Complexity through Translations for Modal Logic with Recursion*

Luca Aceto[1,2], Antonis Achilleos[1], Elli Anastasiadi[1], Adrian Francalanza[3], and Anna Ingólfsdóttir[1]

[1] School of Computer Science, Reykjavik University, Reykjavik, Iceland
[2] Gran Sasso Science Institute, L'Aquila, Italy
[3] Dept. of Computer Science, ICT, University of Malta, Msida, Malta

Abstract

This paper studies the complexity of classical modal logics and of their extension with fixed-point operators, using translations to transfer results across logics. In particular, we show several complexity results for multi-agent logics via translations to and from the μ-calculus and modal logic, which allow us to transfer known upper and lower bounds. We also use these translations to introduce a terminating tableau system for the logics we study, based on Kozen's tableau for the μ-calculus, and the one of Fitting and Massacci for modal logic.

1 Introduction

We introduce a family of multi-modal logics with fixed-point operators that are interpreted on restricted classes of Kripke models. One can consider these logics as extensions of the usual multi-agent logics of knowledge and belief [14] by adding recursion to their syntax or of the μ-calculus [21] by interpreting formulas on different classes of frames and thus giving an epistemic interpretation to the modalities. We define *translations* between these logics, and we demonstrate how one can rely on these translations to prove finite-model theorems, complexity bounds, and tableau termination for each logic in the family.

Modal logic comes in several variations [5]. Some of these, such as multi-modal logics of knowledge and belief [14], are of particular interest to Epistemology and other application areas. Semantically, the classical modal logics used in epistemic (but also other) contexts result from imposing certain restrictions on their models. On the other hand, the modal μ-calculus [21] can be seen as an extension of the smallest normal modal logic **K** with greatest and least fixed-point operators, νX and μX respectively. We explore the situation where one allows both recursion (*i.e.* fixed-point) operators in a multi-modal language and imposes restrictions on the semantic models.

We are interested in the complexity of satisfiability for the resulting logics. Satisfiability for the μ-calculus is known to be EXP-complete [21], while for the modal logics between **K** and **S5** the problem is PSPACE-complete or NP-complete, depending on whether they have Negative Introspection [18, 22]. In the multi-modal case, satisfiability for those modal logics becomes PSPACE-complete, and is EXP-complete with the addition of a common knowledge operator [17].

*This work has been funded by the projects "Open Problems in the Equational Logic of Processes (OPEL)" (grant no. 196050), "Epistemic Logic for Distributed Runtime Monitoring" (grant no. 184940), "Mode(l)s of Verification and Monitorability" (MoVeMent) (grant no 217987) of the Icelandic Research Fund, and the project "Runtime and Equational Verification of Concurrent Programs" (ReVoCoP) (grant 222021) of the Reykjavik University Research Fund.

There is plenty of relevant work on the μ-calculus on restricted frames, mainly in its single-agent form. Alberucci and Facchini examine the alternation hierarchy of the μ-calculus over reflexive, symmetric, and transitive frames in [2]. D'Agostino and Lenzi have studied the μ-calculus on different classes of frames in great detail. In [9], they reduce the μ-calculus over finite transitive frames to first-order logic. In [10], they prove that $\mathbf{S5}^\mu$-satisfiability is NP-complete, and that the two-agent version of $\mathbf{S5}^\mu$ does not have the finite model property. In [11], they consider finite symmetric frames, and they prove that $\mathbf{B}^\mu$-satisfiability is in 2EXP, and EXP-hard. They also examine planar frames in [12], where they show that the alternation hierarchy of the μ-calculus over planar frames is infinite.

Our primary method of proving complexity results is through translations to and from the multi-modal μ-calculus. We show that we can use surprisingly simple translations from modal logics without recursion to the base modal logic $\mathbf{K}_n$, reproving the PSPACE upper bound for these logics (Theorem 8 and Corollary 9). These translations and our constructions to prove their correctness do not generally transfer to the corresponding logics with recursion. We present translations from specific logics to the μ-calculus and back, and we discuss the remaining open cases. We discover, through the properties of our translations, that several behaviors induced on the transitions do not affect the complexity of the satisfiability problem. As a result, we prove that all logics with axioms among D, T, and 4, and the least-fixed-point fragments of logics that also have B, have their satisfiability in EXP, and a matching lower bound for the logics with axioms from D, T, B (Corollaries 15 and 16). Finally, we present tableaux for the discussed logics, based on the ones by Kozen for the μ-calculus [21], and by Fitting and Massacci for modal logic [16, 23]. We give tableau-termination conditions for every logic with a finite model property (Theorem 19).

The addition of recursive operators to modal logic increases expressiveness. An important example is that of *common knowledge* or *common belief*, which can be expressed with a greatest fixed-point thus: $\nu X.(\varphi \wedge \bigwedge_\alpha [\alpha]X)$. But the combination of epistemic logics and fixed-points can potentially express more interesting epistemic concepts. For instance, the formula $\mu X. \bigvee_\alpha ([\alpha]\varphi \vee [\alpha]X)$, in the context of a belief interpretation, can be thought to claim that there is a rumour of φ. It would be interesting to see what other meaningful sentences of epistemic interest one can express using recursion. Furthermore, the family of logics we consider allows each agent to behave according to a different logic. This flexibility allows one to mix different interpretations of modalities, such as a temporal interpretation for one agent and an epistemic interpretation for another. Such logics can even resemble hyper-logics [8] if a set of agents represents different streams, and combinations of epistemic and temporal or hyper-logics have recently been used to express safety and privacy properties of systems [7].

The paper is organized as follows. Section 2 gives the necessary background and an overview of current results. Section 3 defines a class of translations that provide us with several upper and lower bounds, and identifies conditions under which they can be composed. In Section 4 we finally give tableaux for our multi-modal logics with recursion. We conclude in Section 5 with a set of open questions and directions. Omitted proofs can be found in the appendix.

2 Definitions and Background

This section introduces the logics that we study and the necessary background on the complexity of modal logic and the μ-calculus.

2.1 The Multi-Modal Logics with Recursion

We start by defining the syntax of the logics.

Definition 1. *We consider formulas constructed from the following grammar:*

$$\begin{array}{lllllll} \varphi, \psi \in L ::= p & \mid \neg p & \mid \texttt{tt} & \mid \texttt{ff} & \mid X & \mid \varphi \wedge \psi & \mid \varphi \vee \psi \\ \quad\quad \mid \langle \alpha \rangle \varphi & \mid [\alpha]\varphi & \mid \mu X.\varphi & \mid \nu X.\varphi, & & & \end{array}$$

where X comes from a countable set of logical (or fixed-point) variables, LVar*, α from a finite set of agents,* Ag*, and p from a finite set of propositional variables,* PVar*. When* $\text{Ag} = \{\alpha\}$*, $\Box\varphi$ stands for $[\alpha]\varphi$, and $\Diamond\varphi$ for $\langle\alpha\rangle\varphi$. We also write $[A]\varphi$ to mean $\bigwedge_{\alpha\in A}[\alpha]\varphi$ and $\langle A\rangle\varphi$ for $\bigvee_{\alpha\in A}\langle\alpha\rangle\varphi$.*

A formula is closed when every occurrence of a variable X is in the scope of recursive operator νX or μX. Henceforth we consider only closed formulas, unless we specify otherwise.

Moreover, for recursion-free closed formulas we associate the notion of *modal depth*, which is the nesting depth of the modal operators[1]. The modal depth of φ is defined inductively as:

- $md(p) = md(\neg p) = md(\texttt{tt}) = md(\texttt{ff}) = 0$, where $p \in \text{PVar}$,
- $md(\varphi \vee \psi) = md(\varphi \wedge \psi) = max(md(\varphi), md(\psi))$, and
- $md([a]\varphi) = md(\langle a\rangle\varphi) = 1 + md(\varphi)$, where $a \in \text{Ag}$.

We assume that in formulas, each recursion variable X appears in a unique fixed-point formula $\mathsf{fx}(X)$, which is either of the form $\mu X.\varphi$ or $\nu X.\varphi$. If $\mathsf{fx}(X)$ is a least-fixed-point (*resp.* greatest-fixed-point) formula, then X is called a least-fixed-point (*resp.* greatest-fixed-point) variable. We can define a partial order on fixed-point variables, such that $X \leq Y$ iff $\mathsf{fx}(X)$ is a subformula of $\mathsf{fx}(Y)$, and $X < Y$ when $X \leq Y$ and $X \neq Y$. If X is $\leq$-minimal among the free variables of φ, then we define the *closure* of φ to be $cl(\varphi) = cl(\varphi[\mathsf{fx}(X)/X])$, where $\varphi[\psi/X]$ is the usual substitution operation, and if φ is closed, then $cl(\varphi) = \varphi$.

We define $\mathsf{sub}(\varphi)$ as the set of subformulas of φ, and $|\varphi| = |\mathsf{sub}(\varphi)|$ is bounded by the length of φ as a string of symbols. Negation, $\neg\varphi$, and implication, $\varphi \to \psi$, can be defined in the usual way. Then, we define $\overline{\mathsf{sub}}(\varphi) = \mathsf{sub}(\varphi) \cup \{\neg\psi \in L \mid \psi \in \mathsf{sub}(\varphi)\}$.

Semantics We interpret formulas on the states of a *Kripke model.* A Kripke model, or simply model, is a quadruple $M = (W, R, V)$ where W is a nonempty set of states, $R \subseteq W \times \text{Ag} \times W$ is a transition relation, and $V : W \to 2^{\text{PVar}}$ determines on which states a propositional variable is true. (W, R) is called a *frame.* We usually write $(u, v) \in R_\alpha$ or $uR_\alpha v$ instead of $(u, \alpha, v) \in R$, or uRv, when Ag is a singleton $\{\alpha\}$.

Formulas are evaluated in the context of an *environment* $\rho : \text{LVar} \to 2^W$, which gives values to the logical variables. For an environment ρ, variable X, and set $S \subseteq W$, we write $\rho[X \mapsto S]$ for the environment that maps X to S and all $Y \neq X$ to $\rho(Y)$. The semantics for our formulas is given through a function $[\![-]\!]_{\mathcal{M}}$, defined in Table 1. The semantics of $\neg\varphi$ are constructed as usual, where $[\![\neg X, \rho]\!]_{\mathcal{M}} = W \backslash \rho(X)$.

[1]The modal depth of recursive formulas can be either zero, or infinite. However, this is not relevant for the spectrum of this work.

$$[\![\mathtt{tt},\rho]\!] = W, \qquad [\![\mathtt{ff},\rho]\!] = \emptyset, \qquad [\![p,\rho]\!] = \{s \mid p \in V(s)\}, \qquad [\![\neg p,\rho]\!] = W \backslash [\![p,\rho]\!],$$
$$[\![[\alpha]\varphi,\rho]\!] = \left\{s \;\middle|\; \forall t.\ sR_\alpha t \text{ implies } t \in [\![\varphi,\rho]\!]\right\}, \qquad [\![\varphi_1 \wedge \varphi_2,\rho]\!] = [\![\varphi_1,\rho]\!] \cap [\![\varphi_2,\rho]\!],$$
$$[\![\langle\alpha\rangle\varphi,\rho]\!] = \left\{s \;\middle|\; \exists t.\ sR_\alpha t \text{ and } t \in [\![\varphi,\rho]\!]\right\}, \qquad [\![\varphi_1 \vee \varphi_2,\rho]\!] = [\![\varphi_1,\rho]\!] \cup [\![\varphi_2,\rho]\!],$$
$$[\![\mu X.\varphi,\rho]\!] = \bigcap\left\{S \;\middle|\; S \supseteq [\![\varphi,\rho[X \mapsto S]]\!]\right\}, \qquad [\![X,\rho]\!] = \rho(X),$$
$$[\![\nu X.\varphi,\rho]\!] = \bigcup\left\{S \;\middle|\; S \subseteq [\![\varphi,\rho[X \mapsto S]]\!]\right\}.$$

Table 1: Semantics of modal formulas on a model $\mathcal{M} = (W, R, V)$. We omit $\mathcal{M}$ from the notation.

We sometimes use $\mathcal{M}, s \models_\rho \varphi$ for $s \in [\![\varphi,\rho]\!]_\mathcal{M}$, and as the environment has no effect on the semantics of a closed formula φ, we often drop it from the notation and write $\mathcal{M}, s \models \varphi$ or $s \in [\![\varphi]\!]_\mathcal{M}$. If $\mathcal{M}, s \models \varphi$, we say that φ is true, or satisfied, in s. When the particular model does not matter, or is clear from the context, we may omit it.

Depending on how we further restrict our syntax and the model, we can describe several logics. Without further restrictions, the resulting logic is the μ-calculus [21]. The max-fragment (resp. min-fragment) of the μ-calculus is the fragment that only allows the νX (resp. the μX) recursive operator. If $|\textsc{Ag}| = k$ and we allow no recursive operators (or recursion variables), then we have the basic modal logic $\mathbf{K}_k$ (or $\mathbf{K}$, if $k = 1$), and further restrictions on the frames can result in a wide variety of modal logics (see [6]). We give names to the following frame conditions, or frame constraints, for the case where $\textsc{Ag} = \{\alpha\}$. These conditions correspond to the usual axioms for normal modal logics — see [5, 6, 14], which we will revisit in Section 3.

- D**:** R is serial: $\forall s.\exists t.sRt$;
- T**:** R is reflexive: $\forall s.sRs$;
- B**:** R is symmetric: $\forall s, t.(sRt \Rightarrow tRs)$;
- 4**:** R is transitive: $\forall s, t, r.(sRtRr \Rightarrow sRr)$;
- 5**:** R is euclidean: $\forall s, t, r.$ if sRt and sRr, then tRr.

We consider modal logics that are interpreted over models that satisfy a combination of these constraints for each agent. D, which we call Consistency, is a special case of T, called Factivity. Constraint 4 is Positive Introspection and 5 is called Negative Introspection.[2] Given a logic $\mathbf{L}$ and constraint c, $\mathbf{L} + c$ is the logic that is interpreted over all models with frames that satisfy all the constraints of $\mathbf{L}$ and c. The name of a single-agent logic is a combination of the constraints that apply to its frames, including K, if the constraints are among 4 and 5. Therefore, logic $\mathbf{D}$ is $\mathbf{K} + D$, $\mathbf{T}$ is $\mathbf{K} + T$, $\mathbf{B}$ is $\mathbf{K} + B$, $\mathbf{K4} = \mathbf{K} + 4$, $\mathbf{D4} = \mathbf{K} + D + 4 = \mathbf{D} + 4$, and so on. We use the special names $\mathbf{S4}$ for $\mathbf{T4}$ and $\mathbf{S5}$ for $\mathbf{T45}$. We define a (multi-agent) logic $\mathbf{L}$ on $\textsc{Ag}$ as a map from agents to single-agent logics. $\mathbf{L}$ is interpreted on Kripke models of the form (W, R, V), where for every $\alpha \in \textsc{Ag}$, (W, R_α) is a frame for $\mathbf{L}(\alpha)$.

For a logic $\mathbf{L}$, $\mathbf{L}^\mu$ is the logic that results from $\mathbf{L}$ after we allow recursive operators in the syntax — in case they were not allowed in $\mathbf{L}$. Furthermore, if for every $\alpha \in \textsc{Ag}$, $\mathbf{L}(\alpha)$ is the same single-agent logic $\mathbf{L}$, we write $\mathbf{L}$ as $\mathbf{L}_k$, where $|\textsc{Ag}| = k$. Therefore, the μ-calculus is $\mathbf{K}^\mu_k$.

[2] These are names for properties or axioms of a logic. When we refer to these conditions as conditions of a frame or model, we may refer to them with the name of the corresponding relation condition: seriality, reflexivity, symmetry, transitivity, and euclidicity.

From now on, unless we explicitly say otherwise, by a logic, we mean one of the logics we have defined above. We call a formula satisfiable for a logic $\mathbf{L}$, if it is satisfied in some state of a model for $\mathbf{L}$.

Example 1. For a formula φ, we define $Inv(\varphi) = \nu X.(\varphi \wedge [\mathrm{AG}]X)$. $Inv(\varphi)$ asserts that φ is true in *all* reachable states, or, alternatively, it can be read as an assertion that φ is common knowledge. We dually define $Eve(\varphi) = \mu X.(\varphi \vee \langle\mathrm{AG}\rangle X)$, which asserts that φ is true in *some* reachable state.

2.2 Known Results

For logic $\mathbf{L}$, the satisfiability problem for $\mathbf{L}$, or $\mathbf{L}$-satisfiability is the problem that asks, given a formula φ, if φ is satisfiable. Similarly, the model checking problem for $\mathbf{L}$ asks if φ is true at a given state of a given finite model.

Ladner [22] established the classical result of PSPACE-completeness for the satisfiability of **K**, **T**, **D**, **K4**, **D4**, and **S4** and NP-completeness for the satisfiability of **S5**. Halpern and Rêgo later characterized the NP–PSPACE gap for one-action logics by the presence or absence of Negative Introspection [18], resulting in Theorem 1. Later, Rybakov and Shkatov [26] proved the PSPACE-completeness of **B** and **TB**. For formulas with fixed-point operators, D'Agostino and Lenzi in [10] show that satisfiability for single-agent logics with constraint 5 is also NP-complete.

Theorem 1 ([18,22,26]). *If* $\mathbf{L} \in \{\mathbf{K}, \mathbf{T}, \mathbf{D}, \mathbf{B}, \mathbf{TB}, \mathbf{K4}, \mathbf{D4}, \mathbf{S4}\}$, *then* $\mathbf{L}$*-satisfiability is* PSPACE*-complete; and* $\mathbf{L}+5$*-satisfiability and* $(\mathbf{L}+5)^\mu$*-satisfiability is* NP*-complete.*

Theorem 2 ([17]). *If* $k > 1$ *and* $\mathbf{L}$ *has a combination of constraints from* $D, T, 4, 5$ *and no recursive operators, then* $\mathbf{L}_k$*-satisfiability is* PSPACE*-complete.*

Remark 1. We note that Halpern and Moses in [17] only prove these bounds for the cases of $\mathbf{K}_k, \mathbf{T}_k, \mathbf{S4}_k, \mathbf{KD45}_k$, and $\mathbf{S5}_k$; and D'Agostino and Lenzi in [10] only prove the NP-completeness of satisfiability for $\mathbf{S5}^\mu$. However, it is not hard to see that their respective methods also work for the rest of the logics of Theorems 1 and 2. ■

Theorem 3 ([21]). *The satisfiability problem for the* μ*-calculus is* EXP*-complete.*

Theorem 4 ([13]). *The model checking problem for the* μ*-calculus is in* NP ∩ coNP.[3]

Finally we have the following initial known results about the complexity of satisfiability, when we have recursive operators. Theorems 5 and 6 have already been observed in [1].

Theorem 5. *The satisfiability problem for the min- and max-fragments of the* μ*-calculus is* EXP*-complete, even when* $|\mathrm{AG}| = 1$.

Proof sketch. It is known that satisfiability for the min- and max-fragments of the μ-calculus (on one or more action) is EXP-complete. It is in EXP due to Theorem 3, and these fragments suffice [25] to describe the PDL formula that is constructed by the reduction used in [15] to prove EXP-hardness for PDL. Therefore, that reduction can be adjusted to prove that satisfiability for the min- and max-fragments of the μ-calculus is EXP-complete. □

It is not hard to express in logics with both frame constraints and recursion operators that formula φ is common knowledge, with formula $\nu X.\varphi \wedge [\mathrm{AG}]X$. Since validity for $\mathbf{L}_k$ with

[3] In fact, the problem is known to be in UP ∩ coUP [20].

common knowledge (and without recursive operators) and $k > 1$ is EXP-complete [17][4], $\mathbf{L}_k^\mu$ is EXP-hard.

Proposition 6. *Satisfiability for $\mathbf{L}_k^\mu$, where $k > 1$, is EXP-hard.*

3 Complexity through Translations

In this section, we examine **L**-satisfiability. We use formula translations to reduce the satisfiability of one logic to the satisfiability of another. We investigate the properties of these translations and how they compose with each other, and we achieve complexity bounds for several logics.

In the context of this paper, a formula translation from logic $\mathbf{L}_1$ to logic $\mathbf{L}_2$ is a mapping f on formulas such that each formula φ is $\mathbf{L}_1$ -satisfiable if and only if $f(\varphi)$ is $\mathbf{L}_2$ -satisfiable. We only consider translations that can be computed in polynomial time, and therefore, our translations are polynomial-time reductions, transfering complexity bounds between logics.

According to Theorem 3, $\mathbf{K}_k^\mu$-satisfiability is EXP-complete, and therefore for each logic **L**, we aim to connect $\mathbf{K}_k^\mu$ and **L** via a sequence of translations in either direction, to prove complexity bounds for **L**-satisfiability.

3.1 Translating Towards $\mathbf{K}_k$

We begin by presenting translations from logics with more to logics with fewer frame conditions. To this end, we study how taking the closure of a frame under one condition affects any other frame conditions.

3.1.1 Composing Frame Conditions

We now discuss how the conditions for frames affect each other. For example, to construct a transitive frame, one can take the transitive closure of a possibly non-transitive frame. The resulting frame will satisfy condition 4. As we see, taking the closure of a frame under condition x may affect whether that frame maintains condition y, depending on x and y. In the following we observe that one can apply the frame closures in certain orders that preserve the properties one aquires with each application.

Let $F = (W, R)$ be a frame, $\alpha \in A \subseteq$ AG, and x a frame restriction among $T, B, 4, 5$. Then, $\overline{R_\alpha}^x$ is the closure of R_α under x, $\overline{R}^{x,A}$ is defined by $\overline{R}_\beta^{x,A} = \overline{R_\beta}^x$, if $\beta \in A$, and $\overline{R}_\beta^{x,A} = R_\beta$, otherwise. Then, $\overline{F}^{x,A} = (W, \overline{R}^{x,A})$. We make the following observation.

Lemma 7. *Let x be a frame restriction among $D, T, B, 4, 5$, and y a frame restriction among $T, B, 4, 5$, such that $(x, y) \neq (4, B), (5, T), (5, B)$. Then, for every frame F that satisfies x, $\overline{F}^y$ also satisfies x.*

According to Lemma 7, frame conditions are preserved as seen in Figure 1. In Figure 1, an arrow from x to y indicates that property x is preserved though the closure of a frame under y. Dotted red arrows indicate one-way arrows. For convenience, we define $\overline{F}^D = (W, \overline{R}^D)$, where $\overline{R}^D = R \cup \{(a, a) \in W^2 \mid \nexists\, (a, b) \in R\}$.

[4] Similarly to Remark 1, [17] does not explicitly cover all these cases, but the techniques can be adjusted.

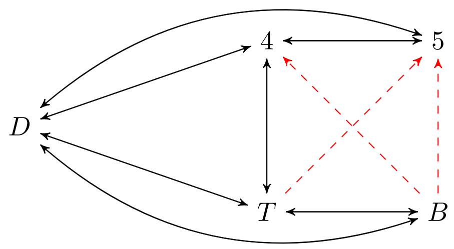

Figure 1: The frame property hierarchy

Remark 2. We note that, in general, not all frame conditions are preserved through all closures under another condition. For example, the accessibility relation $\{(a, b), (b, b)\}$ is euclidean, but its reflexive closure $\{(a, b), (b, b), (a, a)\}$ is not.

There is at least one linear ordering of the frame conditions $D, T, B, 4, 5$, such that all preceding conditions are preserved by closures under the following conditions. We call such an order a closure-preserving order. We use the linear order $D, T, B, 4, 5$ in the rest of the paper.

3.1.2 Modal Logics

We start with translations that map logics without recursive operators to logics with fewer constraints. As mentioned in Subsection 2.2, all of the logics $\mathbf{L} \in \{\mathbf{K}, \mathbf{T}, \mathbf{D}, \mathbf{K4}, \mathbf{D4}, \mathbf{S4}\}$ and $\mathbf{L} + 5$ with one agent have known completeness results, and the complexity of modal logic is well-studied for multi-agent modal logics as well. The missing cases are very few and concern the combination of frame conditions (other than 5) as well as the multi-agent case. However we take this opportunity to present an intuitive introduction to our general translation method. In fact, the translations that we use for logics without recursion are surprisingly straightforward. Each frame condition that we introduced in Section 2 is associated with an axiom for modal logic, such that whenever a model has the condition, every substitution instance of the axiom is satisfied in all worlds of the model (see [5, 6, 14]). We give for each frame condition x and agent α, the axiom ax^x_α:

ax^D_α: $\langle\alpha\rangle\mathtt{tt}$ ax^B_α: $\langle\alpha\rangle[\alpha]p \rightarrow p$ ax^5_α: $\langle\alpha\rangle[\alpha]p \rightarrow [\alpha]p$

ax^T_α: $[\alpha]p \rightarrow p$ ax^4_α: $[\alpha]p \rightarrow [\alpha][\alpha]p$

For each formula φ and $d \geq 0$, let $Inv_d(\varphi) = \bigwedge_{i \leq d} [\mathrm{AG}]^i \varphi$. Our first translations are straightforwardly defined from the above axioms.

Translation 1 (One-step Translation)**.** Let $A \subseteq \mathrm{AG}$ and let x be one of the frame conditions. For every formula φ, let $d = md(\varphi)$ if $x \neq 4$, and $d = md(\varphi)|\varphi|$, if $x = 4$. We define:

$$\mathsf{F}^x_A(\varphi) = \varphi \wedge Inv_d\left(\bigwedge_{\substack{\psi \in \overline{\mathsf{sub}}(\varphi) \\ \alpha \in A}} \mathsf{ax}^x_\alpha[\psi/p] \right).$$

Theorem 8. *Let $A \subseteq \mathrm{AG}$, x be one of the frame conditions, and let $\mathbf{L}_1, \mathbf{L}_2$ be logics without recursion operators, such that $\mathbf{L}_1(\alpha) = \mathbf{L}_2(\alpha) + x$ when $\alpha \in A$, and $\mathbf{L}_2(\alpha)$ otherwise, and $\mathbf{L}_2(\alpha)$*

only includes frame conditions that precede x in the fixed order of frame conditions. Then, φ is $\mathbf{L}_1$-satisfiable if and only if $\mathsf{F}^x_A(\varphi)$ is $\mathbf{L}_2$-satisfiable.

We present here a short proof sketch of this theorem. The full proof of Theorem 8 can be found in Appendix A

Proof sketch. The proof of the "only if" direction is straightforward, as for any agent α with frame condition x, ax^x_α is valid in $\mathbf{L}_1$. Thus the translation holds on any $\mathbf{L}_1$-model that satisfies φ.

The other, and more involved direction requires the construction of an $\mathbf{L}_1$-model for φ from an $\mathbf{L}_2$-model for $\mathsf{F}^x_A(\varphi)$. We make use of the observation that no modal logic formula can describe a model at depth more than the constant d. Therefore, we use the unfolding of the $\mathbf{L}_2$-model, to keep track of the path that one takes to reach a certain state, and that path's length. We then carefully reapply the necessary closures on the accessibility relations and we use induction on its subformulas, to prove that φ is true in the constructed model. We do this as a separate case for each axiom. It is worth noting that we pay special care for the case of $x = 4$ to account for the fact that the translation does not allow the modal depth of the relevant subformulas to decrease with each transition during the induction; and that we needed to include the negations of subformulas of φ in the conjunction of Translation 1, only for the case of $x = 5$. □

Corollary 9. *The satisfiability problem for every logic without fixed-point operators is in* PSPACE*.*

3.1.3 Modal Logics with Recursion

In the remainder of this section we will modify our translations and proof technique, in order to lift our results to logics with fixed-point operators. It is not clear whether the translations of Subsection 3.1.2 can be extended straightforwardly in the case of logics with recursion, by using unbounded invariance Inv, instead of the bounded Inv_d.

Example 2. Let $\varphi_f = \mu X.\Box X$, which requires all paths in the model to be finite, and thus it is not satisfiable in reflexive frames. In Subsection 3.1.2, to translate formulas from reflexive models, we did not need to add the negations of subformulas as conjuncts. In this case, such a translation would give

$$\varphi_t := \varphi_f \wedge Inv((\Box\varphi_f \rightarrow \varphi_f) \wedge (\Box\Box\varphi_f \rightarrow \Box\varphi_f)).$$

Indeed, on reflexive frames, the formulas $\Box\varphi_f \rightarrow \varphi_f$ and $\Box\Box\varphi_f \rightarrow \Box\varphi_f$ are valid, and therefore φ_t is equivalent to φ_f, which is $\mathbf{K}$-satisfiable. This was not an issue in Subsection 3.1.2, as the finiteness of the paths in a model cannot be expressed without recursion.

One would then naturally wonder whether conjoining over $\overline{\mathsf{sub}}(\varphi_f)$ in the translation would make a difference. The answer is affirmative, as the tranlation

$$\varphi_f \wedge Inv\left(\bigwedge_{\psi\in\overline{\mathsf{sub}}(\varphi_f)} \Box\psi \rightarrow \psi\right)$$

would then yield a formula that is not satisfiable. However, our constructions would not work to prove that such a translation preserves satisfiability. For example, consider $\mu X.\Box(p \rightarrow (r \wedge (q \rightarrow X)))$, whose translation is satisfied on a pointed model that satisfies at the same time p and q. We invite the reader to verify the details.

The only case where the approach that we used for the logics without recursion can be applied is for the case of seriality (condition D), as $Inv(\langle\alpha\rangle\mathtt{tt})$ directly ensures the seriality of a model.

Translation 2.

$$\mathsf{F}_A^{D^\mu}(\varphi) = \varphi \wedge Inv\big(\bigwedge_{\alpha\in A}\langle\alpha\rangle\mathtt{tt}\big).$$

Theorem 10. *Let $A \subseteq$ Ag and $|\text{Ag}| = k$, and let $\mathbf{L}$ be a logic, such that $\mathbf{L}(\alpha) = \mathbf{D}$ when $\alpha \in A$, and $\mathbf{K}$ otherwise. Then, φ is $\mathbf{L}$-satisfiable if and only if $\mathsf{F}_A^{D^\mu}(\varphi)$ is $\mathbf{K}_k^\mu$-satisfiable.*

For the cases of reflexivity and transitivity, our simple translations substitute the modal subformulas of a formula to implicitly enforce the corresponding condition.

Translation 3. The operation $\mathsf{F}_A^{T^\mu}(-)$ is defined to be such that

- $\mathsf{F}_A^{T^\mu}([\alpha]\varphi) = [\alpha]\mathsf{F}_A^{T^\mu}(\varphi) \wedge \mathsf{F}_A^{T^\mu}(\varphi)$;
- $\mathsf{F}_A^{T^\mu}(\langle\alpha\rangle\varphi) = \langle\alpha\rangle\mathsf{F}_A^{T^\mu}(\varphi) \vee \mathsf{F}_A^{T^\mu}(\varphi)$;
- and it commutes with all other operations.

Theorem 11. *Let $\emptyset \neq A \subseteq$ Ag, and let $\mathbf{L}_1, \mathbf{L}_2$ be logics, such that $\mathbf{L}_1(\alpha) = \mathbf{L}_2(\alpha) + T$ when $\alpha \in A$, and $\mathbf{L}_2(\alpha)$ otherwise, and $\mathbf{L}_2(\alpha)$ at most includes frame condition D. Then, φ is $\mathbf{L}_1$-satisfiable if and only if $\mathsf{F}_A^{T^\mu}(\varphi)$ is $\mathbf{L}_2$-satisfiable.*

Proof sketch. The "only if" direction is proven by taking the appropriate reflexive closure and showing by induction that the subformulas of φ are preserved. The full proof can be found in Appendix A.2 □

Translation 4. The operation $\mathsf{F}_A^{4^\mu}(-)$ is defined to be such that

- $\mathsf{F}_A^{4^\mu}([\alpha]\psi) = Inv([\alpha](\mathsf{F}_A^{4^\mu}(\psi))$,
- $\mathsf{F}_A^{4^\mu}(\langle\alpha\rangle\psi) = Eve(\langle\alpha\rangle(\mathsf{F}_A^{4^\mu}(\psi))$,
- $\mathsf{F}_A^{4^\mu}(-)$ commutes with all other operations.

Theorem 12. *Let $\emptyset \neq A \subseteq$ Ag, and let $\mathbf{L}_1, \mathbf{L}_2$ be logics, such that $\mathbf{L}_1(\alpha) = \mathbf{L}_2(\alpha) + 4$ when $\alpha \in A$, and $\mathbf{L}_2(\alpha)$ otherwise, and $\mathbf{L}_2(\alpha)$ at most includes frame conditions D, T, B. Then, φ is $\mathbf{L}_1$-satisfiable if and only if $\mathsf{F}_A^{4^\mu}(\varphi)$ is $\mathbf{L}_2$-satisfiable.*

Proof. If $\mathsf{F}_A^{4^\mu}(\varphi)$ is satisfied in a model $M = (W, R, V)$, let $M' = (W, R^+, V)$, where R_α^+ is the transitive closure of R_α, if $\alpha \in A$, and $R_\alpha^+ = R_\alpha$, otherwise. It is now not hard to use induction on ψ to show that for every (possibly open) subformula ψ of φ, for every environment ρ, $[\![\mathsf{F}_A^{4^\mu}(\psi), \rho]\!]_{\mathcal{M}} = [\![\psi, \rho]\!]_{\mathcal{M}'}$. The other direction is more straightforward. □

In order to produce a similar translation for symmetric frames, we needed to use a more intricate type of construction. Moreover, we only prove the correctness of the following translation for formulas without least-fixed-point operators.

Translation 5. The operation $\mathsf{F}_A^{\mathbf{B}^\mu}(-)$ is defined as

$$\mathsf{F}_A^{\mathbf{B}^\mu}(\varphi) = \varphi \wedge Inv\big([\alpha]\langle\alpha\rangle p \wedge \bigwedge_{\psi\in\overline{\mathsf{sub}}(\varphi)}(\psi \rightarrow [\alpha][\alpha](p \rightarrow \psi))\big),$$

where p is a new propositional variable, not occurring in φ.

Theorem 13. *Let $\emptyset \neq A \subseteq \textsc{Ag}$, and let $\mathbf{L}_1, \mathbf{L}_2$ be logics, such that $\mathbf{L}_1(\alpha) = \mathbf{L}_2(\alpha) + \mathbf{B}$ when $\alpha \in A$, and $\mathbf{L}_2(\alpha)$ otherwise, and $\mathbf{L}_2(\alpha)$ at most includes frame conditions D, T. Then, a formula φ that has no μX operators is $\mathbf{L}_1$-satisfiable if and only if $\mathsf{F}_A^{\mathbf{B}^\mu}(\varphi)$ is $\mathbf{L}_2$-satisfiable.*

Proof. The proof can be found in Appendix A.3 ☐

Remark 3. A translation for euclidean frames and for the full syntax on symmetric frames would need different approaches. D'Agostino and Lenzi show in [10] that $\mathbf{S5}_2^\mu$ does not have a finite model property, and their result can be easily extended to any logic $\mathbf{L}$ with fixed-point operators, where there are at least two distinct agents α, β, such that $\mathbf{L}(\alpha)$ and $\mathbf{L}(\beta)$ have constraint B or 5. Therefore, as our constructions for the translations to $\mathbf{K}_k^\mu$ guarantee the finite model property to the corresponding logics, they do not apply to multimodal logics with B or 5.

3.2 Embedding $\mathbf{K}_n^\mu$

In this subsection, we present translations from logics with fewer frame conditions to ones with more conditions. This will allow us to prove EXP-completeness in the following subsection. Let p, q be distinguished propositional variables that do not appear in our formulas. We let $\vec{p}$ range over p, $\neg p$, $p \wedge q$, $p \wedge \neg q$, and $\neg p \wedge q$.

Definition 2 (function *next*)**.** *$next(p \wedge q) = p \wedge \neg q$, $next(p \wedge \neg q) = \neg p \wedge q$, and $next(\neg p \wedge q) = p \wedge q$; and $next(p) = \neg p$ and $next(\neg p) = p$.*

We use a uniform translation from $\mathbf{K}_k^\mu$ to any logic with a combination of conditions D, T, B.

Translation 6. The operation $\mathsf{F}_A^{\mathbf{K}^\mu}(-)$ on formulas is defined such that:

- $\mathsf{F}_A^{\mathbf{K}^\mu}(\langle\alpha\rangle\psi) = \bigwedge_{\vec{p}}(\vec{p} \rightarrow \langle\alpha\rangle(next(\vec{p}) \wedge \mathsf{F}_A^{\mathbf{K}^\mu}(\psi)))$, if $\alpha \in A$;
- $\mathsf{F}_A^{\mathbf{K}^\mu}([\alpha]\psi) = \bigwedge_{\vec{p}}(\vec{p} \rightarrow [\alpha](next(\vec{p}) \rightarrow \mathsf{F}_A^{\mathbf{K}^\mu}(\psi)))$, if $\alpha \in A$;
- $\mathsf{F}_A^{\mathbf{K}^\mu}(-)$ commutes with all other operations.

We note that there are simpler translations for the cases of logics with only D ot T as a constraint, but the $\mathsf{F}_A^{\mathbf{K}^\mu}(-)$ is uniform for all the logics that we consider in this subsection.

Theorem 14. *Let $\emptyset \neq A \subseteq \textsc{Ag}$, $|\textsc{Ag}| = k$, and let $\mathbf{L}$ be such that $\mathbf{L}(\alpha)$ includes only frame conditions from $\mathbf{D}, \mathbf{T}, \mathbf{B}$ when $\alpha \in A$, and $\mathbf{L}(\alpha) = \mathbf{K}$ otherwise. Then, φ is $\mathbf{K}_k^\mu$-satisfiable if and only if $\mathsf{F}_A^{\mathbf{K}^\mu}(\varphi)$ is $\mathbf{L}$-satisfiable.*

Proof. The proof of Theorem 14 can be found in Appendix A.4. It is worth noting that the "if" direction uses the symmetric closure to construct a new model, while the "only if" direction requires the introduction of new states that behave as each original state in the model. ☐

3.3 Complexity results

We observe that our translations all result in formulas of size at most linear with respect to the original. The exceptions are Translations 1 and 5, which have a quadratic cost.

Corollary 15. *If $\mathbf{L}$ only has frame conditions D, T, then its satisfiability problem is* EXP*-complete; if $\mathbf{L}$ only has frame conditions $D, T, 4$, then its satisfiability problem is in* EXP.

Proof. Immediately from Theorems 10, 11, 12, and 14. □

Corollary 16. *If* $\mathbf{L}$ *only has frame conditions* D, T, B*, then*

1. $\mathbf{L}$*-satisfiability is* EXP*-hard; and*
2. *the restriction of* $\mathbf{L}$*-satisfiability on formulas without* μX *operators is* EXP*-complete.*

Proof. Immediately from Theorems 10, 11, 13, and 14. □

4 Tableaux for $\mathbf{L}_k^\mu$

We give a sound and complete tableau system for logic $\mathbf{L}$. Furthermore, if $\mathbf{L}$ has a finite model property, then we give terminating conditions for its tableau. The system that we give in this section is based on Kozen's tableaux for the μ-calculus [21] and the tableaux of Fitting [16] and Massacci [23] for modal logic. We can use Kozen's finite model theorem [21] to help us ensure the termination of the tableau for some of these logics.

Theorem 17 ([21])**.** *There is a computable* $\kappa : \mathbb{N} \to \mathbb{N}$*, such that every* $\mathbf{K}_k^\mu$*-satisfiable formula* φ *is satisfied in a model with at most* $\kappa(|\varphi|)$ *states.*[5]

Corollary 18. *If* $\mathbf{L}$ *only has frame conditions* $D, T, 4$*, then there is a computable* $\kappa : \mathbb{N} \to \mathbb{N}$*, such that every* $\mathbf{L}$*-satisfiable formula* φ *is satisfied in a model with at most* $\kappa(|\varphi|)$ *states.*

Proof. Immediately, from Theorems 17, 10, 11, and 12, and Lemma 7. □

Remark 4. We note that not all modal logics with recursion have a finite model property – see Remark 3.

Intuitively, a tableau attempts to build a model that satisfies the given formula. When it needs to consider two possible cases, it branches, and thus it may generate several branches. Each branch that satisfies certain consistency conditions, which we define below, represents a corresponding model.

Our tableaux use *prefixed formulas*, that is, formulas of the form $\sigma\ \varphi$, where $\sigma \in (\text{AG} \times L)^*$ and $\varphi \in L$; σ is the prefix of φ in that case, and we say that φ is prefixed by σ. We note that we separate the elements of σ with a dot. We say that the prefix σ is α-flat when α has axiom 5 and $\sigma = \sigma'.\alpha\langle\psi\rangle$ for some ψ. Each prefix possibly represents a state in a corresponding model, and a prefixed formula $\sigma\ \varphi$ declares that φ is satisfied in the state represented by σ. As we will see below, the prefixes from $(\text{AG} \times L)^*$ allow us to keep track of the diamond formula that generates a prefix through the tableau rules. For agents with condition 5, this allows us to restrict the generation of new prefixes and avoid certain redundancies, due to the similarity of euclidean binary relations to equivalence relations [18, 24].

The tableau rules that we use appear in Table 2. These include fixed-point and propositional rules, as well as rules that deal with modalities. Depending on the logic that each agent α is based on, a different set of rules applies for α: for rule (d), $\mathbf{L}(\alpha)$ must have condition D; for rule (t), $\mathbf{L}(\alpha)$ must have condition T; for rule (4), $\mathbf{L}(\alpha)$ must have condition 4; for rule (B5), (D5), and (D55), $\mathbf{L}(\alpha)$ must have condition 5; for (b) $\mathbf{L}(\alpha)$ must have condition B; and for (b4) $\mathbf{L}(\alpha)$ must have both B and 4. Rule (or) is the only rule that splits the current tableau branch

[5]It is easy to extract an upper bound of $2^{2^{O(n^3)}}$ for $\kappa_0(n)$ from the tableau in [21], but this bound is not useful to extract a decision procedure. The purpose of this section is not to establish any good upper bound for satisfiability testing, which is done in Section 3.

$$\frac{\sigma\ \pi X.\varphi}{\sigma\ \varphi}\ (\mathsf{fix}) \qquad \frac{\sigma\ X}{\sigma\ \mathsf{fx}(X)}\ (\mathsf{X}) \qquad \frac{\sigma\ \varphi \vee \psi}{\sigma\ \varphi \mid \sigma\ \psi}\ (\mathsf{or}) \qquad \frac{\sigma\ \varphi \wedge \psi}{\begin{array}{c}\sigma\ \varphi\\ \sigma\ \psi\end{array}}\ (\mathsf{and})$$

$$\frac{\sigma\ [\alpha]\varphi}{\sigma.\alpha\langle\psi\rangle\ \varphi}\ (\mathsf{B}) \qquad \frac{\sigma\ \langle\alpha\rangle\varphi}{\sigma.\alpha\langle\varphi\rangle\ \varphi}\ (\mathsf{D}) \qquad \frac{\sigma\ [\alpha]\varphi}{\sigma.\alpha\langle\varphi\rangle\ \varphi}\ (\mathsf{d}) \qquad \frac{\sigma\ [\alpha]\varphi}{\sigma.\alpha\langle\psi\rangle\ [\alpha]\varphi}\ (\mathsf{4})$$

where, for rules (B) and (4), $\sigma.\alpha\langle\psi\rangle$ has already appeared in the branch; and for (D), σ is not α-flat.

$$\frac{\sigma.\alpha\langle\psi\rangle\ [\alpha]\varphi}{\sigma\ [\alpha]\varphi}\ (\mathsf{B5}) \qquad \frac{\sigma.\alpha\langle\psi\rangle\ \langle\alpha\rangle\varphi}{\sigma.\alpha\langle\psi\rangle.\alpha\langle\varphi\rangle\ \varphi}\ (\mathsf{D5}) \qquad \frac{\sigma.\alpha\langle\psi\rangle\ [\alpha]\varphi}{\sigma\ \varphi}\ (\mathsf{b}) \qquad \frac{\sigma\ [\alpha]\varphi}{\sigma\ \varphi}\ (\mathsf{t})$$

$$\frac{\sigma.\alpha\langle\psi\rangle\ [\alpha]\varphi}{\sigma.\alpha\langle\psi'\rangle\ [\alpha]\varphi}\ (\mathsf{B55}) \qquad \frac{\sigma.\alpha\langle\psi\rangle.\alpha\langle\psi'\rangle\ \langle\alpha\rangle\varphi}{\sigma.\alpha\langle\psi\rangle.\alpha\langle\varphi\rangle\ \varphi}\ (\mathsf{D55}) \qquad \frac{\sigma.\alpha\langle\psi\rangle\ [\alpha]\varphi}{\sigma\ [\alpha]\varphi}\ (\mathsf{b4})$$

where, for rule (B55), $\sigma.\alpha\langle\psi'\rangle$ has already appeared in the branch; for rule (D5), σ is not α-flat, and $\sigma\ \langle\alpha\rangle\varphi$ does not appear in the branch; for rule (D55), $\sigma\ \langle\alpha\rangle\varphi$ does not appear in the branch.

Table 2: The tableau rules for $\mathbf{L} = \mathbf{L}_n^\mu$

into two. A tableau branch is propositionally closed when $\sigma\ \mathtt{ff}$ or both $\sigma\ p$ and $\sigma\ \neg p$ appear in the branch for some prefix σ. For each prefix σ that appears in a fixed tableau branch, let $\Phi(\sigma)$ be the set of formulas prefixed by σ in that branch. We use the notation $\sigma \prec \sigma'$ to mean that $\sigma' = \sigma.\sigma''$ for some σ'', in which case σ is an ancestor of σ'.

We define the relation $\xrightarrow{X}$ on prefixed formulas in a tableau branch as $\chi_1 \xrightarrow{X} \chi_2$, if $\frac{\chi_1}{\chi_2}$ is a tableau rule and χ_1 is not of the form $\sigma\ Y$, where $X < Y$; then, $\xrightarrow{X}^+$ is the transitive closure of $\xrightarrow{X}$ and $\xrightarrow{X}^*$ is its reflexive and transitive closure. We can also extend this relation to prefixes, so that $\sigma \xrightarrow{X} \sigma'$, if and only if $\sigma\ \psi \xrightarrow{X} \sigma'\ \psi'$, for some $\psi \in \Phi(\sigma)$ and $\psi' \in \Phi(\sigma')$. If in a branch there is a $\xrightarrow{X}$-sequence where X is a least fixed-point and appears infinitely often, then the branch is called fixed-point-closed. A branch is closed when it is either fixed-point-closed or propositionally closed; if it is not closed, then it is called open.

Now, assume that there is a $\kappa : \mathbb{N} \to \mathbb{N}$, such that every $\mathbf{L}$-satisfiable formula φ is satisfied in a model with at most $\kappa(|\varphi|)$ states. An open tableau branch is called (*resp. locally*) *maximal* when all tableau rules (*resp. the tableau rules that do not produce new prefixes*) have been applied. A branch is called *sufficient* for φ when it is locally maximal and for every $\sigma\ \psi$ in the branch, for which a rule can be applied and has not been applied to $\sigma\ \psi$, $|\sigma| > |\mathrm{AG}| \cdot \kappa(|\varphi|)^{|\varphi|^2} \cdot 2^{2|\varphi|+1}$. A tableau is called maximal when all of its open branches are maximal, and closed when all of its branches are closed. It is called sufficiently closed for φ if it is propositionally closed, or for some least fixed-point variable X, it has a $\xrightarrow{X}$-path, where X appears at least $\kappa(|\varphi|) + 1$ times. A sufficient branch for φ that is not sufficiently closed is called sufficiently open for φ.

A tableau for φ starts from $\varepsilon\ \varphi$ and is built using the tableau rules of Table 2. A tableau proof for φ is a closed tableau for the negation of φ.

Theorem 19 (Soundness, Completeness, and Termination of $\mathbf{L}_k^\mu$-Tableaux)**.** *From the following, the first two are equivalent for any formula $\varphi \in L$ and any logic $\mathbf{L}$. Furthermore, if there is a $\kappa : \mathbb{N} \to \mathbb{N}$, such that every $\mathbf{L}$-satisfiable formula φ is satisfied in a model with at most $\kappa(|\varphi|)$ states, then all the following are equivallent.*

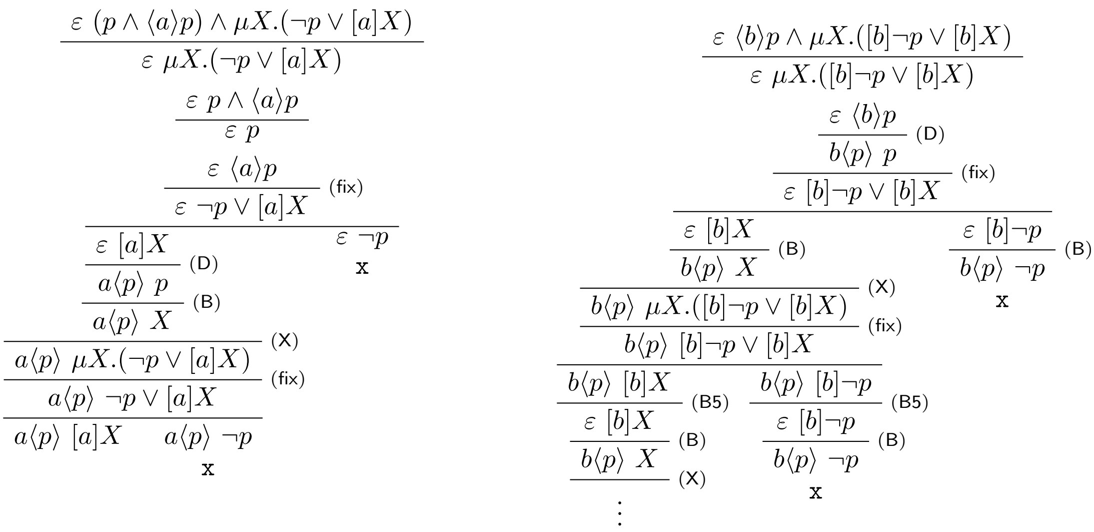

Figure 2: Tableaux for φ_1 and φ_2. The dots represent that the tableau keeps repeating as from the identical node above. The x mark represents a propositionally closed branch.

1. *φ has a maximal* **L***-tableau with an open branch;*
2. *φ is* **L***-satisfiable; and*
3. *φ has an* **L***-tableau with a sufficiently open branch for φ.*

Proof sketch. The direction from 1 to 2 uses the usual model construction, but where one needs to take into account the fixed-point formulas; the direction from 2 to 3 uses techniques and results from [21,27], including Corollary 18; and the direction from 3 to 1 shows how to detect appropriate parts of the branch to repeat until we safely get a maximal branch. The full proof can be found in Appendix B. ☐

Corollary 20. **L***-tableaux are sound and complete for* **L***.*

Example 3. Let $\text{AG} = \{a, b\}$ and $\mathbf{L}$ be a logic, such that $\mathbf{L}(a) = \mathbf{K}^\mu$ and $\mathbf{L}(b) = \mathbf{K5}^\mu$. Let

$$\varphi_1 = (p \wedge \langle a\rangle p) \wedge \mu X.(\neg p \vee [a]X) \qquad \textit{and} \qquad \varphi_2 = \langle b\rangle p \wedge \mu X.([b]\neg p \vee [b])X.$$

As we see in Figure 2, the tableau for φ_1 produces an open branch, while the one for φ_2 has all of its branches closed, the leftmost one due to an infinite $\xrightarrow{X}$-sequence.

5 Conclusions

We studied multi-modal logics with recursion. These logics mix the frame conditions from epistemic modal logic, and the recursion of the μ-calculus. We gave simple translations among these logics that connect their satisfiability problems. This allowed us to offer complexity bounds for satisfiability and to prove certain finite model results. We also presented a sound and complete tableau that has termination guarantees, conditional on a logic's finite model property.

Conjectures and Future Work We currently do not posses full translations for the cases of symmetric and euclidean frames. What is interesting is that we also do not have a counterexample to prove that the translations that we already have, as well as other attempts, are *not* correct. In the case of symmetric frames, we have managed to prove that our construction works for formulas without least-fixed-point operators. A translation for euclidean frames and for the full syntax on symmetric frames is left as future work. We know that we cannot use the same model constructions that preserve the finiteness of the model as in Subection 3.1.3 (see Remark 3).

We do not prove the finite model property on all logics. We note that although it is known that logics with recursion with at least two agents with either B or 5 do not have this property (see 3, [10]), the situation is unclear if there is only one such agent.

We further conjecture that it is not possible to prove EXP-completenes for all the single-agent cases. Specifically, we expect $\mathbf{K4}^{\mu}$-satisfiability to be in PSPACE, similarly to how $\mathbf{K5}^{\mu}$-satisfiability is in NP [10]. As such, we do not expect Translation 6 to be correct for these cases.

The model checking problem for the μ-calculus is an important open problem. The problem does not depend on the frame restrictions of the particular logic, though one may wonder whether additional frame restrictions would help solve the problem more efficiently. We are not aware of a way to use our translations to solve model checking more efficiently.

As, to the best of our knowledge, most of the logics described in this paper have not been explicitly defined before, with notable exceptions such as [3, 9, 10], they also lack any axiomatizations and completeness theorems. We do expect the classical methods from [17,21,22] and others to work out in these cases as well. However it would be interesting to see if there are any unexpected situations that arise.

Given the importance of common knowledge for epistemic logic and the fact that it has been known that common knowledge can be thought of as a (greatest) fixed-point already from [4,19], we consider the logics that we presented to be natural extensions of modal logic. Besides the examples given in Section 2, we are interested in exploring what other natural concepts can be defined with this enlarged language.

References

[1] Luca Aceto, Antonis Achilleos, Adrian Francalanza & Anna Ingólfsdóttir (2020): *The complexity of identifying characteristic formulae. J. Log. Algebraic Methods Program.* 112, p. 100529. Available at `https://doi.org/10.1016/j.jlamp.2020.100529`.

[2] Luca Alberucci & Alessandro Facchini (2009): *The modal μ-calculus hierarchy over restricted classes of transition systems. The Journal of Symbolic Logic* 74(4), p. 1367–1400, doi:10.2178/jsl/1254748696.

[3] Luca Alberucci & Alessandro Facchini (2009): *The modal μ-calculus hierarchy over restricted classes of transition systems. The Journal of Symbolic Logic* 74(4), pp. 1367–1400.

[4] Jon Barwise (1988): *Three views of common knowledge.* In: *Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning About Knowledge*, TARK '88, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp. 365–379. Available at `http://dl.acm.org/citation.cfm?id=1029718.1029753`.

[5] Patrick Blackburn, Johan van Benthem & Frank Wolter (2006): *Handbook of Modal Logic. Studies in Logic and Practical Reasoning* 3, Elsevier Science.

[6] Patrick Blackburn, Maarten de Rijke & Yde Venema (2001): *Modal Logic.* Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, doi:10.1017/cbo9781107050884.

[7] Laura Bozzelli, Bastien Maubert & Sophie Pinchinat (2015): *Unifying hyper and epistemic temporal logics*. In Andrew Pitts, editor: *Foundations of Software Science and Computation Structures*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 167–182.

[8] Michael R. Clarkson & Fred B. Schneider (2010): *Hyperproperties*. *Journal of Computer Security* 18(6), p. 1157–1210, doi:10.3233/JCS-2009-0393.

[9] Giovanna D'Agostino & Giacomo Lenzi (2010): *On the μ-calculus over transitive and finite transitive frames*. *Theoretical Computer Science* 411(50), pp. 4273–4290, doi:https://doi.org/10.1016/j.tcs.2010.09.002. Available at `https://www.sciencedirect.com/science/article/pii/S030439751000469X`.

[10] Giovanna D'Agostino & Giacomo Lenzi (2013): *On modal μ-calculus in S5 and applications*. *Fundamenta Informaticae* 124(4), pp. 465–482.

[11] Giovanna D'Agostino & Giacomo Lenzi (2015): *On the modal μ-Calculus over finite symmetric graphs*. *Mathematica Slovaca* 65(4), pp. 731–746.

[12] Giovanna D'Agostino & Giacomo Lenzi (2018): *The μ-Calculus Alternation Depth Hierarchy is infinite over finite planar graphs*. *Theoretical Computer Science* 737, pp. 40–61, doi:https://doi.org/10.1016/j.tcs.2018.04.009. Available at `https://www.sciencedirect.com/science/article/pii/S0304397518302317`.

[13] E Allen Emerson, Charanjit S Jutla & A Prasad Sistla (2001): *On model checking for the μ-calculus and its fragments*. *Theoretical Computer Science* 258(1-2), pp. 491–522, doi:10.1016/S0304-3975(00)00034-7.

[14] Ronald Fagin, Joseph Y. Halpern, Yoram Moses & Moshe Y. Vardi (1995): *Reasoning About Knowledge*. The MIT Press, doi:10.7551/mitpress/5803.001.0001.

[15] Michael J. Fischer & Richard E. Ladner (1979): *Propositional dynamic logic of regular programs*. *Journal of computer and system sciences* 18(2), pp. 194–211, doi:10.1016/0022-0000(79)90046-1.

[16] Melvin Fitting (1972): *Tableau methods of proof for modal logics*. *Notre Dame Journal of Formal Logic* 13(2), pp. 237–247.

[17] Joseph Y. Halpern & Yoram Moses (1992): *A guide to completeness and complexity for modal logics of knowledge and belief*. *Artificial Intelligence* 54(3), pp. 319–379, doi:10.1016/0004-3702(92)90049-4.

[18] Joseph Y. Halpern & Leandro Chaves Rêgo (2007): *Characterizing the NP-PSPACE gap in the satisfiability problem for modal logic*. *Journal of Logic and Computation* 17(4), pp. 795–806, doi:10.1093/logcom/exm029.

[19] Gilbert Harman (1977): *Review of linguistic behavior by Jonathan Bennett*. *Language* 53, pp. 417–424.

[20] Marcin Jurdziński (1998): *Deciding the winner in parity games is in UP $\cap$ co-UP*. *Information Processing Letters* 68(3), pp. 119–124.

[21] Dexter Kozen (1983): *Results on the propositional μ-calculus*. *Theoretical Computer Science* 27(3), pp. 333–354, doi:10.1016/0304-3975(82)90125-6.

[22] Richard E. Ladner (1977): *The computational complexity of provability in systems of modal propositional logic*. *SIAM Journal on Computing* 6(3), pp. 467–480, doi:10.1137/0206033.

[23] Fabio Massacci (1994): *Strongly analytic tableaux for normal modal logics*. In: *CADE*, pp. 723–737, doi:10.1007/3-540-58156-1_52.

[24] Michael C. Nagle & S. K. Thomason (1985): *The extensions of the modal logic K5*. *Journal of Symbolic Logic* 50(1), pp. 102—-109, doi:10.2307/2273793.

[25] Vaughan R. Pratt (1981): *A decidable mu-calculus: Preliminary report*. In: *22nd Annual Symposium on Foundations of Computer Science (SFCS 1981)*, IEEE, doi:10.1109/sfcs.1981.4.

[26] Mikhail Rybakov & Dmitry Shkatov (2018): *Complexity of finite-variable fragments of propositional modal logics of symmetric frames*. *Logic Journal of the IGPL* 27(1), pp. 60–68, doi:10.1093/jigpal/jzy018.

[27] Wieslaw Zielonka (1998): *Infinite games on finitely coloured graphs with applications to automata on infinite trees*. *Theoretical Computer Science* 200(1-2), pp. 135–183.

Appendix

A Proofs of Section 3

A.1 The proof of theorem 8

Proof. Assume first that $\mathcal{M}, w \models \varphi$, where $\mathcal{M}$ is an $\mathbf{L}_1$-model. For every subformula ψ of φ and $\alpha \in A$, $\mathsf{ax}^x_\alpha[cl(\psi)/p]$ is an instantiation of the axiom ax^x_α, and therefore it holds at all states of $\mathcal{M}$ that ar reachable from w. Thus, $\mathcal{M}, w \models Inv_d\left(\bigwedge_{\substack{\psi\in\mathsf{sub}(\varphi)\\ \alpha\in A}} \mathsf{ax}^x_\alpha[\psi/p]\right)$, which yields that $\mathcal{M}, w \models \mathsf{F}^x_A(\varphi)$.

For the other direction, assume that $\mathcal{M}, w \models \mathsf{F}^x_A(\varphi)$ for some $\mathbf{L}_2$-model $\mathcal{M} = (W, R, V)$ and state w. We assume that for every $\alpha, \beta \in \textsc{Ag}$, if $\alpha \neq \beta$, then $R_\alpha \cap R_\beta = \emptyset$. For each $k \geq 0$, we define $W_k \subseteq W^* \times \mathbb{N}$ in the following way: $W_0 := \{(w, 0)\}$; and $W_{k+1} = W_k \cup \{(pba, k+1) \mid \exists (pb, k) \in W_k, \alpha \in \textsc{Ag}.\ bR_\alpha a\}$. Let $W_\infty = \bigcup_{k=0}^{\infty} W_k$. Let for each $\alpha \in \textsc{Ag}$, $R^u_\alpha = \{((p,k),(pv,k+1)) \in W_\infty \times W_\infty \mid k \geq 0\}$ and $V^u(pv, k) = V(v)$ for all $k \geq 0$. Then, $\mathcal{M}^u = (W_\infty, R^u, V^u)$ is a bisimilar unfolding of $\mathcal{M}$, and therefore for every formula ψ and $(pv, k) \in W_k$, $\mathcal{M}^u, v \models \psi$ if and only if $\mathcal{M}, v \models \psi$. Then, it is not hard to see that for every formula ψ and $(v, k) \in W_k$, $\mathcal{M}^u, (v, k) \models \psi$ if and only if $\mathcal{M}^c, (v, k) \models \psi$, where $\mathcal{M}^c = (W_\infty, R^c, V^u)$, where R^c is the closure of R^u under the conditions of $\mathbf{L}_2$.

Let $d = md(\varphi)$ if $x \neq 4$, and $d = md(\varphi)|\varphi|$, if $x = 4$.

We first handle the case for $x \in \{D, T\}$. Let $\mathcal{M}' = (W', R', V')$, where $W' = W_d$, R' is the (*resp.* D-closure of) the restriction of R^c on W' (*resp.* if $\mathbf{L}_2$ has constraint D), and V' is the restriction of V^u on W'. We observe that $\mathcal{M}'$ remains a $\mathbf{L}_2$-model: removing states only affects condition D, and the D-closure does not affect the other conditions. Furthermore, one can see that, by induction on ψ, for every formula ψ with $md(\psi) \leq d$, and every $v \in W_{d-md(\psi)}$, $\mathcal{M}, v \models \psi$ if and only if $\mathcal{M}', v \models \psi$: propositional cases are straightforward, and modal cases ensure that $md(\psi) > 0$, and therefore $v \in W_{d-1}$, thus the accessible states from v remain the same in $\mathcal{M}$ and in $\mathcal{M}'$. Specifically, $\mathcal{M}', w \models \varphi$.

Let $\mathcal{M}^x = (W', \overline{R'}^x, V')$. It remains to prove that for every subformula ψ of φ, and every $v \in W_{d-md(\psi)}$, $\mathcal{M}', v \models \psi$ if and only if $\mathcal{M}^x, v \models \psi$. We continue by induction on the structure of ψ. The propositional and diamond cases are straightforward, since they are preserved by the introduction of pairs in the accessibility relation. We now consider the case of $\psi = [\alpha]\psi'$. We observe that $md(\psi) > 0$, and therefore $v \in W_{d-1}$. Specifically, if $md(\psi) = e$, then $v \in W_{d-e}$. If $\alpha \notin A$, then the accessible states from v remain the same in $\mathcal{M}'$ and in $\mathcal{M}^x$, and we are done. Therefore, we assume that $\alpha \in A$. If there is some $v\overline{R'_\alpha}^x u$, but not $vR'_\alpha u$, then we take cases for x:

$x = D$ We observe that $md(\langle\alpha\rangle\mathtt{tt}) = 1$, and therefore $\mathcal{M}', v \models \langle\alpha\rangle\mathtt{tt}$, since we observe above that $v \in W_{d-1}$. This contradicts our assumtion that $v\overline{R'_\alpha}^x u$ but not $vR'_\alpha u$, due to definition of $\overline{R'_\alpha}^D$.

$x = T$ We observe that $md([\alpha]\psi' \to \psi') = md([\alpha]\psi') = e$, and therefore $\mathcal{M}', v \models [\alpha]\psi' \to \psi'$, yielding that $\mathcal{M}', v \models \psi'$, and, by the inductive hypothesis, $\mathcal{M}^x, v \models \psi'$. By the definition of the T-closure, $v = u$, and therefore $\mathcal{M}^x, u \models \psi'$.

We now consider the case for $x = B$. We construct model $\mathcal{M}^x$ similarly, and we then prove that for every subformula ψ of φ, and every $(v, d-md(\psi)) \in W_{d-md(\psi)}$, $\mathcal{M}', (v, d-md(\psi)) \models \psi$ if and only if $\mathcal{M}^x, v \models \psi$.

We examine the case for $\psi = [\alpha]\psi'$, where $a \in A$. If $(v, d-e)\overline{R'_\alpha}^x(u,k)$, but not $vR'_\alpha u$, then we see that $k = d - e - 1$, and $uR'_\alpha v$. Then, $\mathcal{M}', u \models \langle\alpha\rangle[\alpha]\psi' \to \psi'$, and therefore $\mathcal{M}', u \models \langle\alpha\rangle[\alpha]\psi'$ and we are done by the inductive hypothesis.

We now consider the case for $x = 4$. For each $(v,k) \in W_\infty$ and $\alpha \in \textsc{Ag}$, let $b_\alpha(v,k) = \{[\alpha]\psi \in \mathsf{sub}(\varphi) \mid \mathcal{M}^u, (v,k) \models [\alpha]\psi\}$, and $d_\alpha(v,k) = \{\langle\alpha\rangle\psi \in \mathsf{sub}(\varphi) \mid \mathcal{M}^u, (v,k) \models \langle\alpha\rangle\psi\}$. Observe that for all $(v,k)R_\alpha(v',k+1)$, where $\alpha \in A$ and $k \leq d$, $b_\alpha(v,k) \subseteq b_\alpha(v',k+1)$ and $d_\alpha(v',k+1) \subseteq d_\alpha(v,k)$. We call a state $(pv,k+1)$ *α-stable*, when $(p,k)R^u_\alpha(pv,k+1)$, and $b_\alpha(p,k) = b_\alpha(pv,k+1)$ and $d_\alpha(p,k) = d_\alpha(pv,k+1)$ for some $\alpha \in A$, and write $(p,k) \triangleright_\alpha (pv,k+1)$. Observe that, by the Pigeonhole Principle, for $\alpha \in A$, in any sequence $(p,k)R_\alpha(pv_1,k+1)R_\alpha \cdots R_\alpha(pv_1\cdots v_l,k+l)$, where $l \geq |\varphi|$, there must be an α-stable state.

Let for each $\alpha \notin A$, $R^4_\alpha = R^u_\alpha$, and for each $\alpha \in A$,

$$R^4_\alpha = \{(a,b) \in R^u \mid a \text{ is not } \alpha\text{-stable}\} \cup \{(a,b) \mid \exists c \triangleright_\alpha a.\ cR^u_\alpha b\}.$$

Observe that if $a \triangleright_\alpha bR^u_\alpha c$, then c is not reachable from $(w,0)$ by R^4.

Let $\mathcal{M}_4 = (W', R^4, V')$ and $\mathcal{M}' = (W', R', V')$, where

$$W' = \{v \in W_d \mid v \text{ is reachable from } (w,0) \text{ by } R^4\},$$

R' is the closure of the restriction of R^4 on W' under the $\mathbf{L}_2$ constraints, and V' is the restriction of V^u on W'. $\mathcal{M}'$ is now a $\mathbf{L}_2$-model. We prove, by induction on ψ, that for every formula $\psi \in \mathsf{sub}(\varphi)$ and every $v = (p,k) \in W'$, where $k \leq d - md(\psi)|\varphi|$, $\mathcal{M}^u, v \models \psi$ if and only if $\mathcal{M}^4, v \models \psi$ if and only if $\mathcal{M}', v \models \psi$. The propositional cases are straightforward for both biimplications.

The modal cases $\psi = \langle\alpha\rangle\psi'$ or $\psi = [\alpha]\psi'$ ensure that $md(\psi) > 0$, and therefore $v \in W_{d-1}$, and there is some $(v,u) \in R^4_\alpha$. We first prove that $\mathcal{M}^u, v \models \psi$ if and only if $\mathcal{M}^4, v \models \psi$. If $\alpha \notin A$ or v is not α-stable, then this case follows from the observation that the accessible states from v in $\mathcal{M}^u$ and in $\mathcal{M}^4$ are the same. If $\alpha \in A$ and $v' \triangleleft v$, then v' is not α-stable, because otherwise $v \notin W'$, as it would not be reachable from $(w,0)$ by R^4 in $\mathcal{M}^u$. From $v' \triangleleft v$, we get that $\mathcal{M}^u, v' \models \psi$; by the inductive hypothesis, for every state $uR^4_\alpha v'$, $\mathcal{M}^u, u \models \psi'$ implies $\mathcal{M}^4, u \models \psi'$, and since v' has the same accessible states in $\mathcal{M}^u$ and in $\mathcal{M}^4$, $\mathcal{M}^u, v' \models \psi$. This completes the induction, and we conclude that $\mathcal{M}^u, v \models \psi$ if and only if $\mathcal{M}^4, v \models \psi$.

To prove the second biimplication, note that we have that for every formula $\psi \in \mathsf{sub}(\varphi)$ and every $v = (pu,k) \in W'$, where $k \leq d - md(\psi)|\varphi|$, $((pu,k),(puu',k')) \in R^4_\alpha$. Furthermore, for every $((pu,k),(puu',k')) \in R^4_\alpha$, we have that $(u,u') \in R_\alpha$. We then see that, since $\mathcal{M}$ is an $\mathbf{L}_2$-model, for every $((pu,k),(puu',k')) \in R^4_\alpha$, it must be the case that $(u,u') \in R_\alpha$. We can then conclude that $\mathcal{M}^4, v \models \psi$ if and only if $\mathcal{M}', v \models \psi$. Specifically, $\mathcal{M}', w \models \varphi$.

Let $\mathcal{M}^x = (W', \overline{R'}^{x,A}, V')$. It is straightforward now to prove that for every subformula ψ of φ, and every $v \in W_{d-md(\psi)}$, $\mathcal{M}', v \models \psi$ if and only if $\mathcal{M}^x, v \models \psi$, and the proof is complete.

We now consider the case for $x = 5$. We construct model $\mathcal{M}^x$ similarly, and we then prove that for every subformula ψ of φ, and every $(v, d-md(\psi)) \in W_{d-md(\psi)}$, $\mathcal{M}', (v, d-md(\psi)) \models \psi$ if and only if $\mathcal{M}^x, v \models \psi$.

We examine the case for $\psi = [\alpha]\psi'$, where $\alpha \in A$. Let $v\overline{R'_\alpha}^x u$, but not $vR'_\alpha u$. It suffices to prove that $\mathcal{M}^u, u \models \psi'$. From our assumption that $v\overline{R'_\alpha}^x u$, but not $vR'_\alpha u$, the euclidean closure condition, and the tree-structure of $\mathcal{M}^u$, we can see that there are some $a, b, c \in W'$, such that $aR'_\alpha \underline{b, c}$, and v is R'_α-reachable from b, and u is R'_α-reachable from c. Observe that for every $\psi \in \overline{\mathsf{sub}}(\varphi)$, $\alpha \in A$ and $v_1, v_2 \in W_d$, $v_1 R^c_\alpha v_2$, $\mathcal{M}^c, v_1 \models \langle\alpha\rangle\psi$ implies $\mathcal{M}^c, v_2 \models \langle\alpha\rangle\psi$, and $\mathcal{M}^c, v_2 \models [\alpha]\psi$ implies $\mathcal{M}^c, v_1 \models [\alpha]\psi$.

Assume that $\mathcal{M}^u, u \not\models \psi'$, to reach a contradiction. We can see that there is some u' that is R'_α-reachable from a and $u' R^u u$, such that $\mathcal{M}^u, u' \models \langle\alpha\rangle\neg\psi'$; in other words, $\mathcal{M}^u, u' \not\models [\alpha]\psi'$, and by the definition of ax^5_α, it must be the case that $\mathcal{M}^u, u' \not\models \langle\alpha\rangle[\alpha]\psi'$. Therefore, $\mathcal{M}^u, u' \models [\alpha]\neg\psi$, and by the definition of ax^5_α, it must be the case that $\mathcal{M}^u, a \models [\alpha]\neg\psi$, or, equivallently, $\mathcal{M}^u, a \models [\alpha]\langle\alpha\rangle\neg\psi'$. In turn, this yields that $\mathcal{M}^u, b \models \langle\alpha\rangle\neg\psi'$, and therefore $\mathcal{M}^u, v \models \neg\psi$, which is a contradiction. ◻

A.2 The proof of Theorem 11

Proof. First we assume that $\mathcal{M}, w \models \varphi$, where $\mathcal{M} = (W, R, V)$ is an $\mathbf{L}_1$-model and $w \in W$. We can easily verify, by induction on every subformula ψ of φ on every environment ρ and state s of W, that $\mathcal{M}, s \models_\rho \psi$ if and only if $\mathcal{M}, s \models_\rho \mathsf{F}^{T^\mu}_A(\psi)$. We then conclude that $\mathcal{M}, w \models \mathsf{F}^{T^\mu}_A(\varphi)$, and thus the translated formula remains satisfiable.

For the converse direction, we assume an $\mathbf{L}_2$-model $\mathcal{M} = (W, R, V)$ and $w \in W$, such that $\mathcal{M}, w \models \mathsf{F}^{T^\mu}_A(\varphi)$. We construct an $\mathbf{L}_1$-model that satisfies φ: let $\mathcal{M}_p = (W, R^p, V^p)$, where for each $\alpha \in A$, $R^p_\alpha = R_\alpha \cup \{(u, u) \mid u \in W\}$, the reflexive closure of R_α, and for each $\alpha \notin A$, $R^p_\alpha = R_\alpha$. It now suffices to prove that for every $v \in W$, environment ρ, and $\psi \in \mathsf{sub}(\varphi)$, $\mathcal{M}, v \models_\rho \mathsf{F}^{T^\mu}_A(\psi)$ if and only if $\mathcal{M}_p, v \models_\rho \psi$. We proceed by induction on ψ. The cases of propositional and recursion variables and boolean connectives are straightforward. The cases of fixed-points are also not hard by using the inductive hypothesis. Finally, the modal cases use the form of the translation to show that boxes are preserved in the reflexive closure, and that no diamonds are introduced. ◻

A.3 The proof of Theorem 13

Proof. First we assume that the μ-free formula φ is satisfied in $\mathcal{M}, w$, where $\mathcal{M} = (W, R, V)$ is an $\mathbf{L}_1$-model and $w \in W$. We assume that $p \notin V(u)$ for every $u \in W$. We construct an $\mathbf{L}_2$-model that satisfies $\mathsf{F}^{B^\mu}_A(\varphi)$. Let $\mathcal{M}_p = (W_u \cup W_p, R^p, V^p)$ and $\mathcal{M}_u = (W_u \cup W_u, R^u, V^u)$, where:

- $\mathcal{M}_u$ is the unfolding of $\mathcal{M}$.
- $W_p = \{u^p \mid u \in W_u\}$ is a set of distinct copies of states from W_u;
- For each $u \in W_u$, $V^p(u) = V(u)$ and $V^p(u_p) = V(u) \cup \{p\}$.
- $R^p_\alpha = R^u_\alpha \cup \{(s, t_p), (t_p, s) \mid (s, t) \in R^u_\alpha\}$, if $\alpha \in A$, and $R^p_\alpha = R^u_\alpha$ otherwise.

This construction is demonstrated Figure 3.

It is easy to see that we have $Inv(\Box\Diamond p)$ in all worlds s of $\mathcal{M}_p$. Namely, for each original world in the unfolding, we have that if there existed any arrow from it, now there is a new world s_p at distance 2 from s, where p holds. Moreover, since $\mathcal{M}_u$ is an unfolding, all other worlds t_p in W_p are at distance strictly larger or smaller than 2 from s since they have been only connected to either s itself or to worlds that are not neighbors of s. Since the world s_p has identical neighbors and propositional variables as s we have that they are bissimilar and thus they satisfy the same formulae. For all worlds $w \in W_p$, we have that they can reach themselves in 2 steps, and for all other worlds they can reach in 2 steps (remember that $\mathcal{M}_u$ is an unfolding

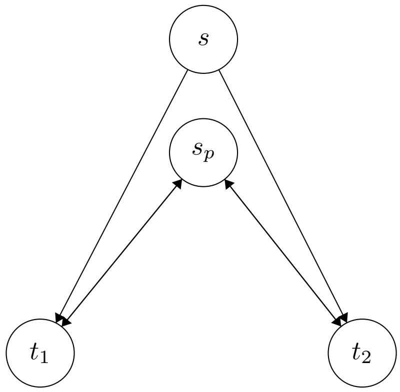

Figure 3: The new model, based on a world s in $\mathcal{M}$, with two neighbors t_1, t_2.

and thus each world has exactly one incoming edge), p does not hold. Thus we can also see that the whole translation of φ is satisfied in $\mathcal{M}'$. Furthermore, it is not hard to see that if we take the reflexive closure of R^p_α for each α such that $\mathbf{L}_2(\alpha)$ has condition T, then each state in the resulting model would be bisimilar to itself in $\mathcal{M}^p$; and that the respective serial closure would not affect the model.

For the converse, we assume an $\mathbf{L}_2(\alpha)$ model $\mathcal{M} = (W, R, V)$ and $w \in W$, such that $\mathcal{M}, w \models \mathsf{F}^{B^\mu}_A(\varphi)$, and we construct an $\mathbf{L}_1$-model $\mathcal{M}'$ that satisfies φ. Let $\mathcal{M}' = (W, R', V)$, where R'_α is the symmetric closure of R_α, if $\alpha \in A$, and $R'_\alpha = R_\alpha$ otherwise. Let ρ be an environment such that for every Y, $\rho(Y) = [\![cl(Y)]\!]$. We prove that for every state $v \in W$ and $\psi \in \mathsf{sub}(\varphi)$, $\mathcal{M}, v \models_\rho \mathsf{F}^{B^\mu}_A(\psi)$ implies that $\mathcal{M}', v \models_\rho \psi$. First, notice that for every $\alpha \in A$ and $vR_\alpha v_1 R_\alpha v_2$, $\mathcal{M}, v \models_\rho \psi$ iff $\mathcal{M}, v_2 \models_\rho \mathsf{F}^{B^\mu}_A(\psi)$. We now proceed by induction on ψ.

- Propositional cases, the case of logical variables, and the case of $\psi = [\alpha]\psi'$, where $\alpha \notin A$, are straightforward.

- For the case of $\psi = \langle\alpha\rangle\psi'$, notice that introducing pairs in the accessibility relation preserves diamonds.

- For the case of $\psi = [\alpha]\psi'$, where $\alpha \in A$, let $(v, v') \in R'_\alpha$. It suffices to prove that $\mathcal{M}', v' \models \psi'$. If $(v, v') \in R_\alpha$, then we are done. If not, then $v' R_\alpha v$, and therefore there is some $(v, v_p) \in R_\alpha$, such that $\mathcal{M}, v_p \models p$. Therefore, $\mathcal{M}, v_p \models_\rho \psi'$, yielding $\mathcal{M}, v' \models_\rho \psi'$.

- For the case of $\psi = \nu Y.\psi'$, notice that $[\![\psi', \rho[Y \mapsto [\![cl(Y)]\!]_\mathcal{M}]]\!]_{\mathcal{M}'} = [\![\psi', \rho]\!]_{\mathcal{M}'} = [\![\psi', \rho]\!]_\mathcal{M}$, by the inductive hypothesis. Therefore, $[\![\psi', \rho]\!]_\mathcal{M} \subseteq [\![\psi', \rho]\!]_{\mathcal{M}'}$, which is what we wanted to prove. □

A.4 Proof of Theorem 14

For the "only if" direction, let $\mathcal{M} = (W, R, V)$ be an unfolded model and $w \in W$ is its root, such that $\mathcal{M}, w \models \varphi$. Variables p and q do not appear in φ, so we can assume that $\mathcal{M}, w \models pq$, and that for each $vR_\alpha v'$, if $\alpha \in A$, then for each $\vec{p}$, $\mathcal{M}, v \models \vec{p}$ implies that $\mathcal{M}, v' \models next(\vec{p})$. Let $\mathcal{M}' = (W, R', V)$, where R is the appropriate closure of R under the conditions of L. We observe that for every $\alpha \in A$, $R_\alpha = R'_\alpha \cap \bigcup_{\vec{p}} [\![\vec{p}]\!] \times [\![next(\vec{p})]\!]$. We prove that for any environment

ρ, any $\psi \in \mathsf{sub}(\varphi)$, and any $v \in W$, $\mathcal{M}, v \models_\rho \psi$ iff $\mathcal{M}', v \models_\rho \mathsf{F}_A^{\mathbf{K}^\mu}(\psi)$. The proof proceeds by induction on ψ. We fix a $v \in W$ and a $\vec{p}$, such that $\mathcal{M}, v \models \vec{p}$.

- The propositional cases and the case of $\psi = X$ are immediate.
- So are the cases of $\psi = \langle\alpha\rangle\psi'$ and $\psi = [\alpha]\psi'$, where $\alpha \notin A$.
- For the case of $\psi = \nu X.\psi'$ or $\psi = \mu X.\psi'$, note that due to the inductive hypothesis, for any $S \subseteq W$, $[\![\psi', \rho[X \mapsto S]]\!]_\mathcal{M} = \left[\!\left[\mathsf{F}_A^{\mathbf{K}^\mu}(\psi'), \rho[X \mapsto S]\right]\!\right]_{\mathcal{M}'}$, and therefore $[\![\psi, \rho]\!]_\mathcal{M} = \left[\!\left[\mathsf{F}_A^{\mathbf{K}^\mu}(\psi), \rho\right]\!\right]_{\mathcal{M}'}$.
- For the case of $\psi = \langle\alpha\rangle\psi'$, where $\alpha \in A$, $\mathcal{M}, v \models_\rho \psi$ iff $\mathcal{M}, v \models_\rho \langle\alpha\rangle\psi'$ iff there is some $(v, v') \in R_\alpha$, $\mathcal{M}, v' \models_\rho \psi'$ iff there is some $(v, v') \in R'_\alpha$, where $\mathcal{M}, v' \models \mathit{next}(\vec{p})$, $\mathcal{M}, v' \models_\rho \psi'$, iff there is some $(v, v') \in R'_\alpha$, $\mathcal{M}, v' \models \mathit{next}(\vec{p}) \wedge \psi'$ iff
$$\mathcal{M}, v' \models \mathsf{F}_A^{\mathbf{K}^\mu}(\langle\alpha\rangle\psi') = \bigwedge_{\vec{p}}(\vec{p} \rightarrow \langle\alpha\rangle(\mathit{next}(\vec{p}) \wedge \mathsf{F}_A^{\mathbf{K}^\mu}(\psi'))).$$
- For the case of $\psi = [\alpha]\psi'$, where $\alpha \in A$, $\mathcal{M}, v \models_\rho \psi$ iff $\mathcal{M}, v \models_\rho [\alpha]\psi'$ iff for every $(v, v') \in R_\alpha$, $\mathcal{M}, v' \models_\rho \psi'$ iff for every $(v, v') \in R'_\alpha$, where $\mathcal{M}, v' \models \mathit{next}(\vec{p})$, $\mathcal{M}, v' \models_\rho \psi'$, iff for every $(v, v') \in R'_\alpha$, $\mathcal{M}, v' \models \mathit{next}(\vec{p}) \rightarrow \psi'$ iff
$$\mathcal{M}, v' \models \mathsf{F}_A^{\mathbf{K}^\mu}([\alpha]\psi') = \bigwedge_{\vec{p}}(\vec{p} \rightarrow [\alpha](\mathit{next}(\vec{p}) \rightarrow \mathsf{F}_A^{\mathbf{K}^\mu}(\psi'))),$$
which completes the proof by induction.

For the "if" direction, let $\mathcal{M} = (W, R, V)$ be an L-model and $w \in W$, such that $\mathcal{M}, w \models \mathsf{F}_A^{\mathbf{K}^\mu}(\varphi)$. Let $\mathcal{M}' = (W, R', V)$, where for every $\alpha \in A$, $R'_\alpha = R_\alpha \cap \bigcup_{\vec{p}} [\![\vec{p}]\!] \times [\![\mathit{next}(\vec{p})]\!]$, and for every $\alpha \notin A$, $R'_\alpha = R_\alpha$. We can prove that for any environment ρ, any $\psi \in \mathsf{sub}(\varphi)$, and any $v \in W$, $\mathcal{M}, v \models_\rho \mathsf{F}_A^{\mathbf{K}^\mu}(\psi)$ iff $\mathcal{M}', v \models_\rho \psi$. The proof proceeds by induction on ψ and is very similar to the "only if" direction. □

B Proofs of Section 4

Here we state and prove certain definitions and lemmata that will be used in the main proof of this Appendix. For an agent $\alpha \in \text{Ag}$ and a state s in a model (W, R, V), let $\mathit{Reach}_\alpha(s)$ be the states that are reachable in W by R_α. We also use $R|_S$ for the restriction of a relation R on a set S.

For a logic $\mathbf{L}$, we call a state s in a model $\mathcal{M} = (W, R, V)$ for $\mathbf{L}$ *flat* when for every $\alpha \in \text{Ag}$ for which $\mathbf{L}(\alpha)$ has constraint 5, there is a set of states W_0, such that:

- $\mathit{Reach}_\alpha(s) = \{s\} \cup W_0$;
- $R_\alpha|_{\mathit{Reach}_\alpha(s)} = E_0 \cup E_1$, where
  $E_0 \subseteq \{s\} \times W_0$ and
  $E_1 = W_0^2$; and
- if $\mathbf{L}(\alpha)$ has constraint T, or $E_0 \neq \emptyset$ and $\mathbf{L}(\alpha)$ has constraint B, then $s \in W_0$.

Lemma 21 ([18,24]). *Every pointed* $\mathbf{L}$*-model is bisimilar to* $\mathbf{L}$*-model whose states are all flat.*

B.1 The proof of Theorem 19

To prove that 1 implies 2, let b be a maximal open branch in the tableau for φ. We construct a $\mathbf{K}_k^\mu$-model $\mathcal{M} = (W, R, V)$ for φ in the following way. Let W be the set of prefixes that appear in the branch, and let, for each $\alpha \in \text{Ag}$,

$$R_\alpha^0 = \{\sigma, \sigma.\alpha\langle\psi\rangle \in W^2\} \cup \left\{ \sigma, \sigma \in W^2 \mid \mathbf{L}(\alpha) \text{ has } \begin{array}{c} \text{reflexive frames, or} \\ \text{serial frames and} \\ \forall\psi.\sigma.\alpha\langle\psi\rangle \notin W^2 \end{array} \right\};$$

R_α^1 is the symmetric closure of R_α^0, if $\mathbf{L}(\alpha)$ has symmetric frames, and it is R_α^0 otherwise; R_α^2 is the euclidean closure of R_α^1, if $\mathbf{L}(\alpha)$ has euclidean frames, and it is R_α^1 otherwise; and finally, R_α is the transitive closure of R_α^2, if $\mathbf{L}(\alpha)$ has transitive frames, and it is R_α^2 otherwise. By Lemma 7, R_a satisfies all the necessary closure conditions. We also set $V(p) = \{\sigma \in W \mid \sigma\ p \text{ appears in the branch }\}$.

It is now possible to prove, by straightforward induction, that for every subformula ψ of φ, if $\sigma\ \psi$ appears in the branch, then for any environment ρ, such that $\{\sigma' \in W \mid \sigma\ \psi \xrightarrow{X}^* \sigma'\ X\} \subseteq \rho(X)$, $\sigma \in [\![\psi, \rho_{\sigma,\psi}]\!]$. The only interesting cases are fixed-point formulas, so let $\psi = \nu X.\psi'$. Let S_X be the set of prefixes of X in the branch. We can immediately see that if $\sigma\ X$ appears in the branch, then so does $\sigma\ \psi'$, and therefore, by the inductive hypothesis, $S_X \subseteq [\![\psi, \rho[X \mapsto S_X]]\!]$. From the semantics in Table 1, $\sigma \in [\![\psi, \rho]\!]$.

On the other hand, if $\psi = \mu X.\psi'$, then we prove that if $\sigma \notin S \subseteq W$, then $S \not\supseteq [\![\psi', \rho[X \mapsto S]]\!]$. Let $\Psi = \{\sigma'\ \chi S \mid \sigma\ \psi' \xrightarrow{X}^* \sigma'\ \chi\ \sigma' \notin S\}$. We know that $\sigma\ \psi' \in \Psi$, so $\Psi \neq \emptyset$. There are no infinite $\xrightarrow{X}$-paths in the branch, so there is some $\sigma'\ \psi' \in \Psi$, such that $\sigma'\ \psi' \not\xrightarrow{X}^+ \sigma''\ \psi'$ for any σ''. Then, we see that $\{\sigma'' \in W \mid \sigma'\ \psi \xrightarrow{X}^* \sigma''\ X\} \subseteq S$, because if $\sigma'\ \psi \xrightarrow{X}^* \sigma''\ X$, then $\sigma'\ \psi' \xrightarrow{X}^* \sigma''\ \psi'$. But then, by the inductive hypothesis, $\sigma' \in [\![\psi', \rho[X \mapsto S]]\!]$, and therefore $S \not\supseteq [\![\psi', \rho[X \mapsto S]]\!]$, which was what we wanted to prove.

To prove that 2 implies 3, let $\mathcal{M} = (W, R, V)$ be a $\mathbf{L}$-model and $w \in W$, such that $\mathcal{M}, w \models \varphi$, and W has at most $\kappa(|\varphi|)$ states. Let the environment ρ be such that $\rho(X) = [\![\mathsf{fx}(X), \rho]\!]_\mathcal{M}$ for every variable X. We say that $\rightarrow \subseteq (W \times L)^2$ is a dependency relation on $\mathcal{M}$ when it satisfies the following conditions:

- if $u \models_\rho \psi_1 \wedge \psi_2$, then $(u, \psi_1 \wedge \psi_2) \rightarrow (u, \psi_1)$ and $(u, \psi_1 \wedge \psi_2) \rightarrow (u, \psi_2)$;
- if $u \models_\rho \psi_1 \vee \psi_2$, then $(u, \psi_1 \vee \psi_2) \rightarrow (u, \psi_1)$ and $u \models_\rho \psi_1$ or $(u, \psi_1 \wedge \psi_2) \rightarrow (u, \psi_2)$ and $u \models_\rho \psi_2$;
- if $u \models_\rho [\alpha]\psi$, then $(u, [\alpha]\psi) \rightarrow (v, \psi)$ for all $v \in W$, such that $uR_\alpha v$;
- if $u \models_\rho \langle\alpha\rangle\psi$, then $(u, \langle\alpha\rangle\psi) \rightarrow (v, \psi)$ for some $v \in W$, such that $uR_\alpha v \models \psi$;
- if $u \models_\rho \mu X.\psi$ or $u \models_\rho \nu X.\psi$, then $(u, \mathsf{fx}(X)) \rightarrow (u, \psi)$; and
- if $u \models_\rho X$, then $(u, X) \rightarrow (u, \mathsf{fx}(X))$.

For each variable X and dependency relation $\rightarrow$, we also define $\xrightarrow{X}$, such that $(w_1, \psi_1) \xrightarrow{X} (w_2, \psi_2)$ whenever $(w_1, \psi_1) \rightarrow (w_2, \psi_2)$ and $\psi_1 \neq Y$ for all variables Y where $\mathsf{fx}(Y)$ is not a subformula of $\mathsf{fx}(X)$. We call a dependency relation $\rightarrow$ lfp-finite, if for every least-fixed-point variable X, X appears finitely many times on every $\xrightarrow{X}$-sequence.

Claim: There is a lfp-finite dependency relation. The claim amounts to the memoryless determinacy of parity games and it can be proven similarly, as in [27]. Thus, we fix such a lfp-finite dependency relation.

The tableau starts with $\varepsilon\ \varphi$ and we can keep expanding this branch to a sufficient one using the tableau rules, such that every prefix is mapped to a state in W, whenever σ is mapped to u, $\mathcal{M}, u \models_\rho \psi$ for every $\psi \in \Phi(\sigma)$, and if $\sigma.\alpha\langle\psi\rangle$ to v, then $uR_\alpha v$; furthermore, this can be done by following the lfp-finite dependency relation. This is by straightforward induction on the application of the tableau rules. A special case are the agents with euclidean accessibility relations, for which we can use Lemma 21. It is not hard to see that in this way we generate a set of branches that are not propositionally closed. Furthermore, since the tableau rule applications follow a lfp-finite dependency relation and W has at most $\kappa(|\varphi|)$ states, it is not hard to see that for every least-fixed-point variable X, on every $\xrightarrow{X}$-path, X appears at most $\kappa(|\varphi|)$ times.

To prove that 3 implies 1, we assume that there is a sufficient open branch b in a tableau for φ and we demonstrate that φ has a maximal tableau with an open branch — specifically, we construct such an open branch from b.

We call a prefix σ a leaf when there is no $\sigma' \neq \sigma$ in the branch, such that $\sigma \prec \sigma'$; we call σ productive when a tableau rule on a formula $\sigma\ \psi$ in the branch can produce a new prefix. We call σ ready if it is of the form $\sigma'.\alpha\langle\psi_1\rangle.\alpha\langle\psi_2\rangle$, or if σ is not α-flat for any $\alpha \in \textsc{Ag}$.

For each $\sigma\ \psi$ in b and each least-fixed-point variable X, let

$$c(\sigma\ \psi, X) = \max\{n \mid \text{there is a } \xrightarrow{X} \begin{array}{l}\text{-path that ends in } \sigma\ \psi,\\ \text{where } X \text{ appears } n \text{ times}\end{array}\}.$$

Since b is not sufficiently closed, always $c(\sigma\ \psi, X) \leq \kappa(|\varphi|)$. We use the notation $\sigma \sim \sigma'$ to mean that $\sigma = \sigma_1.\alpha\langle\psi_1\rangle$, $\sigma = \sigma_2.\alpha\langle\psi_2\rangle$ for some $\sigma_1, \sigma_2, \alpha$, and $\Phi(\sigma) = \Phi(\sigma')$; and the notation $\sigma \equiv \sigma'$ to mean that $\sigma \sim \sigma'$ for every least-fixed-point variable X and $\psi \in \Phi(\sigma)$, $c(\sigma\ \psi, X) = c(\sigma'\ \psi, X)$ and either both are ready, or both σ and σ' are not ready and $\sigma + 1 \sim \sigma_2$, where $\sigma = \sigma_1.\alpha\langle\psi\rangle$ and $\sigma' = \sigma_2.\alpha\langle\psi'\rangle$. We then say that σ and σ' are equivalent.

Every productive leaf σ in the branch has an ancestor $e(\sigma)$ that has more that has a distinct, ready, and $\equiv$-equivalent ancestor; let $s(\sigma)$ be such an ancestor of $e(\sigma)$. We further assume that $e(\sigma)$ is the $\prec$-minimal ancestor of σ with these properties. We note that for any two productive leaves σ_1 and σ_2, if $e(\sigma_1) \prec \sigma_2$, then $e(\sigma_1) = e(\sigma_2)$.

Let b_0 be the branch that results by removing from b all prefixed formulas of the form $\sigma.\alpha\langle\psi_1\rangle\ \psi_2$, where $e(\sigma') \prec \sigma$ for some productive leaf σ'. Observe that b_0 is such that each of its productive leafs has an equivalent ancestor that is not a leaf. To complete the proof, it suffices to show how to extend any branch b_i, where each of its productive leafs σ has an equivalent proper ancestor $s(\sigma)$, to a branch b_{i+1} that preserves this property, and has an increased minimum length of its productive leafs. To form b_{i+1}, simply add to b_i all formulas of the form $\sigma.\sigma'\ \psi$, where $s(\sigma).\sigma'\ \psi$ appears in b_i.

Finally, observe that the finite model property of $\mathbf{L}$ was only used to prove equivalence with the third statement of the theorem. □