Logical Methods in Computer Science
Volume 15, Issue 2, 2019, pp. 2:1–2:45
https://lmcs.episciences.org/

Submitted Jan. 09, 2019
Published Apr. 11, 2019

CANONICAL MODELS AND THE COMPLEXITY OF MODAL TEAM LOGIC

MARTIN LÜCK

Leibniz Universität Hannover, Institut für Theoretische Informatik, Appelstraße 4, 30167 Hannover
e-mail address: lueck@thi.uni-hannover.de

Abstract. We study modal team logic MTL, the team-semantical extension of modal logic ML closed under Boolean negation. Its fragments, such as modal dependence, independence, and inclusion logic, are well-understood. However, due to the unrestricted Boolean negation, the satisfiability problem of full MTL has been notoriously resistant to a complexity theoretical classification.

In our approach, we introduce the notion of canonical models into the team-semantical setting. By construction of such a model, we reduce the satisfiability problem of MTL to simple model checking. Afterwards, we show that this approach is optimal in the sense that MTL-formulas can efficiently enforce canonicity.

Furthermore, to capture these results in terms of complexity, we introduce a non-elementary complexity class, TOWER(poly), and prove that it contains satisfiability and validity of MTL as complete problems. We also prove that the fragments of MTL with bounded modal depth are complete for the levels of the elementary hierarchy (with polynomially many alternations). The respective hardness results hold for both strict or lax semantics of the modal operators and the splitting disjunction, and also over the class of reflexive and transitive frames.

1. Introduction

It is well-known that non-linear quantifier dependencies, such as w depending only on z in the sentence $\forall x\, \exists y\, \forall z\, \exists w\, \varphi$, cannot be expressed in first-order logic. To overcome this restriction, logics of incomplete information such as *independence-friendly logic* [HS89] have been studied. Later, Hodges [Hod97] introduced *team semantics* to provide these logics with a compositional interpretation. The fundamental idea is to not consider single assignments to free variables, but instead whole sets of assignments, called *teams*.

In this vein, Väänänen [Vää07] expressed non-linear quantifier dependencies by the *dependence atom* $=(x_1, \ldots, x_n, y)$, which intuitively states that the values of y in the team functionally depend on those of $x_1, \ldots, x_n$. Logics with numerous other non-classical atoms such as *independence* $\perp$ [GV13], *inclusion* $\subseteq$ and *exclusion* $|$ [Gal12] have been studied since, and manifold connections to scientific areas such as statistics, database theory, physics, cryptography and social choice theory have emerged (see also Abramsky et al. [AKVV16]).

2012 ACM CCS: Theory of computation → Modal and temporal logics; Complexity theory and logic;
Key words and phrases: team semantics, modal logic, complexity, satisfiability.

* This article is an extension of the conference publication [Lüc18b].

Logic	Satisfiability	Validity	References
PDL	NP	NEXPTIME	[LV13, Vir17]
MDL	NEXPTIME	NEXPTIME	[Sev09, Han17]
PIL	NP	NEXPTIME-hard, in Π_2^E	[HKVV15]
MIL	NEXPTIME	Π_2^E-hard	[KMSV17, Han16]
PInc	EXPTIME	co-NP	[HKVV15]
MInc	EXPTIME	co-NEXPTIME-hard	[HKMV15]
PTL	ATIME-ALT(exp, poly)	ATIME-ALT(exp, poly)	[HKLV16, HKVV18]
MTL_k	ATIME-ALT($\exp_{k+1}$, poly)	ATIME-ALT($\exp_{k+1}$, poly)	Theorem 8.1
MTL	TOWER(poly)	TOWER(poly)	Theorem 8.1

Table 1: Complexity landscape of propositional and modal logics of dependence ($*$DL), independence ($*$IL), inclusion ($*$Inc) and team logic ($*$TL). Entries are completeness results unless stated otherwise.

Team semantics have also been adapted to a range of propositional [YV16, HKLV16], modal [Vää08], and temporal logics [KMV15, KMVZ18]. Besides *propositional dependence logic* PDL [YV16] and *modal dependence logic* MDL [Vää08], also propositional and modal logics of independence and inclusion have been studied [KMSV17, HKVV15, HS15, Han17]. Unlike in the first-order setting, the atoms such as the dependence atom range over flat formulas. For example, the instance $=(p_1, \ldots, p_n, \Diamond\mathsf{unsafe})$ of a modal dependence atom may specify that the reachability of an unsafe state is a function of $p_1 \cdots p_n$, but instead of exhibiting the explicit function, the atom only stipulates its existence.

Most team logics lack the Boolean negation, and adding it as a connective $\sim$ usually increases both the expressive power and the complexity tremendously. The respective extensions of propositional and modal logic are called *propositional team logic* PTL [HKLV16, YV17, HKVV18] and *modal team logic* MTL [Mül14, KMSV15]. With $\sim$, these logics can express all the non-classical atoms mentioned above, and in fact are expressively complete for their respective class of models [KMSV15, YV17]. For these reasons, they are both interesting and natural logics.

The expressive power of MTL is well-understood [KMSV15], and a complete axiomatization was presented by the author [Lüc18a]. Yet the complexity of the satisfiability problem has been an open question [Mül14, KMSV15, DKV16, HKMV17]. Recently, certain fragments of MTL with restricted negation were shown ATIME-ALT(exp, poly)-complete using the well-known filtration method [Lüc17]. In the same paper, however, it was shown that no elementary upper bound for full MTL can be established by the same approach, whereas the best known lower bound is ATIME-ALT(exp, poly)-hardness, inherited from propositional team logic [HKVV18].

Contribution. We show that MTL is complete for a non-elementary class we call TOWER(poly), which contains the problems decidable in a runtime that is a tower of nested exponentials of polynomial height. Likewise, we show that the fragments MTL_k of bounded modal depth k are complete for classes we call ATIME-ALT($\exp_{k+1}$, poly) and which corresponds to $(k+1)$-fold exponential runtime and polynomially many alternations. These results fill a long-standing gap in the active field of propositional and modal team logics (see Table 1).

In our approach, we consider so-called *canonical models*. Loosely speaking, a canonical model satisfies every satisfiable formula in some of its submodels, and such models have been long known for, e.g., many systems of modal logic [BRV01]. In Section 4, we adapt this notion for modal logics with team semantics, and prove that such models exist for MTL. This enables us to reduce the satisfiability problem to simple model checking, albeit on models that are of non-elementary size with respect to $|\Phi| + k$, where Φ are the available propositional variables and k is a bound on the modal depth.

Nonetheless, this approach is essentially optimal: In Section 5 and 6, we show that MTL can, in a certain sense, *efficiently enforce* canonical models, that is, with formulas that are of size polynomial in $|\Phi| + k$. In this vein, we then obtain the matching complexity lower bounds in Section 7 and 8, where we encode computations of non-elementary length in such large models.

Finally, in Section 9 we extend the preliminary version of this paper [Lüc18b] and consider restrictions of MTL to specific frame classes, and to so-called *strict* team-semantical connectives.

2. Preliminaries

The length of (the encoding of) x is denoted by $|x|$. We assume the reader to be familiar with alternating Turing machines [CKS81] and basic complexity theory. When a problem is hard or complete for a complexity class, in this paper we are always referring to logspace reductions.

The class ATIME-ALT$(\exp, \text{poly})$ (also known as AEXPTIME(poly)) contains the problems decidable by an alternating Turing machine in time $2^{p(n)}$ with $p(n)$ alternations, where p is a polynomial. We generalize it to capture the *elementary hierarchy* as follows.

Let $\exp_0(n) := n$ and $\exp_{k+1}(n) := 2^{\exp_k(n)}$. A function $f\colon \mathbb{N} \to \mathbb{N}$ is *elementary* if it is computable in time $\mathcal{O}(\exp_k(n))$ for some fixed k. In this paper, we consider the elementary hierarchy with polynomially many alternations:

Definition 2.1. For $k \geq 0$, ATIME-ALT$(\exp_k, \text{poly})$ is the class of problems decidable by an alternating Turing machine with at most $p(n)$ alternations and runtime at most $\exp_k(p(n))$, for a polynomial p.

Note that setting $k = 0$ or $k = 1$ yields the classes PSPACE and ATIME-ALT$(\exp, \text{poly})$, respectively [CKS81]. Schmitz [Sch16] proposed the following non-elementary class that contains ATIME-ALT$(\exp_k, \text{poly})$ for all k.

Definition 2.2 [Sch16]**.** TOWER is the class of problems decidable by a deterministic Turing machine in time (or equivalently, space) $\exp_{f(n)}(1)$ for an elementary function f.

A suitable notion of reduction for this class is the following: An *elementary reduction* from A to B is an elementary function f such that $x \in A \Leftrightarrow f(x) \in B$. $A \leq_{\mathrm{m}}^{\mathrm{elem}} B$ means that there exists an elementary reduction from A to B.

Proposition 2.3 [Sch16]**.** TOWER *is closed under* $\leq_{\mathrm{m}}^{\mathrm{elem}}$.

The next class results from imposing a polynomial bound on the number of exponentials in the definition of TOWER, which leads to a strict subclass.

Definition 2.4. TOWER(poly) is the class of problems that are decided by a deterministic Turing machine in time (or equivalently, space) $\exp_{p(n)}(1)$ for some polynomial p.

The reader may verify that both ATIME-ALT($\exp_k$, poly) and TOWER(poly) are closed under $\leq_{\mathrm{m}}^{\mathrm{P}}$ and $\leq_{\mathrm{m}}^{\log}$. Furthermore, by the time hierarchy theorem, TOWER(poly) $\subsetneq$ TOWER.

To the author's best knowledge, neither has been explicitly considered before. However, candidates for natural complete problems exist. Although not proved complete, several problems in TOWER(poly) are provably non-elementary, such as the satisfiability problem of separated first-order logic [Voi17], the equivalence problem for star-free expressions [SM73], or the first-order theory of finite trees [CH90], to only name a few. We refer the reader also to the survey of Meyer [Mey74].

Another example is the two-variable fragment of first-order team logic, $\mathsf{FO}^2(\sim)$. It is related to MTL in the same fashion as classical two-variable logic FO^2 to ML. By reduction from MTL to $\mathsf{FO}^2(\sim)$, the satisfiability problem of $\mathsf{FO}^2(\sim)$ is TOWER(poly)-complete problems as a corollary of our main result, Theorem 8.1, while its fragments $\mathsf{FO}^2_k(\sim)$ of bounded quantifier rank k are ATIME-ALT($\exp_{k+1}$, poly)-hard [Lüc18c].

Next, we justify why we use only $\leq_{\mathrm{m}}^{\log}$-reductions (or polynomial time reductions in general) in this paper instead of $\leq_{\mathrm{m}}^{\mathrm{elem}}$.

Proposition 2.5. *Every problem that is $\leq_{\mathrm{m}}^{\mathrm{elem}}$-complete for* TOWER(poly) *is also $\leq_{\mathrm{m}}^{\mathrm{elem}}$-complete for* TOWER*.*

Proof. Clearly, TOWER(poly) $\subseteq$ TOWER. For the lower bound, let A be $\leq_{\mathrm{m}}^{\mathrm{elem}}$-complete for TOWER(poly), and let $B \in$ TOWER be arbitrary. B is decidable in time $\exp_{r(n)}(1)$ for some elementary r. Define the set $C := \{x\#0^{r(|x|)} \mid x \in B\}$. First, we show that $C \in$ TOWER(poly). Consider the algorithm that first checks if the input z is of the form $x\#0^*$, computes $r(|x|)$ in elementary time, checks whether $z = x\#0^{r(|x|)}$, and then whether $x \in B$. The first two steps clearly take elementary time in n, where $n := |x\#0^{r(|x|)}|$, and the final step runs in time $\exp_{r(|x|)}(1) \leq \exp_n(1)$.

By assumption, $C \leq_{\mathrm{m}}^{\mathrm{elem}} A$ via an elementary reduction f. But clearly also $B \leq_{\mathrm{m}}^{\mathrm{elem}} C$ by the elementary reduction $g\colon x \mapsto x\#0^{r(|x|)}$. As a consequence, the function $h := f \circ g$ is a reduction from B to A. h is computable in time $\exp_{k_1}(\exp_{k_2}(n)) = \exp_{k_1+k_2}(n)$ for fixed $k_1, k_2 \geq 0$ depending on f and g, and hence again elementary. $\square$

Corollary 2.6. TOWER(poly) *is not closed under $\leq_{\mathrm{m}}^{\mathrm{elem}}$-reductions.*

Proof. Suppose TOWER(poly) is closed under $\leq_{\mathrm{m}}^{\mathrm{elem}}$-reductions, and let A be any problem complete for TOWER(poly) (such A exists; see also our main result, Theorem 8.1). By the previous proposition, then TOWER $\subseteq$ TOWER(poly), contradiction. $\square$

3. Modal team logic

We fix a countably infinite set $\mathcal{PS}$ of propositional symbols. *Modal team logic* MTL, introduced by Müller [Mül14], extends classical modal logic ML. Formulas of classical ML are built following the grammar

$$\alpha ::= \neg\alpha \mid \alpha \wedge \alpha \mid \alpha \vee \alpha \mid \Box\alpha \mid \Diamond\alpha \mid p \mid \top,$$

where $p \in \mathcal{PS}$ and $\top$ is constant truth. MTL extends ML by the grammar

$$\varphi ::= {\sim}\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \Box\varphi \mid \Diamond\varphi \mid \alpha,$$

where α denotes an ML-formula.

The set of propositional variables occurring in a formula $\varphi \in \mathsf{MTL}$ is $\mathsf{Prop}(\varphi)$. We use the common abbreviations $\bot := \neg\top$, $\alpha \to \beta := \neg\alpha \vee \beta$ and $\alpha \leftrightarrow \beta := (\alpha \wedge \beta) \vee (\neg\alpha \wedge \neg\beta)$. For easier distinction, we have classical ML-formulas denoted by $\alpha, \beta, \gamma, \dots$ and reserve $\varphi, \psi, \vartheta, \dots$ for general MTL-formulas.

The *modal depth* $\mathsf{md}(\varphi)$ of a formula φ is recursively defined:

$$\begin{aligned}
\mathsf{md}(p) &:= \mathsf{md}(\top) &&:= 0\\
\mathsf{md}({\sim}\varphi) &:= \mathsf{md}(\neg\varphi) &&:= \mathsf{md}(\varphi)\\
\mathsf{md}(\varphi \wedge \psi) &:= \mathsf{md}(\varphi \vee \psi) &&:= \max\{\mathsf{md}(\varphi), \mathsf{md}(\psi)\}\\
\mathsf{md}(\Diamond\varphi) &:= \mathsf{md}(\Box\varphi) &&:= \mathsf{md}(\varphi) + 1
\end{aligned}$$

ML_k and MTL_k are the fragments of ML and MTL with modal depth $\leq k$, respectively. If the propositions are restricted to a fixed set $\Phi \subseteq \mathcal{PS}$ as well, then the fragment is denoted by ML_k^{Φ}, or MTL_k^{Φ}, respectively.

Let $\Phi \subseteq \mathcal{PS}$ be finite. A *Kripke structure* (over Φ) is a tuple $\mathcal{K} = (W, R, V)$, where W is a set of *worlds* or *points*, (W, R) is a directed graph called *frame*, and $V : \Phi \to \mathfrak{P}(W)$ is the *valuation*, with $\mathfrak{P}(X)$ being the power set of X.

Occasionally, by slight abuse of notation, we use the inverse mapping $V^{-1} : W \to \mathfrak{P}(\Phi)$ defined by $V^{-1}(w) := \{p \in \Phi \mid w \in V(p)\}$ instead of V, i.e., the set of propositions that are true in a given world. If $w \in W$, then $(\mathcal{K}, w)$ is called *pointed structure*. ML is evaluated on pointed structures in the classical Kripke semantics.

By contrast, MTL is evaluated on pairs $(\mathcal{K}, T)$ called *structures with teams*, where $\mathcal{K} = (W, R, V)$ is a Kripke structure and $T \subseteq W$ is called *team* (in $\mathcal{K}$). Every team T has an *image* $RT := \{v \mid w \in T, (w, v) \in R\}$, and for $w \in W$, we simply write Rw instead of $R\{w\}$. R^iT is inductively defined as $R^0T := T$ and $R^{i+1}T := RR^iT$. An *R-successor team* (or simply successor team) of T is a team S such that $S \subseteq RT$ and $T \subseteq R^{-1}S$, where $R^{-1} := \{(v, w) \mid (w, v) \in R\}$. Intuitively, S is formed by picking at least one R-successor of every world in T. The semantics of MTL can now be defined as follows. [1]

$$\begin{aligned}
(\mathcal{K}, T) \vDash \alpha &\Leftrightarrow \forall w \in T : (\mathcal{K}, w) \vDash \alpha \text{ if } \alpha \in \mathsf{ML}\text{, and otherwise as}\\
(\mathcal{K}, T) \vDash {\sim}\psi &\Leftrightarrow (\mathcal{K}, T) \nvDash \psi,\\
(\mathcal{K}, T) \vDash \psi \wedge \theta &\Leftrightarrow (\mathcal{K}, T) \vDash \psi \text{ and } (\mathcal{K}, T) \vDash \theta,\\
(\mathcal{K}, T) \vDash \psi \vee \theta &\Leftrightarrow \exists S, U \subseteq T \text{ such that } T = S \cup U,\ (\mathcal{K}, S) \vDash \psi, \text{ and } (\mathcal{K}, U) \vDash \theta,\\
(\mathcal{K}, T) \vDash \Diamond\psi &\Leftrightarrow (\mathcal{K}, S) \vDash \psi \text{ for some successor team } S \text{ of } T,\\
(\mathcal{K}, T) \vDash \Box\psi &\Leftrightarrow (\mathcal{K}, RT) \vDash \psi.
\end{aligned}$$

We often omit $\mathcal{K}$ and write only $T \vDash \varphi$ (for team semantics) or $w \vDash \alpha$ (for Kripke semantics).

An MTL-formula φ is *satisfiable* if it is true in some structure with team over $\mathsf{Prop}(\varphi)$, which is then called a *model* of φ. Analogously, φ is *valid* if it is true in every structure with team (over $\mathsf{Prop}(\varphi)$). For a logic $\mathcal{L}$, the sets of all satisfiable resp. valid formulas of $\mathcal{L}$ are $\mathsf{SAT}(\mathcal{L})$ and $\mathsf{VAL}(\mathcal{L})$, respectively.

[1] Often, the "atoms" of MTL are restricted to literals $p, \neg p$ instead of ML-formulas α. However, this implies a restriction to formulas in negation normal form, and both definitions are equivalent due to the *flatness* property of ML (cf. [KMSV15, Proposition 2.2]).

In the literature on team semantics, the empty team is usually excluded in the above definition, since most $\sim$-free logics with team semantics have the *empty team property*, i.e., the empty team satisfies every formula [Vää08, KMSV17, HS15]. However, this distinction is unnecessary for MTL: φ is satisfiable iff $\top \vee \varphi$ is satisfied by some non-empty team[2], and φ is satisfied by some non-empty team iff $\sim\bot \wedge \varphi$ is satisfiable.

The modality-free fragment MTL_0 syntactically coincides with *propositional team logic* PTL [HKLV16, HKVV18, YV17]. The usual interpretations of the latter, i.e., sets of Boolean assignments, can easily be represented as teams in Kripke structures. For this reason, we treat PTL and MTL_0 as identical in this article.

Note that the connectives $\vee$, $\rightarrow$ and $\neg$ are not the Boolean disjunction, implication and negation, except on singleton teams, which correspond to Kripke semantics. Using $\wedge$ and $\sim$ however, we can define team-wide Boolean disjunction $\varphi_1 \owedge \varphi_2 := \sim(\sim\varphi_1 \wedge \sim\varphi_2)$ and material implication $\varphi_1 \twoheadrightarrow \varphi_2 := \sim\varphi_1 \owedge \varphi_2$.

The notation $\Box^i\varphi$ is defined via $\Box^0\varphi := \varphi$ and $\Box^{i+1}\varphi := \Box\Box^i\varphi$, and analogously for $\Diamond^i\varphi$. To express that at least one element of a team satisfies $\alpha \in \mathsf{ML}$, we use $\mathsf{E}\alpha := \sim\neg\alpha$.

MTL can express the *(extended) dependence atom* $=\!(\alpha_1, \ldots, \alpha_{n-1}, \alpha_n)$ of (extended) modal dependence logic [Vää08, EHM+13], which states that the truth value of α_n is a function of the truth values of $\alpha_1, \ldots, \alpha_{n-1}$, where $\alpha_1, \ldots, \alpha_n \in \mathsf{ML}$. It is definable in MTL as $\sim\left[\top \vee \sim\left(\bigwedge_{i=1}^{n-1} =\!(\alpha_i) \twoheadrightarrow =\!(\alpha_n)\right)\right]$, where $=\!(\alpha) := \alpha \owedge \neg\alpha$ is the *constancy atom*, stating that the truth value of $\alpha \in \mathsf{ML}$ is constant throughout the team.

The well-known *bisimulation* relation $\rightleftharpoons_k^\Phi$ fundamentally characterizes the expressive power of modal logic [BRV01] and plays a key role in our results.

Definition 3.1. Let $\Phi \subseteq \mathcal{PS}$ and $k \geq 0$. For $i \in \{1, 2\}$, let $(\mathcal{K}_i, w_i)$ be a pointed structure, where $\mathcal{K}_i = (W_i, R_i, V_i)$. Then $(\mathcal{K}_1, w_1)$ and $(\mathcal{K}_2, w_2)$ are (Φ, k)-*bisimilar*, in symbols $(\mathcal{K}_1, w_1) \rightleftharpoons_k^\Phi (\mathcal{K}_2, w_2)$, if
- $\forall p \in \Phi$: $w_1 \in V_1(p) \Leftrightarrow w_2 \in V_2(p)$,
- and if $k > 0$,
  - $\forall v_1 \in R_1 w_1 \colon \exists v_2 \in R_2 w_2 \colon (\mathcal{K}_1, v_1) \rightleftharpoons_{k-1}^\Phi (\mathcal{K}_2, v_2)$ (*forward* condition),
  - $\forall v_2 \in R_2 w_2 \colon \exists v_1 \in R_1 w_1 \colon (\mathcal{K}_1, v_1) \rightleftharpoons_{k-1}^\Phi (\mathcal{K}_2, v_2)$ (*backward* condition).

So-called *characteristic formulas* or *Hintikka formulas* capture the essence of the bisimulation relation in the following sense:

Proposition 3.2 [GO07, Theorem 32]**.** *Let $\Phi \subseteq \mathcal{PS}$ be finite, $k \geq 0$, and let $(\mathcal{K}, w)$ be a pointed structure. Then there is a formula $\zeta \in \mathsf{ML}_k^\Phi$ such that for all pointed structures $(\mathcal{K}', w')$ we have $(\mathcal{K}', w') \vDash \zeta$ if and only if $(\mathcal{K}, w) \rightleftharpoons_k^\Phi (\mathcal{K}', w')$.*

The notion of bisimulation was lifted to team semantics by Hella et al. [HLSV14, KMSV17, KMSV15]:

Definition 3.3. Let $\Phi \subseteq \mathcal{PS}$ and $k \geq 0$. For $i \in \{1, 2\}$, let $(\mathcal{K}_i, T_i)$ be a structure with team. Then $(\mathcal{K}_1, T_1)$ and $(\mathcal{K}_2, T_2)$ are (Φ, k)-*team-bisimilar*, written $(\mathcal{K}_1, T_1) \rightleftharpoons_k^\Phi (\mathcal{K}_2, T_2)$, if
- $\forall w_1 \in T_1 \colon \exists w_2 \in T_2 \colon (\mathcal{K}_1, w_1) \rightleftharpoons_k^\Phi (\mathcal{K}_2, w_2)$,

[2]Note that $\top \vee \varphi$ is not a tautology in general, since $\vee$ is not the Boolean disjunction. Rather, $\top \vee \varphi$ existentially quantifies a subteam where φ holds. In fact, $\top \vee \varphi$ is a tautology if and only if φ holds in the empty team.

- $\forall w_2 \in T_2 \colon \exists w_1 \in T_1 \colon (\mathcal{K}_1, w_1) \rightleftharpoons_k^\Phi (\mathcal{K}_2, w_2)$.

If no confusion can arise, we will also refer to teams T_1, T_2 that are (Φ, k)-team-bisimilar simply as (Φ, k)-bisimilar. Throughout the paper, we will make use of the following characterizations of bisimilarity.

Proposition 3.4. *Let $\Phi \subseteq \mathcal{PS}$ be finite, and $k \geq 0$. For $i \in \{1, 2\}$, let $(\mathcal{K}_i, w_i)$ be a pointed structure, where $\mathcal{K}_i = (W_i, R_i, V_i)$. The following statements are equivalent:*

(1) $\forall \alpha \in \mathsf{ML}_k^\Phi \colon (\mathcal{K}_1, w_1) \vDash \alpha \Leftrightarrow (\mathcal{K}_2, w_2) \vDash \alpha$,
(2) $(\mathcal{K}_1, w_1) \rightleftharpoons_k^\Phi (\mathcal{K}_2, w_2)$,
(3) $(\mathcal{K}_1, \{w_1\}) \rightleftharpoons_k^\Phi (\mathcal{K}_2, \{w_2\})$,
and if $k > 0$,
(4) $(\mathcal{K}_1, w_1) \rightleftharpoons_0^\Phi (\mathcal{K}_2, w_2)$ *and* $(\mathcal{K}_1, R_1 w_1) \rightleftharpoons_{k-1}^\Phi (\mathcal{K}_2, R_2 w_2)$.

Proof. (1) $\Leftrightarrow$ (2) is a standard result ([GO07, Theorem 32]). (2) $\Leftrightarrow$ (3) follows from Definition 3.3. For $k > 0$, we show that (2) + (3) implies (4). Clearly, $(\mathcal{K}_1, w_1) \rightleftharpoons_0^\Phi (\mathcal{K}_2, w_2)$ follows from (2). Due to Hella et al. [HLSV14, Lemma 3.3], (3) implies $(\mathcal{K}_1, R_1 w_1) \rightleftharpoons_{k-1}^\Phi (\mathcal{K}_2, R_2 w_2)$.

Finally, we show (4) $\Rightarrow$ (2). Suppose $(\mathcal{K}_1, w_1) \rightleftharpoons_0^\Phi (\mathcal{K}_2, w_2)$ and $(\mathcal{K}_1, R_1 w_1) \rightleftharpoons_{k-1}^\Phi (\mathcal{K}_2, R_2 w_2)$. Then to show $(\mathcal{K}_1, w_1) \rightleftharpoons_k^\Phi (\mathcal{K}_2, w_2)$, it is sufficient to prove the *forward* and *backward* conditions of Definition 3.1. Suppose $v_1 \in R_1 w_1$. Since $(\mathcal{K}_1, R_1 w_1) \rightleftharpoons_{k-1}^\Phi (\mathcal{K}_2, R_2 w_2)$, by Definition 3.3 there exists $v_2 \in R_2 w_2$ such that $(\mathcal{K}_1, v_1) \rightleftharpoons_{k-1}^\Phi (\mathcal{K}_2, v_2)$, proving the *forward* condition. The *backward* condition is symmetric. ☐

As a consequence, the *forward* and *backward* condition from Definition 3.1 can be equivalently stated in terms of team-bisimilarity of the respective image teams. A similar characterization exists for team-bisimilarity:

Proposition 3.5. *Let $\Phi \subseteq \mathcal{PS}$ be finite, and $k \geq 0$. Let $(\mathcal{K}_i, T_i)$ be a structure with team for $i \in \{1, 2\}$. Then the following statements are equivalent:*

(1) $\forall \alpha \in \mathsf{ML}_k^\Phi \colon (\mathcal{K}_1, T_1) \vDash \alpha \Leftrightarrow (\mathcal{K}_2, T_2) \vDash \alpha$,
(2) $\forall \varphi \in \mathsf{MTL}_k^\Phi \colon (\mathcal{K}_1, T_1) \vDash \varphi \Leftrightarrow (\mathcal{K}_2, T_2) \vDash \varphi$,
(3) $(\mathcal{K}_1, T_1) \rightleftharpoons_k^\Phi (\mathcal{K}_2, T_2)$.

Proof. The above statements are all true if $T_1 = T_2 = \emptyset$, and they are all false if exactly one of the teams is empty, since a team T satisfies the ML-formula $\bot$ precisely if $T = \emptyset$. For this reason, we can assume that both T_1 and T_2 are non-empty.

By Kontinen et al. [KMSV15, Proposition 3.10], for non-empty T_1, T_2 there exists an MTL_k^Φ-formula φ that is true in $(\mathcal{K}_1, T_1)$, but holds in $(\mathcal{K}_2, T_2)$ if and only if $(\mathcal{K}_1, T_1) \rightleftharpoons_k^\Phi (\mathcal{K}_2, T_2)$. This immediately proves (2) $\Rightarrow$ (3). The direction (3) $\Rightarrow$ (2) is due to Kontinen et al. [KMSV15, Proposition 2.8] as well.

Finally, (1) $\Leftrightarrow$ (2) follows from the fact that $\mathsf{ML}_k^\Phi \subseteq \mathsf{MTL}_k^\Phi$, and that conversely every MTL_k^Φ-formula is equivalent to a formula of the form

$$\bigovee_{i=1}^{n} \Big(\alpha_i \wedge \bigwedge_{j=1}^{m_i} \mathsf{E}\beta_{i,j} \Big),$$

where $\{\alpha_1, \ldots, \alpha_n, \beta_{1,1}, \ldots, \beta_{n,m_n}\} \subseteq \mathsf{ML}_k^\Phi$ (see [Lüc18a, Theorem 5.2] or [KMSV15, p. 11]). ☐

Note that the analog of condition 4 in Proposition 3.4 for team bisimulation is not equivalent: It is possible that $(\mathcal{K}_1, T_1) \rightleftharpoons_0^{\Phi} (\mathcal{K}_2, T_2)$ and $(\mathcal{K}_1, R_1T_1) \rightleftharpoons_{k-1}^{\Phi} (\mathcal{K}_2, R_2T_2)$, but $(\mathcal{K}_1, T_1) \not\rightleftharpoons_k^{\Phi} (\mathcal{K}_2, T_2)$.

4. Types and canonical models

Many modal logics admit a "universal" model, also called *canonical model*. The defining property of a canonical model is that it simultaneously witnesses all satisfiable (sets of) formulas in some of its points. These models are a popular tool for proving the completeness of manifold systems of modal logics; for the explicit construction of such a model for ML, consult, e.g., Blackburn et al. [BRV01, Section 4.2].

Unfortunately, any canonical model for ML is necessarily infinite, and consequently impractical for complexity theoretic considerations. Instead, we use so-called (Φ, k)-*canonical models* for finite $\Phi \subseteq \mathcal{PS}$ and $k \in \mathbb{N}$; as the name suggests they are canonical for the fragment ML_k^{Φ}. While these models are finite, by Proposition 3.4 their size is at least the number of equivalence classes of $\rightleftharpoons_k^{\Phi}$. We call the equivalence classes of $\rightleftharpoons_k^{\Phi}$ *types.*

A first issue arises since types are then proper classes, and in team semantics, we need to speak about *sets* of types. For this reason, we begin this section by defining types on proper set-theoretic grounds, by indentifying the type of a point with the set of formulas that are true in it, which is a standard approach in first-order model theory.

4.1. Types.

Definition 4.1. A set $\tau \subseteq \mathsf{ML}_k^{\Phi}$ is a (Φ, k)-*type* if it is satisfiable and for all $\alpha \in \mathsf{ML}_k^{\Phi}$ contains either α or $\neg\alpha$. The (Φ, k)-type of a pointed structure $(\mathcal{K}, w)$ is

$$[\![\mathcal{K}, w]\!]_k^{\Phi} := \{\alpha \in \mathsf{ML}_k^{\Phi} \mid (\mathcal{K}, w) \vDash \alpha\}.$$

The set of all (Φ, k)-types is Δ_k^{Φ}. Given a team T in $\mathcal{K}$, the types in T are

$$[\![\mathcal{K}, T]\!]_k^{\Phi} := \{[\![\mathcal{K}, w]\!]_k^{\Phi} \mid w \in T\}.$$

The following assertions ascertain that the above definition of types properly reflects the bisimulation relation.

Proposition 4.2. *Let $\Phi \subseteq \mathcal{PS}$ and $k \geq 0$. Then*

(1) *The unique (Φ, k)-type satisfied by $(\mathcal{K}, w)$ is $[\![\mathcal{K}, w]\!]_k^{\Phi}$.*
(2) *$(\mathcal{K}, w) \rightleftharpoons_k^{\Phi} (\mathcal{K}', w')$ if and only if $[\![\mathcal{K}, w]\!]_k^{\Phi} = [\![\mathcal{K}', w']\!]_k^{\Phi}$.*
(3) *$(\mathcal{K}, T) \rightleftharpoons_k^{\Phi} (\mathcal{K}', T')$ if and only if $[\![\mathcal{K}, T]\!]_k^{\Phi} = [\![\mathcal{K}', T']\!]_k^{\Phi}$.*

Proof. Property (1) is straightforward: two distinct types τ, τ' satisfied by $(\mathcal{K}, w)$ differ in some $\alpha \in \mathsf{ML}_k^{\Phi}$. But then $(\mathcal{K}, w) \vDash \alpha, \neg\alpha$, contradiction. Property (2) immediately follows from Proposition 3.4. For (3), first consider "$\Rightarrow$". Due to symmetry, we only show that $(\mathcal{K}, T) \rightleftharpoons_k^{\Phi} (\mathcal{K}', T')$ implies $[\![\mathcal{K}, T]\!]_k^{\Phi} \subseteq [\![\mathcal{K}', T']\!]_k^{\Phi}$. Hence suppose $\tau \in [\![\mathcal{K}, T]\!]_k^{\Phi}$. Then there exists $w \in T$ of type $[\![\mathcal{K}, w]\!]_k^{\Phi} = \tau$. By Definition 3.3, there is $w' \in T'$ with $(\mathcal{K}, w) \rightleftharpoons_k^{\Phi} (\mathcal{K}', w')$. Then $[\![\mathcal{K}', w']\!]_k^{\Phi} = \tau \in [\![\mathcal{K}', T']\!]_k^{\Phi}$ by property (2). The direction "$\Leftarrow$" of (3) is shown analogously. $\square$

It is unsurprising that the type of a point w is determined solely by the propositions in w and the types in the image Rw. In other words, all pointed structures of type τ satisfy the same propositions in their roots, viz. $\tau \cap \Phi$, and have the same types contained in their image teams. Regarding the latter, we define $\mathcal{R}\tau := \left\{\tau' \in \Delta_k^\Phi \mid \{\alpha \mid \Box\alpha \in \tau\} \subseteq \tau'\right\}$, given a $(\Phi, k+1)$-type τ. Intuitively, $\mathcal{R}\tau$ is the set of (Φ, k)-types that occur in the image team of a world of type τ.

The following proposition shows that types are indeed uniquely determined by the above constituents:

Proposition 4.3. *Let $\Phi \subseteq \mathcal{PS}$ be finite and $k \geq 0$.*

(1) $[\![w]\!]_k^\Phi \cap \Phi = V^{-1}(w) \cap \Phi$ *and* $[\![Rw]\!]_k^\Phi = \mathcal{R}[\![w]\!]_{k+1}^\Phi$, *for all pointed structures* (W, R, V, w).

(2) *The mapping* $h\colon \tau \mapsto \tau \cap \Phi$ *is a bijection from* Δ_0^Φ *to* $\mathfrak{P}(\Phi)$.

(3) *The mapping* $h\colon \tau \mapsto (\tau \cap \Phi, \mathcal{R}\tau)$ *is a bijection from* Δ_{k+1}^Φ *to* $\mathfrak{P}(\Phi) \times \mathfrak{P}(\Delta_k^\Phi)$.

Proof. See the appendix. $\square$

Lemma 4.4. *Let (W, R, V, w) be a pointed structure.*

(1) *If* $\tau \in \Delta_0^\Phi$, *then* $[\![w]\!]_0^\Phi = \tau$ *if and only if* $V^{-1}(w) = \tau \cap \Phi$.

(2) *If* $\tau \in \Delta_{k+1}^\Phi$, *then* $[\![w]\!]_{k+1}^\Phi = \tau$ *if and only if* $V^{-1}(w) = \tau \cap \Phi$ *and* $[\![Rw]\!]_k^\Phi = \mathcal{R}\tau$.

Proof. The direction "$\Rightarrow$" of 1. and 2. follows directly from Proposition 4.3. Moreover, we prove "$\Leftarrow$" only for statement 2., as the proof is analogous for 1.

Suppose that there are $\tau, \tau' \in \Delta_{k+1}^\Phi$ such that $V^{-1}(w) = \tau \cap \Phi$ and $[\![Rw]\!]_k^\Phi = \mathcal{R}\tau$, but $[\![w]\!]_{k+1}^\Phi = \tau'$. Then, by "$\Rightarrow$", we have $V^{-1}(w) = \tau' \cap \Phi$ and $[\![Rw]\!]_k^\Phi = \mathcal{R}\tau'$ as well. In other words, $\tau \cap \Phi = \tau' \cap \Phi$ and $\mathcal{R}\tau = \mathcal{R}\tau'$. However, since the mapping $h\colon \tau \mapsto (\tau \cap \Phi, \mathcal{R}\tau)$ is bijective according to Proposition 4.3, we have $\tau = \tau' = [\![w]\!]_{k+1}^\Phi$. $\square$

We are now ready to state the formal definition of canonicity by the notion of types:

Definition 4.5. A structure with team $(\mathcal{K}, T)$ is (Φ, k)-*canonical* if $[\![\mathcal{K}, T]\!]_k^\Phi = \Delta_k^\Phi$.

In the following, we often omit Φ and $\mathcal{K}$ and instead write $[\![w]\!]_k$ and $[\![T]\!]_k$, respectively, and simply say that T is (Φ, k)-canonical if $\mathcal{K}$ is clear.

4.2. **Canonical models in team semantics.** It is a standard result that for every Φ and $k \geq 0$ there exists a (Φ, k)-canonical model [BRV01], or in other words, that the logic ML_k^Φ admits canonical models.

We will show that, given a (Φ, k)-canonical model $\mathcal{K}$, every satisfiable MTL_k^Φ-formula can be satisfied in some team of $\mathcal{K}$ as well, despite MTL being significantly more expressive than ML [KMSV15]. In other words, the canonical models for MTL_k^Φ and ML_k^Φ coincide:

Theorem 4.6. *Let $(\mathcal{K}, T)$ be (Φ, k)-canonical and $\varphi \in \mathsf{MTL}_k^\Phi$. Then φ is satisfiable if and only if $(\mathcal{K}, T') \vDash \varphi$ for some $T' \subseteq T$.*

Proof. Assume $(\mathcal{K}, T)$ and φ are as above. As the direction from right to left is trivial, suppose that φ is satisfiable, i.e., has a model $(\hat{\mathcal{K}}, \hat{T})$. As a team in $\mathcal{K}$ that satisfies φ, we define

$$T' := \left\{ w \in T \;\middle|\; [\![\mathcal{K}, w]\!]_k^\Phi \in [\![\hat{\mathcal{K}}, \hat{T}]\!]_k^\Phi \right\}.$$

By Proposition 3.5 and 4.2, it suffices to prove $[\![\hat{\mathcal{K}}, \hat{T}]\!]_k^\Phi = [\![\mathcal{K}, T']\!]_k^\Phi$. Moreover, the direction "$\supseteq$" is clear by definition. As T is (Φ, k)-canonical, for every $\tau \in [\![\hat{\mathcal{K}}, \hat{T}]\!]_k^\Phi$ there exists a world $w \in T$ of type τ. Consequently, $[\![\hat{\mathcal{K}}, \hat{T}]\!]_k^\Phi \subseteq [\![\mathcal{K}, T']\!]_k^\Phi$. $\square$

How large is a (Φ, k)-canonical model at least? The number of types is captured by the function $\exp_k^*$, defined by

$$\exp_0^*(n) := n \qquad \exp_{k+1}^*(n) := n \cdot 2^{\exp_k^*(n)}.$$

Proposition 4.7. $|\Delta_k^\Phi| = \exp_k^*\left(2^{|\Phi|}\right)$ *for all* $k \geq 0$ *and finite* $\Phi \subseteq \mathcal{PS}$*.*

Proof. By induction on k. For the base case $k = 0$, this follows from Proposition 4.3, as there is a bijection between Δ_0^Φ and $\mathfrak{P}(\Phi)$ and $\exp_0^*\left(2^{|\Phi|}\right) = 2^{|\Phi|} = |\Delta_0^\Phi|$.

We proceed with the inductive step, i.e., $k+1$. First note that by induction hypothesis

$$\exp_{k+1}^*\left(2^{|\Phi|}\right) = 2^{|\Phi|} \cdot 2^{\exp_k^*(2^{|\Phi|})} = |\mathfrak{P}(\Phi) \times \mathfrak{P}(\Delta_k^\Phi)|.$$

Again, there exists a bijection from Δ_{k+1}^Φ to $\mathfrak{P}(\Phi) \times \mathfrak{P}(\Delta_k^\Phi)$ by Proposition 4.3. ☐

Next, we present an algorithm that solves the satisfiability and validity problems of MTL_k by computing a canonical model. Let us first explicate this construction in a lemma.

Lemma 4.8. *There is an algorithm that, given* $\Phi \subseteq \mathcal{PS}$ *and* $k \geq 0$*, computes a* (Φ, k)*-canonical model in time polynomial in* $|\Delta_k^\Phi|$*.*

Proof. The idea is to construct sets $L_0 \cup L_1 \cup \cdots \cup L_k$ of worlds in stage-wise manner such that L_i is (Φ, i)-canonical. For L_0, we simply add a world w for each $\Phi' \in \mathfrak{P}(\Phi)$ such that $V^{-1}(w) = \Phi'$. For $i > 0$, we iterate over all $L' \in \mathfrak{P}(L_{i-1})$ and $\Phi' \in \mathfrak{P}(\Phi)$ and insert a new world w into L_i such that L' is the image of w and such that again $V^{-1}(w) = \Phi'$. An inductive argument based on Proposition 3.5 and 4.3 shows that L_i is (Φ, i)-canonical for all $i \in \{0, \ldots, k\}$. As $k \leq |\Delta_k^\Phi|$, and each L_i is constructed in time polynomial in $|\Delta_i^\Phi| \leq |\Delta_k^\Phi|$, the overall runtime is polynomial in $|\Delta_k^\Phi|$. ☐

With the help of a small lemma, we conclude the upper bound for the satisfiability and validity problem of MTL and its fragments.

Lemma 4.9. *For every polynomial* p *there is a polynomial* q *such that*

$$p(\exp_k^*(n)) \leq \exp_k(q((k+1) \cdot n))$$

for all $k \geq 0$ *and* $n \geq 1$*.*

Proof. See the appendix. ☐

Theorem 4.10. $\mathsf{SAT}(\mathsf{MTL}_k)$ *and* $\mathsf{VAL}(\mathsf{MTL}_k)$ *are in* ATIME-ALT$(\exp_{k+1}, \text{poly})$*.*

Proof. Consider the following algorithm. Let $\varphi \in \mathsf{MTL}_k$ be the input, $n := |\varphi|$, and $\Phi := \mathsf{Prop}(\varphi)$. Construct deterministically, as in Lemma 4.8, a (Φ, k)-canonical structure $\mathcal{K} = (W, R, V)$ in time $p(|\Delta_k^\Phi|)$ for a polynomial p.

By a result of Müller [Mül14], the model checking problem of MTL is solvable by an alternating Turing machine that has runtime polynomial in $|\varphi| + |\mathcal{K}|$, and alternations polynomial in $|\varphi|$. We call this algorithm as a subroutine: by Theorem 4.6, φ is satisfiable (resp. valid) if and only if for at least one subteam (resp. all subteams) $T \subseteq W$ we have $(\mathcal{K}, T) \vDash \varphi$. Equivalently, this is the case if and only if $(\mathcal{K}, W)$ satisfies $\top \vee \varphi$ (resp. $\sim(\top \vee \sim\varphi)$).

Let us turn to the overall runtime. $\mathcal{K}$ is constructed in time polynomial in $|\Delta_k^\Phi| = \exp_k^*(2^{|\Phi|}) \leq \exp_{k+1}^*(|\Phi|) \leq \exp_{k+1}^*(n)$. The subsequent model checking runs in time polynomial in $|\mathcal{K}| + n$, and hence polynomial in $\exp_{k+1}^*(n)$ as well. By Lemma 4.9, we obtain a total runtime of $\exp_{k+1}(q((k+2) \cdot n))$ for a polynomial q. ☐

The upper bound for MTL is proved identically, since $k := \mathsf{md}(\varphi)$ is polynomial in $|\varphi|$.

Corollary 4.11. *$\mathsf{SAT}(\mathsf{MTL})$ and $\mathsf{VAL}(\mathsf{MTL})$ are in* TOWER(poly).

The usual definition of a canonical model is a structure that has all (infinite) maximal consistent subsets of a certain class of modal formulas as worlds (see virtually any textbook on modal logic, e.g., [BRV01]). This indeed results in a finite number of worlds in the case of, say, ML_k^Φ (cf. [Cre83, CH96]). Truly finitary constructions of canonical models can be traced back to Fine [Fin75], whose work has been extended towards various other modal systems (e.g., by Moss [Mos07]). Furthermore, Cresswell and Hughes [CH96] used *mini canonical models*, models that are "canonical" only with respect to all subformulas of a fixed ML-formula, which allows them to be finite models with finite sets of formulas as worlds.

All these approaches have in common that they still are non-constructive and intended for completeness proofs. Even computing a "mini canonical model" would not be guaranteed to be feasible enough for MTL: This would require an explicit translation of a given input MTL_k^Φ-formula to a Boolean combination of ML_k^Φ-formulas first (see the proof of Proposition 3.5), and it is open whether there is an elementary translation for every fixed k (cf. [Lüc18a]).

In this light, our approach yields a purely constructive definition of a canonical model (in Lemma 4.8), which can easily be plugged into the algorithms used for the above results, and has optimal runtime up to a polynomial.

5. Scopes and Subteam Quantifiers

Kontinen et al. [KMSV15] proved that MTL is expressively complete up to bisimulation: it can define every property of teams that is (Φ, k)-bisimulation invariant, that is closed under $\rightleftharpoons_k^\Phi$, for some finite Φ and k. Two team properties that fall into this category are in fact (Φ, k)-bisimilarity itself—in the sense that all worlds in a team have the same (Φ, k)-type—as well as (Φ, k)-canonicity. Consequently, these properties are definable by MTL_k^Φ-formulas. However, by a simple counting argument, formulas defining arbitrary team properties require non-elementary size w. r. t. Φ and k.

In this section, we consider a special class of structures, and on these, define k-bisimilarity by a formula χ_k of polynomial size in Φ and k. (From now on, we always assume some finite $\Phi \subseteq \mathcal{PS}$ and omit it in the notation, i.e., we write *k-canonicity*, *k-bisimilarity*, $\rightleftharpoons_k$, and so on.) Afterwards, in Section 6 we devise a formula canon_k of polynomial size that expresses k-canonicity.

5.1. Scopes.

It is natural to implement k-bisimilarity by mutual recursion in the spirit of Proposition 3.4: the $(k+1)$-bisimilarity of two points w, v is expressed in terms of k-team-bisimilarity of Rw and Rv, and conversely, to verify k-team-bisimilarity of Rw and Rv, we proceed analogously to the *forward* and *backward* conditions of Definition 3.1 and reduce the problem to checking k-bisimilarity of pairs of points in Rw and Rv.

MTL-formulas define team properties, but we want to express a *relation* between teams such as Rw and Rv. For this reason, we consider the "marked union" of Rw and Rv as a single team using the following tool. Formally, if $\alpha \in \mathsf{ML}$, then the "conditioned" subteam $T_\alpha \subseteq T$ is defined as

$$T_\alpha := \{\, w \in T \mid w \vDash \alpha \,\}.$$

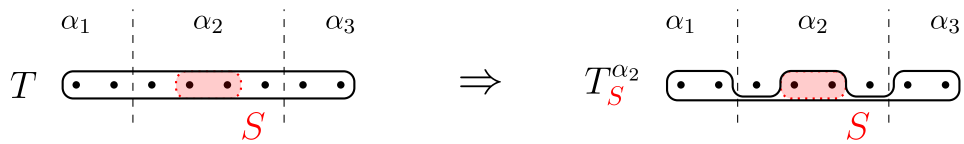

Figure 1: Example of subteam selection in the scope α_2

In the literature, T_α is also written $T \upharpoonright \alpha$ [Gal15, Gal16, Gal18]. The corresponding "decoding" operator

$$\alpha \hookrightarrow \varphi := \neg\alpha \vee (\alpha \wedge \varphi)$$

was introduced by Galliani [Gal15, Gal16, Gal18] as well: $\alpha \hookrightarrow \varphi$ is true in T if and only if $T_\alpha \vDash \varphi$.

Now, instead of defining an n-ary relation on teams, a formula φ can define a unary relation—a team property—parameterized by formulas $\alpha_1, \ldots, \alpha_n \in \mathsf{ML}$. We emphasize this by writing $\varphi(\alpha_1, \ldots, \alpha_n)$.

It will be useful if the "markers" of the constituent teams are invariant under traversing edges in the structure. In that case, we call these formulas *scopes*:

Definition 5.1. Let $\mathcal{K} = (W, R, V)$ be a Kripke structure. A formula $\alpha \in \mathsf{ML}$ is called a *scope (in $\mathcal{K}$)* if $(w, v) \in R$ implies $w \vDash \alpha \Leftrightarrow v \vDash \alpha$. Two scopes α, β are called *disjoint (in $\mathcal{K}$)* if W_α and W_β are disjoint.

To avoid interference, we always assume that scopes are formulas in $\mathsf{ML}_0^{\mathcal{PS}\setminus\Phi}$, i.e., they are always purely propositional and do not contain propositions from Φ.

It is desirable to be able to speak about subteams in a specific scope. If S is a team, let $T^\alpha_S := T_{\neg\alpha} \cup (T_\alpha \cap S)$. For singletons $\{w\}$, we simply write T^α_w instead of $T^\alpha_{\{w\}}$. Intuitively, T^α_S is obtained from T by "shrinking" the subteam T_α down to S without impairing $T \setminus T_\alpha$ (see Figure 1 for an example). Scopes have several desirable properties:

Proposition 5.2. *Let α, β be disjoint scopes and S, U, T teams in a Kripke structure $\mathcal{K} = (W, R, V)$. Then the following laws hold:*

(1) *Distributive laws: $(T \cap S)_\alpha = T_\alpha \cap S = T \cap S_\alpha = T_\alpha \cap S_\alpha$ and $(T \cup S)_\alpha = T_\alpha \cup S_\alpha$.*
(2) *Disjoint selection commutes: $\big(T^\alpha_S\big)^\beta_U = \big(T^\beta_U\big)^\alpha_S$.*
(3) *Disjoint selection is independent: $\big((T^\alpha_S)^\beta_U\big)_\alpha = T_\alpha \cap S$.*
(4) *Image and selection commute: $(RT)_\alpha = \big(R(T_\alpha)\big)_\alpha = R(T_\alpha)$*
(5) *Selection propagates: If $S \subseteq T$, then $R\big(T^\alpha_S\big) = (RT)^\alpha_{RS}$.*

Proof. Straightforward; see the appendix. □

Accordingly, we write R^iT_α instead of $(R^iT)_\alpha$ or $R^i(T_\alpha)$ and $T^{\alpha_1,\alpha_2}_{S_1,S_2}$ for $(T^{\alpha_1}_{S_1})^{\alpha_2}_{S_2}$.

5.2. **Subteam quantifiers.** We refer to the following abbreviations as *subteam quantifiers*, where $\alpha \in \mathsf{ML}$:

$$\exists^{\subseteq}_\alpha\, \varphi := \alpha \vee \varphi \qquad\qquad \forall^{\subseteq}_\alpha\, \varphi := {\sim}\exists^{\subseteq}_\alpha{\sim}\varphi$$

$$\exists^1_\alpha\, \varphi := \exists^{\subseteq}_\alpha \big[\mathsf{E}\alpha \wedge \forall^{\subseteq}_\alpha(\mathsf{E}\alpha \multimap \varphi)\big] \qquad\qquad \forall^1_\alpha\, \varphi := {\sim}\exists^1_\alpha{\sim}\varphi$$

Intuitively, they quantify over subteams $S \subseteq T_\alpha$ or worlds $w \in T_\alpha$ such that T^α_S resp. T^α_w satisfies φ.

Proposition 5.3. *The subteam quantifiers have the following semantics:*

$$T \vDash \exists^{\subseteq}_{\alpha}\varphi \Leftrightarrow \exists S \subseteq T_\alpha \colon T^\alpha_S \vDash \varphi \qquad T \vDash \exists^{1}_{\alpha}\varphi \Leftrightarrow \exists w \in T_\alpha \colon T^\alpha_w \vDash \varphi$$
$$T \vDash \forall^{\subseteq}_{\alpha}\varphi \Leftrightarrow \forall S \subseteq T_\alpha \colon T^\alpha_S \vDash \varphi \qquad T \vDash \forall^{1}_{\alpha}\varphi \Leftrightarrow \forall w \in T_\alpha \colon T^\alpha_w \vDash \varphi$$

Proof. We prove the existential cases, as the other ones work dually.

Let us first consider the "⇒" direction for $\exists^{\subseteq}_{\alpha}$. Accordingly, suppose $T \vDash \exists^{\subseteq}_{\alpha}\varphi$, i.e., $T \vDash \alpha \vee \varphi$. Then there exist $S \subseteq T$ and $U \subseteq T_\alpha$ such that $S \vDash \varphi$ and $T = S \cup U$. Since $U \cap T_{\neg\alpha} = \emptyset$, it holds $T_{\neg\alpha} \subseteq S$. Moreover, $S = (S \cap T_\alpha) \cup (S \cap T_{\neg\alpha}) = ((S \cap T_\alpha) \cap T_\alpha) \cup T_{\neg\alpha} = T^\alpha_{S \cap T_\alpha}$. Consequently, $T^\alpha_{S\cap T_\alpha} \vDash \varphi$ for some set $S \cap T_\alpha \subseteq T_\alpha$.

For "⇐", suppose $T^\alpha_S \vDash \varphi$ for some $S \subseteq T_\alpha$. Then T^α_S and $T \setminus T^\alpha_S$ form a division of T. Since $T \setminus T^\alpha_S = T \setminus (T_{\neg\alpha} \cup (T_\alpha \cap S)) \subseteq T \setminus T_{\neg\alpha} = T_\alpha$, it holds $T \setminus T^\alpha_S \vDash \alpha$. As a consequence, $T \vDash \alpha \vee \varphi$.

We proceed with $\exists^1_\alpha$. For "⇒", suppose that $T \vDash \exists^1_\alpha\varphi$. Then there exists $S \subseteq T_\alpha$ such that $T^\alpha_S \vDash \mathsf{E}\alpha \wedge \forall^{\subseteq}_{\alpha}(\mathsf{E}\alpha \multimap \varphi)$. Since $T^\alpha_S \vDash \mathsf{E}\alpha$, there exists $w \in (T^\alpha_S)_\alpha$. As $\forall^{\subseteq}_{\alpha}$ now applies to $(T^\alpha_S)^\alpha_{\{w\}} = T^\alpha_w$ as well, it follows $T^\alpha_w \vDash \mathsf{E}\alpha \multimap \varphi$, and consequently $T^\alpha_w \vDash \varphi$.

Suppose for "⇐" that $T^\alpha_w \vDash \varphi$ for some $w \in T_\alpha$. Let $S \subseteq T_\alpha$ be arbitrary. If $w \notin S$, then $(T^\alpha_w)^\alpha_S = T^\alpha_\emptyset \nvDash \mathsf{E}\alpha$, and if $w \in S$, then $(T^\alpha_w)^\alpha_S = T^\alpha_w \vDash \varphi$. Therefore, for any $S \subseteq T_\alpha$ it holds $(T^\alpha_w)^\alpha_S \vDash (\mathsf{E}\alpha \multimap \varphi)$, so $T^\alpha_w \vDash \forall^{\subseteq}_{\alpha}(\mathsf{E}\alpha \multimap \varphi)$. Since also $T^\alpha_w \vDash \mathsf{E}\alpha$, it follows $T \vDash \exists^{\subseteq}_{\alpha}\left[\mathsf{E}\alpha \wedge \forall^{\subseteq}_{\alpha}(\mathsf{E}\alpha \multimap \varphi)\right]$. ☐

5.3. Implementing bisimulation.

With scopes and subteam quantifiers at our hands, we have all ingredients to implement k-bisimulation.

$$\chi_0(\alpha,\beta) := (\alpha \vee \beta) \hookrightarrow \bigwedge_{p \in \Phi} {=}(p)$$

$$\chi_{k+1}(\alpha,\beta) := \chi_0(\alpha,\beta) \wedge \Box\chi^*_k(\alpha,\beta)$$

$$\chi^*_k(\alpha,\beta) := (\neg\alpha \wedge \neg\beta) \otimes \Big(\mathsf{E}\alpha \wedge \mathsf{E}\beta \wedge {\sim}\big[(\alpha \otimes \beta) \vee (\mathsf{E}\alpha \wedge \mathsf{E}\beta \wedge {\sim}\exists^1_\alpha\exists^1_\beta\chi_k(\alpha,\beta))\big]\Big)$$

Note that a literal translation of the forward and backward condition would rather result in the formula $\chi^*_k(\alpha,\beta) := \forall^1_\alpha\exists^1_\beta\chi_k(\alpha,\beta) \wedge \forall^1_\beta\exists^1_\alpha\chi_k(\alpha,\beta)$. The more complicated formula shown above however avoids the exponential size that would come with two recursive calls.

Theorem 5.4. *Let $k \geq 0$. For all Kripke structures $\mathcal{K}$, teams T and disjoint scopes α, β in $\mathcal{K}$, and points $w \in T_\alpha$ and $v \in T_\beta$ it holds:*

$$T^{\alpha,\beta}_{w,v} \vDash \chi_k(\alpha,\beta) \text{ if and only if } w \rightleftharpoons_k v,$$
$$T \vDash \chi^*_k(\alpha,\beta) \text{ if and only if } T_\alpha \rightleftharpoons_k T_\beta.$$

Moreover, both $\chi_k(\alpha,\beta)$ and $\chi^*_k(\alpha,\beta)$ are MTL_k-formulas that are constructible in space $\mathcal{O}(\log(k + |\Phi| + |\alpha| + |\beta|))$.

Proof. The idea is to isolate a single point in $z \in T_\alpha \cup T_\beta$ that serves as a *counter-example* against $[\![T_\alpha]\!]_k = [\![T_\beta]\!]_k$ by, say, $[\![z]\!]_k \in [\![T_\beta]\!]_k \setminus [\![T_\alpha]\!]_k$. We erase $T_\beta \setminus \{z\}$ from T using the disjunction $\vee$, as $T_\beta \setminus \{z\} \vDash \alpha \otimes \beta$. The remaining team is exactly T^β_z, in which $\exists^1_\alpha\exists^1_\beta\chi_k(\alpha,\beta)$ fails (see Figure 2). The case $[\![z]\!]_k \in [\![T_\alpha]\!]_k \setminus [\![T_\beta]\!]_k$ is detected analogously.

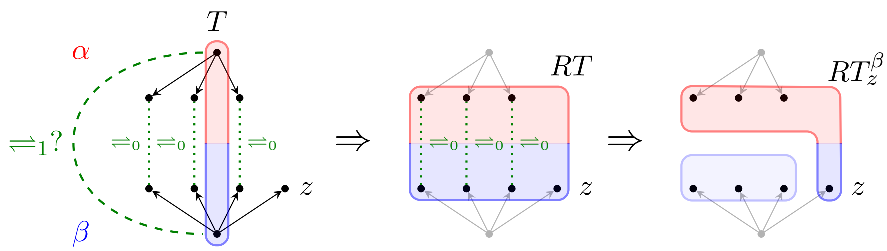

Figure 2: As z violates the *backward* condition, shrinking RT_β leads to a $\rightleftharpoons_0$-free subteam, falsifying $\exists^1_\alpha \exists^1_\beta \chi_0(\alpha, \beta)$.

We proceed with a formal correctness proof by induction on k. Let $\mathcal{K} = (W, R, V)$ as in the theorem. The base case $k = 0$ is straightforward, as no proposition $p \in \Phi$ occurs in α or β. The induction step is split into two parts.

"$\chi_k \Rightarrow \chi_k^*$": Let T be a team and α, β disjoint scopes. Observe that χ_k^* is always true if T_α and T_β are both empty (then $\llbracket T_\alpha \rrbracket_k = \llbracket T_\beta \rrbracket_k$), and that it is always false if exactly one of them is empty (then $\llbracket T_\alpha \rrbracket_k \neq \llbracket T_\beta \rrbracket_k$). Therefore, let $T_\alpha \neq \emptyset$ and $T_\beta \neq \emptyset$. Then $\chi_k^*(\alpha, \beta)$ boils down to $\sim((\alpha \ovee \beta) \vee \mathsf{E}\alpha \wedge \mathsf{E}\beta \wedge {\sim}\exists^1_\alpha \exists^1_\beta \chi_k(\alpha, \beta))$, which we prove equivalent to $\llbracket T_\alpha \rrbracket_k = \llbracket T_\beta \rrbracket_k$.

The first direction is proved by contradiction. Suppose $\llbracket T_\alpha \rrbracket_k = \llbracket T_\beta \rrbracket_k$ but $T \vDash (\alpha \ovee \beta) \vee \mathsf{E}\alpha \wedge \mathsf{E}\beta \wedge {\sim}\exists^1_\alpha \exists^1_\beta \chi_k(\alpha, \beta)$. The disjunction is witnessed by some division $T = S \cup U$, where w.l.o.g. $S \subseteq T_\alpha$ satisfies $\alpha \ovee \beta$, (if $S \subseteq T_\beta$, the proof is symmetric), and $U \vDash \mathsf{E}\alpha \wedge \mathsf{E}\beta \wedge {\sim}\exists^1_\alpha \exists^1_\beta \chi_k(\alpha, \beta)$. Since $T_\alpha \cap T_\beta = \emptyset$, then $T_\beta \subseteq U$, and clearly $T_\beta \subseteq U_\beta$. By the formula, some $w \in U_\alpha$ exists. By assumption that $\llbracket T_\alpha \rrbracket_k = \llbracket T_\beta \rrbracket_k$, U_β must contain a world v of type $\llbracket w \rrbracket_k$ as well. But then $U^{\alpha,\beta}_{w,v} \vDash \chi_k(\alpha, \beta)$ by induction hypothesis, contradiction to $U \vDash {\sim}\exists^1_\alpha \exists^1_\beta \chi_k(\alpha, \beta)$.

For the other direction, suppose $\llbracket T_\alpha \rrbracket_k \neq \llbracket T_\beta \rrbracket_k$. W.l.o.g. there exists $w \in T_\alpha$ such that $\llbracket w \rrbracket_k \notin \llbracket T_\beta \rrbracket_k$. (For $w \in T_\beta$, the proof is again symmetric.) Consider $S := T_\alpha \setminus \{w\}$ and $U := T^\alpha_w$ as a division of T. Then $S \vDash \alpha \ovee \beta$ and $U \vDash \mathsf{E}\alpha \wedge \mathsf{E}\beta$. It remains to show $U \vDash {\sim}\exists^1_\alpha \exists^1_\beta \chi_k(\alpha, \beta)$. However, this is easy to see: $U \vDash \exists^1_\alpha \exists^1_\beta \chi_k(\alpha, \beta)$ if and only if $U \vDash \exists^1_\beta \chi_k(\alpha, \beta)$, but T_β and hence U_β contains no world of type $\llbracket w \rrbracket_k$, so by induction hypothesis U cannot satisfy $\exists^1_\beta \chi_k(\alpha, \beta)$.

"$\chi_k^* \Rightarrow \chi_{k+1}$": We follow Definition 3.1 and Proposition 3.4.

$$
\begin{aligned}
& T^{\alpha,\beta}_{w,v} \vDash \chi_{k+1}(\alpha, \beta) & \\
\Leftrightarrow\ & T^{\alpha,\beta}_{w,v} \vDash \chi_0(\alpha, \beta) \wedge \Box \chi_k^*(\alpha, \beta) & \text{(Definition of } \chi_{k+1}\text{)} \\
\Leftrightarrow\ & w \rightleftharpoons_0 v \text{ and } T^{\alpha,\beta}_{w,v} \vDash \Box \chi_k^*(\alpha, \beta) & \text{(Induction hypothesis)} \\
\Leftrightarrow\ & w \rightleftharpoons_0 v \text{ and } RT^{\alpha,\beta}_{Rw,Rv} \vDash \chi_k^*(\alpha, \beta) & \text{(Proposition 5.2)} \\
\Leftrightarrow\ & w \rightleftharpoons_0 v \text{ and } Rw \rightleftharpoons_k Rv & \text{(Induction hypothesis)} \\
\Leftrightarrow\ & w \rightleftharpoons_{k+1} v. & \text{(Proposition 3.4)}
\end{aligned}
$$

It is routine to check that the formulas are constructible in logarithmic space from α, β, Φ and k, and that $\mathsf{md}(\chi_k) = \mathsf{md}(\chi_k^*) = k$. $\square$

Let us stress that χ_k relies on disjoint scopes to be present in the structure, and it is open whether the property $|[\![T]\!]_k| \leq 1$ is polynomially definable without these. Incidentally, the related property $|[\![Rw]\!]_k| \leq 1$ of points w was recently studied by Hella and Vilander [HV16], and was proven to be expressible in ML, but only by formulas of non-elementary size. However, they proved that it is definable in exponential size in *2-dimensional modal logic* ML^2 (for an introduction to ML^2, see Marx and Venema [MV97]). Roughly speaking, ML^2 is evaluated by traversing over *pairs* of points independently. The relationship between ML^2 and MTL is unclear: Arguably, pairs of points are a special case of teams. But on the other hand, the modalities in MTL do not act on the points in a team independently, as in ML^2, but instead always proceed to a successor team "synchronously". As a consequence, it is also open whether MTL can define one of the above properties by a formula of elementary size.

6. Enforcing a canonical model

In this section, we approach the canonical models of MTL from a lower bound perspective. Here, we devise an MTL_k-formula that is satisfiable but permits *only* k-canonical models.

For $k = 0$, that is propositional team logic, Hannula et al. [HKVV15] defined the PTL-formula

$$\mathsf{max}(X) := \sim \bigvee_{x \in X} {=}(x)$$

and proved that $T \vDash \mathsf{max}(\Phi)$ if and only if T is 0-canonical, i.e., contains all Boolean assignment over Φ. We generalize this for all k, i.e., construct a satisfiable formula canon_k that has only k-canonical models.

6.1. Staircase models.

Our approach is to express k-canonicity by inductively enforcing i-canonical sets of worlds for $i = 0, \ldots, k$ located in different "height" inside the model. For this purpose, we employ distinct scopes $\mathfrak{s}_0, \ldots, \mathfrak{s}_k$ ("*stairs*"), and introduce a specific class of models:

Definition 6.1. Let $k, i \geq 0$ and let $(\mathcal{K}, T)$ be a Kripke structure with team, $\mathcal{K} = (W, R, V)$. Then T is *k-canonical with offset i* if for every $\tau \in \Delta_k$ there exists $w \in T$ with $[\![R^i w]\!]_k = \{\tau\}$. $(\mathcal{K}, T)$ is called *k-staircase* if for all $i \in \{0, \ldots, k\}$ we have that $T_{\mathfrak{s}_i}$ is i-canonical with offset $k - i$.

As an example, a 3-staircase for $\Phi = \emptyset$ is depicted in Figure 3. Observe that it is a *directed forest*, i.e., it is acyclic and all worlds are either *roots* (i.e., without predecessor) or have exactly one predecessor. Moreover, it has bounded *height*, where the height of a directed forest is the greatest number h such that every path traverses at most h edges.

Proposition 6.2. *For each $k \geq 0$, there is a finite k-staircase $(\mathcal{K}, T)$ such that $\mathfrak{s}_0, \ldots, \mathfrak{s}_k$ are disjoint scopes in $\mathcal{K}$, and $\mathcal{K}$ is a directed forest with height at most k and its set of roots being exactly T.*

Proof. See Figure 3. □

Observe that in such a model, $T_{\mathfrak{s}_k}$ is k-canonical with offset 0, which is simply k-canonical:

Corollary 6.3 (Finite tree model property of MTL)**.** *Every satisfiable MTL_k-formula has a finite model $(\mathcal{K}, T)$ such that $\mathcal{K}$ is a directed forest with height at most k and its set of roots being exactly T.*

6.2. **Enforcing canonicity.** In the rest of the section, we illustrate how a k-staircase can be enforced in MTL inductively.

For $\Phi = \emptyset$, the inductive step—obtaining $(k+1)$-canonicity from k-canonicity—is done by the formula $\forall^{\subseteq}_{\alpha} \exists^{1}_{\beta} \Box\chi^{*}_{k}(\alpha, \beta)$. The idea is that this formula states that for every *subteam* $T' \subseteq T_\alpha$ there exists a *point* $w \in T_\beta$ such that $[\![RT']\!]_k = [\![Rw]\!]_k$. Intuitively, every possible set of types is captured as the image of some point in T_β. As a consequence, if T_α is k-canonical with offset 1, then T_β will be $(k+1)$-canonical.

Note that the simpler formula $\Box^k \mathsf{max}(\Phi)$ expresses 0-canonicity of R^kT, but not 0-canonicity of T with offset k (consider, e.g., a singleton T). Instead, we use the formula

$$\mathsf{max}_i := \top \lor (\Diamond^i\top \land \sim\bigvee_{p\in\Phi}(\Diamond^i p \ovee \Diamond^i \neg p)).$$

It states not only that R^iT is 0-canonical, but also that R^iw contains exactly one propositional assignment for each $w \in T$, which together yields 0-canonicity with offset i.

Lemma 6.4. *$T \vDash \mathsf{max}_i$ iff T is 0-canonical with offset i.*

Proof. By the distributive law $\varphi \lor (\psi_1 \ovee \psi_2) \equiv (\varphi \lor \psi_1) \ovee (\varphi \lor \psi_2)$, the duality $\sim(\psi_1 \ovee \psi_2) \equiv \sim\psi_1 \land \sim\psi_2$, and the definition $\mathsf{E}\psi = \sim\neg\psi$,

$$\sim\bigvee_{p\in\Phi}(\Diamond^i p \ovee \Diamond^i\neg p) \equiv \sim\bigovee_{P\subseteq\Phi}\Big(\bigvee_{p\in P}\Diamond^i p \lor \bigvee_{p\in\Phi\setminus P}\Diamond^i\neg p\Big) \equiv \bigwedge_{P\subseteq\Phi}\mathsf{E}\Big(\bigwedge_{p\in P}\Box^i\neg p \land \bigwedge_{p\in\Phi\setminus P}\Box^i p\Big).$$

The rightmost formula now states that for all types $\tau \in \Delta_0$ (each represented by a subset of Φ, cf. Proposition 4.3), there exists a world $w \in T$ such that $[\![R^iw]\!]^{\Phi}_0 \subseteq \{\tau\}$. Likewise, $T \vDash \Diamond^i\top$ iff $R^iw \neq \emptyset$ for every $w \in T$. ☐

Based on this, k-canonicity with offset i is now recursively defined as ρ^i_k:

$$\begin{aligned}
\rho^i_0(\beta) &:= \beta \hookrightarrow \mathsf{max}_i\\
\rho^i_{k+1}(\alpha,\beta) &:= \forall^{\subseteq}_{\alpha}\exists^{\subseteq}_{\beta}\left(\rho^i_0(\beta) \land \Box^i\forall^1_\beta\,\Box\chi^*_k(\alpha,\beta)\right)\\
\mathsf{canon}_k &:= \rho^k_0(\mathfrak{s}_0) \land \bigwedge_{m=1}^{k}\rho^{k-m}_m(\mathfrak{s}_{m-1},\mathfrak{s}_m)
\end{aligned}$$

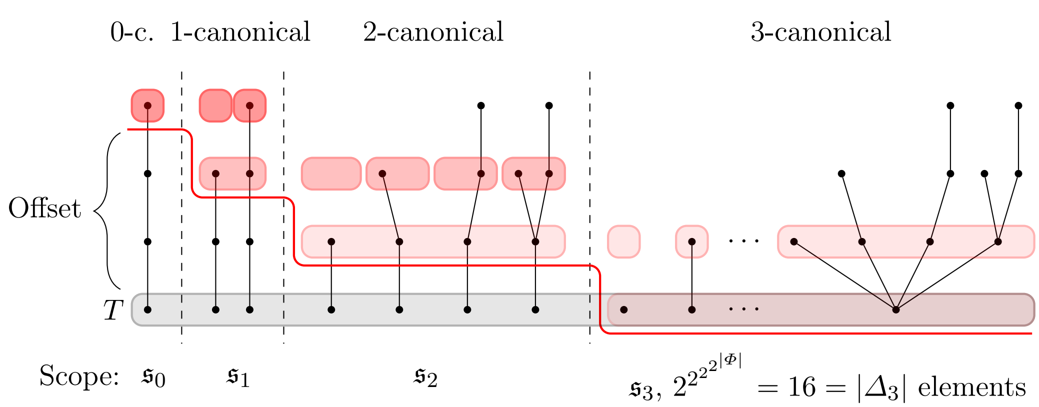

Figure 3: Visualization of the 3-staircase for $\Phi = \emptyset$, where the subteam $T_{\mathfrak{s}_i}$ is i-canonical with offset $3 - i$.

Theorem 6.5. *Let $k \geq 0$ and $\mathcal{K}$ be a structure with disjoint scopes $\mathfrak{s}_0, \dots, \mathfrak{s}_k$. Then $(\mathcal{K}, T) \vDash \mathsf{canon}_k$ if and only if $(\mathcal{K}, T)$ is a k-staircase. Moreover, canon_k is an MTL_k-formula constructible in space $\mathcal{O}(\log(|\Phi| + k))$.*

Proof. Similar to Theorem 5.4, the construction of the above formula in logspace is straightforward. We proceed with the correctness of the formula. Suppose that $\mathfrak{s}_0, \dots, \mathfrak{s}_k$ are disjoint scopes in $\mathcal{K}$. We show the following by induction on $0 \leq i \leq k$: Assuming that T_α is k-canonical with offset $i+1$, it holds that T_β is $(k+1)$-canonical with offset i if and only if $T \vDash \rho^i_{k+1}(\alpha, \beta)$. With the induction basis done in Lemma 6.4, the inductive step is proved by the following equivalence:

$$T_\beta \text{ is } (k+1)\text{-canonical with offset } i$$
$$\Leftrightarrow \forall \tau \in \Delta_{k+1} \colon \exists w \in T_\beta \colon [\![R^i w]\!]_{k+1} = \{\tau\}$$

Using the inverse of the bijection $h \colon \tau \mapsto (\tau \cap \Phi, \mathcal{R}\tau)$ from Proposition 4.3, we can equivalently quantify over $\mathfrak{P}(\Delta_k)$ and $\mathfrak{P}(\Phi)$:

$$\Leftrightarrow \forall \Delta' \subseteq \Delta_k \colon \forall \Phi' \subseteq \Phi \colon \exists w \in T_\beta \colon [\![R^i w]\!]_{k+1} = \{h^{-1}(\Phi', \Delta')\}$$
$$\Leftrightarrow \forall \Delta' \subseteq \Delta_k \colon \forall \Phi' \subseteq \Phi \colon \exists w \in T_\beta \colon R^i w \neq \emptyset \text{ and } \forall v \in R^i w \colon [\![v]\!]_{k+1} = h^{-1}(\Phi', \Delta')$$

By Lemma 4.4, $V^{-1}(v) = \Phi'$ and $[\![Rv]\!]_k = \Delta'$ is equivalent to $[\![v]\!]_{k+1} = h^{-1}(\Phi', \Delta')$:

$$\Leftrightarrow \forall \Delta' \subseteq \Delta_k \colon \forall \Phi' \subseteq \Phi \colon \exists w \in T_\beta \colon R^i w \neq \emptyset$$
$$\text{and } \forall v \in R^i w \colon V^{-1}(v) = \Phi' \text{ and } [\![Rv]\!]_k = \Delta'$$

Again by Proposition 4.3, $h \colon \tau \mapsto \tau \cap \Phi$ is a bijection from Δ_0 to $\mathfrak{P}(\Phi)$:

$$\Leftrightarrow \forall \Delta' \subseteq \Delta_k \colon \forall \tau_0 \in \Delta_0 \colon \exists w \in T_\beta \colon R^i w \neq \emptyset$$
$$\text{and } \forall v \in R^i w \colon V^{-1}(v) = \tau_0 \cap \Phi \text{ and } [\![Rv]\!]_k = \Delta'$$

Once more by Lemma 4.4:

$$\Leftrightarrow \forall \Delta' \subseteq \Delta_k \colon \forall \tau_0 \in \Delta_0 \colon \exists w \in T_\beta \colon R^i w \neq \emptyset$$
$$\text{and } \forall v \in R^i w \colon [\![v]\!]_0 = \tau_0 \text{ and } [\![Rv]\!]_k = \Delta'$$
$$\Leftrightarrow \forall \Delta' \subseteq \Delta_k \colon \forall \tau_0 \in \Delta_0 \colon \exists w \in T_\beta \colon [\![R^i w]\!]_0 = \{\tau_0\} \text{ and } \forall v \in R^i w \colon [\![Rv]\!]_k = \Delta'$$

Since T_α is assumed k-canonical with offset $i+1$, for every $\tau' \in \Delta_k$ there exists $u \in T_\alpha$ such that $[\![R^{i+1} u]\!]_k = \{\tau'\}$. Accordingly, for every *set* $\Delta' \subseteq \Delta_k$ there exists $S \subseteq T_\alpha$ such that $[\![R^{i+1} S]\!]_k = \Delta'$:

$$\Leftrightarrow \forall S \subseteq T_\alpha \colon \forall \tau_0 \in \Delta_0 \colon \exists w \in T_\beta \colon [\![R^i w]\!]_0 = \{\tau_0\} \text{ and } \forall v \in R^i w \colon [\![Rv]\!]_k = [\![R^{i+1} S]\!]_k$$

For each S, gather the respective w in a team $U \subseteq T_\beta$:

$$\Leftrightarrow \forall S \subseteq T_\alpha \colon \exists U \subseteq T_\beta \colon \big(\forall \tau_0 \in \Delta_0 \colon \exists w \in U \colon [\![R^i w]\!]_0 = \{\tau_0\}\big)$$
$$\text{and } \forall v \in R^i U \colon [\![Rv]\!]_k = [\![R^{i+1} S]\!]_k$$
$$\Leftrightarrow \forall S \subseteq T_\alpha : \exists U \subseteq T_\beta \colon U \text{ is 0-canonical with offset } i$$
$$\text{and } \forall v \in R^i U \colon [\![Rv]\!]_k = [\![R^{i+1} S]\!]_k$$

By the base case $k = 0$, and since $U = (T^{\alpha,\beta}_{S,U})_\beta$:

$$\Leftrightarrow \forall S \subseteq T_\alpha \colon \exists U \subseteq T_\beta \colon T^{\alpha,\beta}_{S,U} \vDash \rho^i_0(\beta) \text{ and } \forall v \in R^i U \colon [\![Rv]\!]_k = [\![R^{i+1} S]\!]_k$$

By Theorem 5.4:

$$\Leftrightarrow \forall S \subseteq T_\alpha \colon \exists U \subseteq T_\beta \colon T^{\alpha,\beta}_{S,U} \vDash \rho^i_0(\beta) \text{ and } \forall v \in R^i U \colon (R^{i+1}T)^{\alpha,\beta}_{R^{i+1}S,Rv} \vDash \chi^*_k(\alpha,\beta)$$

By Proposition 5.2 (5.):

$$\Leftrightarrow \forall S \subseteq T_\alpha \colon \exists U \subseteq T_\beta \colon T^{\alpha,\beta}_{S,U} \vDash \rho^i_0(\beta) \text{ and } \forall v \in R^i U \colon (R^i T)^{\alpha,\beta}_{R^i S,v} \vDash \Box\chi^*_k(\alpha,\beta)$$

By Proposition 5.3 applied to $(R^iT)^{\alpha,\beta}_{R^iS,R^iU}$:

$$\Leftrightarrow \forall S \subseteq T_\alpha \colon \exists U \subseteq T_\beta \colon T^{\alpha,\beta}_{S,U} \vDash \rho^i_0(\beta) \text{ and } (R^iT)^{\alpha,\beta}_{R^iS,R^iU} \vDash \forall^1_\beta \Box\chi^*_k(\alpha,\beta)$$

Again by Proposition 5.2 (5.) and Proposition 5.3:

$$\Leftrightarrow \forall S \subseteq T_\alpha \colon \exists U \subseteq T_\beta \colon T^{\alpha,\beta}_{S,U} \vDash \rho^i_0(\beta) \text{ and } R^i\left(T^{\alpha,\beta}_{S,U}\right) \vDash \forall^1_\beta \Box\chi^*_k(\alpha,\beta)$$

$$\Leftrightarrow \forall S \subseteq T_\alpha \colon \exists U \subseteq T_\beta \colon T^{\alpha,\beta}_{S,U} \vDash \rho^i_0(\beta) \wedge \Box^i \forall^1_\beta \Box\chi^*_k(\alpha,\beta)$$

$$\Leftrightarrow T \vDash \forall^{\subseteq}_\alpha \exists^{\subseteq}_\beta \left(\rho^i_0(\beta) \wedge \Box^i \forall^1_\beta \Box\chi^*_k(\alpha,\beta)\right)$$

$$\Leftrightarrow T \vDash \rho^i_{k+1}(\alpha,\beta).$$

□

6.3. **Enforcing scopes.** As the next step, we lift the restriction of the $\mathfrak{s}_i$ being scopes *a priori*. In a sense, this condition is definable in MTL as well. For this, let $\Psi \subseteq \mathcal{PS}$ be disjoint from Φ. Then the formula below ensures that Ψ is a set of disjoint scopes "up to height k".

$$\mathsf{scopes}_k(\Psi) := \bigwedge_{\substack{x,y\in\Psi\\ x\neq y}} \neg(x \wedge y) \wedge \bigwedge_{i=1}^{k} \Big((x \wedge \Box^i x) \vee (\neg x \wedge \Box^i \neg x)\Big).$$

The definition up to height k is sufficient for our purposes, which follows from the next lemma.

Lemma 6.6. *If $\varphi \in \mathsf{MTL}_k$, then φ is satisfiable if and only if $\varphi \wedge \Box^{k+1}\bot$ is satisfiable.*

Proof. As the direction from right to left is trivial, suppose φ is satisfiable. By Corollary 6.3, it then has a model $(\mathcal{K},T)$ that is a directed forest of height at most k. But then $(\mathcal{K},T) \vDash \Box^{k+1}\bot$, since $R^{k+1}T = \emptyset$ and $(\mathcal{K},\emptyset)$ satisfies all ML-formulas, including $\bot$. □

Theorem 6.7. $\mathsf{canon}_k \wedge \mathsf{scopes}_k(\{\mathfrak{s}_0,\ldots,\mathfrak{s}_k\}) \wedge \Box^{k+1}\bot$ *is satisfiable, but has only k-staircases as models.*

Proof. By combining Proposition 6.2, Theorem 6.5 and Lemma 6.6, the formula is satisfiable. Since in every model $(\mathcal{K},T)$ the propositions $\mathfrak{s}_0,\ldots,\mathfrak{s}_k$ must be disjoint scopes due to $\Box^{k+1}\bot$ and scopes_k, we can apply Theorem 6.5. □

As for bisimilarity, it is open whether (Φ,k)-canonicity can be defined in MTL^Φ_k *efficiently* without restricting the models to those with scopes. Note that the results of this section alone do not imply that the brute force algorithm given in Theorem 4.10 is optimal, as there could possibly be a satisfiability algorithm that does not need to construct a model. To show proper complexity theoretic hardness, we need to encode non-elementary computations in such models, to which we will proceed in the next sections.

7. Defining an order on types

In the previous section, we enforced k-canonicity with a formula, i.e., such that $|\Delta_k|$ different types are contained in the team. In order to encode computations of length $|\Delta_k|$, we additionally need to be able to talk about an ordering of Δ_k.

Let us call any finite strict linear ordering simply an *order*. We specify an order $\prec_k$ on Δ_k, and analogously to team bisimilarity, an order $\prec_k^*$ on $\mathfrak{P}(\Delta_k)$. To begin with, let us first agree on some arbitrary order $<$ on Φ, say, $p_1 < p_2 < \cdots < p_{|\Phi|}$. Furthermore, if $\sqsubset$ is some order on X, then the *lexicographic order* $\sqsubset^*$ on $\mathfrak{P}(X)$ is defined by

$$X_1 \sqsubset^* X_2 \text{ iff } \exists x \in X_2 \setminus X_1 \text{ such that } \forall x' \in X \colon (x \sqsubset x') \Rightarrow (x' \in X_1 \Leftrightarrow x' \in X_2).$$

For example, let $X = \{0,1\}$ and $0 \sqsubset 1$. Then $\emptyset \sqsubset^* \{0\} \sqsubset^* \{1\} \sqsubset^* \{0,1\}$. The order $\prec_k$ depends on the propositions true in a world, and otherwise recursively on the lexicographic order of the image team:

$$\begin{aligned}
\tau \prec_0 \tau' &\Leftrightarrow \tau \cap \Phi <^* \tau' \cap \Phi,\\
\tau \prec_{k+1} \tau' &\Leftrightarrow \tau \cap \Phi <^* \tau' \cap \Phi \text{ or } (\tau \cap \Phi = \tau' \cap \Phi \text{ and } \mathcal{R}\tau \ \prec_k^* \ \mathcal{R}\tau').
\end{aligned}$$

It is easy to verify by induction that $\prec_k$ and $\prec_k^*$ are orders on Δ_k and $\mathfrak{P}(\Delta_k)$, respectively.

The next step is to prove that $\prec_k$ and $\prec_k^*$ are (efficiently) definable in MTL_k. For this, we pursue the same approach as for χ_k and χ_k^* in Section 5, and show that $\prec_k$ and $\prec_k^*$ are definable in formulas ζ_k and ζ_k^* in a mutually recursive fashion. Since order is a binary relation, the formulas below are once more parameterized by two scopes.

$$\zeta_0(\alpha,\beta) := \bigvee_{p\in\Phi} \Big[(\alpha \hookrightarrow \neg p) \wedge (\beta \hookrightarrow p) \wedge \bigwedge_{\substack{q\in\Phi\\ q<p}} (\alpha \vee \beta) \hookrightarrow {=}(q)\Big]$$

$$\zeta_{k+1}(\alpha,\beta) := \zeta_0(\alpha,\beta) \ \otimes\ \chi_0(\alpha,\beta) \wedge \Box\zeta_k^*(\alpha,\beta)$$

$$\begin{aligned}
\zeta_k^*(\alpha,\beta) := {}& \exists^1_{\mathfrak{s}_k} \left(\exists^1_\beta \chi_k(\mathfrak{s}_k,\beta)\right) \wedge \left({\sim}\exists^1_\alpha \chi_k(\mathfrak{s}_k,\alpha)\right)\\
&\wedge \Big(\left(\chi_k^*(\alpha,\beta) \wedge (\alpha \vee \beta)\right) \vee \left(\forall^1_{\alpha\vee\beta} {\sim}\zeta_k(\mathfrak{s}_k, \alpha\vee\beta)\right)\Big)
\end{aligned}$$

Note that we make use of the scopes $\mathfrak{s}_0,\dots,\mathfrak{s}_k$ in the formula, and in the following we restrict ourselves to k-staircase models. Moreover, in the subformula $\zeta_k(\mathfrak{s}_k, \alpha\vee\beta)$, we use the fact that $\alpha \vee \beta$ is a scope whenever α, β are scopes.

We require the next lemma for the correctness of ζ_k and ζ_k^*. Intuitively, it states that MTL_k is invariant under substitution of "locally equivalent" ML-formulas.

Lemma 7.1. *Let $\alpha, \beta \in \mathsf{ML}$ and $\varphi \in \mathsf{MTL}_k$. Let T be a team such that $R^iT \vDash \alpha \leftrightarrow \beta$ for all $i \in \{0,\dots,k\}$. Then $T \vDash \varphi$ if and only if $T \vDash \mathrm{Sub}(\varphi,\alpha,\beta)$, where $\mathrm{Sub}(\varphi,\alpha,\beta)$ is the formula obtained from φ by substituting every occurrence of α with β.*

Proof. By straightforward induction; see the appendix. □

The following theorem states that in the class of k-staircase models (see the previous section) ζ_k and ζ_k^* define the required orders.

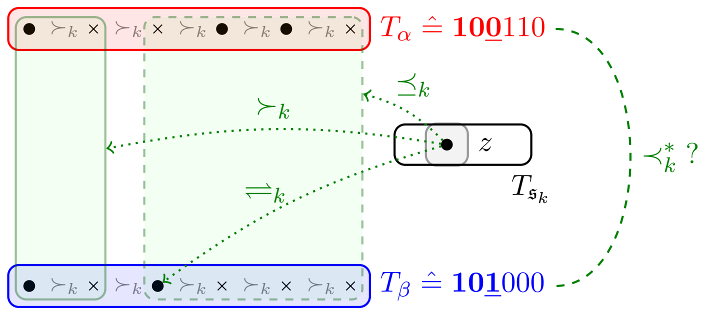

Figure 4: The pivot $z \in T_{\mathfrak{s}_k}$ determines that $[\![T_\alpha]\!]_k \prec_k^* [\![T_\beta]\!]_k$. The subteam of $T_{\alpha\vee\beta}$ of worlds $\prec_k$-greater than z must satisfy $\chi_k^*(\alpha,\beta)$.

Theorem 7.2. *Let $k \geq 0$, and let $(\mathcal{K}, T)$ be a k-staircase with disjoint scopes $\alpha, \beta, \mathfrak{s}_0, \ldots, \mathfrak{s}_k$. If $w \in T_\alpha$ and $v \in T_\beta$, then*

$$T_{w,v}^{\alpha,\beta} \vDash \zeta_k(\alpha,\beta) \quad \textit{if and only if} \quad [\![w]\!]_k \prec_k [\![v]\!]_k,$$
$$T \vDash \zeta_k^*(\alpha,\beta) \quad \textit{if and only if} \quad [\![T_\alpha]\!]_k \prec_k^* [\![T_\beta]\!]_k.$$

Furthermore, both $\zeta_k(\alpha,\beta)$ and $\zeta_k^*(\alpha,\beta)$ are MTL_k-formulas that are constructible in space $\mathcal{O}(\log(k + |\Phi| + |\alpha| + |\beta|))$.

We first give a rough idea of the proof, and after a series of required lemmas fully prove the theorem. The definition of ζ_{k+1} simply follows the definition of $\prec_{k+1}$. Furthermore, the formula ζ_k^* implements the lexicographic order $\prec_k^*$ as follows. As shown in Figure 4, we first choose some $z \in T_{\mathfrak{s}_k}$ that acts as an *pivot* to determine if $[\![T_\alpha]\!]_k \prec_k^* [\![T_\beta]\!]$, in the sense that it is the $\prec_k$-maximal type in which T_α and T_β differ.[3] The first line of ζ_k^* indeed expresses that $[\![z]\!]_k \in [\![T_\beta]\!]_k \setminus [\![T_\alpha]\!]_k$.

The disjunction in the second line intuitively states that we then can "split off" the subteam of $T_\alpha \cup T_\beta$ consisting of the elements $\prec_k$-greater than z (the solid green area in Figure 4), while χ_k^* ensures that they agree on the contained types (this reflects the part after the quantifier in the definition of $\sqsubset^*$). To achieve this, the subformula $\forall^1_{\alpha\vee\beta}{\sim}\zeta_k(\mathfrak{s}_k, \alpha\vee\beta)$ stipulates that any "remaining" elements from $T_\alpha \cup T_\beta$ possess only types not $\prec_k$-greater than $[\![z]\!]_k$ (the dashed green area in the figure).

Here, Lemma 7.1 is applied, as it ensures that after processing $\forall^1_{\alpha\vee\beta}$ the formula $\zeta_k(\mathfrak{s}_k, \alpha\vee\beta)$ in fact behaves as either $\zeta_k(\mathfrak{s}_k, \alpha)$ or $\zeta_k(\mathfrak{s}_k, \beta)$; and hence behaves correctly by induction hypothesis.

Definition 7.3. Let $k \geq 0$. Let α, β be disjoint scopes and T a team in a Kripke structure. Then α and β are called *$\prec_k$-comparable in T* if for all $w \in T_\alpha, v \in T_\beta$

$$T_{w,v}^{\alpha,\beta} \vDash \zeta_k(\alpha,\beta) \text{ iff } [\![w]\!]_k \prec_k [\![v]\!]_k \text{ and}$$
$$T_{w,v}^{\alpha,\beta} \vDash \zeta_k(\beta,\alpha) \text{ iff } [\![v]\!]_k \prec_k [\![w]\!]_k.$$

Likewise, α and β are *$\prec_k^*$-comparable in T* if

$$T \vDash \zeta_k^*(\alpha,\beta) \text{ iff } [\![T_\alpha]\!]_k \prec_k^* [\![T_\beta]\!]_k \text{ and}$$
$$T \vDash \zeta_k^*(\beta,\alpha) \text{ iff } [\![T_\beta]\!]_k \prec_k^* [\![T_\alpha]\!]_k.$$

[3] Since the pivot is selected from $T_{\mathfrak{s}_k}$, at this point it is crucial that the underlying structure is a k-staircase.

The next lemma shows that the correctness of $\prec_k^*$ follows from that of $\prec_k$.

Lemma 7.4. *Suppose that $(\mathcal{K}, T)$ is a k-staircase with disjoint scopes $\alpha, \beta, \mathfrak{s}_0, \ldots, \mathfrak{s}_k$. If both α and β are $\prec_k$-comparable to $\mathfrak{s}_k$ in all subteams S of the form $T_{\mathfrak{s}_0} \cup \cdots \cup T_{\mathfrak{s}_{k-1}} \subseteq S \subseteq T$, then α and β are $\prec_k^*$-comparable in T.*

Proof. Assuming $\mathcal{K}, T, \alpha, \beta, \mathfrak{s}_0, \ldots, \mathfrak{s}_k$ as above, the proof is split into the following claims.

Claim (a)**.** *The disjoint scopes $\alpha \vee \beta$ and $\mathfrak{s}_k$ are $\prec_k$-comparable in any team S that satisfies $T_{\mathfrak{s}_0} \cup \cdots \cup T_{\mathfrak{s}_{k-1}} \subseteq S \subseteq T$.*

Proof of claim. Let $w \in S_{\alpha\vee\beta}$ and $v \in S_{\mathfrak{s}_k}$. W.l.o.g. $w \in S_\alpha$ (the case $w \in S_\beta$ works analogously). Then

$$
\begin{aligned}
& S^{\alpha\vee\beta,\mathfrak{s}_k}_{w,v} \vDash \zeta_k(\alpha \vee \beta, \mathfrak{s}_k) \\
\Leftrightarrow\ & S^{\alpha,\beta,\mathfrak{s}_k}_{w,\emptyset,v} \vDash \zeta_k(\alpha \vee \beta, \mathfrak{s}_k) && \text{(Since } S^{\alpha\vee\beta,\mathfrak{s}_k}_{w,v} = S^{\alpha,\beta,\mathfrak{s}_k}_{w,\emptyset,v}\text{)} \\
\Leftrightarrow\ & S^{\alpha,\beta,\mathfrak{s}_k}_{w,\emptyset,v} \vDash \zeta_k(\alpha, \mathfrak{s}_k) && \text{(By Lemma 7.1, as } \textstyle\bigcup_{i=0}^k R^i S^{\alpha,\beta,\mathfrak{s}_k}_{w,\emptyset,v} \vDash \alpha \leftrightarrow (\alpha \vee \beta)\text{)} \\
\Leftrightarrow\ & [\![w]\!]_k \prec_k [\![v]\!]_k. && \text{(By assumption of the lemma)}
\end{aligned}
$$

The case $\zeta_k(\mathfrak{s}_k, \alpha \vee \beta)$ is symmetric. ◁

For the remaining proof, we omit the subscript k when referring to types and $\prec$. Furthermore, for all $\tau \in \Delta_k$, let $[\![T]\!]^\tau$ denote the restriction of $[\![T]\!]$ to types τ' such that $\tau' \succ \tau$. Intuitively, these types are the "more significant positions" for the lexicographic ordering. In the next claim, we essentially show that the second line in the definition of $\zeta_k^*(\alpha, \beta)$ can be expressed as a statement of the form $[\![T_\alpha]\!]^\tau = [\![T_\beta]\!]^\tau$.

Claim (b)**.** *Let T be a team and $\tau \in \Delta_k$. Then $[\![T_\alpha]\!]^\tau = [\![T_\beta]\!]^\tau$ if and only if there exists $S \subseteq T_{\alpha\vee\beta}$ such that $[\![S_\alpha]\!] = [\![S_\beta]\!]$ and $[\![w]\!] \not\succ \tau$ for all $w \in T_{\alpha\vee\beta} \setminus S$.*

Proof of claim. "$\Rightarrow$": Let $S := \{v \in T_{\alpha\vee\beta} \mid [\![v]\!] \succ \tau\}$. Then $[\![S_\alpha]\!] = [\![T_\alpha]\!]^\tau = [\![T_\beta]\!]^\tau = [\![S_\beta]\!]$. Moreover, for every $w \in T_{\alpha\vee\beta} \setminus S$ clearly $[\![w]\!] \not\succ \tau$ holds.

"$\Leftarrow$": Assume that S exists as stated in the claim. By symmetry, we only prove $[\![T_\alpha]\!]^\tau \subseteq [\![T_\beta]\!]^\tau$. Consequently, let $w \in T_\alpha$ such that $[\![w]\!] \in [\![T_\alpha]\!]^\tau$. Then $[\![w]\!] \succ \tau$ by definition. But then $w \notin T_{\alpha\vee\beta} \setminus S$. However, we have $w \in T_\alpha$, hence $w \in T_{\alpha\vee\beta}$, which only leaves the possibility $w \in S$. Combining $w \in S$ and $w \in T_\alpha$ yields $w \in S_\alpha$, which by assumption also implies $[\![w]\!] \in [\![S_\beta]\!]$. As $[\![S_\beta]\!] \subseteq [\![T_\beta]\!]$ and $[\![w]\!] \succ \tau$, the membership $[\![w]\!] \in [\![T_\beta]\!]^\tau$ follows. ◁

Claim (c)**.** *α and β are $\prec_k^*$-comparable in T.*

Proof of claim. Due to symmetry, we prove only that $T \vDash \zeta_k^*(\alpha, \beta)$ iff $[\![T_\alpha]\!]_k \prec_k^* [\![T_\beta]\!]_k$.

$$
\begin{aligned}
& [\![T_\alpha]\!] \prec^* [\![T_\beta]\!] \\
\Leftrightarrow\ & \exists \tau \in [\![T_\beta]\!] \setminus [\![T_\alpha]\!] \colon \forall \tau' \in \Delta, \tau \prec \tau' \colon \tau' \in [\![T_\alpha]\!] \Leftrightarrow \tau' \in [\![T_\beta]\!] && \text{(Definition of } \prec_k^*\text{)} \\
\Leftrightarrow\ & \exists \tau \in [\![T_\beta]\!] \setminus [\![T_\alpha]\!] \colon [\![T_\alpha]\!]^\tau = [\![T_\beta]\!]^\tau && \text{(Definition of } [\![\cdot]\!]^\tau\text{)}
\end{aligned}
$$

Since $T_{\mathfrak{s}_k}$ is k-canonical, for every $\tau \in \Delta$ there exists $z \in T_{\mathfrak{s}_k}$ of type τ:

$$
\begin{aligned}
\Leftrightarrow\ & \exists z \in T_{\mathfrak{s}_k} \colon [\![T_\alpha]\!]^{[\![z]\!]} = [\![T_\beta]\!]^{[\![z]\!]} \text{ and } [\![z]\!] \in [\![T_\beta]\!] \setminus [\![T_\alpha]\!] \\
\Leftrightarrow\ & \exists z \in T_{\mathfrak{s}_k} \colon [\![T_\alpha]\!]^{[\![z]\!]} = [\![T_\beta]\!]^{[\![z]\!]} \text{ and } \exists x \in T_\beta \colon [\![z]\!] = [\![x]\!] \text{ and } \nexists y \in T_\alpha \colon [\![z]\!] = [\![y]\!]
\end{aligned}
$$

As α, β and $\mathfrak{s}_k$ are disjoint, we have $T_\alpha = O_\alpha$, where $O := T_z^{\mathfrak{s}_k}$, and likewise $T_\beta = O_\beta$:

$$\Leftrightarrow \exists z \in T_{\mathfrak{s}_k} \colon [\![O_\alpha]\!]^{[\![z]\!]} = [\![O_\beta]\!]^{[\![z]\!]} \text{ and } \exists x \in O_\beta \colon [\![z]\!] = [\![x]\!] \text{ and } \nexists y \in O_\alpha \colon [\![z]\!] = [\![y]\!]$$

$$\Leftrightarrow \exists z \in T_{\mathfrak{s}_k} \colon \exists x \in O_\beta \colon [\![z]\!] = [\![x]\!] \text{ and } \nexists y \in O_\alpha \colon [\![z]\!] = [\![y]\!]$$
$$\text{and } \exists S \subseteq O_{\alpha\vee\beta} \colon [\![S_\alpha]\!] = [\![S_\beta]\!] \text{ and } \forall w \in O_{\alpha\vee\beta} \setminus S \colon [\![z]\!] \not\prec [\![w]\!] \qquad \text{(by Claim (b))}$$

Clearly S is a subteam of $O_{\alpha\vee\beta}$ if and only if it is a subteam of O and satisfies $\alpha \vee \beta$:

$$\Leftrightarrow \exists z \in T_{\mathfrak{s}_k} \colon \exists x \in O_\beta \colon [\![z]\!] = [\![x]\!] \text{ and } \nexists y \in O_\alpha \colon [\![z]\!] = [\![y]\!]$$
$$\text{and } \exists S \subseteq O \colon [\![S_\alpha]\!] = [\![S_\beta]\!] \text{ and } S \vDash \alpha \vee \beta \text{ and } \forall w \in O_{\alpha\vee\beta} \setminus S \colon [\![z]\!] \not\prec [\![w]\!]$$

Letting $U = O \setminus S$, we have $O_{\alpha\vee\beta} \setminus S = U_{\alpha\vee\beta}$:

$$\Leftrightarrow \exists z \in T_{\mathfrak{s}_k} \colon \exists x \in O_\beta \colon [\![z]\!] = [\![x]\!] \text{ and } \nexists y \in O_\alpha \colon [\![z]\!] = [\![y]\!] \text{ and } \exists S \subseteq O \colon$$
$$[\![S_\alpha]\!] = [\![S_\beta]\!] \text{ and } S \vDash \alpha \vee \beta \text{ and } \exists U \subseteq O \colon U = O \setminus S \text{ and } \forall w \in U_{\alpha\vee\beta} \colon [\![z]\!] \not\prec [\![w]\!]$$

Clearly, the property $\forall w \in U_{\alpha\vee\beta} : [\![z]\!] \not\prec [\![w]\!]$ is preserved when taking subteams of U. Hence, $U = O \setminus S$ satisfies it if and only if some (not necessarily proper) superteam U' of $O \setminus S$ does:

$$\Leftrightarrow \exists z \in T_{\mathfrak{s}_k} \colon \exists x \in O_\beta \colon [\![z]\!] = [\![x]\!] \text{ and } \nexists y \in O_\alpha \colon [\![z]\!] = [\![y]\!]$$
$$\text{and } \exists S \subseteq O \colon [\![S_\alpha]\!] = [\![S_\beta]\!] \text{ and } S \vDash \alpha \vee \beta$$
$$\text{and } \exists U' \subseteq O \colon U' \supseteq O \setminus S \text{ and } \forall w \in U'_{\alpha\vee\beta} \colon [\![z]\!] \not\prec [\![w]\!]$$

By Theorem 5.4:

$$\Leftrightarrow \exists z \in T_{\mathfrak{s}_k} \colon O \vDash (\exists^1_\beta \chi_k(\mathfrak{s}, \beta)) \wedge ({\sim}\exists^1_\alpha \chi_k(\mathfrak{s}, \alpha)) \text{ and } \exists S \subseteq O \colon$$
$$S \vDash (\alpha \vee \beta) \wedge \chi_k^*(\alpha, \beta) \text{ and } \exists U' \subseteq O \colon U' \supseteq O \setminus S \text{ and } \forall w \in U'_{\alpha\vee\beta} \colon [\![z]\!] \not\prec [\![w]\!]$$

Note that $T_{\mathfrak{s}_0}, \ldots, T_{\mathfrak{s}_{k-1}}$ are retained in O. Moreover, $S \subseteq O_{\alpha\vee\beta}$, which implies that they are still subteams of $O \setminus S$ and hence of U'. But by Claim (a), $\alpha \vee \beta$ and $\mathfrak{s}_k$ are then $\prec_k$-comparable scopes in U' and we can replace $[\![z]\!] \not\prec [\![w]\!]$:

$$\Leftrightarrow \exists z \in T_{\mathfrak{s}_k} \colon O \vDash (\exists^1_\beta \chi_k(\mathfrak{s}, \beta)) \wedge ({\sim}\exists^1_\alpha \chi_k(\mathfrak{s}, \alpha))$$
$$\text{and } \exists S \subseteq O \colon S \vDash (\alpha \vee \beta) \wedge \chi_k^*(\alpha, \beta)$$
$$\text{and } \exists U' \subseteq O \colon U' \supseteq O \setminus S \text{ and } \forall w \in U'_{\alpha\vee\beta} \colon (U')_w^{\alpha\vee\beta} \vDash {\sim}\zeta_k(\mathfrak{s}_k, \alpha \vee \beta)$$

Recalling that $O = T_z^{\mathfrak{s}_k}$, and by Proposition 5.3, we obtain:

$$\Leftrightarrow \exists z \in T_{\mathfrak{s}_k} \colon T_z^{\mathfrak{s}_k} \vDash (\exists^1_\beta \chi_k(\mathfrak{s}, \beta)) \wedge ({\sim}\exists^1_\alpha \chi_k(\mathfrak{s}, \alpha))$$
$$\text{and } \exists S \subseteq T_z^{\mathfrak{s}_k} \colon S \vDash (\alpha \vee \beta) \wedge \chi_k^*(\alpha, \beta)$$
$$\text{and } \exists U' \subseteq T_z^{\mathfrak{s}_k} \colon U' \supseteq T_z^{\mathfrak{s}_k} \setminus S \text{ and } U' \vDash \forall^1_{\alpha\vee\beta} {\sim}\zeta_k(\mathfrak{s}_k, \alpha \vee \beta)$$
$$\Leftrightarrow T \vDash \exists^1_{\mathfrak{s}_k} (\exists^1_\beta \chi_k(\mathfrak{s}, \beta)) \wedge ({\sim}\exists^1_\alpha \chi_k(\mathfrak{s}, \alpha))$$
$$\wedge \big((\alpha \vee \beta) \wedge \chi_k^*(\alpha, \beta)\big) \vee \big(\forall^1_{\alpha\vee\beta} {\sim}\zeta_k(\mathfrak{s}_k, \alpha \vee \beta)\big)$$
$$\Leftrightarrow T \vDash \zeta^*(\alpha, \beta). \qquad \triangleleft \quad \square$$

In the next lemma, we prove the converse direction of Lemma 7.4.

Lemma 7.5. *Let $k > 0$, and let $(\mathcal{K}, T)$ be a k-staircase with disjoint scopes $\alpha, \beta, \mathfrak{s}_0, \ldots, \mathfrak{s}_{k-1}$. Then α and β are $\prec_k$-comparable in every subteam S of T that contains $T_{\mathfrak{s}_0} \cup \cdots \cup T_{\mathfrak{s}_{k-1}}$.*

Proof. The proof is by induction on k. Disjoint scopes α and β are always $\prec_0$-comparable, which can be easily seen in ζ_0. For the inductive step to $k+1$, assume $(\mathcal{K}, T)$ and S as above, and let $\mathcal{K} = (W, R, V)$. Let $O := S^{\alpha,\beta}_{w,v}$ with $w \in S_\alpha, v \in S_\beta$ arbitrary.

Claim (a). *α and β are $\prec^*_k$-comparable in RO.*

Proof of claim. In the inductive step, now $\mathfrak{s}_0, \ldots, \mathfrak{s}_k, \alpha, \beta$ are disjoint scopes. Additionally, $(\mathcal{K}, RT)$ is a k-staircase. In particular, in the induction step α and β are disjoint from $\mathfrak{s}_k$. For this reason, $(\mathcal{K}, RO)$ is a k-staircase as well, as $(RO)_{\mathfrak{s}_0 \vee \cdots \vee \mathfrak{s}_k} = (RT)_{\mathfrak{s}_0 \vee \cdots \vee \mathfrak{s}_k}$.

Hence, by induction hypothesis, for every team U such that $RO_{\mathfrak{s}_0} \cup \cdots \cup RO_{\mathfrak{s}_{k-1}} \subseteq U \subseteq RO$, we obtain that $\mathfrak{s}_k$ and α are $\prec_k$-comparable in U, as well as $\mathfrak{s}_k$ and β. Consequently, we can apply Lemma 7.4, which proves the claim. ◁

We proceed with the induction step. Again by symmetry, we only show that $O \vDash \zeta_{k+1}(\alpha, \beta)$ iff $[\![w]\!]_{k+1} \prec_{k+1} [\![v]\!]_{k+1}$. We distinguish three cases w. r. t. $\prec_0$:

- If $[\![w]\!]_0 \prec_0 [\![v]\!]_0$, then $O \vDash \zeta_0(\alpha, \beta)$ by the induction basis. As the former implies $[\![w]\!]_{k+1} \prec_{k+1} [\![v]\!]_{k+1}$ and the latter $O \vDash \zeta_{k+1}(\alpha, \beta)$, the equivalence holds.
- If $[\![w]\!]_0 \succ_0 [\![v]\!]_0$, then $[\![w]\!]_{k+1} \not\prec_{k+1} [\![v]\!]_{k+1}$. Moreover, $O \nvDash \zeta_0(\alpha, \beta)$ by induction basis. Additionally, $O \nvDash \chi_0(\alpha, \beta)$ by Theorem 5.4. Consequently, both sides of the equivalence are false.
- If $[\![w]\!]_0 = [\![v]\!]_0$, then $O \vDash \chi_0(\alpha, \beta)$ by Theorem 5.4, but $O \nvDash \zeta_0(\alpha, \beta)$ by induction basis. Consequently, $O \vDash \zeta_{k+1}(\alpha, \beta)$ iff $O \vDash \Box\zeta^*_k(\alpha, \beta)$. Also, $[\![w]\!]_{k+1} \prec_{k+1} [\![v]\!]_{k+1}$ iff $\mathcal{R}[\![w]\!]_{k+1} \prec^*_k \mathcal{R}[\![v]\!]_{k+1}$. The following equivalence concludes the proof:

$$\begin{aligned} & \mathcal{R}[\![w]\!]_{k+1} \prec^*_k \mathcal{R}[\![v]\!]_{k+1} \\ \Leftrightarrow\; & [\![Rw]\!]_k \prec^*_k [\![Rv]\!]_k && \text{(By Proposition 4.3)} \\ \Leftrightarrow\; & RO \vDash \zeta^*_k(\alpha, \beta) && \text{(By Claim (a))} \\ \Leftrightarrow\; & O \vDash \Box\zeta^*_k(\alpha, \beta). \end{aligned}$$ □

With the above lemmas we are now in the position to prove Theorem 7.2:

Proof of Theorem 7.2. First, it is straightforward to construct ζ_k and ζ^*_k in logarithmic space. For the correctness, let $(\mathcal{K}, T)$ be a model with disjoint scopes $\alpha, \beta, \mathfrak{s}_0, \ldots, \mathfrak{s}_k$ as in the theorem. By Lemma 7.5 it immediately follows that α and β are $\prec_k$-comparable in T. The second part, that α and β are $\prec^*_k$-comparable in T, follows from the combination of Lemma 7.4 and 7.5. □

8. Encoding non-elementary computations

We combine all the previous sections and extend Theorem 4.10 and Corollary 4.11 by their matching lower bounds:

Theorem 8.1.

- $\mathsf{SAT}(\mathsf{MTL})$ *and* $\mathsf{VAL}(\mathsf{MTL})$ *are complete for* TOWER(poly).
- *If* $k \geq 0$*, then* $\mathsf{SAT}(\mathsf{MTL}_k)$ *and* $\mathsf{VAL}(\mathsf{MTL}_k)$ *are complete for* ATIME-ALT($\exp_{k+1}$, poly).

The above complexity classes are complement-closed, and additionally MTL and MTL_k are syntactically closed under negation. For this reason, it suffices to prove the hardness of $\mathsf{SAT}(\mathsf{MTL})$ and $\mathsf{SAT}(\mathsf{MTL}_k)$, respectively. Moreover, the case $k = 0$ is equivalent to $\mathsf{SAT}(\mathsf{PTL})$

being ATIME-ALT(exp, poly)-hard, which was proven by Hannula et al. [HKVV18]. Their reduction also works in logarithmic space. Consequently, the result boils down to the following lemma:

Lemma 8.2.

- *If* $L \in \mathrm{TOWER(poly)}$*, then* $L \leq_{\mathrm{m}}^{\log} \mathsf{SAT}(\mathsf{MTL})$*.*
- *If* $k \geq 1$ *and* $L \in \mathrm{ATIME\text{-}ALT}(\exp_{k+1}, \mathrm{poly})$*, then* $L \leq_{\mathrm{m}}^{\log} \mathsf{SAT}(\mathsf{MTL}_k)$*.*

We devise for each L a reduction $x \mapsto \varphi_x$ such that φ_x is a formula that is satisfiable if and only if $x \in L$. By assumption, there exists a single-tape alternating Turing machine M that decides L (for $L \in \mathrm{TOWER(poly)}$, w.l.o.g. M is alternating as well).

Let M have states Q, which is the disjoint union of $Q_\exists$ (*existential states*), $Q_\forall$ (*universal states*), Q_{acc} (*accepting states*) and Q_{rej} (*rejecting states*). Also, Q contains some initial state q_0. Let M have a finite tape alphabet Γ with blank symbol $\flat \in \Gamma$, and a transition relation δ.

We design φ_x in a fashion that forces its models $(\mathcal{K}, T)$ to encode an accepting computation of M on x. Let us call any legal sequence of configurations of M (not necessarily starting with the initial configuration) a *run*. Then, similarly as in Cook's theorem [Coo71], we encode runs as square "grids" with a vertical "time" coordinate and a horizontal "space" coordinate in the model, i.e., each row of the grid represents a configuration of M.

W.l.o.g. M never leaves the input to the left, and there exists N that is an upper bound on both the length of a configuration and the runtime of M. Formally, a run of M is then a function $C\colon \{1, \ldots, N\}^2 \to \Gamma \cup (Q \times \Gamma)$, Here, $C(i,j) = c$ for $c \in \Gamma$ means that the i-th configuration (i.e., after M performed $i-1$ transitions) contains the symbol c in its j-th cell. The same holds if $C(i,j) = (q,c)$ for $(q,c) \in Q \times \Gamma$, but then additionally the machine is in the state q with its head visiting the j-th cell in the i-th configuration. As an example, for a run C from M's initial configuration we have $C(1,1) = (q_0, x_1)$, $C(1,i) = x_i$ for $2 \leq i \leq n$, and $C(1,i) = \flat$ for $n < i \leq N$.

Due to the semantics of MTL, such a run must be encoded in $(\mathcal{K}, T)$ very carefully. We let the team T contain N^2 worlds $w_{i,j}$ in which the respective value of $C(i,j)$ is encoded as a propositional assignment. However, we cannot simply pursue the standard approach of assembling a large $N \times N$-grid in the edge relation R in order to compare successive configurations; by Corollary 6.3, we cannot force the model to contain R-paths longer than $|\varphi_x|$. Instead, to define grid neighborship, we let $w_{i,j}$ encode i and j in its *type*. More precisely, we use the linear order $\prec_k$ on Δ_k we defined with the MTL_k-formula ζ_k in the previous section. Then, instead of using $\Box$ and $\Diamond$, we examine the grid by letting ζ_k judge whether a given pair of worlds is deemed (horizontally or vertically) adjacent.

8.1. **Encoding runs in a team.** Next, we discuss how runs $C\colon \{1, \ldots, N\}^2 \to \Gamma \cup (Q \times \Gamma)$ are encoded in T. Given a world $w \in T$, we partition the image Rw with two special propositions $\mathfrak{t} \notin \Phi$ ("timestep") and $\mathfrak{p} \notin \Phi$ ("position"). Then we assign to w the pair $\ell(w) := (i,j)$ such that $[\![(Rw)_{\mathfrak{t}}]\!]_{k-1}$ is the i-th element, and $[\![(Rw)_{\mathfrak{p}}]\!]_{k-1}$ is the j-th element in the order $\prec_{k-1}^*$. We call the pair $\ell(w)$ the *location* of w (in the grid).

Accordingly, we fix $N := |\mathfrak{P}(\Delta_{k-1}^{\Phi})|$. For the case of fixed k, M has runtime bounded by $\exp_{k+1}(g(n))$ for a polynomial g. Then taking $\Phi := \{p_1, \ldots, p_{g(n)}\}$ yields a sufficiently large coordinate space, as

$$\exp_{k+1}(g(n)) = \exp_{k+1}(|\Phi|) = 2^{\exp_{k-1}(2^{|\Phi|})} \leq 2^{\exp^*_{k-1}(2^{|\Phi|})} = 2^{|\Delta^\Phi_{k-1}|} = N$$

by Proposition 4.7. For runtime $\exp_{g(n)}(1)$ of M, we let $\Phi := \emptyset$ and precompute $k := g(|x|)+1$, but otherwise proceed identically.

Next, let Ξ be a constant set of propositions disjoint from Φ that encodes the range of C via some bijection $c\colon \Xi \to \Gamma \cup (Q \times \Gamma)$. If a world w satisfies exactly one proposition p of those in Ξ, then by slight abuse of notation we write $c(w)$ instead of $c(p)$. Intuitively, $c(w) \in \Gamma \cup (Q \times \Gamma)$ is the content of the grid cell represented by w.

Using ℓ and c, the function C can be encoded into a team T as follows. First, a team T is called *grid* if every point in T satisfies exactly one proposition in Ξ, and if every location $(i,j) \in \{1,\ldots,N\}^2$ occurs as $\ell(w)$ for some point $w \in T$. Moreover, a grid T is called *pre-tableau* if for every location (i,j) *and* every element $p \in \Xi$ there is some world $w \in T$ such that $\ell(w) = (i,j)$ and $w \vDash p$. Finally, a grid T is a *tableau* if any two elements $w, w' \in T$ with $\ell(w) = \ell(w')$ also agree on Ξ, i.e., $c(w) = c(w')$.

Let us motivate the above definitions. Clearly, the definition of a *grid* T means that T captures the whole domain of C, and that c is well-defined on the level of *points*. If T is additionally a *tableau*, then c is also well-defined on the level of *locations*. In other words, a tableau T induces a function $C_T\colon \{1,\ldots,N\}^2 \to \Gamma \cup (Q \times \Gamma)$ via $C_\alpha(i,j) := c(w)$, where $w \in T$ is arbitrary such that $\ell(w) = (i,j)$.

A *pre-tableau* can be seen as the union of all possible C. In particular, given any pre-tableau, the definition ensures that arbitrary tableaus can be obtained from it by the means of subteam quantification $\exists^{\subseteq}$ (cf. p. 12).

A tableau T is *legal* if C_T is a run of M, i.e., if every row is a configuration of M, and if every pair of two successive rows represents a valid δ-transition.

The idea of the reduction is now to capture the alternating computation of M by nesting polynomially many quantifications (via $\exists^{\subseteq}$ and $\forall^{\subseteq}$) of legal tableaus, of which each one is the continuation of the computation of the previous one.

8.2. **Accessing two components of locations.** An discussed earlier, we choose to represent a location (i,j) in a point w as a pair (Δ', Δ'') by stipulating that $\Delta' = [\![(Rw)_{\mathfrak{t}}]\!]_{k-1}$ and $\Delta'' = [\![(Rw)_{\mathfrak{p}}]\!]_{k-1}$. To access the two components of a encoded location independently, we introduce the operator

$$|^{\alpha}_{\mathfrak{q}}\, \psi := (\alpha \wedge \neg\mathfrak{q}) \vee ((\alpha \hookrightarrow \mathfrak{q}) \wedge \psi),$$

where $\mathfrak{q} \in \{\mathfrak{t}, \mathfrak{p}\}$ and $\alpha \in \mathsf{ML}$. It is easy to check that $T \vDash |^{\alpha}_{\mathfrak{q}}\, \psi$ iff $T^{\alpha}_{T_{\mathfrak{q}}} \vDash \psi$.

In order to *compare* the locations of grid cells, for each component $\mathfrak{q} \in \{\mathfrak{t}, \mathfrak{p}\}$ we define the following formulas: $\psi^{\mathfrak{q}}_{\prec}(\alpha,\beta)$ tests whether the location in T_α is less than the one in T_β w.r.t. its $\mathfrak{q}$-component (assuming singleton teams T_α and T_β). Analogously, $\psi^{\mathfrak{q}}_{\equiv}(\alpha,\beta)$ checks for equality of the respective component:

$$\psi^{\mathfrak{q}}_{\prec}(\alpha,\beta) := \Box\, |^{\alpha}_{\mathfrak{q}}|^{\beta}_{\mathfrak{q}} \zeta^*_{k-1}(\alpha,\beta)$$
$$\psi^{\mathfrak{q}}_{\equiv}(\alpha,\beta) := \Box\, |^{\alpha}_{\mathfrak{q}}|^{\beta}_{\mathfrak{q}} \chi^*_{k-1}(\alpha,\beta)$$

For this purpose, $\psi^{\mathfrak{q}}_{\prec}$ is built upon the formula ζ^*_{k-1} from Theorem 7.2, while $\psi^{\mathfrak{q}}_{\equiv}$ checks for equality with the help of χ^*_{k-1} from Theorem 5.4.

Claim (a)**.** *Let $\mathcal{K}$ be a structure with a team T and disjoint scopes α and β. Suppose $w \in T_\alpha$ and $v \in T_\beta$, where $\ell(w) = (i_w, j_w)$ and $\ell(v) = (i_v, j_v)$. Then:*

$$T^{\alpha,\beta}_{w,v} \vDash \psi^{\mathfrak{t}}_{\equiv}(\alpha, \beta) \;\Leftrightarrow\; i_w = i_v$$
$$T^{\alpha,\beta}_{w,v} \vDash \psi^{\mathfrak{p}}_{\equiv}(\alpha, \beta) \;\Leftrightarrow\; j_w = j_v$$

Moreover, if $\alpha, \beta, \mathfrak{s}_0, \ldots, \mathfrak{s}_k$ are disjoint scopes in $\mathcal{K}$ and $(\mathcal{K}, T)$ is a k-staircase, then:

$$T^{\alpha,\beta}_{w,v} \vDash \psi^{\mathfrak{t}}_{\prec}(\alpha, \beta) \;\Leftrightarrow\; i_w < i_v$$
$$T^{\alpha,\beta}_{w,v} \vDash \psi^{\mathfrak{p}}_{\prec}(\alpha, \beta) \;\Leftrightarrow\; j_w < j_v$$

Proof of claim. Let us begin with $\psi^{\mathfrak{t}}_{\equiv}$ ($\psi^{\mathfrak{p}}_{\equiv}$ works identically):

$$\begin{aligned}
i_w = i_v \;&\Leftrightarrow \llbracket (Rw)_{\mathfrak{t}} \rrbracket_{k-1} = \llbracket (Rv)_{\mathfrak{t}} \rrbracket_{k-1} && \text{(By Definition)}\\
&\Leftrightarrow RT^{\alpha,\beta}_{(Rw)_{\mathfrak{t}},(Rv)_{\mathfrak{t}}} \vDash \chi^*_{k-1}(\alpha, \beta) && \text{(By Theorem 5.4)}\\
&\Leftrightarrow \left(RT^{\alpha,\beta}_{Rw,Rv}\right)^{\alpha,\beta}_{RT_{\mathfrak{t}},RT_{\mathfrak{t}}} \vDash \chi^*_{k-1}(\alpha, \beta) &&\\
&\Leftrightarrow RT^{\alpha,\beta}_{Rw,Rv} \vDash |^{\alpha}_{\mathfrak{t}}|^{\beta}_{\mathfrak{t}} \chi^*_{k-1}(\alpha, \beta) &&\\
&\Leftrightarrow T^{\alpha,\beta}_{w,v} \vDash \Box\, |^{\alpha}_{\mathfrak{t}}|^{\beta}_{\mathfrak{t}} \chi^*_{k-1}(\alpha, \beta) && \text{(Proposition 5.2)}
\end{aligned}$$

Similarly for $\psi^{\mathfrak{t}}_{\prec}$ ($\psi^{\mathfrak{p}}_{\prec}$ again works identically):

$$\begin{aligned}
i_w < i_v \;&\Leftrightarrow \llbracket (Rw)_{\mathfrak{t}} \rrbracket_{k-1} \prec^*_{k-1} \llbracket (Rv)_{\mathfrak{t}} \rrbracket_{k-1} && \text{(By Definition)}\\
&\Leftrightarrow RT^{\alpha,\beta}_{(Rw)_{\mathfrak{t}},(Rv)_{\mathfrak{t}}} \vDash \zeta^*_{k-1}(\alpha, \beta) && \text{(By Theorem 7.2)}\\
&\Leftrightarrow \left(RT^{\alpha,\beta}_{Rw,Rv}\right)^{\alpha,\beta}_{T_{\mathfrak{t}},T_{\mathfrak{t}}} \vDash \zeta^*_{k-1}(\alpha, \beta) &&\\
&\Leftrightarrow RT^{\alpha,\beta}_{Rw,Rv} \vDash |^{\alpha}_{\mathfrak{t}}|^{\beta}_{\mathfrak{t}} \zeta^*_{k-1}(\alpha, \beta) &&\\
&\Leftrightarrow T^{\alpha,\beta}_{w,v} \vDash \Box\, |^{\alpha}_{\mathfrak{t}}|^{\beta}_{\mathfrak{t}} \zeta^*_{k-1}(\alpha, \beta) && \text{(Proposition 5.2)} \triangleleft
\end{aligned}$$

8.3. **Defining grids, pre-tableaus, and tableaus.** Next, we aim at constructing formulas that check whether a given team is a grid, pre-tableau, or a tableau, respectively.

First, to check that every location $(i, j) \in \{1, \ldots, N\}^2$ of the grid occurs as $\ell(w)$ of some $w \in T$, we quantify over all corresponding pairs $(\Delta', \Delta'') \in \mathfrak{P}(\Delta_{k-1})^2$. To cover all these sets of types we can quantify, for instance, over the images of all points of $T_{\mathfrak{s}_k}$. However, as subteam quantifiers $\exists^{\subseteq}, \exists^1, \forall^{\subseteq}, \forall^1$ cannot pick *two* subteams from the same scope, we enforce a k-canonical copy $\mathfrak{s}'_k$ of $\mathfrak{s}_k$ in the spirit of Theorem 6.5:

$$\mathsf{canon}' := \rho^k_0(\mathfrak{s}_0) \wedge \bigwedge_{m=1}^{k} \rho^{k-m}_m(\mathfrak{s}_{m-1}, \mathfrak{s}_m) \wedge \rho^0_k(\mathfrak{s}_{k-1}, \mathfrak{s}'_k)$$

Claim (b)**.** *If $\mathfrak{s}_0, \ldots, \mathfrak{s}_k, \mathfrak{s}'_k$ are disjoint scopes in $\mathcal{K}$, then $(\mathcal{K}, T) \vDash \mathsf{canon}'$ if and only if $(\mathcal{K}, T)$ is a k-staircase and $T_{\mathfrak{s}'_k}$ is k-canonical. Moreover, $\mathsf{canon}' \wedge \mathsf{scopes}_k(\{\mathfrak{s}_0, \ldots, \mathfrak{s}_k, \mathfrak{s}'_k\}) \wedge \Box^{k+1}\bot$ is satisfiable, but is only satisfied by k-staircases $(\mathcal{K}, T)$ in which both $T_{\mathfrak{s}_k}$ and $T_{\mathfrak{s}'_k}$ are k-canonical. Furthermore, both formulas are constructible in space $\mathcal{O}(\log(|\Phi| + k))$.*

Proof of claim. Proven similarly to Theorem 6.5 and 6.7. $\triangleleft$

The next formula checks whether a given team is a grid. More precisely, the subformula ψ_{pair} compares the $\mathfrak{t}$-component of the selected location in α to the image of the world quantified in $\mathfrak{s}_k$, and its $\mathfrak{p}$-component to $\mathfrak{s}'_k$, respectively. That every world satisfies exactly one element of Ξ is guaranteed by ψ_{grid} as well.

$$\psi_{\text{grid}}(\alpha) := \Big(\alpha \hookrightarrow \bigvee_{e\in\Xi} e \wedge \bigwedge_{\substack{e'\in\Xi\\ e'\neq e}} \neg e'\Big) \wedge \forall^1_{\mathfrak{s}_k} \forall^1_{\mathfrak{s}'_k} \exists^1_\alpha \, \psi_{\text{pair}}(\alpha)$$

$$\psi_{\text{pair}}(\alpha) := \Box \left[\left(|^\alpha_{\mathfrak{t}} \; \chi^*_{k-1}(\mathfrak{s}_k, \alpha)\right) \wedge \left(|^\alpha_{\mathfrak{p}} \; \chi^*_{k-1}(\mathfrak{s}'_k, \alpha)\right)\right]$$

In the following and all subsequent claims, we always assume that T is a team in a Kripke structure $\mathcal{K}$ such that $(\mathcal{K}, T)$ satisfies $\mathsf{canon}' \wedge \Box^{k+1}\bot$. Moreover, all stated scopes are always assumed pairwise disjoint in $\mathcal{K}$ (as we can enforce this later in the reduction with $\mathsf{scopes}_k(\cdots)$).

Claim (c). *$T \vDash \psi_{grid}(\alpha)$ if and only if T_α is a grid.*

Proof of claim. Clearly $T \vDash \alpha \hookrightarrow \bigvee_{e\in\Xi} e \wedge \bigwedge_{e'\in\Xi, e'\neq e} \neg e'$ if and only if every world $w \in T_\alpha$ satisfies exactly one element of Ξ. Consequently, for the proof it remains to show the following equivalence:

$$\forall (i,j) \in \{1,\ldots,N\}^2 \colon \exists w \in T_\alpha \colon \ell(w) = (i,j)$$
$$\Leftrightarrow \forall \Delta', \Delta'' \subseteq \Delta_{k-1} \colon \exists w \in T_\alpha : [\![(Rw)_{\mathfrak{t}}]\!]_{k-1} = \Delta' \text{ and } [\![(Rw)_{\mathfrak{p}}]\!]_{k-1} = \Delta''$$

By k-canonicity of $\mathfrak{s}_k, \mathfrak{s}'_k$ due to Claim (b):

$$\Leftrightarrow \forall v \in T_{\mathfrak{s}_k}, v' \in T_{\mathfrak{s}'_k} : \exists w \in T_\alpha \colon [\![(Rw)_{\mathfrak{t}}]\!]_{k-1} = [\![Rv]\!]_{k-1} \text{ and } [\![(Rw)_{\mathfrak{p}}]\!]_{k-1} = [\![Rv']\!]_{k-1}$$

By Theorem 5.4:

$$\Leftrightarrow \forall v \in T_{\mathfrak{s}_k}, v' \in T_{\mathfrak{s}'_k} : \exists w \in T_\alpha \colon RT^{\alpha,\mathfrak{s}_k,\mathfrak{s}'_k}_{(Rw)_{\mathfrak{t}},Rv,Rv'} \vDash \chi^*_{k-1}(\mathfrak{s}_k,\alpha)$$
$$\text{and } RT^{\alpha,\mathfrak{s}_k,\mathfrak{s}'_k}_{(Rw)_{\mathfrak{p}},Rv,Rv'} \vDash \chi^*_{k-1}(\mathfrak{s}'_k,\alpha)$$
$$\Leftrightarrow \forall v \in T_{\mathfrak{s}_k}, v' \in T_{\mathfrak{s}'_k} : \exists w \in T_\alpha : \left(RT^{\alpha,\mathfrak{s}_k,\mathfrak{s}'_k}_{Rw,Rv,Rv'}\right)^\alpha_{RT_{\mathfrak{t}}} \vDash \chi^*_{k-1}(\mathfrak{s}_k,\alpha)$$
$$\text{and } \left(RT^{\alpha,\mathfrak{s}_k,\mathfrak{s}'_k}_{Rw,Rv,Rv'}\right)^\alpha_{RT_{\mathfrak{p}}} \vDash \chi^*_{k-1}(\mathfrak{s}'_k,\alpha)$$
$$\Leftrightarrow \forall v \in T_{\mathfrak{s}_k}, v' \in T_{\mathfrak{s}'_k} : \exists w \in T_\alpha \colon RT^{\alpha,\mathfrak{s}_k,\mathfrak{s}'_k}_{Rw,Rv,Rv'} \vDash \; |^\alpha_{\mathfrak{t}}\chi^*_{k-1}(\mathfrak{s}_k,\alpha) \wedge \; |^\alpha_{\mathfrak{p}}\chi^*_{k-1}(\mathfrak{s}'_k,\alpha)$$

By Proposition 5.2:

$$\Leftrightarrow \forall v \in T_{\mathfrak{s}_k}, v' \in T_{\mathfrak{s}'_k} : \exists w \in T_\alpha \colon T^{\alpha,\mathfrak{s}_k,\mathfrak{s}'_k}_{w,v,v'} \vDash \Box \, |^\alpha_{\mathfrak{t}}\chi^*_{k-1}(\mathfrak{s}_k,\alpha) \wedge \; |^\alpha_{\mathfrak{p}}\chi^*_{k-1}(\mathfrak{s}'_k,\alpha)$$

By Proposition 5.3:

$$\Leftrightarrow T \vDash \forall^1_{\mathfrak{s}_k} \forall^1_{\mathfrak{s}'_k} \exists^1_\alpha \Box \, |^\alpha_{\mathfrak{t}}\chi^*_{k-1}(\mathfrak{s}_k,\alpha) \wedge \; |^\alpha_{\mathfrak{p}}\chi^*_{k-1}(\mathfrak{s}'_k,\alpha)$$
$$\Leftrightarrow T \vDash \forall^1_{\mathfrak{s}_k} \forall^1_{\mathfrak{s}'_k} \exists^1_\alpha \, \psi_{\text{pair}}(\alpha) \qquad \triangleleft$$

With slight modifications it is straightforward to define pre-tableaus:

$$\psi_{\text{pre-tableau}}(\alpha) := \psi_{\text{grid}}(\alpha) \wedge \forall^1_{\mathfrak{s}_k} \forall^1_{\mathfrak{s}'_k} \bigwedge_{e\in\Xi} \exists^1_\alpha \big(\psi_{\text{pair}}(\alpha) \wedge (\alpha \hookrightarrow e)\big)$$

Claim (d). *$T \vDash \psi_{pre\text{-}tableau}(\alpha)$ if and only if T_α is a pre-tableau.*

Proof of claim. Proven similarly to Claim (c). ◁

The other special case of a grid, that is, a *tableau*, requires a more elaborate approach to define in MTL. The difference to a grid or pre-tableau is that we have to quantify over all *pairs* (w, w') of points in T, and check that they agree on Ξ if $\ell(w) = \ell(w')$. However, as discussed before, while $\forall^1$ can quantify over all points in a team, it cannot quantify over pairs.

As a workaround, we consider not only a tableau T_α, but also a *second* tableau that acts as a copy of T_α. Formally, for grids T_α, T_β, let $T_\alpha \approx T_\beta$ denote that for all pairs $(w, w') \in T_\alpha \times T_\beta$ it holds that $\ell(w) = \ell(w')$ implies $c(w) = c(w')$.

As $\approx$ is symmetric and transitive, $T_\alpha \approx T_\beta$ in fact implies both $T_\alpha \approx T_\alpha$ and $T_\beta \approx T_\beta$, and hence that both T_α and T_β are tableaus such that $C_{T_\alpha} = C_{T_\beta}$, where $C_{T_\alpha}, C_{T_\beta} \colon \{1, \dots, N\}^2 \to \Gamma \cup (Q \times \Gamma)$ are the induced runs as discussed on p. 25.

$$\psi_{\text{tableau}}(\alpha) := \psi_{\text{grid}}(\alpha) \wedge \exists^{\subseteq}_{\gamma_0} \psi_{\text{grid}}(\gamma_0) \wedge \psi_{\approx}(\alpha, \gamma_0)$$

$$\psi_{\approx}(\alpha, \beta) := \forall^1_\alpha \forall^1_\beta \left(\left(\psi^{\mathfrak{t}}_{\equiv}(\alpha, \beta) \wedge \psi^{\mathfrak{p}}_{\equiv}(\alpha, \beta) \right) \multimap \bigovee_{e \in \Xi} ((\alpha \vee \beta) \hookrightarrow e) \right)$$

In the following claim (and in the subsequent ones), we use the scopes $\gamma_0, \gamma_1, \gamma_2, \dots$ as "auxiliary pre-tableaus". Later, we will also use them as domains to quantify extra locations or rows from. (The index of γ_i is incremented whenever necessary to avoid quantifying from the same scope twice.) For this reason, from now on we always assume, for sufficiently large i, that T_{γ_i} is a pre-tableau. This can be later enforced in the reduction with $\psi_{\text{pre-tableau}}(\gamma_i)$.

Claim (e)**.**
(1) *$T \vDash \psi_{tableau}(\alpha)$ if and only if T_α is a tableau.*
(2) *For grids T_α, T_β, it holds $T \vDash \psi_{\approx}(\alpha, \beta)$ if and only if $T_\alpha \approx T_\beta$.*

Proof of claim. (2) follows straightforwardly from Claim (a). Let us consider (1). As ψ_{tableau} implies ψ_{grid}, and by Claim (c), we can assume that T_α is a grid.

Suppose that the formula is true. Then there exists $S \subseteq T_{\gamma_0}$ such that $T^\alpha_S \vDash \psi_{\text{grid}}(\gamma_0)$. By Claim (c), then S is a grid as well. Moreover, $T_\alpha \approx S$ by (2). As argued above, this implies that T_α (and S) is a tableau.

For the other direction, suppose that T_α is a tableau. Then it defines a function C_{T_α}. Since T_{γ_0} is a pre-tableau, we can pick a subteam S of it that contains for each $(i, j) \in \{1, \dots, N\}^2$ exactly those worlds w with $\ell(w) = (i, j)$ such that $c(w) = C_{T_\alpha}(i, j)$. Then $T_\alpha \approx S$, and ψ_{tableau} is true, with the quantifier $\exists^{\subseteq}_{\gamma_0}$ witnessed by S. ◁

8.4. **From tableaus to runs.** To ascertain that a tableau contains a run of M, we have to check whether each row indeed is a configuration of M—in other words, exactly one cell of each row contains a pair $(q, a) \in Q \times \Gamma$—and whether consecutive configurations obey the transition relation δ of M.

For this, in the spirit of Cook's theorem [Coo71] it suffices to consider all *legal windows* in the grid, i.e., cells that are adjacent as follows, where $e_1, \dots, e_6 \in \Gamma \cup (Q \times \Gamma)$:

e_1	e_2	e_3
e_4	e_5	e_6

If, say, $(q, a, q', a', R) \in \delta$—$M$ switches to state q' from q, replacing a on the tape by a', and moves to the right—then the windows obtained by setting $e_1 = e_4 = b$, $e_2 = (q, a)$, $e_5 = a'$, $e_3 = b'$, $e_6 = (q', b')$ are legal for all $b, b' \in \Gamma$. Using this scheme, δ is completely represented by a constant finite set $\mathsf{win} \subseteq \Xi^6$ of tuples $(e_1, \ldots, e_6)$ that represent the allowed windows in a run of M.

Let us next explain how adjacency of cells is expressed. Suppose that two points $w \in T_\alpha$ and $v \in T_\beta$ are given. That v is the immediate ($\mathfrak{t}$- or $\mathfrak{p}$-)successor of w then means that no element of the order exists between them. Simultaneously, w and v have to agree on the other component of their location, which is expressed by the first conjunct below. If $\mathfrak{q} \in \{\mathfrak{t}, \mathfrak{p}\}$ and $\overline{\mathfrak{q}} \in \{\mathfrak{t}, \mathfrak{p}\} \setminus \{\mathfrak{q}\}$, we define:

$$\psi^{\mathfrak{q}}_{\mathrm{succ}}(\alpha, \beta) := \psi^{\overline{\mathfrak{q}}}_{\equiv}(\alpha, \beta) \wedge \psi^{\mathfrak{q}}_{\prec}(\alpha, \beta) \wedge {\sim}\exists^1_{\gamma_0} \left(\psi^{\mathfrak{q}}_{\prec}(\alpha, \gamma_0) \wedge \psi^{\mathfrak{q}}_{\prec}(\gamma_0, \beta)\right)$$

Claim (f). *If $w \in T_\alpha$ and $v \in T_\beta$, then:*

$$T^{\alpha,\beta}_{w,v} \vDash \psi^{\mathfrak{t}}_{succ}(\alpha, \beta) \Leftrightarrow \exists i, j \in \{1, \ldots, N\} \colon \ell(w) = (i, j) \text{ and } \ell(v) = (i + 1, j)$$

$$T^{\alpha,\beta}_{w,v} \vDash \psi^{\mathfrak{p}}_{succ}(\alpha, \beta) \Leftrightarrow \exists i, j \in \{1, \ldots, N\} \colon \ell(w) = (i, j) \text{ and } \ell(v) = (i, j + 1)$$

Proof of claim. Let us consider only $\mathfrak{q} = \mathfrak{t}$, as the case $\mathfrak{q} = \mathfrak{p}$ is proven analogously. Assume that the formula $\psi^{\mathfrak{t}}_{\mathrm{succ}}(\alpha, \beta)$ is true in $T^{\alpha,\beta}_{w,v}$. By Claim (a), $\psi^{\mathfrak{p}}_{\equiv}$ holds if and only if there is a unique j such that $\ell(w) = (i, j)$ and $\ell(v) = (i', j)$, for some i, i'; in other words, if w and v agree on their $\mathfrak{p}$-component.

Next, consider the sets $\Delta_w := [\![(Rw)_{\mathfrak{t}}]\!]_{k-1}$ and $\Delta_v := [\![(Rv)_{\mathfrak{t}}]\!]_{k-1}$ which correspond to the $\mathfrak{t}$-components of $\ell(w)$ and $\ell(v)$. Suppose that Δ_w is the i-th element of $\prec^*_{k-1}$. By $\psi^{\mathfrak{t}}_{\prec}$ and Claim (a), then clearly Δ_v is the i'-th element for some $i' > i$.

Suppose for the sake of contradiction that also $i' > i+1$, and let then instead $\Delta' \subseteq \Delta_{k-1}$ be the $(i+1)$-th element of $\prec^*_{k-1}$. As T_{γ_0} is a pre-tableau, it contains a world z such that $\ell(z) = (i + 1, j)$. But then $\psi^{\mathfrak{q}}_{\prec}(\alpha, \gamma_0) \wedge \psi^{\mathfrak{q}}_{\prec}(\gamma_0, \beta)$ is true in $T^{\alpha,\beta,\gamma_0}_{w,v,z}$, contradiction to $\psi^{\mathfrak{t}}_{\mathrm{succ}}$. Consequently, $i' = i + 1$. The direction from right to left is shown similarly. $\triangleleft$

To check all windows in the tableau T_α, we need to simultaneously quantify elements from *six* tableaus $T_{\gamma_1}, \ldots, T_{\gamma_6}$ that are copies of T_α. For this purpose, we define

$$\exists^{\approx\alpha}_{\gamma_i} \varphi := \exists^{\subseteq}_{\gamma_i} \psi_{\mathrm{grid}}(\gamma_i) \wedge \psi_{\approx}(\alpha, \gamma_i) \wedge \varphi.$$

Intuitively, under the premise that T_{γ_i} is a pre-tableau and T_α is a tableau, it "copies" the tableau T_α into T_{γ_i} by shrinking T_{γ_i} accordingly. This is proven analogously to Claim (e). The next formula states that the picked points are arranged as in the picture:

$$\begin{aligned}\psi_{\mathrm{window}}(\gamma_1, \ldots, \gamma_6) := \psi^{\mathfrak{t}}_{\mathrm{succ}}(\gamma_1, \gamma_4) \wedge \psi^{\mathfrak{t}}_{\mathrm{succ}}(\gamma_2, \gamma_5) \wedge \psi^{\mathfrak{t}}_{\mathrm{succ}}(\gamma_3, \gamma_6) \wedge \\ \psi^{\mathfrak{p}}_{\mathrm{succ}}(\gamma_1, \gamma_2) \wedge \psi^{\mathfrak{p}}_{\mathrm{succ}}(\gamma_2, \gamma_3)\end{aligned}$$

T_{γ_1}	T_{γ_2}	T_{γ_3}
T_{γ_4}	T_{γ_5}	T_{γ_6}

The formula defining legal tableaus follows.

$$\psi_{\mathrm{legal}}(\alpha) := \psi_{\mathrm{tableau}}(\alpha) \wedge \exists^{\approx\alpha}_{\gamma_1} \cdots \exists^{\approx\alpha}_{\gamma_6} \vartheta_1 \wedge \vartheta_2 \wedge \vartheta_3$$

We check that at most cell per row contains a state of M:

$$\vartheta_1 := \forall^1_{\gamma_1}\forall^1_{\gamma_2}\Big(\psi^{\mathfrak{t}}_{\equiv}(\gamma_1,\gamma_2)\wedge\psi^{\mathfrak{p}}_{\prec}(\gamma_1,\gamma_2)\Big) \mathrel{-\!\!\triangleright} \bigwedge_{(q_1,a_1),(q_2,a_2)\in Q\times\Gamma} \sim\big((\gamma_1\hookrightarrow c^{-1}(q_1,a_1))\wedge(\gamma_2\hookrightarrow c^{-1}(q_2,a_2))\big)\Big)$$

We also check that every row contains some state. For this, $\forall^1_{\gamma_1}$ fixes some row and $\exists^1_{\gamma_2}\psi^{\mathfrak{t}}_{\equiv}(\gamma_1,\gamma_2)$ searches that particular row for a state:

$$\vartheta_2 := \forall^1_{\gamma_1}\exists^1_{\gamma_2}\ \psi^{\mathfrak{t}}_{\equiv}(\gamma_1,\gamma_2)\wedge\bigovee_{(q,a)\in Q\times\Gamma}(\gamma_2\hookrightarrow c^{-1}(q,a))$$

Finally, every window must obey the transition relation:

$$\vartheta_3 := \forall^1_{\gamma_1}\cdots\forall^1_{\gamma_6}\ \Big(\psi_{\text{window}}(\gamma_1,\ldots,\gamma_6) \mathrel{-\!\!\triangleright} \bigovee_{(e_1,\ldots,e_6)\in\text{win}}\ \bigwedge_{i=1}^{6}(\gamma_i\hookrightarrow e_i)\Big)$$

Claim (g)**.** *$T\vDash\psi_{legal}(\alpha)$ iff T_α is a legal tableau, i.e., iff C_{T_α} exists and is a run of M.*

Proof of claim. Suppose that the formula holds. We show that T_α is a legal tableau; the other direction is proven similarly.

Due to Claim (e), there are tableaus $S_1\subseteq T_{\gamma_1}$, …, $S_6\subseteq T_{\gamma_6}$ that are copies of T_α such that $\vartheta_1\wedge\vartheta_2\wedge\vartheta_2$ holds in $T^{\gamma_1,\ldots,\gamma_6}_{S_1,\ldots,S_6}$.

Due to Claim (a), the subformula ϑ_1 ensures the following: For all $w\in S_1, w'\in S_2$, $\ell(w)=(i,j)$, $\ell(w')=(i',j')$, if $i=i'$ and $j<j'$ hold, then it is not the case that both $c(w)=(q,a)$ and $c(w')=(q',a')$ for any state symbols $q,q'\in Q$. Since $C_{S_1}=C_{S_2}=C_{T_\alpha}$, this is precisely the case if each row of C_{T_α} contains at most one state symbol.

Conversely, again by Claim (a), the subformula ϑ_2 states that for every cell $w\in S_1$ there is another cell $w'\in S_2$ in the same row that carries a state symbol: in other words, every row of C_{T_α} contains at least one state symbol.

Finally, ϑ_3 relies on Claim (f) and states for every choice of singletons $w_1,\ldots,w_6$ in $S_1,\ldots,S_6$, assuming that they are arranged as a window, that there exists a tuple $(e_1,\ldots,e_6)\in\text{win}$ such that $w_i\in S_i$ satisfies $c(w_i)=e_i$. As we showed that C_{T_α} contains in each row a configuration of M, this implies that C_{T_α} exists and is a run of M. ◁

8.5. **From runs to a computation.** To encode the initial configuration on input $x=x_1\cdots x_n$ in a tableau, we access the first n cells of the first row and assign the respective letter of x, as well as the initial state, to the first cell. Moreover, we assign $\flat$ to all other cells in that row. For each $\mathfrak{q}\in\{\mathfrak{t},\mathfrak{p}\}$, we can check whether the location of a point in T_α is minimal in its $\mathfrak{q}$-component:

$$\psi^{\mathfrak{q}}_{\min}(\alpha) := \sim\exists^1_{\gamma_0}\psi^{\mathfrak{q}}_{\prec}(\gamma_0,\alpha)$$

This enables us to fix the first row of the configuration:

$$\begin{aligned}\psi_{\text{input}}(\alpha) := \exists^{\approx\alpha}_{\gamma_1}\cdots\exists^{\approx\alpha}_{\gamma_{n+1}}\ \exists^1_{\gamma_1}\cdots\exists^1_{\gamma_n}\ &\psi^{\mathfrak{t}}_{\min}(\gamma_1)\wedge\psi^{\mathfrak{p}}_{\min}(\gamma_1)\wedge\big(\gamma_1\hookrightarrow c^{-1}(q_0,x_1)\big)\\ &\bigwedge_{i=2}^{n}\psi^{\mathfrak{p}}_{\text{succ}}(\gamma_{i-1},\gamma_i)\ \wedge\big(\gamma_i\hookrightarrow c^{-1}(x_i)\big)\\ &\quad\wedge\forall^1_{\gamma_{n+1}}\Big(\big(\psi^{\mathfrak{t}}_{\equiv}(\gamma_n,\gamma_{n+1})\wedge\psi^{\mathfrak{p}}_{\prec}(\gamma_n,\gamma_{n+1})\big) \mathrel{-\!\!\triangleright} \big(\gamma_{n+1}\hookrightarrow c^{-1}(\flat)\big)\Big)\end{aligned}$$

Claim (h)**.** *Let T_α be a tableau. Then $T \vDash \psi_{input}(\alpha)$ if and only if*
(1) $C_{T_\alpha}(1,1) = (q_0, x_1)$,
(2) $C_{T_\alpha}(1,i) = x_i$ *for* $2 \leq i \leq n$,
(3) $C_{T_\alpha}(1,i) = \flat$ *for* $n < i \leq N$.

Proof of claim. Suppose that the formula holds. After processing the quantifiers $\exists^{\approx\alpha}_{\gamma_1} \cdots \exists^{\approx\alpha}_{\gamma_{n+1}}$, for all $m \in \{1, \ldots, n+1\}$ the team T_{γ_m} is a tableau such that $C_{T_{\gamma_m}} = C_{T_\alpha}$. (Obviously this requires these teams to be pre-tableaus beforehand.) For this reason, we can freely replace $C_{T_\alpha}(i,j)$ with $C_{T_{\gamma_m}}(i,j)$ when proving the properties (1)–(3).

In the second line of the formula, we make sure that $c(w) = (q_0, x_1)$ holds for least one point $w \in C_{T_{\gamma_1}}$ of location $\ell(w) = (1,1)$. That $\ell(w) = (1,1)$ holds follows from Claim (a), $\psi^{\mathfrak{q}}_{\min}$, and the assumption that T_{γ_0} is a pre-tableau (which it still is after processing $\exists^{\approx\alpha}_{\gamma_1} \cdots \exists^{\approx\alpha}_{\gamma_{n+1}}$). In particular, $C_{T_{\gamma_1}}(1,1) = (q_0, x_1)$.

The third line works similarly: for $2 \leq i \leq n$, it assigns x_i to $C_{T_{\gamma_i}}(1,i)$ and hence to $C_{T_\alpha}(1,i)$. Note that $\psi^{\mathfrak{p}}_{\text{succ}}$ also preserves the position in "$\mathfrak{p}$-direction", i.e., it is not necessary to repeat it for every cell of the first row. Finally, the last two lines state that every other location $(1,j')$ with $j' > n$ contains $\flat$. The other direction is again similar. ◁

Until now, we ignored the fact that M (polynomially often) alternates. To simulate this, we alternatingly quantify polynomially many tableaus, each containing a part of the computation of M. Each of these tableaus possesses a *tail configuration*, which is the configuration where M either accepts, rejects, or alternates. Formally, a number $i \in \{1, \ldots, N\}$ is a *tail index* of C if there exists j such that either
(1) $C(i,j)$ has an accepting or rejecting state,
(2) or $C(i,j)$ has an existential state and and there are $i' < i$ and j' with a universal state in $C(i',j')$,
(3) or $C(i,j)$ has a universal state and there are $i' < i$ and j' with an existential state in $C(i',j')$.

The least such i is called *first tail index*, and the corresponding configuration is the *first tail configuration*. The idea is that we can split the computation of M into multiple tableaus if any tableau (except the initial one) contains a run that continues from the previous tableau's first tail configuration.

We formalize the above as follows. Assume that T_α is a tableau, and that T_β marks a single row i by being a singleton $\{w\}$ with $\ell(w) = (i,j)$ for some j. Then the formula $\psi_{\text{tail}}(\alpha,\beta)$ below will be true if and only if the i-th row of C_{T_α} is a tail configuration. With

$$Q'\text{-}\mathsf{state}(\beta) := \bigovee_{(q,a)\in Q'\times\Gamma} (\beta \hookrightarrow c^{-1}(q,a)),$$

we check if a given singleton $T_\beta = \{w\}$ encodes an accepting, rejecting, existential, universal, or any state by setting Q' to Q_{acc}, Q_{rej}, $Q_\exists$, $Q_\forall$ or Q, respectively. We define ψ_{tail}:

$$\begin{aligned}
\psi_{\text{tail}}(\alpha,\beta) &:= \exists^{\approx\alpha}_{\gamma_0} \exists^1_\alpha \psi^{\mathfrak{t}}_{\equiv}(\alpha,\beta) \wedge Q\text{-}\mathsf{state}(\alpha) \wedge \Big[Q_{\text{acc}}\text{-}\mathsf{state}(\alpha) \varovee Q_{\text{rej}}\text{-}\mathsf{state}(\alpha) \varovee \\
&\qquad \exists^1_{\gamma_0}\Big(\psi^{\mathfrak{t}}_{\prec}(\gamma_0,\alpha) \wedge \big(Q_\exists\text{-}\mathsf{state}(\alpha) \wedge Q_\forall\text{-}\mathsf{state}(\gamma_0)\big) \varovee \big(Q_\forall\text{-}\mathsf{state}(\alpha) \wedge Q_\exists\text{-}\mathsf{state}(\gamma_0)\big)\Big)\Big] \\
\psi_{\text{first-tail}}(\alpha,\beta) &:= \psi_{\text{tail}}(\alpha,\beta) \wedge {\sim}\exists^1_{\gamma_1}\Big(\psi^{\mathfrak{t}}_{\prec}(\gamma_1,\beta) \wedge \psi_{\text{tail}}(\alpha,\gamma_1)\Big)
\end{aligned}$$

Claim (i)**.** *Suppose that T_α is a tableau, $T_\beta = \{w\}$, and $\ell(w) = (i, j)$. Then $T \vDash \psi_{\mathit{tail}}(\alpha, \beta)$ if and only if i is a tail index of C_{T_α}. Moreover, $T \vDash \psi_{\mathit{first\text{-}tail}}(\alpha, \beta)$ if and only if i is the first tail index of C_{T_α}.*

Proof of claim. Since T_{γ_1} is a pre-tableau and hence contains all locations in rows $i' < i$, it is easy to see that the proof for $\psi_{\text{first-tail}}$ boils down to that of ψ_{tail}. Consequently, let us consider ψ_{tail}.

First, due to $\exists^{\approx\alpha}_{\gamma_0}$, we can assume that T_{γ_0} is a tableau that is a copy of T_α, i.e., $C_{T_\alpha} = C_{T_{\gamma_0}}$. Here, it is required for the inner quantification in the definition of a tail index.

The first line of the formula reduces T_α to a singleton that is (due to $\psi^{\mathfrak{t}}_{\equiv}$) in row i. Furthermore, it carries a state q of M due to $Q\text{-}\mathsf{state}(\alpha)$. The further examination of this state will determine if i is a tail index. Now, q is exactly one of accepting, rejecting, existential, or universal. If $q \in Q_{\text{acc}} \cup Q_{\text{rej}}$, then i is a tail index by definition.

Otherwise we quantify over the states q' of all (copies of) earlier rows in T_α, using $\exists^1_{\gamma_0} \psi^{\mathfrak{t}}_{\prec}(\gamma_0, \alpha)$, and search for a universal state if q is existential and vice versa, which as well, if it exists, proves by definition that i is a tail index. ◁

Formally, given a run C of M that has a tail configuration, C *accepts* if the state q in its first tail configuration is in Q_{acc}, C *rejects* if that q is in Q_{rej}, and C *alternates* otherwise. That a run of the form C_{T_α} accepts or rejects is expressed by

$$\psi_{\text{acc}}(\alpha) := \exists^{\approx\alpha}_{\gamma_2} \exists^1_{\gamma_2} Q_{\text{acc}}\text{-}\mathsf{state}(\gamma_2) \wedge \psi_{\text{first-tail}}(\alpha, \gamma_2),$$
$$\psi_{\text{rej}}(\alpha) := \exists^{\approx\alpha}_{\gamma_2} \exists^1_{\gamma_2} Q_{\text{rej}}\text{-}\mathsf{state}(\gamma_2) \wedge \psi_{\text{first-tail}}(\alpha, \gamma_2).$$

In this formula, first the tableau T_α is copied to T_{γ_2} to extract with $\exists^1_{\gamma_2}$ the world carrying an accepting/rejecting state, while $\psi_{\text{first-tail}}(\alpha, \gamma_2)$ ensures that no alternation or rejecting/accepting state occurs at some earlier point in C_{T_α}.

If the first tail configuration of the run contains an alternation, and if the run was existentially quantified, then it should be continued in a universally quantified tableau, and vice versa. The following formula expresses, given two tableaus T_α, T_β, that C_{T_β} is a *continuation* of C_{T_α}, i.e., that the first configuration of C_{T_β} equals the first tail configuration of C_{T_α}. In other words, if i is the first tail index of C_{T_α}, then $C_{T_\alpha}(i, j) = C_{T_\beta}(1, j)$ for all $j \in \{1, \dots, N\}$.

$$\psi_{\text{cont}}(\alpha, \beta) := \exists^1_{\gamma_2} \psi_{\text{first-tail}}(\alpha, \gamma_2) \wedge \forall^1_\alpha \forall^1_\beta$$
$$\Big[\Big(\psi^{\mathfrak{t}}_{\min}(\beta) \wedge \psi^{\mathfrak{t}}_{\equiv}(\alpha, \gamma_2) \wedge \psi^{\mathfrak{p}}_{\equiv}(\alpha, \beta)\Big) \rightarrow \Big(\bigotimes_{e \in \Xi} (\alpha \vee \beta) \hookrightarrow e\Big)\Big]$$

The above formula first obtains the first tail index i of C_{T_α} and stores it in a singleton $y \in T_{\gamma_2}$. Then for all worlds $w \in T_\alpha$ and $v \in T_\beta$, where v is $\mathfrak{t}$-minimal (i.e., in the first row) and w is in the same row as y, and which additionally agree on their $\mathfrak{p}$-component, the third line states that w and v agree on Ξ. Altogether, the i-th row of C_{T_α} and the first row of C_{T_β} then have to coincide.

M performs at most $r(n) - 1$ alternations for some polynomial r. Then we require $r = r(n)$ tableaus, which we call $\alpha_1, \dots, \alpha_r$. In the following, the formula $\psi_{\text{run},i}$ describes the behaviour of the i-th run, i.e., the part of the computation after $i - 1$ alternations. W.l.o.g. r is even and $q_0 \in Q_\exists$. We may then define the final run by

$$\psi_{\text{run},r} := \forall^{\subseteq}_{\alpha_r} \Big[\Big(\psi_{\text{legal}}(\alpha_r) \wedge \psi_{\text{cont}}(\alpha_{r-1}, \alpha_r)\Big) \rightarrow \Big({\sim}\psi_{\text{rej}}(\alpha_r) \wedge \psi_{\text{acc}}(\alpha_r)\Big)\Big].$$

For $1 < i < r$ and even i, let

$$\psi_{\text{run},i} := \forall^{\subseteq}_{\alpha_i} \left[\Big(\psi_{\text{legal}}(\alpha_i) \wedge \psi_{\text{cont}}(\alpha_{i-1}, \alpha_i)\Big) \twoheadrightarrow \Big({\sim}\psi_{\text{rej}}(\alpha_i) \wedge \big(\psi_{\text{acc}}(\alpha_i) \owedge \psi_{\text{run},i+1}\big)\Big)\right]$$

and for $1 < i < r$ and odd i

$$\psi_{\text{run},i} := \exists^{\subseteq}_{\alpha_i} \Big[\psi_{\text{legal}}(\alpha_i) \wedge \psi_{\text{cont}}(\alpha_{i-1}, \alpha_i) \wedge {\sim}\psi_{\text{rej}}(\alpha_i) \wedge \Big(\psi_{\text{acc}}(\alpha_i) \owedge \psi_{\text{run},i+1}\Big)\Big].$$

Analogously, the initial run is described by

$$\psi_{\text{run},1} := \exists^{\subseteq}_{\alpha_1} \Big(\psi_{\text{legal}}(\alpha_1) \wedge \psi_{\text{input}}(\alpha_1) \wedge {\sim}\psi_{\text{rej}}(\alpha_1) \wedge \Big(\psi_{\text{acc}}(\alpha_1) \owedge \psi_{\text{run},2}\Big)\Big)$$

We are now in the position to state the full reduction. Let us gather all relevant scopes in the set $\Psi \subseteq \mathcal{PS}$:

$$\Psi := \{\mathfrak{s}_i \mid 0 \leq i \leq k\} \cup \{\mathfrak{s}'_k\} \ \cup \ \{\gamma_i \mid 0 \leq i \leq n+1\} \ \cup \ \{\alpha_i \mid 1 \leq i \leq r\}$$

The scopes that accommodate pre-tableaus are

$$\Psi' := \{\gamma_i \mid 0 \leq i \leq n+1\} \cup \{\alpha_i \mid 1 \leq i \leq r\}.$$

W.l.o.g. $n \geq 5$, as $\gamma_1, \ldots, \gamma_6$ are always required in the construction. The reduction now maps x to

$$\varphi_x := \mathsf{canon}' \wedge \mathsf{scopes}_k(\Psi) \wedge \bigwedge_{p \in \Psi'} \psi_{\text{pre-tableau}}(p) \wedge \psi_{\text{run},1}.$$

It is easy to see that this formula is an MTL_k-formula that is logspace-constructible from x and k, where k itself is either constant or a polynomial in $|x|$ and hence logspace-computable. By Lemma 6.6, φ_x is satisfiable if and only if $\varphi_x \wedge \Box^{k+1}\bot$ is satisfiable. For this reason, we conclude the reduction with the following proof.

Proof of Lemma 8.2. It remains to argue that $\varphi_x \wedge \Box^{k+1}\bot$ is satisfiable if and only if M accepts x. For the sake of simplicity, assume $r = 2$. The cases $r > 2$ are proven analogously.

"$\Rightarrow$": Suppose $(\mathcal{K}, T) \vDash \varphi_x \wedge \Box^{k+1}\bot$. Similarly as in Theorem 6.7, the $p \in \Psi$ are disjoint scopes due to $\mathsf{scopes}_k(\Psi)$. Moreover, by canon' and Claim (b), $(\mathcal{K}, T)$ is then a k-staircase in which $T_{\mathfrak{s}_k}$ and $T_{\mathfrak{s}'_k}$ both are k-canonical teams. Due to Claim (d) and the large conjunction in φ_x, $T_{\alpha_1}, T_{\alpha_2}, T_{\gamma_1}, \ldots, T_{\gamma_{n+1}}$ are then pre-tableaus.

As the formula $\psi_{\text{run},1}$ holds, by Claim (g) and (h), T_{α_1} has a subteam S_1 that is a legal tableau and starts with M's initial configuration on x. In particular, C_{S_1} exists. Moreover, either ψ_{acc} holds (i.e., C_{S_1} and hence M is accepting) or $\psi_{\text{run},2}$ holds (i.e., if C_{S_1} alternates). Consider the latter case. Then for all legal tableaus $S_2 \subseteq T_{\alpha_2}$ such that C_{S_2} is a continuation of C_{S_1} it holds that C_{S_2} is accepting. However, as T_{α_2} is a pre-tableau, every run is of the form C_{S_2} for some $S_2 \subseteq T_{\alpha_2}$. Consequently, M accepts x.

"$\Leftarrow$": Suppose M accepts x. First of all, due to Claim (b), the formula $\mathsf{canon}' \wedge \mathsf{scopes}_k(\{\mathfrak{s}_0, \ldots, \mathfrak{s}_k, \mathfrak{s}'_k\}) \wedge \Box^{k+1}\bot$ has a model $(\mathcal{K}, T)$. Moreover, we can freely add a pre-tableau T_p for each $p \in \Psi$ to satisfy the large conjunction in φ_x. By labeling the propositions in Ψ correctly (as disjoint scopes), we ensure that $\mathsf{scopes}_k(\Psi)$ holds as well.

It remains to demonstrate $T \vDash \psi_{\text{run},1}$. As M accepts x, there exists a run C_1 starting from M's initial configuration such that either C_1 accepts, or, for all runs C_2 continuing C_1, C_2 accepts.

Since T_{α_1} is a pre-tableau, it also contains a subteam S_1 such that S_1 is a legal tableau and $C_{S_1} = C_1$. We choose S_1 as witness for $\exists^{\subseteq}_{\alpha_1}$. If C_1 itself accepts, then $\psi_{\text{acc}}(\alpha_1)$ and hence $\psi_{\text{run},1}$ is satisfied. Otherwise we consider $\psi_{\text{run},2}$. Suppose that $S_2 \subseteq T_{\alpha_2}$ is picked as a subteam by $\forall^{\subseteq}_{\alpha_2}$. If it forms a legal tableau and C_{S_2} is a continuation of C_1, then C_2 must be accepting since M accepts x by assumption. But this implies that $\psi_{\mathrm{acc}}(\alpha_2)$ is true for any such S_2. Consequently, $\psi_{\mathrm{run},2}$ and hence $\psi_{\mathrm{run},1}$ is true. $\square$

9. Hardness under strict semantics and on restricted frame classes

9.1. **Lax and strict semantics.** In this section, we further generalize the hardness result of the previous section.

Team-semantical connectives can be evaluated either in so-called *standard* or *lax semantics*, or alternatively in *strict semantics*. In Section 3, we defined MTL with lax semantics. In strict semantics, the connectives $\vee$ and $\Diamond$ are replaced by their counterparts $\vee_s$ and $\Diamond_s$:

$$(\mathcal{K}, T) \vDash \psi \vee_s \theta \Leftrightarrow \exists S, U \subseteq T \text{ such that } T = S \cup U,\ S \cap U = \emptyset,\ (\mathcal{K}, S) \vDash \psi, \text{ and } (\mathcal{K}, U) \vDash \theta,$$

$$(\mathcal{K}, T) \vDash \Diamond_s \psi \quad \Leftrightarrow\ (\mathcal{K}, S) \vDash \psi \text{ for some strict successor team } S \text{ of } T,$$

where a *strict successor team* of T is a successor team $S \subseteq RT$ for which there exists a surjective $f : T \to S$ satisfying $f(w) \in Rw$ for all $w \in T$. Intuitively, in the lax disjunction the teams of the splitting may overlap, while in the strict disjunction they are disjoint. Likewise, a lax successor team may contain multiple successor of any $w \in T$, while in a strict successor team we pick exactly one successor for each $w \in T$.

An MTL-formula φ is *downward closed* if $(\mathcal{K}, T) \vDash \varphi$ implies $(\mathcal{K}, S) \vDash \varphi$ for all $S \subseteq T$. For example, every ML-formula is downward closed, as is the constancy atom $=\!(\alpha) = \alpha \otimes \neg\alpha$ or generally any monotone Boolean combination of ML-formulas. On such formulas, strict and lax semantics are equivalent:

Proposition 9.1. *Let $\varphi, \psi \in \mathsf{MTL}$ such that φ is downward closed. Then $\varphi \vee \psi \equiv \varphi \vee_s \psi$ and $\Diamond\varphi \equiv \Diamond_s\varphi$.*

Proof. Clearly $\varphi \vee_s \psi$ entails $\varphi \vee \psi$ and $\Diamond_s\varphi$ entails $\Diamond\varphi$. If conversely $T \vDash \varphi \vee \psi$ via subteams $S, U \subseteq T$ such that $S \cup U = T$, $S \vDash \varphi$ and $U \vDash \psi$, then we instead split T into the subteams U and $T \setminus U$. Since $T \setminus U \subseteq S$ and φ is downward closed, this proves $T \vDash \varphi \vee_s \psi$.

Likewise, suppose $T \vDash \Diamond\varphi$ via some successor team S of T. Assuming the axiom of choice, there is some function $f : T \to S$ such that $f(w) \in Rw$ for each $w \in T$. The team $\{f(w) \mid w \in T\} \subseteq S$ is now a strict successor team of T and satisfies φ due to downward closure. $\square$

Due to Proposition 9.1, the distinction between strict and lax semantics was traditionally unnecessary for many team logics such as the original *dependence logic* [Vää07, Vää08], as it has only downward closed formulas. The distinction between strict and lax semantics was first made in the context of first-order team logic by Galliani [Gal12]. It has some interesting consequences, for instance first-order inclusion logic in strict semantics is as expressive as existential second-order logic [GHK13] (see also Hannula and Kontinen [HK15]).

With modal team logic, strict semantics was studied, e.g., by Hella et al. [HS15, HKMV15, HKMV17]. In the works that explicitly study strict semantics, the underlying (first-order or modal) team logic was enriched by not downward closed constructs such as the *inclusion atom* $\subseteq$ or *exclusion atom* $|$, or the *independence atom* $\perp$.

In this article, where we consider team-wide negation $\sim$ as part of the logic, the distinction between strict and lax semantics becomes apparent already for simple formulas such as

$\mathsf{E}\top \vee \mathsf{E}\top \not\equiv \mathsf{E}\top \vee_s \mathsf{E}\top$, where the former defines non-emptiness, but the latter means that the team contains at least two points.

We prove that our hardness results also hold in strict semantics. Let the logics $\mathsf{MTL}(\vee_s, \Box)$ and $\mathsf{MTL}_k(\vee_s, \Box)$ be defined like MTL and MTL_k, but with $\vee_s$ instead of $\vee$ and without $\Diamond$ and $\Diamond_s$ (i.e., only using the modality $\Box$).

Theorem 9.2. *$\mathsf{SAT}(\mathcal{L})$ and $\mathsf{VAL}(\mathcal{L})$ are hard for* TOWER(poly) *if $\mathcal{L} = \mathsf{MTL}(\vee_s, \Box)$, and hard for* ATIME-ALT$(\exp_{k+1}, \text{poly})$ *if $\mathcal{L} = \mathsf{MTL}_k(\vee_s, \Box)$ and $k \geq 0$.*

Proof. An analysis of the proof of Lemma 8.2 yields that the MTL-formula φ_x produced in the reduction can be easily adapted to strict semantics. First, observe that $\Diamond$ occurs only in the subformula max_i, which is by Proposition 9.1 equivalent to

$$\top \vee_s \left(\neg\Box^i\bot \wedge \sim \bigvee_{p\in\Phi}{}_s (\neg\Box^i p \varovee \neg\Box^i\neg p)\right),$$

since $\Diamond\alpha \equiv \neg\Box\neg\alpha$ and $\neg\Box^i p \varovee \neg\Box^i\neg p$ is a downward closed formula. A quick check reveals that all other instances of $\vee$ in φ_x are subject to Proposition 9.1 as well, except of the occurrence in the second line of ζ_k^*. Here, the critical part of the correctness proof is the choice of the subteam U' in Claim (c) of Lemma 7.4. In strict semantics, the only possibility becomes $U' = U = O \setminus S$, for which the proof works identically. Finally, for the case $k = 0$, a similar check of the proof for PTL [HKVV18, Theorem 4.9] reveals that there also every $\vee$ can be replaced by $\vee_s$ due to Proposition 9.1. $\Box$

Note that the corresponding upper bound via the construction of a canonical model (viz. Theorem 4.6) does not apply to strict semantics. The reason for this is the failure of Proposition 3.5: In strict semantics, MTL_k-formulas are not invariant under k-team-bisimulation in general.

As an example, consider the formula $\varphi := \mathsf{E}\top \vee_s \mathsf{E}\top$. It states that the team contains at least two points. However, for every finite $\Phi \subseteq \mathcal{PS}$ and $k \geq 0$ it is easy to find a team T of two points and a singleton S that is (Φ, k)-bisimilar to it, while $T \vDash \varphi$ and $S \nvDash \varphi$.

A possible approach could be to define a bisimulation relation that respects the multiplicity of types in a team, and to define a corresponding canonical model, but this is beyond the scope of this paper.

9.2. **Restricted frame classes.** A natural restriction in the context of modal logic is to focus on a specific subclass of Kripke frames, which is useful for instance for modeling belief or temporal systems. (For an introduction to frame classes, consider, e.g., Fitting [Fit07].) Let $F = (W, R)$ denote a frame. Prominent frame classes include

- K: all frames,
- D: serial frames ($w \in W \Rightarrow Rw \neq \emptyset$),
- T: reflexive frames ($w \in W \Rightarrow w \in Rw$),
- $\mathsf{K4}$: transitive frames ($u \in Rv, v \in Rw, w \in W \Rightarrow u \in Rw$),
- $\mathsf{D4}$: serial and transitive frames,
- $\mathsf{S4}$: reflexive and transitive frames.

In this section, we consider these classes from a complexity theoretic perspective, and show that the lower bounds of MTL hold when restricted to these classes. Given a frame class $\mathcal{F}$ and a fragment $\mathcal{L}$ of MTL, let $\mathsf{SAT}(\mathcal{L}, \mathcal{F})$ denote the set of all $\mathcal{L}$-formulas that are satisfied in a model (W, R, V, T) where (W, R) is a frame in $\mathcal{F}$. Define $\mathsf{VAL}(\mathcal{L}, \mathcal{F})$ analogously.

We prove the team-semantical analog of Ladner's theorem, which states that classical modal satisfiability and validity are PSPACE-hard problem for any frame class between $\mathsf{S4}$ and K [Lad77, Theorem 3.1]. Note that this includes all the frame classes stated above.

Theorem 9.3. *Let $\mathcal{F}$ be a frame class such that $\mathsf{S4} \subseteq \mathcal{F} \subseteq \mathsf{K}$. Then $\mathsf{SAT}(\mathsf{MTL}, \mathcal{F})$ and $\mathsf{VAL}(\mathsf{MTL}, \mathcal{F})$ are hard for* TOWER(poly)*, and $\mathsf{SAT}(\mathsf{MTL}_k, \mathcal{F})$ and $\mathsf{VAL}(\mathsf{MTL}_k, \mathcal{F})$ are hard for* ATIME-ALT$(\exp_{k+1}, \text{poly})$*, for $k \geq 0$.*

Proof. We give the proof for $\mathsf{SAT}(\mathsf{MTL}_k) \leq_{\mathrm{m}}^{\log} \mathsf{SAT}(\mathsf{MTL}_k, \mathcal{F})$. Let $\varphi \in \mathsf{MTL}_k$. The idea is to introduce new propositions $\ell_0, \ldots, \ell_k \notin \mathsf{Prop}(\varphi)$ that mark the layers of different height in a structure, and to modify the formula such that all edges except between consecutive layers i and $i+1$ are ignored. (Here, we make the assumption that $\mathcal{K}$ is a acyclic, which relies on Corollary 6.3 and hence indirectly on Proposition 3.5).

Given a $\Phi \cup \{\ell_0, \ldots, \ell_k\}$-structure $\mathcal{K} = (W, R, V)$, let $\mathcal{K}^\circ := (W, R^\circ, V)$ be the structure where only such edges are retained, i.e.,

$$R^\circ = R \cap \bigcup_{i=0}^{k-1} (V(\ell_i) \times V(\ell_{i+1})).$$

On the side of formulas, the reduction is $\varphi \mapsto \ell_0 \wedge \varphi^0$, where φ^i is inductively as follows. The non-modal connectives are ignored, i.e., $p^i := p$ for $p \in \Phi$, $(\psi \wedge \theta)^i := \psi^i \wedge \theta^i$, $(\sim\psi)^i := \sim\psi^i$, $(\psi \vee \theta)^i := \psi^i \vee \theta^i$. For the modalities, let $(\Diamond\psi)^i := \Diamond(\ell_{i+1} \wedge \psi^{i+1})$ and $(\Box\psi^i) := \Box(\ell_{i+1} \hookrightarrow \psi^{i+1})$. Intuitively, φ^i is meant to be evaluated in layer i, and we make sure that successor teams always are contained in the next layer $i+1$.

For the correctness of the reduction, we will first show the following claim.

Claim. *For all $i \in \{0, \ldots, k\}$ and $T \subseteq V(\ell_i)$, it holds that $(\mathcal{K}, T) \vDash \varphi^i$ iff $(\mathcal{K}^\circ, T) \vDash \varphi$.*

Proof of claim. This is proved by a straightforward induction on the formula size:

- Atomic propositions are clear. The Boolean connectives and splitting follow straightforwardly from the induction hypothesis (as subteams of T are again in $V(\ell_i)$).
- Let $\varphi = \Diamond\psi$. Suppose $(\mathcal{K}, T) \vDash \varphi^i$, i.e., $(\mathcal{K}, S) \vDash \ell_{i+1} \wedge \psi^{i+1}$ for some R-successor team S of T. Then by induction hypothesis $(\mathcal{K}^\circ, S) \vDash \psi$, as $S \subseteq V(\ell_{i+1})$. S is an R°-successor team of T as well, since $(w, v) \in R \Leftrightarrow (w, v) \in R^\circ$ for every $(w, v) \in V(\ell_i) \times V(\ell_{i+1})$. This proves $(\mathcal{K}^\circ, T) \vDash \varphi$.

  Conversely, if $(\mathcal{K}^\circ, T) \vDash \varphi$, then $(\mathcal{K}^\circ, S) \vDash \psi$ for some R°-successor team S of T. However, any R°-successor team of T is a subset of $V(\ell_{i+1})$. As a consequence, $(\mathcal{K}, S) \vDash \ell_{i+1}$. Moreover, by induction hypothesis, $(\mathcal{K}, S) \vDash \psi^{i+1}$. This yields $(\mathcal{K}, T) \vDash \varphi^i$, since S is trivially also a R-successor team of T.
- Let $\varphi = \Box\psi$. Then $(\mathcal{K}, T) \vDash \varphi^i$ iff $(\mathcal{K}, RT) \vDash (\ell_{i+1} \hookrightarrow \psi^{i+1})$ iff $(\mathcal{K}, RT \cap V(\ell_{i+1})) \vDash \psi^{i+1}$ iff $(\mathcal{K}^\circ, RT \cap V(\ell_{i+1}) \vDash \psi$ by induction hypothesis. It remains to show that $R^\circ T = RT \cap V(\ell_{i+1})$. Clearly, $R^\circ T \subseteq RT$ and $R^\circ T \subseteq V(\ell_{i+1})$, since $R^\circ \subseteq R$, $R^\circ V(\ell_i) \subseteq V(\ell_{i+1})$, and $T \subseteq V(\ell_i)$. Conversely, if $w \in RT \cap V(\ell_{i+1})$, then $(v, w) \in R$ for some $v \in T$. As $(v, w) \in V(\ell_i) \times V(\ell_{i+1})$, then $(v, w) \in R^\circ$, hence $w \in R^\circ T$. $\lhd$

Now, due to the above claim, if $\ell_0 \wedge \varphi^0$ is satisfiable, then clearly φ is as well. It remains to show that $\ell_0 \wedge \varphi^0$ has a reflexive and transitive model if φ is satisfiable. Suppose that the latter is satisfied in a Φ-structure $(\mathcal{K}, T)$. By Corollary 6.3, we may assume that $(\mathcal{K}, T)$ is a forest of height k with the set of roots being T. Then we label the new propositions ℓ_i such that $V(\ell_i) = R^i T$, i.e., $V(\ell_0) = T$, $V(\ell_1) = RT$ and so on. As $\mathcal{K}$ is a forest, note that the sets $T, RT, R^2T, \ldots$ are pairwise disjoint. In other words, every world in $\mathcal{K}$ has a unique distance $0 \leq i \leq k$ from T and hence exactly one ℓ_i labeled. This is required for the next part of the proof.

Let now R^* be the reflexive transitive closure of R. It remains to show $(R^*)^\circ = R$, since then we can again apply the previously proved claim and are done. It is easy to see that $R \subseteq (R^*)^\circ$, since for every $(w, v) \in R$ there is some i such that $w \in R^iT = V(\ell_i)$, consequently $(w, v) \in R^iT \times R^{i+1}T = V(\ell_i) \times V(\ell_{i+1})$. For the other direction, suppose $(w, v) \in (R^*)^\circ$. By definition of $(R^*)^\circ$, there is i such that $w \in V(\ell_i)$, $v \in V(\ell_{i+1})$, and v is reachable from w by some R-path $(u_0, \ldots, u_n)$ where $w = u_0$ and $v = u_n$. But since $u_0 \in R^iT$, for all m it holds $u_m \in R^{i+m}T = V(\ell_{i+m})$. As $V(\ell_{i+n}) \cap V(\ell_{i+1}) = \emptyset$ for $n \neq 1$, we conclude $n = 1$, so $(w, v) \in R$. □

10. Conclusion

Theorem 8.1 settles the complexity of MTL and proves that its satisfiability and validity problems are complete for the non-elementary complexity class TOWER(poly). Moreover, the fragments MTL_k are proved complete for ATIME-ALT($\exp_{k+1}$, poly), the levels of the elementary hierarchy with polynomially many alternations.

In our approach, we developed a notion of (k-)canonical models for modal logics with team semantics. We showed that such models exist for MTL and MTL_k, and that logspace-computable MTL_k-formulas exist that are satisfiable, but only have k-canonical models.

Our lower bounds carry over to two-variable first-order team logic $\mathsf{FO}^2(\sim)$ and its fragment $\mathsf{FO}^2_k(\sim)$ of bounded quantifier rank k as well [Lüc18c]. While the former is TOWER(poly)-complete, the latter is ATIME-ALT($\exp_{k+1}$, poly)-hard. However, no matching upper bound for the satisfiability problem of $\mathsf{FO}^2_k(\sim)$ exists.

In the final section, we considered variants of the satisfiability problem for MTL. We showed that it is as hard as the original problem when MTL is interpreted in strict semantics, and in fact for $\Diamond$-free formulas with $\vee$ being interpreted either lax or strict. Also, any restriction of the satisfiability problem to a frame class that includes at least the reflexive-transitive frames is as hard as the original problem.

In future research, it could be useful to further generalize the concept of canonical models to other logics with team semantics. Do logics such as $\mathsf{FO}^2_k(\sim)$ permit a canonical model in the spirit of k-canonical models for MTL_k, and does this yield a tight upper bound on the complexity of their satisfiability problem? How do MTL_k and $\mathsf{FO}^2_k(\sim)$ differ in terms of succinctness?

Other obvious open questions are the upper bounds for Theorem 9.2 and 9.3, and also the combination of the above aspects, e.g., does the lower bound still hold in strict semantics on reflexive-transitive frames? To solve these issues, the model theory of modal team logic has to be refined. For example, what is the analog of Proposition 3.5 for strict semantics?

Acknowledgements. The author wishes to thank Heribert Vollmer, Irena Schindler and Arne Meier, as well as the anonymous referee, for numerous helpful comments and suggestions.

References

[AKVV16] Samson Abramsky, Juha Kontinen, Jouko Väänänen, and Heribert Vollmer, editors. *Dependence Logic, Theory and Applications*. Springer, 2016.

[BRV01] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge University Press, New York, NY, USA, 2001.

[CH90] Kevin J. Compton and C. Ward Henson. A Uniform Method for Proving Lower Bounds on the Computational Complexity of Logical Theories. *Ann. Pure Appl. Logic*, 48(1):1–79, 1990.

[CH96] M. J. Cresswell and G. E. Hughes. *A New Introduction to Modal Logic*. Routledge, 1996.

[CKS81] Ashok K. Chandra, Dexter C. Kozen, and Larry J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, January 1981.

[Coo71] Stephen A. Cook. The Complexity of Theorem-Proving Procedures. In *Proc. of the 3rd Annual ACM Sym. on Theory of Computing*, pages 151–158, 1971.

[Cre83] M. J. Cresswell. KM and the finite model property. *Notre Dame Journal of Formal Logic*, 24(3):323–327, 1983.

[DKV16] Arnaud Durand, Juha Kontinen, and Heribert Vollmer. Expressivity and Complexity of Dependence Logic. In *Dependence Logic, Theory and Applications*, pages 5–32. Springer, 2016.

[EHM+13] Johannes Ebbing, Lauri Hella, Arne Meier, Julian-Steffen Müller, Jonni Virtema, and Heribert Vollmer. Extended modal dependence logic. In *WoLLIC*, volume 8071 of *Lecture Notes in Computer Science*, pages 126–137. Springer, 2013.

[Fin75] Kit Fine. Normal forms in modal logic. *Notre Dame Journal of Formal Logic*, 16(2):229–237, 1975.

[Fit07] Melvin Fitting. 2 Modal proof theory. In Johan Van Benthem Patrick Blackburn and Frank Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical Reasoning*, pages 85 – 138. Elsevier, 2007.

[Gal12] Pietro Galliani. Inclusion and exclusion dependencies in team semantics - On some logics of imperfect information. *Ann. Pure Appl. Logic*, 163(1):68–84, 2012.

[Gal15] Pietro Galliani. Upwards closed dependencies in team semantics. *Information and Computation*, 245:124–135, December 2015.

[Gal16] Pietro Galliani. On strongly first-order dependencies. In *Dependence Logic*, pages 53–71. Springer, 2016.

[Gal18] Pietro Galliani. Safe Dependency Atoms and Possibility Operators in Team Semantics. In *GandALF*, volume 277 of *EPTCS*, pages 58–72, 2018.

[GHK13] Pietro Galliani, Miika Hannula, and Juha Kontinen. Hierarchies in independence logic. In *CSL*, volume 23 of *LIPIcs*, pages 263–280. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.

[GO07] Valentin Goranko and Martin Otto. 5 Model theory of modal logic. In Johan Van Benthem Patrick Blackburn and Frank Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical Reasoning*, pages 249 – 329. Elsevier, 2007.

[GV13] Erich Grädel and Jouko A. Väänänen. Dependence and Independence. *Studia Logica*, 101(2):399–410, 2013.

[Han16] Miika Hannula. Validity and Entailment in Modal and Propositional Dependence Logics. *CoRR*, abs/1608.04301, 2016.

[Han17] Miika Hannula. Validity and Entailment in Modal and Propositional Dependence Logics. In *CSL*, volume 82 of *LIPIcs*, pages 28:1–28:17. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.

[HK15] Miika Hannula and Juha Kontinen. Hierarchies in independence and inclusion logic with strict semantics. *J. Log. Comput.*, 25(3):879–897, 2015.

[HKLV16] Miika Hannula, Juha Kontinen, Martin Lück, and Jonni Virtema. On Quantified Propositional Logics and the Exponential Time Hierarchy. In *GandALF*, volume 226 of *EPTCS*, pages 198–212, 2016.

[HKMV15] Lauri Hella, Antti Kuusisto, Arne Meier, and Heribert Vollmer. Modal Inclusion Logic: Being Lax is Simpler than Being Strict. In *MFCS*, volume 9234 of *LNCS*, pages 281–292. Springer, 2015.

[HKMV17] Lauri Hella, Antti Kuusisto, Arne Meier, and Jonni Virtema. Model Checking and Validity in Propositional and Modal Inclusion Logics. In *MFCS*, volume 83 of *LIPIcs*, pages 32:1–32:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.

[HKVV15] Miika Hannula, Juha Kontinen, Jonni Virtema, and Heribert Vollmer. Complexity of Propositional Independence and Inclusion Logic. In *MFCS*, volume 9234 of *LNCS*, pages 269–280. Springer, 2015.

[HKVV18] Miika Hannula, Juha Kontinen, Jonni Virtema, and Heribert Vollmer. Complexity of Propositional Logics in Team Semantic. *ACM Trans. Comput. Log.*, 19(1):2:1–2:14, 2018.

[HLSV14] Lauri Hella, Kerkko Luosto, Katsuhiko Sano, and Jonni Virtema. The Expressive Power of Modal Dependence Logic. In *Advances in Modal Logic*, pages 294–312. College Publications, 2014.

[Hod97] Wilfrid Hodges. Compositional Semantics for a Language of Imperfect Information. *Logic Journal of the IGPL*, 5(4):539–563, 1997.

[HS89] Jaakko Hintikka and Gabriel Sandu. Informational Independence as a Semantical Phenomenon. In Jens Erik Fenstad, Ivan T. Frolov, and Risto Hilpinen, editors, *Logic, Methodology and Philosophy of Science VIII*, volume 126 of *Studies in Logic and the Foundations of Mathematics*, pages 571 – 589. Elsevier, 1989.

[HS15] Lauri Hella and Johanna Stumpf. The expressive power of modal logic with inclusion atoms. In *GandALF*, volume 193 of *EPTCS*, pages 129–143, 2015.

[HV16] Lauri Hella and Miikka Vilander. The succinctness of first-order logic over modal logic via a formula size game. In *Advances in Modal Logic*, pages 401–419. College Publications, 2016.

[KMSV15] Juha Kontinen, Julian-Steffen Müller, Henning Schnoor, and Heribert Vollmer. A Van Benthem Theorem for Modal Team Semantics. In *CSL*, volume 41 of *LIPIcs*, pages 277–291. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.

[KMSV17] Juha Kontinen, Julian-Steffen Müller, Henning Schnoor, and Heribert Vollmer. Modal independence logic. *J. Log. Comput.*, 27(5):1333–1352, 2017.

[KMV15] Andreas Krebs, Arne Meier, and Jonni Virtema. A Team Based Variant of CTL. In *TIME*, pages 140–149. IEEE Computer Society, 2015.

[KMVZ18] Andreas Krebs, Arne Meier, Jonni Virtema, and Martin Zimmermann. Team semantics for the specification and verification of hyperproperties. In *MFCS*, volume 117 of *LIPIcs*, pages 10:1–10:16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018.

[Lad77] Richard E. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM J. Comput.*, 6(3):467–480, 1977.

[Lüc17] Martin Lück. The Power of the Filtration Technique for Modal Logics with Team Semantics. In *CSL*, volume 82 of *LIPIcs*, pages 31:1–31:20. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.

[Lüc18a] Martin Lück. Axiomatizations of team logics. *Ann. Pure Appl. Logic*, 169(9):928–969, 2018.

[Lüc18b] Martin Lück. Canonical Models and the Complexity of Modal Team Logic. In *CSL*, volume 119 of *LIPIcs*, pages 30:1–30:23. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018.

[Lüc18c] Martin Lück. On the Complexity of Team Logic and Its Two-Variable Fragment. In *MFCS*, volume 117 of *LIPIcs*, pages 27:1–27:22. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018.

[LV13] Peter Lohmann and Heribert Vollmer. Complexity Results for Modal Dependence Logic. *Studia Logica*, 101(2):343–366, 2013.

[Mey74] Albert R Meyer. The inherent computational complexity of theories of ordered sets. In *Proc. Int. Congress of Mathematicians*, volume 19722, page 481, 1974.

[Mos07] Lawrence S. Moss. Finite Models Constructed from Canonical Formulars. *J. Philosophical Logic*, 36(6):605–640, 2007.

[Mül14] Julian-Steffen Müller. *Satisfiability and model checking in team based logics*. PhD thesis, University of Hanover, 2014.

[MV97] Maarten Marx and Yde Venema. Multi-dimensional modal logic. In *Multi-Dimensional Modal Logic*, pages 1–9. Springer, 1997.

[Sch16] Sylvain Schmitz. Complexity Hierarchies beyond Elementary. *TOCT*, 8(1):3:1–3:36, 2016.

[Sev09] Merlijn Sevenster. Model-theoretic and Computational Properties of Modal Dependence Logic. *J. Log. Comput.*, 19(6):1157–1173, 2009.

[SM73] Larry J. Stockmeyer and Albert R. Meyer. Word Problems Requiring Exponential Time: Preliminary Report. In *Proc. of the 5th Annual ACM Sym. on Theory of Computing*, pages 1–9, 1973.

[Vää07] Jouko Väänänen. *Dependence logic: A New Approach to Independence Friendly Logic*. Number 70 in London Mathematical Society student texts. Cambridge University Press, Cambridge ; New York, 2007.

[Vää08] Jouko Väänänen. Modal dependence logic. *New perspectives on games and interaction*, 4:237–254, 2008.

[Vir17] Jonni Virtema. Complexity of validity for propositional dependence logics. *Inf. Comput.*, 253:224–236, 2017.

[Voi17] Marco Voigt. A fine-grained hierarchy of hard problems in the separated fragment. In *LICS*, pages 1–12. IEEE Computer Society, 2017.

[YV16] Fan Yang and Jouko Väänänen. Propositional Logics of Dependence. *Ann. Pure Appl. Logic*, 167(7):557–589, 2016.

[YV17] Fan Yang and Jouko Väänänen. Propositional team logics. *Ann. Pure Appl. Logic*, 168(7):1406–1441, 2017.

Appendix A. Proof details

In the appendix, we include several propositions that have straightforward but lengthy proofs.

Proofs of Section 4.

Proposition 4.3. *Let $\Phi \subseteq \mathcal{PS}$ be finite and $k \geq 0$.*

(1) $[\![w]\!]_k^\Phi \cap \Phi = V^{-1}(w) \cap \Phi$ *and* $[\![Rw]\!]_k^\Phi = \mathcal{R}[\![w]\!]_{k+1}^\Phi$*, for all pointed structures* (W, R, V, w).
(2) *The mapping* $h\colon \tau \mapsto \tau \cap \Phi$ *is a bijection from* Δ_0^Φ *to* $\mathfrak{P}(\Phi)$.
(3) *The mapping* $h\colon \tau \mapsto (\tau \cap \Phi, \mathcal{R}\tau)$ *is a bijection from* Δ_{k+1}^Φ *to* $\mathfrak{P}(\Phi) \times \mathfrak{P}(\Delta_k^\Phi)$.

- *Proof of (1).* Assume $(W, R, V, v), \Phi \subseteq \mathcal{PS}$ and $k \geq 0$ as above. For all $p \in \Phi$, clearly $p \in [\![w]\!]_k^\Phi$ iff $w \vDash p$ iff $p \in V^{-1}(w)$. Next, we show that $[\![Rw]\!]_k^\Phi = \mathcal{R}[\![w]\!]_{k+1}^\Phi$. Let $\tau = [\![w]\!]_{k+1}^\Phi$, and recall that $\mathcal{R}\tau = \{\tau' \in \Delta_k^\Phi \mid \{\alpha \mid \Box\alpha \in \tau\} \subseteq \tau'\}$. To prove $[\![Rw]\!]_k^\Phi \subseteq \mathcal{R}\tau$, let $\tau' \in [\![Rw]\!]_k^\Phi$ be arbitrary. Then $[\![v]\!]_k^\Phi = \tau'$ for some $v \in Rw$. Now, for all $\alpha \in \mathsf{ML}_k^\Phi$, $\Box\alpha \in \tau$ implies $w \vDash \Box\alpha$. In particular, $v \vDash \alpha$, i.e., $\alpha \in \tau'$. Hence, $\{\alpha \mid \Box\alpha \in \tau\} \subseteq \tau'$, which implies $\tau' \in \mathcal{R}\tau$.

  For the converse direction, $\mathcal{R}\tau \subseteq [\![Rw]\!]_k^\Phi$, let $\tau' \in \mathcal{R}\tau$ be arbitrary. By definition, $\{\alpha \mid \Box\alpha \in \tau\} \subseteq \tau'$. Since τ' is a k-type, it has a model $(\mathcal{K}', v')$, and due to Proposition 4.2, $[\![\mathcal{K}', v']\!]_k^\Phi = \tau'$. By Proposition 3.2, there is a formula $\zeta \in \mathsf{ML}_k^\Phi$ such that $(\mathcal{K}'', v'') \vDash \zeta$ if and only if $(\mathcal{K}', v') \rightleftharpoons_k^\Phi (\mathcal{K}'', v'')$. As τ is a $(k+1)$-type, either $\Diamond\zeta \in \tau$ or $\neg\Diamond\zeta \in \tau$.

  First, suppose $\neg\Diamond\zeta \in \tau$. Then $\Box\neg\zeta \in \tau$, hence $\neg\zeta \in \tau'$ by definition of τ'. But as $(\mathcal{K}', v') \vDash \tau'$, then both $(\mathcal{K}', v') \nvDash \zeta$ and $(\mathcal{K}', v') \vDash \zeta$, as $(\mathcal{K}', v') \rightleftharpoons_k^\Phi (\mathcal{K}', v')$. Contradiction, therefore $\Diamond\zeta \in \tau$. Consequently, w has an R-successor v such that $v \vDash \zeta$, i.e., $\tau' = [\![v]\!]_k^\Phi \in [\![Rw]\!]_k^\Phi$.
- *Proof that h in (2) and (3) is injective.* Let $\tau, \tau' \in \Delta_k^\Phi$ be arbitrary. Let $(\mathcal{K}, w) = (W, R, V, w)$ be of type τ, and $(\mathcal{K}', w') = (W', R', V', w')$ of type τ'. We first consider (2) and demonstrate that $h\colon \tau \mapsto \tau \cap \Phi$ injective. This follows from (1), as $\tau \cap \Phi = \tau' \cap \Phi$ implies $V^{-1}(w) = \tau \cap \Phi = \tau' \cap \Phi = V'^{-1}(w')$, i.e., $(\mathcal{K}, w) \rightleftharpoons_0^\Phi (\mathcal{K}', w')$. By Proposition 4.2, then $\tau = \tau'$.

  For (3), let $k > 0$, and additionally suppose $\mathcal{R}\tau = \mathcal{R}\tau'$. Again by (1), we have $[\![\mathcal{K}, Rw]\!]_{k-1}^\Phi = \mathcal{R}\tau = \mathcal{R}\tau' = [\![\mathcal{K}', R'w']\!]_{k-1}^\Phi$. By Proposition 4.2, $(\mathcal{K}, Rw) \rightleftharpoons_{k-1}^\Phi (\mathcal{K}', R'w')$ follows. Since $(\mathcal{K}, w) \rightleftharpoons_0^\Phi (\mathcal{K}', w')$ holds as before, $(\mathcal{K}, w) \rightleftharpoons_k^\Phi (\mathcal{K}', w')$ by Proposition 3.4. By Proposition 4.2, $\tau = [\![\mathcal{K}, w]\!]_k^\Phi = [\![\mathcal{K}', w']\!]_k^\Phi = \tau'$.
- *Proof that h in (2) and (3) is surjective.* First, consider (2). We have to show that, for all $\Phi' \subseteq \Phi$, there exists a type $\tau \in \Delta_0^\Phi$ such that $\tau \cap \Phi = \Phi'$. Likewise, for (3) we have to show that for all $k \geq 0$, $\Phi' \subseteq \Phi$ and $\Delta' \subseteq \Delta_k^\Phi$, there exists a type $\tau \in \Delta_{k+1}^\Phi$ such that $\tau \cap \Phi = \Phi'$ and $\mathcal{R}\tau = \Delta'$. We show the second statement, as the first one is shown analogously. The following model $(\mathcal{K}, w)$ witnesses that there exists $\tau \in \Delta_{k+1}$ such that $\tau \cap \Phi = \Phi'$ and $\mathcal{R}\tau = \Delta'$. First, recall that each $\tau' \in \Delta'$ has a model $(\mathcal{N}_{\tau'}, v_{\tau'})$ such that, by Proposition 4.2, $[\![\mathcal{N}_{\tau'}, v_{\tau'}]\!]_k^\Phi = \tau'$. Define $\mathcal{K}$ as the disjoint union of all $\mathcal{N}_\tau$ and of a distinct point w, and let $V^{-1}(w) = \Phi'$. By (1), then $[\![w]\!]_{k+1}^\Phi \cap \Phi = \Phi'$. Moreover, let $Rw = \{v_{\tau'} \mid \tau' \in \Delta'\}$. Again due to (1), $\mathcal{R}[\![w]\!]_{k+1}^\Phi = [\![Rw]\!]_k^\Phi$. By definition, $[\![Rw]\!]_k^\Phi = [\![\{v_{\tau'} \mid \tau' \in \Delta'\}]\!]_k^\Phi = \{[\![v_{\tau'}]\!]_k^\Phi \mid \tau' \in \Delta'\} = \Delta'$. ◻

Lemma 4.9. *For every polynomial p there is a polynomial q such that*

$$p(\exp_k^*(n)) \leq \exp_k(q((k+1)\cdot n))$$

for all $k \geq 0$ and $n \geq 1$.

We require the following inequalities.

Lemma A.1. *Let $n, k, c \geq 0$. Then $c + \exp_k(n) \leq \exp_k(c+n)$. If also $n \geq 1$, then $c \cdot \exp_k(n) \leq \exp_k(cn)$.*

Proof. Induction on k, where $k = 0$ is trivial. For $k \geq 1$,

$$\begin{aligned} c + \exp_{k+1}(n) &= c + 2^{\exp_k(n)} \leq 2^c \cdot 2^{\exp_k(n)} && \text{(As } c + a \leq 2^c \cdot a \text{ for } c \geq 0, a \geq 1\text{)} \\ &= 2^{c+\exp_k(n)} \leq 2^{\exp_k(c+n)} && \text{(Induction hypothesis)} \\ &= \exp_{k+1}(c+n). \end{aligned}$$

For the product, the cases $c = 0, 1$ are trivial. For $c \geq 2$,

$$\begin{aligned} c \cdot \exp_{k+1}(n) &\leq 2^{c-1} \cdot 2^{\exp_k(n)} && \text{(Since } c \geq 2 \text{ implies } c \leq 2^{c-1}\text{)} \\ &= 2^{c-1+\exp_k(n)} \leq 2^{\exp_k(c-1+n)} && \text{(By + case)} \\ &\leq 2^{\exp_k(cn)} = \exp_{k+1}(cn). && \text{(As } (c-1)+n \leq cn \text{ for } c, n \geq 1\text{)} \quad \square \end{aligned}$$

Recall that $\exp_0^*(n) := n$ and $\exp_{k+1}^*(n) := n \cdot 2^{\exp_k^*(n)}$.

Lemma A.2. *Let $n, k \geq 0$. Then $\exp_k^*(n) \leq \exp_k((k+1)\cdot n)$.*

Proof. Induction on k. For $k = 0$, $\exp_0^*(n) = n = \exp_0((0+1)\cdot n)$. For the inductive step,

$$\begin{aligned} \exp_{k+1}^*(n) &= n \cdot 2^{\exp_k^*(n)} \leq 2^n \cdot 2^{\exp_k^*(n)} = 2^{n+\exp_k^*(n)} \\ &\leq 2^{n+\exp_k((k+1)n)} && \text{(Induction hypothesis)} \\ &\leq 2^{\exp_k(n+(k+1)n)} = \exp_{k+1}((k+2)n) && \text{(Lemma A.1)} \quad \square \end{aligned}$$

The next inequality states that a polynomial can be "pulled inside" $\exp_k$:

Lemma A.3. *For every polynomial p there is a polynomial q such that $p(\exp_k(n)) \leq \exp_k(q(n)))$ for all $k \geq 0, n \geq 1$.*

Proof. For every polynomial p there are integers $c, d \geq 1$ such that $p(n) \leq cn^d$ for all $n \geq 1$. Let $q(n) := cdn^d + c$. Then the case $k = 0$ is clear. For $k \geq 1$ and $n \geq 1$,

$$\begin{aligned} p(\exp_k(n)) &\leq c \cdot \exp_k(n)^d \leq 2^c \cdot \left(2^{\exp_{k-1}(n)}\right)^d = 2^{c+d\cdot\exp_{k-1}(n)} \\ &\leq 2^{q(\exp_{k-1}(n))} && \text{(As } q(n) \geq c + dn\text{)} \\ &\leq 2^{\exp_{k-1}(q(n))} = \exp_k(q(n)). && \text{(Lemma A.1)} \quad \square \end{aligned}$$

Finally, we combine both lemmas:

Proof of 4.9. Let p be a polynomial as above. W.l.o.g. p is non-decreasing. Then by Lemma A.2, $p(\exp_k^*(n)) \leq p\big(\exp_k((k+1)\cdot n)\big)$. Moreover, due to Lemma A.3, there is a polynomial q such that $p\big(\exp_k((k+1)\cdot n)\big) \leq \exp_k\big(q((k+1)\cdot n)\big)$. $\square$

Proofs of Section 5.

Proposition 5.2. *Let α, β be disjoint scopes and S, U, T teams in a Kripke structure $\mathcal{K} = (W, R, V)$. Then the following laws hold:*

(1) *Distributive laws:* $(T \cap S)_\alpha = T_\alpha \cap S = T \cap S_\alpha = T_\alpha \cap S_\alpha$ *and* $(T \cup S)_\alpha = T_\alpha \cup S_\alpha$.
(2) *Disjoint selection commutes:* $\big(T^\alpha_S\big)^\beta_U = \big(T^\beta_U\big)^\alpha_S$.
(3) *Disjoint selection is independent:* $\big((T^\alpha_S)^\beta_U\big)_\alpha = T_\alpha \cap S$.
(4) *Image and selection commute:* $(RT)_\alpha = \big(R(T_\alpha)\big)_\alpha = R(T_\alpha)$
(5) *Selection propagates: If* $S \subseteq T$, *then* $R\big(T^\alpha_S\big) = (RT)^\alpha_{RS}$.

Proof. (1) Observe that $X_\alpha = X \cap W_\alpha$. Hence, for the union $(T \cup S)_\alpha = (T \cup S) \cap W_\alpha = (T \cap W_\alpha) \cup (S \cap W_\alpha) = T_\alpha \cup S_\alpha$ holds. For the intersection, likewise $(T \cap S) \cap W_\alpha = (T \cap W_\alpha) \cap S = T \cap (W_\alpha \cap S) = (T \cap W_\alpha) \cap (S \cap W_\alpha)$.

(2) Proved in the following equation. We use the fact that $X_{\gamma \wedge \gamma'} = (X_\gamma)_{\gamma'} = (X_{\gamma'})_\gamma = X_{\gamma' \wedge \gamma}$ for all teams X and scopes γ, γ.

$$\big(T^\alpha_S\big)^\beta_U$$
$$= \big(T_{\neg\alpha} \cup (T_\alpha \cap S)\big)_{\neg\beta} \cup \Big(\big(T_{\neg\alpha} \cup (T_\alpha \cap S)\big)_\beta \cap U\Big)$$

Distributing all scopes according to (1):

$$= T_{\neg\alpha\wedge\neg\beta} \cup \big(T_{\alpha\wedge\neg\beta} \cap S_{\neg\beta}\big) \cup \big(T_{\neg\alpha\wedge\beta} \cap U\big) \cup \big(T_{\alpha\wedge\beta} \cap S_\beta \cap U\big)$$

Replace U by $U_{\neg\alpha}/U_\alpha$ due to the intersection law of (1):

$$= T_{\neg\alpha\wedge\neg\beta} \cup \big(T_{\alpha\wedge\neg\beta} \cap S_{\neg\beta}\big) \cup \big(T_{\neg\alpha\wedge\beta} \cap U_{\neg\alpha}\big) \cup \big(T_{\alpha\wedge\beta} \cap S_\beta \cap U_\alpha\big)$$

Likewise, replace $S_{\neg\beta}/S_\beta$ by S:

$$= T_{\neg\alpha\wedge\neg\beta} \cup \big(T_{\alpha\wedge\neg\beta} \cap S\big) \cup \big(T_{\neg\alpha\wedge\beta} \cap U_{\neg\alpha}\big) \cup \big(T_{\alpha\wedge\beta} \cap S \cap U_\alpha\big)$$

Reverse distribution of scopes:

$$= \big(T_{\neg\beta} \cup (T_\beta \cap U)\big)_{\neg\alpha} \cup \Big(\big(T_{\neg\beta} \cup (T_\beta \cap U)\big)_\alpha \cap S\Big)$$
$$= \big(T^\beta_U\big)^\alpha_S.$$

(3) By definition and application of (2), $\big((T^\alpha_S)^\beta_U\big)_\alpha$ equals

$$\Big[\big(T_{\neg\beta} \cup (T_\beta \cap U)\big)_{\neg\alpha} \cup \Big(\big(T_{\neg\beta} \cup (T_\beta \cap U)\big)_\alpha \cap S\Big)\Big]_\alpha$$
$$= \big(T_{\neg\beta} \cup (T_\beta \cap U)\big)_{\neg\alpha\wedge\alpha} \cup \Big(\big(T_{\neg\beta} \cup (T_\beta \cap U)\big)_\alpha \cap S\Big)_\alpha$$
$$= \emptyset \cup \Big(\big(T_{\neg\beta} \cup (T_\beta \cap U)\big)_\alpha \cap S_\alpha\Big)$$
$$= \big(T_{\neg\beta\wedge\alpha} \cap S_\alpha\big) \cup \big(T_{\beta\wedge\alpha} \cap U_\alpha \cap S_\alpha\big)$$

Since α and β are disjoint:

$$= \big(T_\alpha \cap S_\alpha\big) \cup (\emptyset \cap U_\alpha \cap S_\alpha) = T_\alpha \cap S.$$

(4) $(RT)_\alpha \subseteq \big(R(T_\alpha)\big)_\alpha$: Suppose $v \in (RT)_\alpha$. Then $v \in Rw$ for some $w \in T$. Moreover, $w \in T_\alpha$, since α is a scope. Hence $v \in R(T_\alpha)$. As $v \vDash \alpha$, $v \in \big(R(T_\alpha)\big)_\alpha$ follows.
$\big(R(T_\alpha)\big)_\alpha \subseteq R(T_\alpha)$: Obvious.

$R(T_\alpha) \subseteq (RT)_\alpha$: Again, let $v \in R(T_\alpha)$ be arbitrary. Then $v \in Rw$ for some $w \in T_\alpha$. Hence $v \in RT$. Since $v \vDash \alpha$ follows from $w \vDash \alpha$, we conclude $v \in (RT)_\alpha$.

(5) For "$\subseteq$", suppose $v \in R(T^\alpha_S)$, i.e., $v \in Rw$ for some $w \in T^\alpha_S$. In particular, $v \in RT$. If $w \nvDash \alpha$, then $v \in RT_{\neg\alpha}$ and trivially $v \in (RT)^\alpha_{RS}$. If $w \vDash \alpha$, then necessarily $w \in S$. Moreover, $v \vDash \alpha$. Consequently, $v \in RS_\alpha \cap RT_\alpha$, hence $v \in (RT)^\alpha_{RS}$.

For "$\supseteq$", suppose $v \in (RT)^\alpha_{RS} = RT_{\neg\alpha} \cup (RT_\alpha \cap RS)$.

If $v \in RT_{\neg\alpha}$, then by (4) $v \in Rw$ for some $w \in T_{\neg\alpha}$. In particular, $w \in T^\alpha_S$, hence $v \in R(T^\alpha_S)$.

If $v \in RT_\alpha \cap RS$, then by (1) $v \in RS_\alpha$. By (4) $v \in R(S_\alpha)$, in other words, $v \in Rw$ for some $w \in S_\alpha$. As $S \subseteq T$, then $w \in S_\alpha \cap T$, and in fact $w \in T_\alpha \cap S$ due to (1) Consequently, $w \in T^\alpha_S$ and $v \in R(T^\alpha_S)$. □

Proofs of Section 7.

Lemma 7.1. *Let $\alpha, \beta \in \mathsf{ML}$ and $\varphi \in \mathsf{MTL}_k$. Let T be a team such that $R^iT \vDash \alpha \leftrightarrow \beta$ for all $i \in \{0, \dots, k\}$. Then $T \vDash \varphi$ if and only if $T \vDash \mathrm{Sub}(\varphi, \alpha, \beta)$, where $\mathrm{Sub}(\varphi, \alpha, \beta)$ is the formula obtained from φ by substituting every occurrence of α with β.*

Proof. Proof by induction on k and the syntax on φ. W.l.o.g. α occurs in φ. If $\varphi = \alpha$, then $\mathrm{Sub}(\varphi, \alpha, \beta) = \beta$, in which case the proof boils down to showing $T \vDash \alpha \Leftrightarrow T \vDash \beta$. However, this easily follows from $T \vDash \alpha \leftrightarrow \beta$ by the semantics for classical ML-formulas.

Otherwise, α is a proper subformula of φ. We distinguish the following cases.

- $\varphi = \neg\gamma$: Then $\mathrm{Sub}(\neg\gamma, \alpha, \beta) = \neg\mathrm{Sub}(\gamma, \alpha, \beta)$, and

$$\begin{aligned}
&T \vDash \mathrm{Sub}(\varphi, \alpha, \beta)\\
\Leftrightarrow\ &T \vDash \neg\mathrm{Sub}(\gamma, \alpha, \beta)\\
\Leftrightarrow\ &\forall w \in T\colon \{w\} \vDash \neg\mathrm{Sub}(\gamma, \alpha, \beta)\\
\Leftrightarrow\ &\forall w \in T\colon \{w\} \vDash \neg\gamma && \text{(Induction hypothesis, as } \{w\}, Rw, \dots \vDash \alpha \leftrightarrow \beta)\\
\Leftrightarrow\ &T \vDash \neg\gamma\\
\Leftrightarrow\ &T \vDash \varphi
\end{aligned}$$

- $\varphi = {\sim}\psi$: By induction hypothesis, $T \vDash \mathrm{Sub}(\varphi, \alpha, \beta)$ iff $T \vDash {\sim}\mathrm{Sub}(\psi, \alpha, \beta)$ iff $T \vDash {\sim}\psi$.
- $\varphi = \psi \wedge \theta$: Proved similarly to $\sim$.
- $\varphi = \psi \vee \theta$: First note that $\mathrm{Sub}(\psi \vee \theta, \alpha, \beta) = \mathrm{Sub}(\psi, \alpha, \beta) \vee \mathrm{Sub}(\theta, \alpha, \beta)$. Then:

$$\begin{aligned}
&T \vDash \mathrm{Sub}(\varphi, \alpha, \beta)\\
\Leftrightarrow\ &T \vDash \mathrm{Sub}(\psi, \alpha, \beta) \vee \mathrm{Sub}(\theta, \alpha, \beta)\\
\Leftrightarrow\ &\exists S, U\colon T = S \cup U, S \vDash \mathrm{Sub}(\psi, \alpha, \beta), U \vDash \mathrm{Sub}(\theta, \alpha, \beta)
\end{aligned}$$

  By induction hypothesis, since $S, U, RS, RU, \dots \vDash \alpha \leftrightarrow \beta$:

$$\begin{aligned}
\Leftrightarrow\ &\exists S, U\colon T = S \cup U, S \vDash \psi, U \vDash \theta\\
\Leftrightarrow\ &T \vDash \varphi
\end{aligned}$$

- $\varphi = \Box\psi$: We have $\mathrm{Sub}(\Box\psi, \alpha, \beta) = \Box\mathrm{Sub}(\psi, \alpha, \beta)$, hence

$$\begin{aligned}
&T \vDash \mathrm{Sub}(\varphi, \alpha, \beta)\\
\Leftrightarrow\ &T \vDash \Box\mathrm{Sub}(\psi, \alpha, \beta)\\
\Leftrightarrow\ &RT \vDash \mathrm{Sub}(\psi, \alpha, \beta).
\end{aligned}$$

However, since $\psi \in \mathsf{MTL}_{k-1}$ and $RT, \ldots, R^{k-1}(RT) \vDash \alpha \leftrightarrow \beta$ holds by assumption, we obtain by induction hypothesis:

$$\begin{aligned} &\Leftrightarrow RT \vDash \psi \\ &\Leftrightarrow T \vDash \varphi \end{aligned}$$

- $\varphi = \Diamond\psi$: As before, $\mathrm{Sub}(\Diamond\psi, \alpha, \beta) = \Diamond\mathrm{Sub}(\psi, \alpha, \beta)$. Then:

$$\begin{aligned} &T \vDash \mathrm{Sub}(\varphi, \alpha, \beta) \\ \Leftrightarrow\ &T \vDash \Diamond\mathrm{Sub}(\psi, \alpha, \beta) \\ \Leftrightarrow\ &\exists S \subseteq RT, T \subseteq R^{-1}S \colon S \vDash \mathrm{Sub}(\psi, \alpha, \beta) \end{aligned}$$

Note that $S, RS, \ldots, R^{k-1}S$ are subteams of $RT, \ldots, R^k T$, respectively. For this reason, the teams $S, RS, \ldots, R^{k-1}S$ satisfy $\alpha \leftrightarrow \beta$ as well. As also $\psi \in \mathsf{MTL}_{k-1}$ holds, we obtain by induction hypothesis:

$$\begin{aligned} \Leftrightarrow\ &\exists S \subseteq RT, T \subseteq R^{-1}S \colon S \vDash \psi \\ \Leftrightarrow\ &T \vDash \varphi \end{aligned}$$

□