiv_{\mathcal{S}_{\mathrm{S5}}} \beta$ and $d(\beta') \leq 1 + 8 \cdot \log_2 m$. Now, consider the formulae

$$\psi' = \big(\alpha_2'(\bot) \wedge \neg\beta'\big) \vee \big(\alpha_2'(\top) \wedge \beta'\big) \quad \text{and} \quad \varphi' = \big(\alpha_1'(\bot) \wedge \neg\Diamond\psi'\big) \vee \big(\alpha_1'(\top) \wedge \Diamond\psi'\big)\ .$$

Then $\psi' \equiv_{\mathcal{S}_{\mathrm{S5}}} \alpha_2(\beta)$, since x does not occur under a $\Diamond$ in $\alpha_2(x)$, hence β is interpreted in the initial world in both formulae. But also $\varphi' \equiv_{\mathcal{S}_{\mathrm{S5}}} \big(\alpha_1(\bot) \wedge \neg\Diamond\psi'\big) \vee \big(\alpha_1(\top) \wedge \Diamond\psi'\big) \equiv_{\mathcal{S}_{\mathrm{S5}}} \alpha_1(\Diamond\psi')$ since, whether or not $\Diamond\psi'$ holds in a particular world, does not depend on the choice of the world (see (4)). Hence $\varphi' \equiv_{\mathcal{S}_{\mathrm{S5}}} \alpha_1(\Diamond\psi') \equiv_{\mathcal{S}_{\mathrm{S5}}} \alpha_1(\Diamond\alpha_2(\beta)) = \varphi$. Similar to the first case, one can show

$$\begin{aligned} d(\psi') &= \max\big\{2 + d\big(\alpha_2'(x)\big),\ 3 + d\big(\beta'\big)\big\} \leq 4 + 8 \cdot \log_2 m \quad \text{and} \\ d(\varphi') &= \max\big\{2 + d\big(\alpha_1'(y)\big),\ 4 + d\big(\psi'\big)\big\} \\ &\leq \max\{2 + 8 \cdot \log_2 m,\ 8 + 8 \cdot \log_2 m\} \leq 8 \cdot (1 + \log_2 m)\,, \end{aligned}$$

which completes the second case and hence the inductive proof. ◀

Let $\varphi \in \mathrm{ML}[\leftrightarrow]$. By the previous lemma, there exists an $\mathrm{ML}[\leftrightarrow]$-formula ψ that is equivalent to φ over $\mathcal{S}_{\mathrm{S5}}$ and has depth $d(\psi) \leq 8 \cdot (1 + \log_2 \|\varphi\|)$. Let φ' be obtained from ψ by replacing each sub-formula $\alpha \leftrightarrow \beta$ by $(\alpha \wedge \beta) \vee (\neg\alpha \wedge \neg\beta)$. Then φ' belongs to ML, is equivalent to φ over $\mathcal{S}_{\mathrm{S5}}$, and the depth increases at most by a factor of three, i.e., $d(\varphi') \leq 3 \cdot d(\psi) \leq 3 \cdot 8 \cdot (1 + \log_2 \|\varphi\|) = 24 \cdot (1 + \log_2 \|\varphi\|)$. Since all operators in φ' are at most binary, it follows that

$$|\varphi'| \leq 2^{d(\varphi')+1} \leq 2^{24 \cdot (1 + \log_2 \|\varphi\|)+1} \leq c \cdot \|\varphi\|^{c'} \leq c \cdot |\varphi|^{c'} \quad \text{for some constants } c, c' > 0.$$

Hence we verified the following claim.

▶ **Lemma 16.** ML[$\leftrightarrow$] *has polynomial translations wrt.* $\mathcal{S}_{\mathrm{S5}}$ *in* ML.

Therefore, the modal logic S5 has only one "succinctness class".

▶ **Theorem 17.** *Let* F *be a finite set of Boolean functions. Then* ML[F] *has polynomial translations wrt.* $\mathcal{S}_{\mathrm{S5}}$ *in* ML.

Proof. Since $\leftrightarrow$ is non-locally-monotone, it follows by Theorem 11 that ML[F] has polynomial translations wrt. $\mathcal{S}_{\mathrm{S5}}$ in ML[$\leftrightarrow$] and therefore in ML by Lemma 16 above. ◀

4 Conclusion

This paper considers the question whether or not the use of additional Boolean functions allows for more succinct formulae. For many logics, elimination of locally monotone functions is possible with polynomial size increase; for arbitrary functions, this holds if we allow bi-implication to appear in the resulting formula. Regarding propositional logic, it is known that also bi-implication can be eliminated with polynomial size increase. The same applies for modal logic if we restrict the class of Kripke structures to equivalence relations. When considering all reflexive Kripke structures however, this is no longer the case – bi-implication cannot be eliminated in modal logic without introducing an exponential size increase when considering a class of Kripke structures that contains all reflexive structures. It remains open, where exactly the change from polynomial to exponential size increase occurs, e.g., whether bi-implication can be eliminated with polynomial size increase when considering all reflexive and symmetric or all reflexive and transitive Kripke structures.

References

1. M. Adler and N. Immerman. An *n!* lower bound on formula size. *ACM Trans. Comput. Log.*, 4(3):296–314, 2003.
2. K. Etessami, M. Vardi, and T. Wilke. First-order logic with two variables and unary temporal logic. *Inf. Comput.*, 179(2):279–295, 2002.
3. T. French, W. van der Hoek, P. Iliev, and B.P. Kooi. On the succinctness of some modal logics. *Artif. Intell.*, 197:56–85, 2013.
4. S.B. Gashkov and I.S. Sergeev. О значении работ В. М. Храпченко. Прикладная дискретная математика, 2020(48):109–124, 2020. doi:10.17223/20710410/48/10.
5. M. Grohe and N. Schweikardt. The succinctness of first-order logic on linear orders. *Log. Methods Comput. Sci.*, 1(1), 2005.
6. L. Hella and M. Vilander. Formula size games for modal logic and μ-calculus. *J. Log. Comput.*, 29(8):1311–1344, 2019.
7. S. Jukna. *Boolean Function Complexity - Advances and Frontiers*, volume 27 of *Algorithms and combinatorics*. Springer, 2012.
8. C Lutz. Complexity and succinctness of public announcement logic. In *AAMAS'06*, pages 137–143. ACM, 2006.
9. C. Lutz, U. Sattler, and F. Wolter. Modal logic and the two-variable fragment. In *CSL'01*, Lecture Notes in Computer Science vol. 2142, pages 247–261. Springer, 2001.
10. N. Markey. Temporal logic with past is exponentially more succinct. *Bull. EATCS*, 79:122–128, 2003.
11. V.R. Pratt. The effect of basis on size of Boolean expressions. In *16th Annual Symposium on Foundations of Computer Science*, pages 119–121, 1975. doi:10.1109/SFCS.1975.29.
12. J.E. Savage. *The Complexity of Computing*. Wiley, 1976.
13. P.M. Spira. On time-hardware tradeoffs for Boolean functions. In *Proc. of 4th Hawaii Intern. Symp. on System Sciences*, pages 525–527, 1971.

14 H. van Ditmarsch, Jie Fan, W. van der Hoek, and P. Iliev. Some exponential lower bounds on formula-size in modal logic. In Rajeev Goré, Barteld P. Kooi, and Agi Kurucz, editors, *Advances in Modal Logic 2014*, pages 139–157. College Publications, 2014.

15 T. Wilke. Ctl^{+} is exponentially more succinct than CTL. In *FSTTCS'99*, Lecture Notes in Computer Science vol. 1738, pages 110–121. Springer, 1999.

16 S.V. Yablonskii and V.P. Kozyrev. Математические бопросы кибернетики. Информационные материалы, Академия наук СССР научный совет по комплексной проблеме "кибернетика", 19a:3–15, 1968.