Modal Logics and Local Quantifiers: A Zoo in the Elementary Hierarchy

Raul Fervari[1] and Alessio Mansutti[2](✉)

[1] CONICET and Universidad Nacional de Córdoba, Córdoba, Argentina
rfervari@unc.edu.ar
[2] Department of Computer Science, University of Oxford, Oxford, UK
alessio.mansutti@cs.ox.ac.uk

Abstract. We study a family of *modal logics* interpreted on tree-like structures, and featuring *local quantifiers* $\exists^k p$ that bind the proposition p to worlds that are accessible from the current one in at most k steps. We consider a first-order and a second-order semantics for the quantifiers, which enables us to relate several well-known formalisms, such as *hybrid logics*, S5Q and *graded modal logic*. To better stress these connections, we explore fragments of our logics, called herein *round-bounded* fragments. Depending on whether first or second-order semantics is considered, these fragments populate the hierarchy 2NExp $\subset$ 3NExp $\subset \cdots$ or the hierarchy 2AExp$_{pol}$ $\subset$ 3AExp$_{pol}$ $\subset \cdots$, respectively. For formulae up-to modal depth k, the complexity improves by one exponential.

1 Introduction

From a traditional perspective, *modal logics* [10] are formalisms to reason about different modes of truth. However, another view consists of seeing these logics as computationally well-behaved fragments of *first-order logic* and *second-order logic* (see e.g., [1] for a discussion). Some examples of well-known modal logics with a good balance between expressivity and computational complexity are *graded modal logic* (GML) [5,28], whose satisfiability problem is PSpace-complete; and the *temporal logics* LTL, CTL and CTL* whose satisfiability problems are complete for PSpace, Exp and 2Exp, respectively [31,19,25].

A family of logics that elude this nice computational picture is that made of modal logics enriched with first-order or second-order propositional quantifiers $\exists p$, which update the set of worlds of a Kripke structure that satisfy the propositional symbol p. The literature of modal logics featuring quantification over propositional symbols can be traced back to [12,26,18]. All these works show that, in spite of the simplicity of the principle, propositional quantification leads to undecidability very quickly. One of the few exceptions is the logic S5Q, i.e. S5 enriched with *second-order propositional quantifiers*, which enjoys an exponential-size small model property, and is thus decidable [22,18]. Here, the success in finding a well-behaved framework for propositional quantification is due to the fact that S5 has a very restricted class of models. In modern literature, the family of *hybrid logics* [2] is one of the most relevant approaches offering *first-order propositional quantification*. Most hybrid logics provide *operators* $\downarrow i$ that binds the current world to the proposition i, and $@_i$ that allows to jump to the world bound to i. This form of quantification is very expressive, and leads to undecidability over standard Kripke structures [3]. To regain decidability, one can restrict the logic to syntactical fragments that avoid the quantification patters $\Box\downarrow$ and $\Diamond\downarrow\Diamond$, or restrict the interpretation to models in which each world has at most two successors [14]. Again, one can also simply consider $\mathsf{S5}$ models: the hybrid logic with $\downarrow$ and $@$ on $\mathsf{S5}$ is known to admit an NExp-complete satisfiability problem [30].

Recent works shed new lights on the role of propositional quantifiers. From a model theoretical perspective, a revision about the different forms of propositional quantification has been put forward in [9]. Novel algebraic insights on $\mathsf{S5}$ with propositional quantification have been discovered in [17]. From a computational perspective, [6] shows that second-order propositional quantification is enough to obtain Tower-complete (hence, non-elementary decidable, [29]) logics on tree-like structures. This last result is of interest, as the second-order logic $\mathsf{QCTL}^t_{\mathbf{X}}$ considered in [6] subsumes several other modal logics with forms of quantification "in disguise", such as the aforementioned GML, as well as *modal separation logics* [16], *ambient logics* [13] and *team logics* [21]. However, when translated into $\mathsf{QCTL}^t_{\mathbf{X}}$, the good computational properties of these logics are lost, and the Tower-hardness of $\mathsf{QCTL}^t_{\mathbf{X}}$ prevents us to grasp the real capabilities of their (often restricted) form of propositional quantifications.

Contributions. The overall message of [6] is that the computational power of propositional quantification in the context of modal logic deserves to be better understood. Driven by this message, we investigate from a unified perspective a family of logics interpreted on tree-like models, featuring a very intuitive form of propositional quantification: the *local quantifier* $\exists^k p$, with $k \geq 1$ integer, that binds the propositional symbol p to world(s) occurring within distance k from the current point of evaluation. More precisely, we look at two families of modal logics: the family $\mathsf{ML}(\exists^1_{\mathit{FO}}), \mathsf{ML}(\exists^2_{\mathit{FO}}), \cdots$, where $\mathsf{ML}(\exists^k_{\mathit{FO}})$ extends the basic modal logic ML with the *first-order* local quantifier $\exists^k p$ binding p to exactly one world occurring within distance k of the current world; and the family $\mathsf{ML}(\exists^1_{\mathit{SO}}), \mathsf{ML}(\exists^2_{\mathit{SO}}), \cdots$, where $\mathsf{ML}(\exists^k_{\mathit{SO}})$ extends ML with the *second-order* local quantifier $\exists^k p$ binding p to a set of worlds occurring within distance k.

As previously mentioned, in introducing these logics our aim is to better understand the similarities and differences between the various modal logics featuring propositional quantification, especially when it comes to their complexity. This analysis cannot be done using Tower-complete logics like $\mathsf{QCTL}^t_{\mathbf{X}}$, as finer complexity classes are required. In this sense, it is worth to notice that our framework features the logic $\mathsf{ML}(\exists^{\infty}_{\mathit{SO}})$, whose quantifier $\exists^{\infty} p$ binds p to arbitrary worlds reachable from the current one. This is exactly the logic $\mathsf{QCTL}^t_{\mathbf{X}}$. Because of this connection and of similarities with other frameworks, e.g. [7], we argue that even if we restrict ourselves to quantifiers $\exists^k$ with small k, the complexity does not improve. In fact, $\mathsf{ML}(\exists^2_{\mathit{FO}})$ is already Tower-complete, although we defer this result to an extended version of the paper, due to the lack of space.

Consequently, to pursue our goal of a fine-grained analysis of the computational power of propositional quantification in modal logic, in this paper we focus on a syntactical restriction for $\mathsf{ML}(\exists^k_{\mathit{FO}})$ and $\mathsf{ML}(\exists^k_{\mathit{SO}})$ where the local quantifiers are *round-bounded* (Sec. 2). Roughly speaking, under the round-bounded condition, $\mathsf{ML}(\exists^k_{\mathit{FO}})$ and $\mathsf{ML}(\exists^k_{\mathit{SO}})$ formulae can be split into parts having k nested modalities. Quantifiers belonging to one part of the formula do not interact with quantifiers from other parts of the formula. The following results are established.

Theorem 1. *The sat. problem for round-bounded* $\mathsf{ML}(\exists^k_{\mathit{FO}})$ *is* $(k{+}1)$NEXP*-complete. It is* kNEXP*-complete for formulae of* $\mathsf{ML}(\exists^k_{\mathit{FO}})$ *of modal depth* k*.*

Theorem 2. *The sat. problem for round-bounded* $\mathsf{ML}(\exists^k_{\mathit{SO}})$ *is* $(k{+}1)$AEXP$_{pol}$*-complete. It is* kAEXP$_{pol}$*-complete for formulae of* $\mathsf{ML}(\exists^k_{\mathit{SO}})$ *of modal depth* k*.*

Here and along the paper, given natural numbers $k, n \geq 1$, we write $\mathfrak{t}$ for the tetration function inductively defined as $\mathfrak{t}(0, n) \stackrel{\text{def}}{=} n$ and $\mathfrak{t}(k, n) = 2^{\mathfrak{t}(k-1,n)}$. Intuitively, $\mathfrak{t}(k, n)$ defines a tower of exponentials of height k. Then, kNEXP is the class of all problems decidable by a non-deterministic Turing machine running in time $\mathfrak{t}(k, f(n))$, for some polynomial f, on each input of length n; whereas kAEXP$_{pol}$ is the class of all problems decidable with an alternating Turing machine [15] in time $\mathfrak{t}(k, f(n))$ and performing at most $g(n)$ alternations, for some polynomials f, g, on each input of length n. For all $k \geq 1$, kNEXP $\subseteq$ kAEXP$_{pol}$ $\subseteq$ TOWER, as we recall that TOWER is the class of all problems decidable with a Turing machine running in time $\mathfrak{t}(g(n), f(n))$ for some polynomial f and elementary function g, on each input of length n [29]. The lower bounds of Thms. 1 and 2 are established by reduction from suitable tiling problems (Sec. 3). The upper bounds are established by designing a quantifier elimination procedure that yields a $(k+1)$EXPSPACE small-model property for round-bounded $\mathsf{ML}(\exists^k_{\mathit{SO}})$, and a kEXPSPACE small-model property for the set of formulae of $\mathsf{ML}(\exists^k_{\mathit{SO}})$ of modal depth k (Sec. 4). The round-bounded condition does not change the set of formulae of $\mathsf{ML}(\exists^1_{\mathit{FO}})$ and $\mathsf{ML}(\exists^1_{\mathit{SO}})$, and thus, as a corollary, we characterise the complexity of these logics:

Corollary 1. *(I) The sat. problem for* $\mathsf{ML}(\exists^1_{\mathit{FO}})$ *is* 2NEXP*-complete.*
(II) The sat. problem for $\mathsf{ML}(\exists^1_{\mathit{SO}})$ *is* 2AEXP$_{pol}$*-complete.*

As promised, our framework yields a refined analysis on the power of propositional quantification in modal logic, which we compare to previous known results in Sec. 2. Quite surprisingly, we show that, on tree-like models, modal logic enriched with propositional quantifiers is as expressive as graded modal logic. Moreover, we establish that $\mathsf{S5Q}$ is AEXP$_{pol}$-complete (refining the previous results from [22,18]), and that hybrid logic with $\downarrow$ and @ on trees is TOWER-complete.

2 Preliminaries

The symbol $\mathbb{N}$ (resp. $\mathbb{N}_+$) denotes the set of *natural numbers* including (resp. excluding) zero, $\overline{\mathbb{N}}$ denotes the set $\mathbb{N} \cup \{\infty\}$, where $n < \infty$, $\infty + n = \infty$ and $n \bmod \infty = n$ for all $n \in \mathbb{N}$, and $\overline{\mathbb{N}}_+ \stackrel{\text{def}}{=} \overline{\mathbb{N}} \setminus \{0\}$. We write $|S| \in \overline{\mathbb{N}}$ for the size of a set S. Finally, let $\mathsf{AP} = \{p, q, r, \dots\}$ be a countable set of *atomic propositions*.

Kripke structures. A *Kripke structure* is a triple $\mathcal{K} = (\mathcal{W}, R, \mathcal{V})$ where $\mathcal{W}$ is a non-empty set of *worlds*, $\mathcal{V}\colon \mathsf{AP} \to 2^{\mathcal{W}}$ is a *valuation*, and $R \subseteq \mathcal{W} \times \mathcal{W}$ is a binary *accessibility relation*. A *Kripke-style forest* is a Kripke structure whose accessibility relation R is such that its inverse R^{-1} is functional and acyclic. In particular, the graph described by $\mathcal{K}$ is a collection of disjoint trees, where R encodes the child relation. We write $R(w)$ for the set of *children* of w, i.e. $\{w' \in \mathcal{W} : (w, w') \in R\}$. For $i \in \mathbb{N}$, R^i is the *i-th composition* of R: R^0 is the identity map on $\mathcal{W}$, and $R^{i+1} \stackrel{\text{def}}{=} \{(w, w') \in \mathcal{W} \times \mathcal{W} : (w, w'') \in R^i$ and $(w'', w') \in R$, for some $w'' \in \mathcal{W}\}$. For $n, m \in \overline{\mathbb{N}}$, $R^{[n,m]} \stackrel{\text{def}}{=} \bigcup_{j=n}^{m} R^j$, and $R^* \stackrel{\text{def}}{=} R^{[0,\infty]}$ is the *Kleene closure* of R. For $\mathcal{W}' \subseteq \mathcal{W}$, $\mathcal{V}[p \leftarrow \mathcal{W}']$ is the valuation obtained from $\mathcal{V}$ by updating to $\mathcal{W}'$ the set assigned to $p \in \mathsf{AP}$. A *pointed forest* $(\mathcal{K}, w)$ is a Kripke-style finite forest $\mathcal{K}$ together with one of its worlds w.

Modal logic with local quantifiers. For $k \in \overline{\mathbb{N}}_+$ written in unary, we introduce the modal logic $\mathsf{ML}(\exists^k)$, whose formulae φ, ψ, χ, etc., are from the grammar below:

$$\varphi, \psi \ := \ \top \ \mid \ p \ \mid \ \varphi \wedge \psi \ \mid \ \neg\varphi \ \mid \ \Diamond\varphi \ \mid \ \exists^k p\, \varphi, \qquad \text{where } p \in \mathsf{AP}.$$

We call $\exists^k p$ a *local (existential) quantifier*. We are interested in two interpretations for the logic $\mathsf{ML}(\exists^k)$, one where the local quantifier $\exists^k p$ performs a first-order quantification, and one where it performs a second-order one. For simplicity, $\mathsf{ML}(\exists^k_{\mathit{FO}})$ (resp. $\mathsf{ML}(\exists^k_{\mathit{SO}})$) stands for $\mathsf{ML}(\exists^k)$ interpreted under first-order (resp. second-order) semantics. The basic modal logic ML is obtained by removing the constructor $\exists^k p\, \varphi$ from the grammar.

Let $(\mathcal{K}, w)$ be a pointed forest, where $\mathcal{K} = (\mathcal{W}, R, \mathcal{V})$. For formulae of $\mathsf{ML}(\exists^k_{\mathit{FO}})$, the satisfaction relation $\models$ is defined as follows (Boolean cases are omitted):

$$\mathcal{K}, w \models p \quad \Leftrightarrow w \in \mathcal{V}(p); \quad \mathcal{K}, w \models \Diamond\varphi \ \Leftrightarrow \ \text{there is } w' \in R(w) \text{ s.t. } \mathcal{K}, w' \models \varphi;$$
$$\mathcal{K}, w \models \exists^k p\, \varphi \Leftrightarrow \text{there is } w' \in R^{[0,k]}(w) \text{ such that } (\mathcal{W}, R, \mathcal{V}[p \leftarrow \{w'\}]), w \models \varphi.$$

An atomic proposition p is said to be a *nominal* for $(\mathcal{K}, w)$ whenever $|\mathcal{V}(p)| = 1$. Additionally, p is *i-local* whenever $\mathcal{V}(p) \subseteq R^i(w)$. In particular, the first-order quantification $\exists^k p\, \varphi$ leads to φ being evaluated in a pointed forest where p is an *i-local nominal* for some $i \in [0, k]$. Given a nominal p, we call $w \in \mathcal{V}(p)$ the *world corresponding to* p, and often denote it by w_p.

For formulae of the *second-order logic* $\mathsf{ML}(\exists^k_{\mathit{SO}})$, the interpretation of the ML fragment remains as for $\mathsf{ML}(\exists^k_{\mathit{FO}})$, whereas we reinterpret the local quantifier as:

$$\mathcal{K}, w \models \exists^k p\, \varphi \Leftrightarrow \text{there is a set } \mathcal{W}' \subseteq R^{[0,k]}(w) \text{ s.t. } (\mathcal{W}, R, \mathcal{V}[p \leftarrow \mathcal{W}']), w \models \varphi.$$

The contradiction $\bot$ and connectives $\vee$, $\Rightarrow$ and $\Leftrightarrow$ are defined as usual. Below, let φ and ψ be two formulae of $\mathsf{ML}(\exists^k)$. The *local universal quantifier* $\forall^k p\, \varphi$ and the modality $\Box\varphi$ are defined as $\neg\exists^k p\, \neg\varphi$ and $\neg\Diamond\neg\varphi$, respectively. We define $\Diamond^0\varphi \stackrel{\text{def}}{=} \varphi$, and given $i \in \mathbb{N}$, $\Diamond^{i+1}\varphi \stackrel{\text{def}}{=} \Diamond^i\Diamond\varphi$. Similarly, $\Box^i\varphi \stackrel{\text{def}}{=} \neg\Diamond^i\neg\varphi$. We write $@^i_p\varphi$ for $\Diamond^i(p \wedge \varphi)$. If p is a nominal, the formula $@^i_p\varphi$ states that p is i-local, and that its corresponding world satisfies φ. We define $\oplus^0\varphi \stackrel{\text{def}}{=} \varphi$ and $\boxplus^0\varphi \stackrel{\text{def}}{=} \varphi$, and given $i \in \mathbb{N}$, $\oplus^{i+1}\varphi \stackrel{\text{def}}{=} \varphi \vee \Diamond \oplus^i \varphi$ and $\boxplus^{i+1}\varphi \stackrel{\text{def}}{=} \varphi \wedge \Box \boxplus^i \varphi$. We use the operator precedence $\{\neg, \Diamond, \Box, \exists^k, \forall^k, @^i_p\} < \{\wedge, \vee\} < \{\Rightarrow, \Leftrightarrow\}$, and sometimes write ":" after a local quantifier with the intuitive meaning that the formula on the right of ":" should be enclosed in brackets, e.g. $\exists^2 p : \varphi \wedge \psi$ abbreviates $\exists^2 p\,(\varphi \wedge \psi)$. Given $i \in \mathbb{N}$, we write $\varphi[\psi \leftarrow_i \chi]$ for the formula obtained from φ by simultaneously substituting with χ each occurrence of the formula ψ appearing under the scope of exactly i nested modalities.

The *length* of φ, denoted with $|\varphi|$, is the number of symbols needed to represent φ. The *modal depth* $\mathsf{md}(\varphi)$ of φ is the maximal number of nested modalities occurring in φ. We write $\mathsf{bp}(\varphi)$ for the set of *bound propositions* of φ, i.e. propositions p that occur in a quantifier $\exists^k p$ inside φ. We say that φ is *well-quantified* whenever each subformula $\exists^k p\,\psi$ of φ quantifies on a different $p \in \mathsf{AP}$, and every occurrence of p in ψ appears under the scope of at most k modalities. One can translate every formula into a well-quantified one at no cost: atomic propositions can be renamed, and occurrences of a quantified atomic proposition that are under the scope of more than k modalities can be replaced with $\bot$.

We write $\varphi \equiv_{\mathit{FO}} \psi$ (resp. $\varphi \equiv_{\mathit{SO}} \psi$) whenever φ and ψ are equivalent under their first-order (resp. second-order) semantics, i.e. they are satisfied by the same pointed forests. When clear from the context or true under both semantics, we drop the subscripts and write $\varphi \equiv \psi$. Notice that $\exists^k p\,\varphi \equiv \exists^{k+1} p\,(\varphi \wedge \Box^{k+1}\neg p)$, and thus $\mathsf{ML}(\exists^k)$ is a syntactical fragment of $\mathsf{ML}(\exists^{k+1})$, and it is able to express all the local quantifiers $\exists^1 p, \dots, \exists^k p$.

Round-bounded fragment. As discussed in Sec. 1, in this paper we focus on a syntactical restriction for $\mathsf{ML}(\exists^k)$ where the local quantifiers are *round-bounded*. The round-bounded formulae of $\mathsf{ML}(\exists^k)$ are those generated from the symbol φ_0^k of the grammar below ($j \in \mathbb{N}$):

$$\varphi_j^k, \psi_j^k := \top \mid p \mid \varphi_j^k \wedge \psi_j^k \mid \neg\varphi_j^k \mid \Diamond\varphi_{j+1}^k \mid \exists^{k-(j \bmod k)} p\,\varphi_j^k, \text{ where } p \in \mathsf{AP}.$$

In a round-bounded formula of $\mathsf{ML}(\exists^k)$, quantifiers appearing under the scope of j modalities are restricted to $\exists^{k-(j \bmod k)}$, e.g. $\exists^3 p\,\Diamond\exists^2 q\,\Diamond\exists^1 r\,\Diamond\exists^3 p\,\varphi$ is a round-bounded formula of $\mathsf{ML}(\exists^3)$, provided that φ is also in this fragment, whereas $\exists^3 p\,\Diamond\exists^3 q\,\varphi$ is not round-bounded. The round-bounded condition does not change the set of formulae of $\mathsf{ML}(\exists^1)$ and $\mathsf{ML}(\exists^\infty)$. Besides, every formula of $\mathsf{ML}(\exists^\infty)$ of modal depth k is equivalent to a round-bounded formula of $\mathsf{ML}(\exists^k)$, of similar size, since given a formula φ of $\mathsf{ML}(\exists^\infty)$, we have $\exists^\infty p\,\varphi \equiv \exists^{\mathsf{md}(\varphi)} p\,\varphi$.

Our framework of local quantifiers enables us to derive connections with other modal logics featuring some form of quantification, which we now briefly discuss.

Graded modal logic. A logic that has been shown related to different forms of quantification is the *graded modal logic* GML [5], that extends ML with modalities $\Diamond_{\geq \ell}$ ($\ell \in \mathbb{N}$), with semantics: $\mathcal{K}, w \models \Diamond_{\geq \ell}\varphi \Leftrightarrow |\{w' \in R(w) \mid \mathcal{K}, w' \models \varphi\}| \geq \ell$. GML has a tree model property, i.e., each of its satisfiable formulae is satisfied by a pointed forest. Then, by syntactically replacing each $\Diamond_{\geq \ell}\varphi$ occurring in a GML formula by $\exists^1 \mathtt{x}_1, \dots, \mathtt{x}_\ell : (\bigwedge_{i=0}^{\ell} \bigwedge_{j=i+1}^{\ell} @^1_{\mathtt{x}_i} \neg\mathtt{x}_j) \wedge \Box((\bigvee_{i=0}^{\ell} \mathtt{x}_i) \Rightarrow \varphi)$, one shows that GML embeds in $\mathsf{ML}(\exists^1_{\mathit{FO}})$. At this point, it is worth noting that, for all $k \in \mathbb{N}_+$, $\mathsf{ML}(\exists^k_{\mathit{FO}})$ can be embedded into $\mathsf{ML}(\exists^k_{\mathit{SO}})$ by replacing, in a well-quantified formula of $\mathsf{ML}(\exists^k_{\mathit{FO}})$, each occurrence of $\exists^k p\,\varphi$ with the $\mathsf{ML}(\exists^k_{\mathit{SO}})$ formula

$\exists^k p : \varphi \land \mathtt{uniq}_k(p)$, where $\mathtt{uniq}_k(p) \stackrel{\text{def}}{=} \oplus^k p \land \forall^k q : \oplus^k(p \land q) \Rightarrow \boxplus^k(p \Rightarrow q)$ states that there is at most one world satisfying p that is reachable from the current one in at most k steps. Hence, $\mathsf{ML}(\exists^k_{SO})$ captures GML, and in fact the converse also holds, as we discover when proving Thm. 2. The corollary below is established.

Corollary 2. *For* $k \in \overline{\mathbb{N}}_+$, $\mathsf{ML}(\exists^k_{FO})$, $\mathsf{ML}(\exists^k_{SO})$ *and* GML *are equally expressive.*

This result is surprising, as it implies that $\mathsf{QCTL}^t_{\mathbf{X}}$ from [6] is as expressive as GML, and that in the context of modal logics, second-order propositional quantifiers do not yield any additional expressive power compared to first-order ones.

Connections with $\mathsf{S5Q}$*.* The sat. problem of $\mathsf{S5Q}$ [18,22] is equireducible to the sat. problem for formulae of $\mathsf{ML}(\exists^1_{SO})$ of modal depth 1. Briefly, any satisfiable formula of $\mathsf{S5Q}$ is satisfied by a Kripke structure $(\mathcal{W}, R, \mathcal{V})$ where $R = \mathcal{W} \times \mathcal{W}$, and $\mathsf{S5Q}$ enriches ML with quantifiers $\exists p$ which, by virtue of the relation R, are essentially the quantifiers $\exists^1 p$ from $\mathsf{ML}(\exists^1_{SO})$. We can simulate the models of $\mathsf{S5Q}$ by using a pointed forest $(\mathcal{K}, w)$ with accessibility relation R' such that $R'(w) = \mathcal{W}$. The current world of the $\mathsf{S5Q}$ model is simulated with a 1-local nominal $\mathtt{x}$ for $(\mathcal{K}, w)$. Then, the translation τ from $\mathsf{S5Q}$ to $\mathsf{ML}(\exists^1_{SO})$ is simple: $\tau(\Diamond\varphi) = \exists^1 \mathtt{x} : \Diamond \mathtt{x} \land \mathtt{uniq}_1(\mathtt{x}) \land \tau(\varphi)$, binding the nominal $\mathtt{x}$ to a new world; $\tau(p) = @^1_{\mathtt{x}} p$, and otherwise τ is homomorphic. A similar translation can be given from formulae of $\mathsf{ML}(\exists^1_{SO})$ with modal depth 1 to $\mathsf{S5Q}$. Following Thm. 2, this allows us to characterise the complexity of $\mathsf{S5Q}$ left open in [18].

Corollary 3. *The sat. problem for* $\mathsf{S5Q}$ *is* AExp$_{pol}$*-complete.*

Connections with hybrid logics. Hybrid logics [3] are among the most studied modal logics featuring first-order propositional quantification. Given a set of nominals $\mathsf{NOM} \subseteq \mathsf{AP}$, the hybrid logic $\mathsf{HL}(\downarrow, @)$ extends ML with the *binder* $\downarrow i$ and the satisfaction operator $@_i$ (where $i \in \mathsf{NOM}$), having the semantics below:

$$(\mathcal{W}, R, \mathcal{V}), w \models \downarrow i.\varphi \;\Leftrightarrow\; (\mathcal{W}, R, \mathcal{V}[i \leftarrow \{w\}]), w \models \varphi;$$
$$(\mathcal{W}, R, \mathcal{V}), w \models @_i\varphi \;\Leftrightarrow\; (\mathcal{W}, R, \mathcal{V}), w_i \models \varphi, \text{ where } \mathcal{V}(i) = \{w_i\}.$$

$\mathsf{ML}(\exists^k_{FO})$ embeds in $\mathsf{HL}(\downarrow, @)$ by replacing with $\downarrow i.\oplus^k \downarrow p.@_i\varphi$ each occurrence of $\exists^k p\, \varphi$ appearing in an $\mathsf{ML}(\exists^k_{FO})$ formula. This translation is (only) exponential in k, and so by uniform reduction for all $k \in \mathbb{N}_+$, and by Rabin's theorem [27] for the upper bound, Thm. 1 implies the following result.

Corollary 4. *The sat. problem for* $\mathsf{HL}(\downarrow, @)$ *on forests is* Tower*-complete.*

3 Lower bounds for $\mathsf{ML}(\exists^k_{FO})$ and $\mathsf{ML}(\exists^k_{SO})$

In this section, we establish the lower bounds of Thms. 1 and 2, which follow by reduction from the *k-exp alternating multi-tiling problem*. While we will introduce this problem in due time, the main difficulty in establishing the reduction is defining, for all $k, n \in \mathbb{N}_+$ given in unary, a formula *type*(k, n) that, whenever satisfied by a pointed forest $(\mathcal{K}, w)$, forces w to have $\mathfrak{t}(k, n)$ children, each of

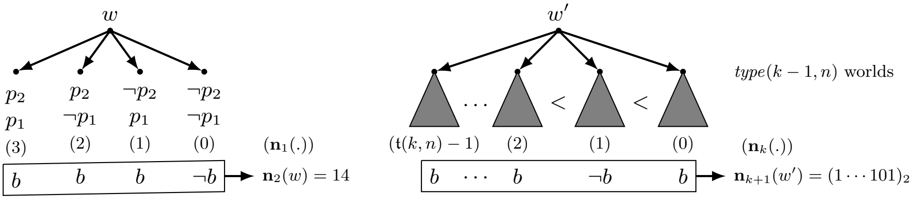

Fig. 1: Two worlds w and w' satisfying $type(1, 2)$ and $type(k, n)$, respectively.

them encoding a different number in $[0, \mathfrak{t}(k, n) - 1]$. To establish Thms. 1 and 2, it is essential that $type(k, n)$ is of size polynomial in k and n, has modal depth k, it is in $\mathsf{ML}(\exists^1_{FO})$ for $k = 1$, and is in round-bounded $\mathsf{ML}(\exists^{k-1}_{FO})$ for all $k \geq 2$. The formula $type(k, n)$ is inspired by the homonymous formula defined in [6] to show that $\mathsf{QCTL}^t_{\mathbf{X}}$ is TOWER-hard, and later adapted in [7] to modal separation logics. With respect to both these works, our definition of $type(k, n)$ poses two serious challenges. First, [6,7] rely on second-order quantification, whereas we only use first-order. Second, in [6,7] the formula $type(k, n)$ is of size exponential in k, whereas our formula is of polynomial size. To achieve both improvements, we rely on a novel gadget that simulates binary addition with carry.

Numeric encoding. First of all, let us define how numbers are encoded by worlds of a pointed forest, following the presentation of [6]. Fix $n + 1$ distinct atomic propositions $p_1, \dots, p_n, b$, and consider a Kripke-style forest $\mathcal{K} = (\mathcal{W}, R, \mathcal{V})$. Given $j \in [1, k]$ and $w \in \mathcal{W}$, we write $\mathbf{n}_j(w)$ for the number in $[0, \mathfrak{t}(j, n) - 1]$ encoded by w. For $j = 1$, we represent $\mathbf{n}_1(w) \in [0, 2^n - 1]$ by using the truth values of the propositions $p_1, \dots, p_n$, where the proposition p_i is responsible for the i-th least significant bit of the number. That is, $\mathbf{n}_1(w) \stackrel{\text{def}}{=} \sum\{2^{i-1} : i \in [1, n] \text{ and } w \in \mathcal{V}(p_i)\}$. For $j > 1$, the number $\mathbf{n}_j(w)$ is represented by the binary encoding of the truth values of the atomic proposition b on the children of w, where a child $w' \in R(w)$ with $\mathbf{n}_{j-1}(w') = i$ from $[0, \mathfrak{t}(j - 1, n) - 1]$ is responsible for the $(i + 1)$-th least significant bit of the number encoded by w. Formally, $\mathbf{n}_j(w) \stackrel{\text{def}}{=} \sum\{2^i : \mathbf{n}_{j-1}(w') = i \text{ and } w' \in \mathcal{V}(b), \text{ for some } w' \in R(w)\}$.

With respect to this encoding of numbers, the forthcoming formula $type(k, n)$ shall satisfy the specification given by the lemma below, which guarantees that in a pointed forest $(\mathcal{K}, w)$ satisfying $type(k, n)$, the numbers encoded by the children of w span all over $[0, \mathfrak{t}(k, n) - 1]$. This is illustrated in Fig. 1.

Lemma 1. *A pointed forest $(\mathcal{K}, w)$, with $\mathcal{K} = (\mathcal{W}, R, \mathcal{V})$, satisfies $type(k, n)$ iff*
1. *for all $i \in [0, \mathfrak{t}(k, n) - 1]$ there is exactly one world $w' \in R(w)$ s.t. $\mathbf{n}_k(w') = i$;*
2. *if $k > 1$, then for every $w' \in R(w)$, $\mathcal{K}, w' \models type(k - 1, n)$.*

Addition with carry. In defining $type(k, n)$, the main challenge lies in how to express the condition (1) of Lemma 1. In [6,7], this boils down to the definition of formulae that express (in)equalities between the numbers encoded by distinct $w_1, w_2 \in R(w)$, e.g. $\mathbf{n}_k(w_1) < \mathbf{n}_k(w_2)$ or $\mathbf{n}_k(w_1) = \mathbf{n}_k(w_2) + 1$. Unfortunately, these formulae are tree-recursive on k, meaning that multiple (possibly negated) occurrences of the inequalities for the case $k - 1$ are required to

Formula:	Expected Semantics:	Assumptions:
$\mathit{0}_j$	$\mathbf{n}_j(w) = 0$	The world w is the current world, which is assumed to satisfy $\mathit{type}(j,n)$. The world w_p corresponds to the i-local nominal $p \in \{\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c}\}$, and is assumed to satisfy $\mathit{type}(k-i,n)$.
$\mathit{1}_j$	$\mathbf{n}_j(w) = 1$	
$\mathcal{E}_j$	$\mathbf{n}_j(w) = \mathfrak{t}(j,n) - 1$	
$\mathit{add}^i_k(\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c})$	$+_{k-i+1}(w_\mathtt{x}, w_\mathtt{y}, w_\mathtt{z}, w_\mathtt{c})$	

Fig. 2: Auxiliary formulae used in the definition of $\mathit{type}(k,n)$, where $i = k = 1$ or $i < k$.

define the inequalities for the case k. Overall, this induces an exponential blow-up on $|\mathit{type}(k,n)|$. To avoid this blow-up, instead of relying on these inequalities we consider a quaternary relation $+_k(w_1,w_2,w_3,w_4)$ that holds whenever $\mathbf{n}_k(w_1)+\mathbf{n}_k(w_2) = \mathbf{n}_k(w_3)$ and $\mathbf{n}_k(w_4)$ represents the sequence of carries needed to perform $\mathbf{n}_k(w_1)+\mathbf{n}_k(w_2)$ in binary, on $\mathfrak{t}(k-1,n)$ bits. For instance, for 4-bits numbers $\mathbf{n}_1(w_1) = 3 = (0011)_2$, $\mathbf{n}_1(w_2) = 5 = (0101)_2$, $\mathbf{n}_1(w_3) = 8 = (1000)_2$ and $\mathbf{n}_1(w_4) = 14 = (1110)_2$, the tuple (w_1,w_2,w_3,w_4) is in $+_1$, as

$$\begin{array}{ll} \color{gray}{1\,1\,1\,0} & : w_4 \text{ (sequence of carries of the sum)} \\ 0\,0\,1\,1\,+ & : w_1 \\ \underline{0\,1\,0\,1} & : w_2 \\ 1\,0\,0\,0 & : w_3 \end{array}$$

corresponds to the table for the binary addition with carry of $3 + 5 = 8$. By looking at the elementary algorithm for addition, a direct characterisation of $+_k$ is as follows. Let $\mathbf{n}_k(w_1) = (x_m \dots x_1)_2$, $\mathbf{n}_k(w_2) = (y_m \dots y_1)_2$, $\mathbf{n}_k(w_3) = (z_m \dots z_1)_2$, $\mathbf{n}_k(w_4) = (c_m \dots c_1)_2$, where $m = \mathfrak{t}(k-1,n)$, and x_i, y_i, z_i and c_i are the i-th least significant digits in the binary encoding of $\mathbf{n}_k(w_1)$, $\mathbf{n}_k(w_2)$, $\mathbf{n}_k(w_3)$, $\mathbf{n}_k(w_4)$, respectively. Then, $+_k(w_1,w_2,w_3,w_4)$ holds if and only if

$$\begin{array}{ll} \textbf{A.} & c_1 = 0 \text{ and at most one among } c_m,\ x_m \text{ and } y_m \text{ is } 1, \\ \textbf{B.} & \text{for every } i \in [2,m],\ c_i = \mathit{maj}(x_{i-1}, y_{i-1}, c_{i-1}), \\ \textbf{C.} & \text{for every } i \in [1,m],\ z_i = (x_i \oplus y_i) \oplus c_i, \end{array} \tag{†}$$

where $\mathit{maj}(\varphi,\psi,\chi) \stackrel{\text{def}}{=} (\varphi\wedge\psi)\vee(\varphi\wedge\chi)\vee(\psi\wedge\chi)$ and $\varphi\oplus\psi \stackrel{\text{def}}{=} (\varphi\vee\psi)\wedge\neg(\varphi\wedge\psi)$ are the standard Boolean functions *majority* and *exclusive or*, respectively. When it comes to capturing $+_k$ with an $\mathsf{ML}(\exists^k_{\mathit{FO}})$ formula, the key property is that the conditions (A), (B) and (C) can be checked with first-order quantification, by going through the binary encodings of $\mathbf{n}_k(w_1), \mathbf{n}_k(w_2), \mathbf{n}_k(w_3)$ and $\mathbf{n}_k(w_4)$ bit by bit, as one would do to check if an addition with carry was performed correctly.

A schema for $\mathit{type}(k,n)$. We move to the definition of $\mathit{type}(k,n)$. In view of its specification given in Lemma 1, the formula is defined recursively on k. For simplicity, we extend $\mathit{type}(k,n)$ to $k = 0$, and define it as $\top$. To express the condition (1) of Lemma 1, we rely on the auxiliary formulae presented in Fig. 2, which we later define. For $k,n \in \mathbb{N}_+$, we define $\mathit{type}(k,n)$ as:

$$\begin{aligned} &\Box\mathit{type}(k-1,n) \wedge \Diamond\mathit{0}_k \wedge \Diamond\mathit{1}_k \wedge \Diamond\mathcal{E}_k \wedge \\ &\quad \forall^1\mathtt{x}\,\forall^1\mathtt{y}\,(\Diamond\mathtt{y} \wedge @^1_\mathtt{x}\neg\mathtt{y} \Rightarrow \exists^1\mathtt{z}\,\exists^1\mathtt{c} : \Diamond\mathtt{c} \wedge @^1_\mathtt{z}\neg\mathit{0}_k \wedge (\mathit{add}^1_k(\mathtt{x},\mathtt{z},\mathtt{y},\mathtt{c}) \vee \mathit{add}^1_k(\mathtt{y},\mathtt{z},\mathtt{x},\mathtt{c}))). \end{aligned}$$

Whereas the first conjunct of $\mathit{type}(k,n)$ clearly encodes the condition (2) of Lemma 1, the remaining part of the formula forces the condition (1) by saying that the current world w has three children encoding the numbers 0, 1 and

$\mathfrak{t}(k,n)-1$, respectively, and that for every two children $w_{\mathtt{x}}, w_{\mathtt{y}}$ of w, if $w_{\mathtt{x}} \neq w_{\mathtt{y}}$ (subformula $\Diamond \mathtt{y} \wedge @^1_{\mathtt{x}}\neg\mathtt{y}$) then there is a child $w_{\mathtt{z}}$ of w such that $\mathbf{n}_k(w_{\mathtt{z}}) \neq 0$, and $\mathbf{n}_k(w_{\mathtt{x}}) + \mathbf{n}_k(w_{\mathtt{z}}) = \mathbf{n}_k(w_{\mathtt{y}})$ or $\mathbf{n}_k(w_{\mathtt{y}}) + \mathbf{n}_k(w_{\mathtt{z}}) = \mathbf{n}_k(w_{\mathtt{x}})$. Hence, in combination with $\Diamond \mathit{0}_k$, $\Diamond \mathit{1}_k$ and $\Diamond \mathcal{E}_k$, the last conjunct of $\mathit{type}(k,n)$ not only states that distinct children of w must encode different numbers, but also that every number of $[0, \mathfrak{t}(k,n)-1]$ must be encoded by some child of w.

To effectively construct $\mathit{type}(k,n)$, what is left is to define the formulae in Fig. 2. Given how the numbers $\mathbf{n}_k(.)$ are encoded, the definitions of $\mathit{0}_k$, $\mathit{1}_k$ and $\mathcal{E}_k$ are simple. For the case $k=1$, we define $\mathit{0}_1 \stackrel{\text{def}}{=} \bigwedge_{j=1}^{n} \neg p_j$, $\mathit{1}_1 \stackrel{\text{def}}{=} (p_1 \wedge \bigwedge_{j=2}^{n} \neg p_j)$ and $\mathcal{E}_1 \stackrel{\text{def}}{=} \bigwedge_{j=1}^{n} p_j$. For $k \geq 2$, we define instead: $\mathit{0}_k \stackrel{\text{def}}{=} \Box\neg b$, $\mathit{1}_k \stackrel{\text{def}}{=} \Box(b \Rightarrow \mathit{0}_{k-1})$, and $\mathcal{E}_k \stackrel{\text{def}}{=} \Box b$. The main difficulty lies in how to define add^i_k, which requires a recursive definition. Below, we consider three cases. First, we consider the base case $i=k=1$ and define add^1_1 by only using the local quantifiers $\exists^1$. Afterwards, we consider the case $1 \leq i < k-1$ and define the formula add^i_k by using local quantifiers $\exists^1, \ldots, \exists^{k-1}$. This formula relies on the definition of add^{i+1}_k, which we assume to be defined by inductive reasoning. Lastly, we consider the only remaining case of $i = k-1$, and define add^{k-1}_k by using quantifiers $\exists^{k-1}$ and $\exists^1$, and without relying on the definition of add^1_1. This case is left for last as it is somewhat more involved than the other two cases, and some ingenuity is required to define add^{k-1}_k without relying on the local quantifiers $\exists^k$. The ad-hoc treatment of this case is however fundamental, as it leads to $\mathit{type}(k,n)$ being a round-bounded formula of the logic $\mathsf{ML}(\exists^{k-1}_{\mathit{FO}})$, for every $k \geq 2$.

Case: $i = k = 1$. Recall that the numbers $\mathbf{n}_1(.)$ are encoded using the truth values of $p_1, \ldots, p_n \in \mathsf{AP}$. Then, add^1_1 simply follows the constraints (†) of $+_1$:

$$\begin{aligned} \mathit{add}^1_1(\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c}) \stackrel{\text{def}}{=}\ & @^1_{\mathtt{c}}\neg p_1 \wedge \textstyle\bigwedge_{q \in \{\mathtt{x},\mathtt{y},\mathtt{c}\}} \big(@^1_q p_n \Rightarrow \bigwedge_{r \in \{\mathtt{x},\mathtt{y},\mathtt{c}\}\setminus\{q\}} @^1_r \neg p_n\big) && \text{(A)}\\ & \wedge \textstyle\bigwedge_{i=2}^{n} \big(@^1_{\mathtt{c}} p_i \Leftrightarrow \mathit{maj}(@^1_{\mathtt{x}} p_{i-1}, @^1_{\mathtt{y}} p_{i-1}, @^1_{\mathtt{c}} p_{i-1})\big) && \text{(B)}\\ & \wedge \textstyle\bigwedge_{i=1}^{n} \big(@^1_{\mathtt{z}} p_i \Leftrightarrow \big((@^1_{\mathtt{x}} p_i \oplus @^1_{\mathtt{y}} p_i) \oplus @^1_{\mathtt{c}} p_i\big)\big) && \text{(C)} \end{aligned}$$

Case: $1 \leq i < k-1$. To define add^i_k, we assume by inductive reasoning that the formula add^{i+1}_k is correctly defined, following its specification in Fig. 2. We specialise add^{i+1}_k to define the two auxiliary formulae below:

$$\begin{aligned} \mathit{eq}^{i+1}_k(\mathtt{x},\mathtt{y}) &\stackrel{\text{def}}{=} \exists^{i+1}\mathtt{z},\mathtt{c} : \Diamond^{i+1}\mathtt{c} \wedge @^{i+1}_{\mathtt{z}} \mathit{0}_{k-i} \wedge \mathit{add}^{i+1}_k(\mathtt{y},\mathtt{z},\mathtt{x},\mathtt{c});\\ \mathit{succ}^{i+1}_k(\mathtt{x},\mathtt{y}) &\stackrel{\text{def}}{=} \exists^{i+1}\mathtt{z},\mathtt{c} : \Diamond^{i+1}\mathtt{c} \wedge @^{i+1}_{\mathtt{z}} \mathit{1}_{k-i} \wedge \mathit{add}^{i+1}_k(\mathtt{y},\mathtt{z},\mathtt{x},\mathtt{c}). \end{aligned}$$

Given $\mathtt{x}$ and $\mathtt{y}$ be two $(i+1)$-local nominals for $(\mathcal{K}, w)$, with corresponding worlds $w_{\mathtt{x}}$ and $w_{\mathtt{y}}$, if $\mathcal{K}, w' \models \mathit{type}(k-i, n)$ for some $w' \in R^i(w)$, then:

- $\mathcal{K}, w \models \mathit{eq}^{i+1}_k(\mathtt{x},\mathtt{y})$ if and only if $\mathbf{n}_{k-i}(w_{\mathtt{x}}) = \mathbf{n}_{k-i}(w_{\mathtt{y}})$;
- $\mathcal{K}, w \models \mathit{succ}^{i+1}_k(\mathtt{x},\mathtt{y})$ if and only if $\mathbf{n}_{k-i}(w_{\mathtt{x}}) = \mathbf{n}_{k-i}(w_{\mathtt{y}}) + 1$.

Notice that the semantics of succ^{i+1}_k and eq^{i+1}_k is given under the hypothesis that a world in $R^i(w)$ satisfies $\mathit{type}(k-i, n)$. This extra hypothesis ensures that the local quantifiers $\exists^{i+1}\mathtt{z}$ and $\exists^{i+1}\mathtt{c}$ used to define succ^{i+1}_k and eq^{i+1}_k quantify over a set of worlds encoding all the numbers in $[0, \mathfrak{t}(k-(i+1), n)-1]$, so that no possible addition with carry is missing. In defining $add^i_k(\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c})$, this hypothesis is clearly satisfied, as the worlds corresponding to the i-local nominals $\mathtt{x}$, $\mathtt{y}$, $\mathtt{z}$ and $\mathtt{c}$ are assumed to satisfy $type(k-i,n)$.

By relying on $succ^{i+1}_k$ and eq^{i+1}_k, we define $add^i_k(\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c})$ again by following the characterisation (†) of $+_{k-i+1}$, as shown below (where $X \stackrel{\text{def}}{=} \{\overline{\mathtt{x}},\overline{\mathtt{y}},\overline{\mathtt{c}}\}$):

$$\forall^{i+1}\overline{\mathtt{x}},\overline{\mathtt{y}},\overline{\mathtt{z}},\overline{\mathtt{c}},\mathtt{g} : @^i_{\mathtt{x}}\Diamond\overline{\mathtt{x}} \wedge @^i_{\mathtt{y}}\Diamond\overline{\mathtt{y}} \wedge @^i_{\mathtt{z}}\Diamond\overline{\mathtt{z}} \wedge @^i_{\mathtt{c}}(\Diamond\overline{\mathtt{c}} \wedge \Diamond\mathtt{g}) \Rightarrow$$

(A): $@^{i+1}_{\overline{\mathtt{c}}}(\mathit{0}_{k-i} \Rightarrow \neg b) \wedge \big((\bigwedge_{q\in X} @^{i+1}_q \mathcal{E}_{k-i}) \Rightarrow \bigwedge_{q\in X}\big(@^{i+1}_q b \Rightarrow \bigwedge_{r\in X\setminus\{q\}} @^{i+1}_r \neg b\big)\big)$

(B): $\wedge \big(eq^{i+1}_k(\overline{\mathtt{x}},\overline{\mathtt{y}}) \wedge eq^{i+1}_k(\overline{\mathtt{y}},\overline{\mathtt{c}}) \wedge succ^{i+1}_k(\mathtt{g},\overline{\mathtt{c}}) \Rightarrow \big(@^{i+1}_{\mathtt{g}} b \Leftrightarrow maj(@^{i+1}_{\overline{\mathtt{x}}} b, @^{i+1}_{\overline{\mathtt{y}}} b, @^{i+1}_{\overline{\mathtt{c}}} b)\big)\big)$

(C): $\wedge \big(eq^{i+1}_k(\overline{\mathtt{x}},\overline{\mathtt{y}}) \wedge eq^{i+1}_k(\overline{\mathtt{y}},\overline{\mathtt{z}}) \wedge eq^{i+1}_k(\overline{\mathtt{z}},\overline{\mathtt{c}}) \Rightarrow \big(@^{i+1}_{\overline{\mathtt{z}}} b \Leftrightarrow ((@^{i+1}_{\overline{\mathtt{x}}} b \oplus @^{i+1}_{\overline{\mathtt{y}}} b) \oplus @^{i+1}_{\overline{\mathtt{c}}} b)\big)\big)$.

The first line of add^i_k binds the propositions $\overline{\mathtt{x}}$, $\overline{\mathtt{y}}$, $\overline{\mathtt{z}}$, and $\overline{\mathtt{c}}$ and $\mathtt{g}$ to children of $\mathtt{x}$, $\mathtt{y}$, $\mathtt{z}$ and $\mathtt{c}$, respectively. Afterwards, the formula follows closely the constraints in (†). For instance, the last conjunct characterises the condition (C) by saying that whenever we consider children $w_{\overline{\mathtt{x}}}$, $w_{\overline{\mathtt{y}}}$, $w_{\overline{\mathtt{z}}}$ and $w_{\overline{\mathtt{c}}}$ of $w_{\mathtt{x}}$, $w_{\mathtt{y}}$, $w_{\mathtt{z}}$ and $w_{\mathtt{c}}$ respectively, if $j = \mathbf{n}_{k-i}(w_{\overline{\mathtt{x}}}) = \mathbf{n}_{k-u}(w_{\overline{\mathtt{y}}}) = \mathbf{n}_{k-i}(w_{\overline{\mathtt{z}}}) = \mathbf{n}_{k-i}(w_{\overline{\mathtt{c}}})$ for some $j \in \mathbb{N}$, then $\mathbf{n}_2(w_{\mathtt{z}})[j] = ((\mathbf{n}_2(w_{\mathtt{x}})[j] \oplus \mathbf{n}_2(w_{\mathtt{y}})[j]) \oplus \mathbf{n}_2(w_{\mathtt{c}})[j])$, where $\mathbf{n}_2(w)[j]$ is the $(j+1)$-th least significant digit of the number encoded by a world w.

Case: $i = k-1$. To complete the definition of add^i_k, what is left is to define add^{k-1}_k by only using quantifiers $\exists^{k-1}$ and $\exists^1$. Below, the worlds $w_{\mathtt{x}}$, $w_{\mathtt{y}}$, $w_{\mathtt{z}}$ and $w_{\mathtt{c}}$, corresponding to the $(k-1)$-local nominals $\mathtt{x}$, $\mathtt{y}$, $\mathtt{z}$ and $\mathtt{c}$, satisfy $type(1,n)$, and so accordingly with $\mathbf{n}_2(.)$ they encode a number by looking at the value of the proposition b in their children, which themselves encode a number $\mathbf{n}_1(.)$. To properly define $add^{k-1}_k(\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c})$, we rely on the fact that these children encode n-bits numbers, with n given in unary. Then, instead of employing a quantifier $\exists^k$ to refer to one of these children, we can rely on $n+1$ local quantifiers $\exists^{k-1}$ to copy the values of $p_1,\dots,p_n$ and b of a child directly on its parent. For instance, to check if $w_{\mathtt{x}}$ and $w_{\mathtt{y}}$ have children encoding the same numbers and equisatisfying b, one can follow the steps below, also sketched in Fig. 3:

1. using $\exists^{k-1}$, we quantify over fresh propositional symbols $r^{\mathtt{v}}_1,\dots,r^{\mathtt{v}}_n$ and $q_{\mathtt{v}}$, with $\mathtt{v} \in \{\mathtt{x},\mathtt{y}\}$, to modify the truth of these symbols on $w_{\mathtt{x}}$ and $w_{\mathtt{y}}$;
2. using $@^{k-1}_{\mathtt{x}}$, we move the evaluation point to $w_{\mathtt{x}}$. We check that the truth of the propositions $r^{\mathtt{x}}_1,\dots,r^{\mathtt{x}}_n,q_{\mathtt{x}}$ on $w_{\mathtt{x}}$ is mirroring the truth of $p_1,\dots,p_n,b$ on a child of $w_{\mathtt{x}}$. For this, we rely on the formula $copy((r^{\mathtt{x}}_1,\dots,r^{\mathtt{x}}_n),q_{\mathtt{x}})$ that, for an n-tuple of atomic propositions $\mathbf{r} = (r_1,\dots,r_n)$ and $q \in \mathsf{AP}$, is defined as: $copy(\mathbf{r},q) \stackrel{\text{def}}{=} \exists^1\mathtt{u} : \Diamond\mathtt{u} \wedge (q \Leftrightarrow @^1_{\mathtt{u}} b) \wedge \bigwedge_{i=1}^n (r_i \Leftrightarrow @^1_{\mathtt{u}} p_i)$. This step is also done (in parallel) for $w_{\mathtt{y}}$, by relying on $copy((r^{\mathtt{y}}_1,\dots,r^{\mathtt{y}}_n),q_{\mathtt{y}})$;
3. with respect to the initial point of evaluation w, we check that the truth of the propositions $r^{\mathtt{x}}_1,\dots,r^{\mathtt{x}}_n,q_{\mathtt{x}}$ on $w_{\mathtt{x}}$ corresponds to the truth of $r^{\mathtt{y}}_1,\dots,r^{\mathtt{y}}_n,q_{\mathtt{y}}$ on $w_{\mathtt{y}}$, i.e. $@^{k-1}_{\mathtt{x}} q_{\mathtt{x}} \Leftrightarrow @^{k-1}_{\mathtt{y}} q_{\mathtt{y}}$ and $@^{k-1}_{\mathtt{x}} r^{\mathtt{x}}_i \Leftrightarrow @^{k-1}_{\mathtt{y}} r^{\mathtt{y}}_i$, for all $i \in [1,n]$.

This idea of copying information about children of $w_{\mathtt{x}}$, $w_{\mathtt{y}}$, $w_{\mathtt{z}}$ and $w_{\mathtt{c}}$ directly in these four worlds is at the base of our definition of add^{k-1}_k, which we now formalise. Similarly to $\mathbf{n}_1(.)$, for an n-tuple of symbols $\mathbf{r} = (r_1,\dots,r_n)$, $\mathbf{n}_{\mathbf{r}}(w) \stackrel{\text{def}}{=} \sum\{2^{i-1} : i \in [1,n], w \in \mathcal{V}(r_i)\}$ stands for the n-bits number encoded by the world w by looking at the truth values of $r_1,\dots,r_n$. Given a second n-tuple of atomic

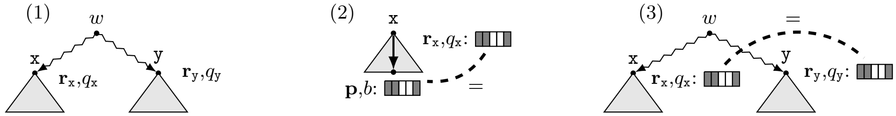

Fig. 3: Steps to check if two children of $w_{\mathtt{x}}$ and $w_{\mathtt{y}}$ encoding the same $\mathbf{n}_1(.)$ equisatisfy b.

propositions $\mathbf{s} = (s_1, \dots, s_n)$, we introduce the formulae $succ(\mathbf{r}@\mathtt{x}, \mathbf{s}@\mathtt{y}) \stackrel{\text{def}}{=} \bigvee_{i=1}^{n}\big(@_{\mathtt{x}}^{k-1} r_i \wedge @_{\mathtt{y}}^{k-1}\neg s_i \wedge \bigwedge_{j=1}^{i-1}(@_{\mathtt{x}}^{k-1}\neg r_j \wedge @_{\mathtt{y}}^{k-1} s_j) \wedge \bigwedge_{j=i+1}^{n}(@_{\mathtt{x}}^{k-1} r_j \Leftrightarrow @_{\mathtt{y}}^{k-1} s_j)\big)$ and $eq(\mathbf{r}@\mathtt{x}, \mathbf{s}@\mathtt{y}) \stackrel{\text{def}}{=} \bigwedge_{i=1}^{n}(@_{\mathtt{x}}^{k-1} r_i \Leftrightarrow @_{\mathtt{y}}^{k-1} s_i)$, having the following semantics:

- $\mathcal{K}, w \models eq(\mathbf{r}@\mathtt{x}, \mathbf{s}@\mathtt{y})$ if and only if $\mathbf{n}_{\mathbf{r}}(w_{\mathtt{x}}) = \mathbf{n}_{\mathbf{s}}(w_{\mathtt{y}})$; and
- $\mathcal{K}, w \models succ(\mathbf{r}@\mathtt{x}, \mathbf{s}@\mathtt{y})$ if and only if $\mathbf{n}_{\mathbf{r}}(w_{\mathtt{x}}) = \mathbf{n}_{\mathbf{s}}(w_{\mathtt{y}}){+}1$.

The correctness of $succ(\mathbf{r}@\mathtt{x}, \mathbf{s}@\mathtt{y})$ follows from standard arithmetical properties: for two n-bits numbers $\mathbf{a}$ and $\mathbf{b}$ represented as binary bit vectors with most significant digit first, $\mathbf{a} = \mathbf{b} + 1$ holds iff $\mathbf{a} = \mathbf{c10}$ and $\mathbf{b} = \mathbf{c01}$ hold for a prefix $\mathbf{c} \in \{0,1\}^*$ and bit vectors of same length $\mathbf{0} \in \{0\}^*$ and $\mathbf{1} \in \{1\}^*$.

The definition of $add_k^{k-1}(\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c})$ is given below, where $X \stackrel{\text{def}}{=} \{\mathtt{x},\mathtt{y},\mathtt{c}\}$ and for $\mathtt{v} \in \{\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c},\mathtt{g}\}$, $\mathbf{r}_{\mathtt{v}} \stackrel{\text{def}}{=} (r_1^{\mathtt{v}}, \dots, r_n^{\mathtt{v}})$ and $\forall^{k-1}\mathbf{r}_{\mathtt{v}}$ is short for $\forall^{k-1} r_1^{\mathtt{v}} \dots \forall^{k-1} r_n^{\mathtt{v}}$.

$$\begin{aligned}
&\forall^{k-1}\mathbf{r}_{\mathtt{x}}, q_{\mathtt{x}}, \mathbf{r}_{\mathtt{y}}, q_{\mathtt{y}}, \mathbf{r}_{\mathtt{z}}, q_{\mathtt{z}}, \mathbf{r}_{\mathtt{c}}, q_{\mathtt{c}}, \mathbf{r}_{\mathtt{g}}, q_{\mathtt{g}} : \textstyle\bigwedge_{\mathtt{v}\in\{\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c}\}} @_{\mathtt{v}}^{k-1} copy(\mathbf{r}_{\mathtt{v}}, q_{\mathtt{v}}) \wedge @_{\mathtt{c}}^{k-1} copy(\mathbf{r}_{\mathtt{g}}, q_{\mathtt{g}}) \Rightarrow \\
&(\mathrm{A}):\quad @_{\mathtt{c}}^{k-1}\Box(o_1 \Rightarrow \neg b) \wedge \textstyle\bigwedge_{q\in X} @_{q}^{k-1}\big(\Diamond(\mathcal{E}_1 \wedge b) \Rightarrow \bigwedge_{r \in X\setminus\{q\}} @_{r}^{k-1}\Box(\mathcal{E}_1 \Rightarrow \neg b)\big) \\
&(\mathrm{B}):\ \wedge \big(eq(\mathbf{r}_{\mathtt{x}}@\mathtt{x}, \mathbf{r}_{\mathtt{y}}@\mathtt{y}) \wedge eq(\mathbf{r}_{\mathtt{y}}@\mathtt{y}, \mathbf{r}_{\mathtt{c}}@\mathtt{c}) \wedge succ(\mathbf{r}_{\mathtt{g}}@\mathtt{c}, \mathbf{r}_{\mathtt{c}}@\mathtt{c}) \\
&\qquad\qquad\qquad \Rightarrow \big(@_{\mathtt{c}}^{k-1} q_{\mathtt{g}} \Leftrightarrow maj(@_{\mathtt{x}}^{k-1} q_{\mathtt{x}}, @_{\mathtt{y}}^{k-1} q_{\mathtt{y}}, @_{\mathtt{c}}^{k-1} q_{\mathtt{c}})\big)\big) \\
&(\mathrm{C}):\ \wedge \big(eq(\mathbf{r}_{\mathtt{x}}@\mathtt{x}, \mathbf{r}_{\mathtt{y}}@\mathtt{y}) \wedge eq(\mathbf{r}_{\mathtt{y}}@\mathtt{y}, \mathbf{r}_{\mathtt{z}}@\mathtt{z}) \wedge eq(\mathbf{r}_{\mathtt{z}}@\mathtt{z}, \mathbf{r}_{\mathtt{c}}@\mathtt{c}) \\
&\qquad\qquad\qquad \Rightarrow \big(@_{\mathtt{z}}^{k-1} q_{\mathtt{z}} \Leftrightarrow \big((@_{\mathtt{x}}^{k-1} q_{\mathtt{x}} \oplus @_{\mathtt{y}}^{k-1} q_{\mathtt{y}}) \oplus @_{\mathtt{c}}^{k-1} q_{\mathtt{c}}\big)\big)\big).
\end{aligned}$$

Notice that this formula first quantifies over fresh atomic propositions $\mathbf{r}_{\mathtt{v}}$ and $q_{\mathtt{v}}$, with $\mathtt{v} \in \{\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c},\mathtt{g}\} \subseteq \mathsf{AP}$, so that the worlds $w_{\mathtt{x}}$, $w_{\mathtt{y}}$, $w_{\mathtt{z}}$, $w_{\mathtt{c}}$ copy the truth of $p_1, \dots, p_n$ and b of some of their children w.r.t. the fresh atomic propositions (see subformula $\bigwedge_{\mathtt{v}\in\{\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c}\}} @_{\mathtt{v}}^{k-1} copy(\mathbf{r}_{\mathtt{v}}, q_{\mathtt{v}}) \wedge @_{\mathtt{c}}^{k-1} copy(\mathbf{r}_{\mathtt{g}}, q_{\mathtt{g}})$). Afterwards, the formula follows very closely the constraints (†) of $+_2$.

By induction on i, we show that add_k^i respects the specification from Fig. 2.

Lemma 2. *Let $(\mathcal{K}, w)$ be a pointed forest, and $\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c}$ be four i-local nominals for $(\mathcal{K}, w)$, with corresponding worlds $w_{\mathtt{x}}$, $w_{\mathtt{y}}$, $w_{\mathtt{z}}$ and $w_{\mathtt{c}}$. If $\mathcal{K}, w_p \models type(k{-}i, n)$ for every $p \in \{\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c}\}$, then $\mathcal{K}, w \models add_k^i(\mathtt{x},\mathtt{y},\mathtt{z},\mathtt{c})$ iff $+_{k-i+1}(w_{\mathtt{x}}, w_{\mathtt{y}}, w_{\mathtt{z}}, w_{\mathtt{c}})$.*

Making add_k^i polynomial. At this stage, add_k^i $(i < k-1)$ has size exponential in k, as it is recursively defined using multiple occurrences of add_k^{i+1} (appearing inside eq_k^{i+1} and $succ_k^{i+1}$). However, all these occurrences have the same polarity, i.e. they all appear positively in the antecedents of the implications for the conditions (B) or (C). This property allows us to rely on a *recursion trick* by Fisher and Rabin [20] to obtain a polynomial size formulation of add_k^i. In a nutshell, given a first-order formula $\varphi(\mathbf{x})$ free in the tuple of variables $\mathbf{x}$, the trick consists in rewriting $\psi \stackrel{\text{def}}{=} \varphi(\mathbf{y}) \wedge \varphi(\mathbf{z})$ as $\forall \mathbf{x} : (\mathbf{x} = \mathbf{y} \vee \mathbf{x} = \mathbf{z}) \Rightarrow \varphi(\mathbf{x})$, so that the size of ψ becomes only $|\varphi(\mathbf{x})|$ plus a constant, instead of being roughly twice $|\varphi(\mathbf{x})|$. In a similar way, one can treat arbitrary formulae, as long as all occurrences of $\varphi(\mathbf{x})$ have the same polarity, as it is the case of add_k^{i+1}. The (simple) manipulation of the formula add_k^i using this trick directly leads to a definition of $\mathit{type}(k,n)$ of size polynomial in k and n.

Multi-tiling. The definition of $\mathit{type}(k,n)$ provides the key technical step required to show the lower bounds of Thms. 1 and 2. Using this formula, both theorems can be proved by suitable reductions from the k-exp alternating multi-tiling problem (kAMTP), as we now briefly discuss.

A *multi-tiling system* $\mathcal{P}$ is a tuple $(\mathcal{T}, \mathcal{T}_0, \mathcal{T}_{\text{acc}}, \mathcal{H}, \mathcal{V}, \mathcal{M}, n)$ where $\mathcal{T}$ is a finite set of *tile types*, $\mathcal{T}_0, \mathcal{T}_{\text{acc}} \subseteq \mathcal{T}$ are sets of *initial* and *accepting* tiles, respectively, $n \in \mathbb{N}_+$ (written in unary) is the *dimension* of the system, and $\mathcal{H}, \mathcal{V}, \mathcal{M} \subseteq \mathcal{T} \times \mathcal{T}$ are the horizontal, vertical and multi-tiling matching relations, respectively.

Fix $k \in \mathbb{N}_+$. We write $\widehat{\Sigma}$ for the set of words of length $\mathfrak{t}(k,n)$ over an alphabet Σ. The *initial row* $I(f)$ of a map $f\colon [0, \mathfrak{t}(k,n)-1]^2 \to \mathcal{T}$ is the word $f(0,0), f(0,1), \ldots, f(0, \mathfrak{t}(k,n)-1)$ from $\widehat{\mathcal{T}}$. A *tiling* for the *grid* $[0, \mathfrak{t}(k,n)-1]^2$ is a tuple $(f_1, f_2, \ldots, f_n)$ such that, for all $\ell \in [1,n]$, the following conditions hold:

maps. $f_\ell\colon [0, \mathfrak{t}(k,n)-1]^2 \to \mathcal{T}$ assigns a tile type to each position of the grid;
init & acc. $I(f_\ell) \in \widehat{\mathcal{T}_0}$, and $f_n(\mathfrak{t}(k,n)-1, j) \in \mathcal{T}_{\text{acc}}$ for some $0 \leq j < \mathfrak{t}(k,n)$;
hori. $(f_\ell(i,j), f_\ell(i+1,j)) \in \mathcal{H}$, for every $i \in [0, \mathfrak{t}(k,n)-2]$ and $0 \leq j < \mathfrak{t}(k,n)$;
vert. $(f_\ell(i,j), f_\ell(i,j+1)) \in \mathcal{V}$, for every $j \in [0, \mathfrak{t}(k,n)-2]$ and $0 \leq i < \mathfrak{t}(k,n)$;
multi. if $\ell < n$ then $(f_\ell(i,j), f_{\ell+1}(i,j)) \in \mathcal{M}$ for every $0 \leq i,j < \mathfrak{t}(k,n)$.

The kAMTP takes as input $\mathcal{P}$ and a *quantifier prefix* $\mathbf{Q} = (Q_1, \cdots, Q_n) \in \{\exists, \forall\}^n$, and accepts whenever the statement "$Q_1 w_1 \in \widehat{\mathcal{T}_0} \ldots Q_n w_n \in \widehat{\mathcal{T}_0}$: there is a tiling $(f_1, \ldots, f_n)$ of $[0, \mathfrak{t}(k,n)-1]^2$ s.t. $I(f_\ell) = w_\ell$ for all $\ell \in [1,n]$" is true.

The AExp$_{pol}$-completeness of kAMTP for $k = 1$ can be traced back to [11]. The proof therein is independent from the size of the grid, and can be easily adapted to show kAExp$_{pol}$-completeness for arbitrary k (see [24] for a self-contained presentation). The problem is kNExp-complete if we fix $\mathbf{Q}$ to only contain existential quantifiers. For the lower bound of Thm. 1, we reduce kAMTP on instances with $\mathbf{Q} \in \{\exists\}^n$ to the sat. problem of $\mathsf{ML}(\exists^k_{\mathit{FO}})$, so that the translation produces a formula of $\mathsf{ML}(\exists^1_{\mathit{FO}})$ of modal depth 1 for the case $k = 1$, and otherwise a round-bounded formula from $\mathsf{ML}(\exists^{k-1}_{\mathit{SO}})$ of modal depth k. For Thm. 2 we get a similar reduction, from instances of the kAMTP with arbitrary $\mathbf{Q}$ to $\mathsf{ML}(\exists^k_{\mathit{SO}})$.

The first step is to define an $\mathsf{ML}(\exists^k_{\mathit{FO}})$ formula $\mathit{grid}(k,n)$ that, when satisfied by a pointed forest $(\mathcal{K}, w)$, forces the children of w to encode every position in the grid $[0, \mathfrak{t}(k,n)-1]^2$, together with a formula $\mathit{tiling}(k, \mathcal{P})$ that characterises the various tiling conditions. Fortunately, both these formulae can be defined as in [7], modulo very minor changes. Briefly, each child w' of w shall encode a different pair of numbers $(\mathbf{n}^{\mathcal{H}}_k(w'), \mathbf{n}^{\mathcal{V}}_k(w'))$ representing a position in the grid. The number of bits required to represent $\mathbf{n}^{\mathcal{H}}_k(w')$ and $\mathbf{n}^{\mathcal{V}}_k(w')$ is the same as $\mathbf{n}_k(.)$, which allows us to define $\mathit{grid}(k,n)$ by slightly updating $\mathit{type}(k,n)$. In particular, $\mathbf{n}^{\mathcal{H}}_k(w')$ and $\mathbf{n}^{\mathcal{V}}_k(w')$ can be encoded requiring w' to satisfy $\mathit{type}(k-1,n)$, and by using fresh symbols $p_1^{\mathcal{H}}, \dots, p_n^{\mathcal{H}}, b^{\mathcal{H}}$ and $p_1^{\mathcal{V}}, \dots, p_n^{\mathcal{V}}, b^{\mathcal{V}}$ to encode $(\mathbf{n}_k^{\mathcal{H}}(w'), \mathbf{n}_k^{\mathcal{V}}(w'))$. For $k = 1$, the horizontal position is $\mathbf{n}_1^{\mathcal{H}}(w') \stackrel{\text{def}}{=} \{2^{i-1} : i \in [1, n] \text{ and } w' \in \mathcal{V}(p_i^{\mathcal{H}})\}$. For $k \geq 2$, $\mathbf{n}_k^{\mathcal{H}}(w') \stackrel{\text{def}}{=} \sum\{2^i : \exists w'' \in R(w') \text{ s.t. } \mathbf{n}_{k-1}(w'') = i \text{ and } w'' \in \mathcal{V}(b^{\mathcal{H}})\}$. The vertical position $\mathbf{n}_k^{\mathcal{V}}(w')$ is defined in a similar way. Notice that, in the case of $k \geq 2$, $\mathbf{n}_k^{\mathcal{H}}(w')$ and $\mathbf{n}_k^{\mathcal{V}}(w')$ are defined in terms of $\mathbf{n}_{k-1}(w'')$, and thus using the $\mathfrak{t}(k-1, n)$ children of w'. For $\mathit{tiling}(k, \mathcal{P})$, we see each tile type $t \in \mathcal{T}$ as an atomic proposition, and consider n distinct copies $t^{(1)}, \dots, t^{(n)} \in \mathsf{AP}$ of it, so that the maps $f_1, \dots, f_n$ can be encoded using just the set of worlds forced by $\mathit{grid}(k, n)$. In particular, for every $i \in [1, n]$, each child w' shall satisfy exactly one proposition in $\{t^{(i)} : t \in \mathcal{T}\}$, encoding the fact that $f_i(\mathbf{n}_k^{\mathcal{H}}(w'), \mathbf{n}_k^{\mathcal{V}}(w')) = t$.

Following the above specification, the toolkit of formulae in Fig. 2 can be easily adapted to express properties of the horizontal and vertical positions encoded by a world, leading to the definition of $\mathit{grid}(k, n)$ and $\mathit{tiling}(k, \mathcal{P})$. For instance, given $G \in \{\mathcal{H}, \mathcal{V}\}$ and $\varphi \in \{\mathit{0}_k, \mathit{1}_k, \mathcal{E}_k\}$ we define the formula φ^G as follows: for $k = 1$ we set $\varphi^G \stackrel{\text{def}}{=} \varphi[p_i \leftarrow_0 p_i^G : i \in [1, n]]$, and for $k \geq 2$ we set $\varphi^G \stackrel{\text{def}}{=} \varphi[b \leftarrow_1 b^G]$. Then, w' satisfies the formula $\mathit{1}_k^{\mathcal{H}} \land \mathit{0}_k^{\mathcal{V}}$ whenever $(\mathbf{n}_k^{\mathcal{H}}(w'), \mathbf{n}_k^{\mathcal{V}}(w')) = (1, 0)$.

Lemma 3. *The $\mathsf{ML}(\exists^k_{\mathit{FO}})$ formula $\mathit{grid}(k, n) \land \mathit{tiling}(k, \mathcal{P})$ is satisfiable if and only if kAMTP accepts on input $(\mathcal{P}, \mathbf{Q})$, with $\mathbf{Q} \in \{\exists\}^n$.*

For the lower bound of Thm. 2, it remains to show how to capture in $\mathsf{ML}(\exists^k_{\mathit{SO}})$ the arbitrary prefixes of quantification $\mathbf{Q} = (Q_1, \dots, Q_n)$ of kAMTP. Compared to [6,7], novel machinery is required to perform this step. As $\mathsf{ML}(\exists^k_{\mathit{SO}})$ captures $\mathsf{ML}(\exists^k_{\mathit{FO}})$, we now see $\mathit{grid}(k, n)$ and $\mathit{tiling}(k, \mathcal{P})$ as formulae of $\mathsf{ML}(\exists^k_{\mathit{SO}})$. For each tile type $t \in \mathcal{T}$, we consider an additional set of copies $t^{(n+1)}, \dots, t^{(2n)} \in \mathsf{AP}$. We also define $\mathbf{t}^{(i)} \stackrel{\text{def}}{=} (t_1^{(i)}, \dots, t_r^{(i)})$, where $\mathcal{T} = \{t_1, \dots, t_r\}$. We use the propositions in $\mathbf{t}^{(n+i)}$ to simulate the quantifier Q_i, which we recall quantifies over the possible initial rows $I(f_i) \in \widehat{\mathcal{T}_0}$ of the map f_i. If $Q_i = \exists$, we simulate this form of quantification with the following shortcut, parametric on φ:

$$E_i(\varphi) \stackrel{\text{def}}{=} \exists^1 \mathbf{t}^{(n+i)} : \varphi \land \Box(\mathit{0}_k^{\mathcal{H}} \Rightarrow \textstyle\bigvee_{t \in \mathcal{T}_0}(t^{(n+i)} \land \bigwedge_{s \in \mathcal{T} \setminus \{t\}} \neg s^{(n+i)})).$$

Here, the last conjunct states that each world encoding a position $(0, j)$ of the grid, for some $j \in [0, \mathfrak{t}(k, n) - 1]$, satisfies exactly one proposition $t^{(n+i)}$ with $t \in \mathcal{T}_0$. For $Q_i = \forall$, we just define $A_i(\varphi) \stackrel{\text{def}}{=} \neg E_i(\neg\varphi)$. Then, the prefix of quantification $\mathbf{Q}$ is captured by $\mathbf{Q}(\varphi) \stackrel{\text{def}}{=} Q_1(Q_2(\dots Q_n(\varphi)))$, where $Q_i(\varphi) \stackrel{\text{def}}{=} E_i(\varphi)$ if $Q_i = \exists$, else $Q_i(\varphi) \stackrel{\text{def}}{=} A_i(\varphi)$. In deciding whether $\mathcal{K}, w \models \mathbf{Q}(\varphi)$ holds for a pointed forest $(\mathcal{K}, w)$ satisfying $\mathit{grid}(k, n)$, the satisfaction of φ is checked w.r.t. a model where each world encoding a position $(0, j)$ of the grid satisfies exactly one $t^{(n+i)}$ with $t \in \mathcal{T}_0$, for all $i \in [1, n]$. In terms of tilings, this corresponds to having set the initial row $I(f_i) \in \widehat{\mathcal{T}_0}$ of each of the maps f_i. We now want to tile the remaining part of the grid by finding a suitable instantiation for φ. To do so, we quantify over all $\mathbf{t}^{(1)}, \dots \mathbf{t}^{(n)}$, searching for an arrangement of these propositions that satisfies $\mathit{tiling}(k, \mathcal{P})$ and such that, on worlds encoding a position $(0, j)$ of the grid, the satisfaction of propositions in $\mathbf{t}^{(i)}$ mirrors the satisfaction of the corresponding propositions in $\mathbf{t}^{(n+i)}$. In formula:

$$\overline{\mathit{tiling}}(k, \mathcal{P}) \stackrel{\text{def}}{=} \exists^1 \mathbf{t}^{(1)}, \dots, \mathbf{t}^{(n)} : \mathit{tiling}(k, \mathcal{P}) \land \Box(\mathit{0}_k^{\mathcal{H}} \Rightarrow \textstyle\bigwedge_{i=1}^n \bigvee_{t \in \mathcal{T}}(t^{(i)} \Leftrightarrow t^{(n+i)})).$$

Lemma 4. *The* $\mathsf{ML}(\exists^k_{SO})$ *formula* $grid(k,n) \wedge \mathbf{Q}(\overline{tiling}(k,\mathcal{P}))$ *is satisfiable if and only if* kAMTP *accepts on input* $(\mathcal{P},\mathbf{Q})$.

Round-boundedness. In defining $type(k,n)$, we made sure to respect the following round-boundedness condition: $type(1,n)$ has modal depth 1 and belongs to $\mathsf{ML}(\exists^1_{FO})$, whereas for every $k \geq 2$, $type(k,n)$ is a round-bounded formula of $\mathsf{ML}(\exists^{k-1}_{FO})$ of modal depth k. The same holds for $grid(k,n)$, $tiling(k,\mathcal{P})$ and $\mathbf{Q}(\overline{tiling}(k,\mathcal{P}))$. Then, Lemmas 3 and 4 imply the lower bounds of Thms. 1 and 2.

4 Upper bounds via a small-model property for $\mathsf{ML}(\exists^k_{SO})$

In this section, we establish the following small model property.

Proposition 1. *Each satisfiable round-bounded formula* φ *in* $\mathsf{ML}(\exists^k_{SO})$ *is satisfied by a pointed forest with* $\mathfrak{t}(k{+}1,\mathcal{O}(|\varphi|))$ *worlds. Each satisfiable* φ *in* $\mathsf{ML}(\exists^k_{SO})$ *with* $\mathsf{md}(\varphi) \leq k$ *is satisfied by a pointed forest with* $\mathfrak{t}(k,\mathcal{O}(|\varphi|^3))$ *worlds.*

As the logic $\mathsf{ML}(\exists^k_{SO})$ captures $\mathsf{ML}(\exists^k_{FO})$, Prop. 1 transfers to the latter logic. With this result at hand, the upper bounds of Thm. 1 and Thm. 2 easily follow. Consider a round-bounded formula φ of either $\mathsf{ML}(\exists^k_{SO})$ of $\mathsf{ML}(\exists^k_{FO})$ (the arguments for a formula of modal depth k are similar). First, we guess a pointed forest $(\mathcal{K},w)$ with bounds as in Prop. 1. This can be done in $(k{+}1)$NExp. Then, we check whether $(\mathcal{K},w)$ satisfies φ. For $\mathsf{ML}(\exists^k_{SO})$, by seeing this logic as a fragment of *monadic second-order logic*, this can be done in polynomial time in the sizes of $(\mathcal{K},w)$ and φ by using an alternating Turing machine that performs $|\varphi|$ many alternations. As $(\mathcal{K},w)$ has $(k{+}1)$-exponential size with respect to $|\varphi|$, the whole algorithm runs in $(k{+}1)\mathrm{AExp}_{pol}$. For $\mathsf{ML}(\exists^k_{FO})$, we rely on the fact that there is a deterministic algorithm for the model checking problem of *first-order logic* that runs in time $\mathcal{O}(|\varphi| \cdot M^{|\varphi|})$ where M is the size of the model. From the bounds on $(\mathcal{K},w)$ we conclude that the procedure for $\mathsf{ML}(\exists^k_{FO})$ is in $(k{+}1)$NExp.

Prop. 1 is shown through a *quantifier elimination (QE) procedure* that translates every formula of $\mathsf{ML}(\exists^k_{SO})$ into an equivalent formula from GML, establishing Cor. 2 as a by-product. Without loss of generality, in this section we extend $\mathsf{ML}(\exists^k_{SO})$ with graded modalities $\Diamond_{\geq j}\varphi$, with $j \in \mathbb{N}$ given in unary, and see the modality $\Diamond$ as a shortcut for $\Diamond_{\geq 1}$. Recall that a GML formula $\Diamond_{\geq j}\varphi$ can be represented with an $\mathsf{ML}(\exists^k_{SO})$ formula of size $\mathcal{O}(j+|\varphi|)$ (Sec. 2).

Parameters of a formula. Fig. 4 introduces a set of parameters for a $\mathsf{ML}(\exists^k_{SO})$ formula φ, which we rely on to establish Prop. 1. For instance, for $\varphi = (p \vee \Diamond_{\geq 3} r) \wedge (q \vee \Diamond_{\geq 5}\Diamond_{\geq 2} q)$ we have $\mathsf{ap}(1,\varphi) = \{r\}$, $\mathsf{gsf}(0,\varphi) = \{\Diamond_{\geq 3} r, \Diamond_{\geq 5}\Diamond_{\geq 2} q\}$, $\mathsf{msf}(1,\varphi) = \{r, \Diamond_{\geq 2} q\}$, $\mathsf{gsf}(1,\varphi) = \{\Diamond_{\geq 2} q\}$, $\mathsf{gr}(0,\varphi) = 5$ and $\mathsf{bd}(0,\varphi) = 8$. Note that every GML formula φ is a Boolean combination of formulae from $\mathsf{ap}(0,\varphi) \cup \mathsf{gsf}(0,\varphi)$, and for every $d \in \mathbb{N}$, $\mathsf{bd}(d,\varphi) \leq \mathsf{gr}(d,\varphi) \cdot |\mathsf{msf}(d+1,\varphi)|$.

For a set of formulae $\Phi = \{\varphi_1,\dots,\varphi_n\}$, we define $\mathcal{C}(\Phi)$ to be the set of all complete conjunctions of possibly negated formulae of Φ. Formally, $\mathcal{C}(\Phi) \overset{\text{def}}{=} \{\gamma_1 \wedge \dots \wedge \gamma_n : \text{for all } i \in [1,n], \gamma_i \in \{\varphi_i, \neg\varphi_i\}\}$, and we fix $\mathcal{C}(\varnothing) = \{\top\}$. Given $\mathsf{P} \subseteq_{fin} \mathsf{AP}$ we refer to the formulae in $\mathcal{C}(\mathsf{P})$ as $\boldsymbol{\rho}_1, \boldsymbol{\rho}_2, \cdots$.

$\mathsf{ap}(d, \varphi)$	: set of atomic propositions of φ in the scope of exactly d graded modalities.
$\mathsf{gsf}(d, \varphi)$	: set of *subformulae* $\Diamond_{\geq j}\psi$ of φ, in the scope of exactly d graded modalities.
$\mathsf{msf}(d, \varphi)$	: set of *maximal subformulae* of φ in the scope of d graded modalities: $\mathsf{msf}(0, \varphi) = \{\varphi\}$, and $\psi \in \mathsf{msf}(d+1, \varphi)$ iff $\Diamond_{\geq j}\psi \in \mathsf{gsf}(d, \varphi)$ for some $j \in \mathbb{N}$.
$\mathsf{gr}(d, \varphi)$	: largest $j \in \mathbb{N}$ such that either $j = 0$ or $\Diamond_{\geq j}\psi \in \mathsf{gsf}(d, \varphi)$, for some ψ.
$\mathsf{bd}(d, \varphi)$	: for $d = 0$ and let $\mathsf{gsf}(0, \varphi) = \{\Diamond_{\geq j_1}\psi_1, \dots, \Diamond_{\geq j_n}\psi_n\}$, $\mathsf{bd}(0, \varphi) \stackrel{\text{def}}{=} j_1 + \dots + j_n$. For $d \geq 1$, $\mathsf{bd}(d, \varphi) \stackrel{\text{def}}{=} \max\{\mathsf{bd}(d-1, \psi) : \psi \in \mathsf{msf}(1, \varphi)\}$.

Fig. 4: Parameters of an $\mathsf{ML}(\exists^k)$ formula φ ($d \in \mathbb{N}$).

Normal forms. We introduce a set of normal forms that are used by our QE procedure. An $\mathsf{ML}(\exists^k_{\mathit{SO}})$ formula φ is in *prenex normal form* if it is of the form $Q_1 p_1 Q_2 p_2 \dots Q_n p_n \psi$ where $Q_i \in \{\exists^k, \forall^k\}$ and ψ is in GML. If ψ is instead in $\mathsf{ML}(\exists^k_{\mathit{SO}})$ but all quantifiers are under the scope of at least k modalities, we say that φ is in *prenex normal form up to* k. An $\mathsf{ML}(\exists^k_{\mathit{SO}})$ formula φ is in *prenex round-bounded (p.r.b.) form* if φ is round-bounded and, for all $i \in \mathbb{N}$, all formulae in $\mathsf{msf}(i \cdot k, \varphi)$ are in prenex normal form up to k. E.g., given a p.r.b. formula ψ in $\mathsf{ML}(\exists^2_{\mathit{SO}})$, $\exists^2 p\, \exists^2 q \Diamond\Diamond\exists^2 r\, \psi$ is in p.r.b. form, while $\exists^2 p \Diamond \exists^1 q \Diamond \exists^2 r\, \psi$ is not. Thanks to the equivalences below one can translate each round-bounded formula φ of $\mathsf{ML}(\exists^k_{\mathit{SO}})$ into an equivalent well-quantified p.r.b. formula of size $\mathcal{O}(|\varphi|)$:

$$\Diamond\exists^{k-1} p\, \varphi \equiv \exists^k p \Diamond\varphi, \qquad \Box\exists^{k-1} p\, \varphi \equiv_{\mathit{SO}} \exists^k p \Box\varphi, \qquad \text{for } k \geq 2. \qquad (\ddagger)$$

Similarly, every φ in $\mathsf{ML}(\exists^k_{\mathit{SO}})$ of modal depth at most k can be translated into a well-quantified prenex formula of $\mathsf{ML}(\exists^k_{\mathit{SO}})$ having size $\mathcal{O}(|\varphi|)$. Notice that the second equivalence in (‡) only holds on pointed forests and for the logic $\mathsf{ML}(\exists^k_{\mathit{SO}})$. It does not hold for arbitrary Kripke structures, nor for $\mathsf{ML}(\exists^k_{\mathit{FO}})$.

Our QE procedure translates each formula of $\mathsf{ML}(\exists^k_{\mathit{SO}})$ into a GML formula in *disjoint normal form* (called *good formulae* in [23, Def. 8.5]) for which it is easy to estimate bounds on the size of the smallest satisfying pointed forest, if any. We say that a set $\{\varphi_1, \dots, \varphi_n\}$ of formulae in GML is a *disjoint set over* $\mathsf{P} \subseteq_{\mathit{fin}} \mathsf{AP}$ whenever for all $i, j \in [1, n]$, we have $\varphi_i = \boldsymbol{\rho}_i \wedge \gamma_i$ and $\varphi_j = \boldsymbol{\rho}_j \wedge \gamma_j$, where $\boldsymbol{\rho}_i, \boldsymbol{\rho}_j \in \mathcal{C}(\mathsf{P})$, $\mathsf{ap}(0, \gamma_i) \cap \mathsf{P} = \mathsf{ap}(0, \gamma_j) \cap \mathsf{P} = \varnothing$, and either $\gamma_i \equiv \gamma_j$ or $(\gamma_i \wedge \gamma_j) \equiv \perp$. By taking $\boldsymbol{\rho}_i$ and $\boldsymbol{\rho}_j$ up-to commutativity and associativity of $\wedge$, a disjoint set over P is also a disjoint set over any $\mathsf{P}' \subset \mathsf{P}$. We say that φ is in *disjoint normal form* (DisjNF) if for every $d \in [0, \mathsf{md}(\varphi)]$, $\mathsf{msf}(d, \varphi)$ is a disjoint set over $\varnothing$.

Proposition 2 ([23], Lemma 8.7). *Each satisfiable* GML *formula* φ *in DisjNF is satisfied by a pointed forest with at most* $(\max_{d \in \mathbb{N}}(\mathsf{bd}(d, \varphi)) + 1)^{\mathsf{md}(\varphi)}$ *worlds.*

To translate a well-quantified p.r.b. formula φ from $\mathsf{ML}(\exists^k_{\mathit{SO}})$ into a GML formula in DisjNF, we consider the largest $i \in \mathbb{N}$ for which $\mathsf{msf}(i \cdot k, \varphi)$ is non-empty, and inductively translate, for each $j = i, i-1, \cdots, 0$, all formulae in $\mathsf{msf}(j \cdot k, \varphi)$ into equivalent ones in GML. At each of these $i + 1$ rounds, the following two steps are applied at most k times:

1. Let $\ell = \min\{r \in \mathbb{N}_+ :$ all formulae of $\mathsf{msf}(j \cdot k, \varphi)$ are in $\mathsf{ML}(\exists^r_{\mathit{SO}})\}$. We update all $\psi \in \mathsf{msf}(j \cdot k, \varphi)$ so that $\mathsf{msf}(\ell, \psi)$ becomes a disjoint set over $\mathsf{bp}(\psi)$.
2. By manipulating all quantified propositions of the formulae in $\mathsf{msf}(\ell, \psi)$, we translate ψ into a formula of either GML (if $\ell = 1$) or $\mathsf{ML}(\exists^{\ell-1}_{\mathit{SO}})$ (if $\ell \geq 2$).

At the end of the round, $\mathsf{msf}(j \cdot k, \varphi)$ solely contains GML formulae in DisjNF, and the next round considers the set $\mathsf{msf}((j-1) \cdot k, \varphi)$, that now contains $\mathsf{ML}(\exists^k_{\mathit{SO}})$ formulae in prenex normal form. The QE procedure has thus three key steps, which we now formalise: (I) manipulating a formula φ so that $\mathsf{msf}(j, \varphi)$ becomes a disjoint set, (II) eliminating the quantifier $\exists^1$ obtaining a formula from GML, and (III) reducing the elimination of $\exists^\ell$ to the elimination of $\exists^{\ell-1}$ (for $\ell \geq 2$).

Step (I): making a single set disjoint. Let $j \in \mathbb{N}_+$ and $\mathsf{P} \subseteq_{\mathit{fin}} \mathsf{AP}$. We show how to transform a GML formula φ into an equivalent formula ψ such that $\mathsf{msf}(j, \psi)$ is a disjoint set over P. Two strategies are possible, which will be combined and carefully chosen in order to obtain the bounds required by Prop. 1.

The *first strategy* considers the set $\mathcal{S} \stackrel{\text{def}}{=} \mathcal{C}(\mathsf{P} \cup \mathsf{ap}(j, \varphi) \cup \mathsf{gsf}(j, \varphi))$, which is disjoint over P (and so over $\varnothing$), and rewrites φ into an equivalent formula ψ with $\mathsf{msf}(j, \psi) \subseteq \mathcal{S}$. Consider $\gamma \in \mathsf{msf}(j, \varphi)$. By definition of $\mathcal{C}(.)$, $\bigvee_{\chi \in \mathcal{S}} \chi$ is a tautology, and since γ is a Boolean combination of formulae in $\mathsf{ap}(j, \varphi) \cup \mathsf{gsf}(j, \varphi)$, for all $\chi \in \mathcal{S}$ the formula $\gamma \wedge \chi$ is equivalent to either $\bot$ or χ. Then, $\gamma \equiv \bigvee_{\chi \in T} \chi$ for some $T \subseteq \mathcal{S}$. Notice that $\gamma \in \mathsf{msf}(j, \varphi)$ holds if and only if $\Diamond_{\geq i} \gamma \in \mathsf{gsf}(j-1, \varphi)$, for some $i \in \mathbb{N}$. By relying on the equivalence of GML

$$\Diamond_{\geq i}(\chi_1 \vee \chi_2) \equiv \bigvee_{i = i_1 + i_2} (\Diamond_{\geq i_1} \chi_1 \wedge \Diamond_{\geq i_2} \chi_2), \qquad \text{whenever } \chi_1 \wedge \chi_2 \equiv \bot,$$

we rewrite $\Diamond_{\geq i} \gamma$ into a Boolean combination of formulae $\Diamond_{\geq i'} \chi$ with $i' \leq i$ and $\chi \in T \subseteq \mathcal{S}$. These steps are applied to all the formulae in $\mathsf{msf}(j, \varphi)$.

The *second strategy* is as follows: for each $\gamma \in \mathsf{msf}(j, \varphi)$ and $\boldsymbol{\rho} \in \mathcal{C}(\mathsf{P})$, let $\gamma_{\boldsymbol{\rho}} \stackrel{\text{def}}{=} \gamma[p \leftarrow_0 v : v \in \{\top, \bot\}, p \in \mathsf{P}, \text{ and } v = \top \text{ iff } p \text{ occurs positively in } \boldsymbol{\rho}]$. Notice that $\mathsf{ap}(0, \gamma_{\boldsymbol{\rho}}) \cap \mathsf{P} = \varnothing$. As $\boldsymbol{\rho}$ gives a polarity to all propositions in P, we have $\boldsymbol{\rho} \wedge \gamma \equiv \boldsymbol{\rho} \wedge \gamma_{\boldsymbol{\rho}}$. Set $\mathcal{T} \stackrel{\text{def}}{=} \mathcal{C}(\{\gamma_{\boldsymbol{\rho}} : \gamma \in \mathsf{msf}(j, \varphi), \boldsymbol{\rho} \in \mathcal{C}(\mathsf{P})\})$. Consider $\mathcal{S}' \stackrel{\text{def}}{=} \mathcal{C}(\mathsf{P} \cup \mathcal{T})$, which is a disjoint set over P, and replay the arguments used for $\mathcal{S}$ in the first strategy to rewrite φ into an equivalent formula ψ with $\mathsf{msf}(j, \psi) \subseteq \mathcal{S}'$.

While both strategies keep most of the parameters of Fig. 4 unchanged (one exception being $\mathsf{ap}(j, \psi) \subseteq \mathsf{ap}(j, \varphi) \cup \mathsf{P}$), they yield profoundly different bounds on the size of $\mathsf{msf}(j, \psi)$. Because of the definition of $\mathcal{S}$, from the first strategy we obtain $|\mathsf{msf}(j, \psi)| \leq 2^{|\mathsf{P}| + |\mathsf{ap}(j, \varphi)| + |\mathsf{gsf}(j, \varphi)|}$, where we highlight the exponential dependence on $|\mathsf{gsf}(j, \varphi)|$, and thus on the number of outermost graded modalities appearing in formulae of $\mathsf{msf}(j, \varphi)$. From the definition of $\mathcal{S}'$, the second strategy yields $|\mathsf{msf}(j, \psi)| \leq 2^{|\mathsf{P}| + 2^{|\mathsf{P}|} \cdot |\mathsf{msf}(j, \varphi)|}$. Here, $|\mathsf{msf}(j, \psi)|$ does not depend on $\mathsf{gsf}(j, \varphi)$, but it is doubly exponential in $|\mathsf{P}|$. Remarkably, in both strategies $\mathsf{gsf}(j, \psi) \subseteq \mathsf{gsf}(j, \varphi)$, thus if $\mathsf{msf}(j+1, \varphi)$ is a disjoint set over $\varnothing$, so is $\mathsf{msf}(j+1, \psi)$. This property is essential, as it allows us to bring the full formula in DisjNF.

Step (II): eliminating $\exists^1$. Given a well-quantified formula $\varphi = \exists^1 p \, \varphi'$, where φ' is in GML and $\mathsf{msf}(1, \varphi)$ is a disjoint set over P, and $p \in \mathsf{P}$, it is quite easy to eliminate the quantifier $\exists^1 p$ and produce a formula ψ in GML equivalent to φ and such that $\mathsf{msf}(1, \psi)$ is a disjoint set over $\mathsf{P} \setminus \{p\}$. We sketch here the main points. First, from standard axioms of propositional calculus and by distributing $\exists^1 p$ over $\vee$, we obtain a representation of φ as a disjunction of formulae of the form $\exists^1 p \, (\boldsymbol{\rho} \wedge \gamma)$ with $\boldsymbol{\rho} \in \mathcal{C}(\mathsf{ap}(0, \varphi))$ and $\gamma \in \mathcal{C}(\mathsf{gsf}(0, \varphi))$. We eliminate the quantifier $\exists^1$ from every such disjunct $\exists^1 p\,(\boldsymbol{\rho} \wedge \gamma)$. Below, let χ be an arbitrary formula with $p \notin \mathsf{ap}(0, \chi)$. First, using the equivalences $\exists^1 p\,(p \wedge \chi) \equiv_{\mathit{SO}} \exists^1 p\, \chi$ and $\exists^1 p\,(\neg p \wedge \chi) \equiv_{\mathit{SO}} \exists^1 p\, \chi$, we get rid of the occurrences of p in $\boldsymbol{\rho}$, obtaining a formula $\boldsymbol{\rho}' \in \mathcal{C}(\mathsf{ap}(0, \varphi) \setminus \{p\})$. Next, we remove p from γ thanks to the equivalences:

$$\exists^1 p\, : \Diamond_{\geq i}(p \wedge \chi) \wedge \Diamond_{\geq j}(\neg p \wedge \chi) \;\equiv_{\mathit{SO}}\; \Diamond_{\geq i+j}\chi;$$
$$\exists^1 p\, : \neg\Diamond_{\geq i}(p \wedge \chi) \wedge \neg\Diamond_{\geq j}(\neg p \wedge \chi) \;\equiv_{\mathit{SO}}\; \neg\Diamond_{\geq i+j-1}\chi.$$

We obtain a GML formula γ' such that $\exists^1 p\,(\boldsymbol{\rho} \wedge \gamma) \equiv_{\mathit{SO}} \boldsymbol{\rho}' \wedge \gamma'$. Size-wise, Step (II) preserves all the parameters of Fig. 4 except $\mathsf{gr}(0, \psi) \leq 2 \cdot \mathsf{gr}(0, \varphi)$.

Step (III): from $\exists^{k+1}$ to $\exists^k$. Consider a well-quantified $\mathsf{ML}(\exists^k_{\mathit{SO}})$ formula φ' having all quantifiers appearing outside the scope of graded modalities, and with the set $\mathsf{msf}(k+1, \varphi')$ disjoint over P. Given $p \in \mathsf{P}$, we translate $\varphi \stackrel{\text{def}}{=} \exists^{k+1} p\, \varphi'$ into an equivalent well-quantified $\mathsf{ML}(\exists^k_{\mathit{SO}})$ formula ψ having all quantifiers outside the scope of graded modalities, and with the set $\mathsf{msf}(k+1, \psi)$ disjoint over $\mathsf{P} \setminus \{p\}$. This is done by replacing $\exists^{k+1} p$ with multiple $\exists^k$. The first step is to single out the occurrences of p under the scope of $k{+}1$ modalities by replacing them with a fresh symbol $\widetilde{p}$ and splitting $\exists^{k+1} p$ into $\exists^k p$ and $\exists^{k+1} \widetilde{p}$. We get $\varphi \equiv_{\mathit{SO}} \exists^k p\; \exists^{k+1} \widetilde{p}\, \varphi''$ where $\varphi'' = \varphi'[p \leftarrow_{k+1} \widetilde{p}]$. Let $\mathsf{gsf}(k, \varphi'') = \{\Diamond_{\geq k_1}\chi_1, \dots, \Diamond_{\geq k_n}\chi_n\}$. From the properties of φ', no proposition from $\mathsf{bp}(\varphi'')$ appears in the GML formulae $\chi_1, \dots, \chi_n$. Using fresh propositions $q_1, \dots, q_n$, we rewrite φ as

$$\exists^k p\; \exists^{k+1} \widetilde{p}\; \exists^k q_1, \dots, q_n\, : \varphi''[\Diamond_{\geq k_i}\chi_i \leftarrow_k q_i : 1 \leq i \leq n] \wedge \Box^k \textstyle\bigwedge_{i=1}^n (q_i \Leftrightarrow \Diamond_{\geq k_i}\chi_i).$$

Essentially, the subformula $\Box^k \bigwedge_{i=1}^n (q_i \Leftrightarrow \Diamond_{\geq k_i}\chi_i)$ constraints each q_i to be true in exactly those worlds satisfying $\Diamond_{\geq k_i}\chi_i$. This allows us to replace with q_i all occurrences of $\Diamond_{\geq k_i}\chi_i$ appearing in φ'' under the scope of k modalities (first conjunct of the formula above), without changing the semantics of φ. By definition, $\varphi''[\Diamond_{\geq k_i}\chi_i \leftarrow_k q_i : 1 \leq i \leq n]$ has modal depth at most k, and thus the proposition $\widetilde{p}$ does not occur in it. We reorder the existential prefix of the formula and, by distributing $\exists^{k+1}\widetilde{p}$, conclude that φ is equivalent to:

$$\exists^k p, q_1, \dots, q_n\, : \varphi''[\Diamond_{\geq k_i}\chi_i \leftarrow_k q_i : 1 \leq i \leq n] \wedge \exists^{k+1}\widetilde{p}\,\Box^k \textstyle\bigwedge_{i=1}^n (q_i \Leftrightarrow \Diamond_{\geq k_i}\chi_i).$$

Lastly, we eliminate $\exists^{k+1}\widetilde{p}$, obtaining the aforementioned $\mathsf{ML}(\exists^k_{\mathit{SO}})$ formula ψ. Using the second equivalence in (‡), we rewrite $\exists^{k+1}\widetilde{p}\,\Box^k \bigwedge_{i=1}^n (q_i \Leftrightarrow \Diamond_{\geq k_i}\chi_i)$ into $\Box^k \exists^1 \widetilde{p} \bigwedge_{i=1}^n (q_i \Leftrightarrow \Diamond_{\geq k_i}\chi_i)$. Since $\{\chi_1, \dots, \chi_n\}$ is a set of formulae form GML that is disjoint over $(\mathsf{P} \setminus \{p\}) \cup \{\widetilde{p}\}$, by applying Step (II) one computes a formula ψ' in GML equivalent to $\exists^1 \widetilde{p} \bigwedge_{i=1}^n (q_i \Leftrightarrow \Diamond_{\geq k_i}\chi_i)$ and such that $\mathsf{msf}(1, \psi')$ is a disjoint set over $\mathsf{P} \setminus \{p\}$. Then, the (output) formula ψ is defined as follows:

$$\psi \stackrel{\text{def}}{=} \exists^k p, q_1, \dots, q_n\, : \varphi''[\Diamond_{\geq k_i}\chi_i \leftarrow_k q_i : 1 \leq i \leq n] \wedge \Box^k \psi'.$$

Down to GML*, inductively.* The manipulation we just described yield the crucial inductive argument that allows us to translate any well-quantified prenex formula of $\mathsf{ML}(\exists^k_{\mathit{SO}})$ into a formula of GML. Inductively on k, consider a well-quantified formula $\varphi = Q_1 p_1 \dots Q_n p_n \varphi'$ where each $Q_i \in \{\exists^k, \forall^k\}$, the formula φ' is in GML and $\mathsf{msf}(k, \varphi)$ is a disjoint set over $\{p_1, \dots, p_n\}$. If $k = 1$, we repeatedly apply Step (II) to translate φ into a GML formula. If $k \geq 2$, starting from p_n down to p_1, we apply Step (III) to translate φ into a well-quantified prenex formula χ from $\mathsf{ML}(\exists^{k-1}_{SO})$. Afterwards, we rely on the first strategy of Step (I) to make the set $\mathsf{msf}(k-1,\chi)$ disjoint over $\mathsf{bp}(\chi)$, and inductively obtain a GML formula ψ equivalent to φ. For a sake of conciseness, let $|\varphi|_k \stackrel{\text{def}}{=} \max(k, |\bigcup_{i\in[0,k]} \mathsf{ap}(i,\varphi)|, \max_{i<k} \mathsf{gr}(i,\varphi))$. Fundamentally, the formula ψ has the same modal depth as φ, and for every $i \in [0, k-1]$ it satisfies:

$$\mathsf{gr}(i,\psi) \le \mathfrak{t}(k-1, 2^{8\cdot|\varphi|_k} \cdot |\mathsf{msf}(k,\varphi)|); \quad \mathsf{msf}(i,\psi) \le \mathfrak{t}(k-1, 2^{8\cdot|\varphi|_k} \cdot |\mathsf{msf}(k,\varphi)|).$$

With these bounds at hand, Prop. 1 follows from Steps (I)–(III) and Prop. 2. First, consider the case of a well-quantified prenex formula φ in $\mathsf{ML}(\exists^k)$ of modal depth k. Using the first strategy from Step (I), we translate φ into an equivalent formula ψ such that the set $\mathsf{msf}(k,\psi)$ is disjoint over $\mathsf{bp}(\varphi)$ and has size exponential in $|\varphi|$. We apply the inductive argument discussed above, and translate ψ into a GML formula χ in DisjNF with $\mathsf{md}(\chi) \le \mathsf{md}(\varphi)$ and $\mathsf{bd}(d,\chi) \le \mathsf{gr}(d,\chi) \cdot |\mathsf{msf}(d+1,\chi)|) \le \mathfrak{t}(k, \mathcal{O}(|\varphi|^2))$ for all $d \in \mathbb{N}$. By Prop. 2, whenever satisfiable, φ is satisfied by a pointed forest with at most $\mathfrak{t}(k, \mathcal{O}(|\varphi|^3))$ worlds. The case of general p.r.b. formulae of $\mathsf{ML}(\exists^k_{SO})$ is similar, but we need to appeal to the second strategy of Step (I) to stop the chain of exponential blow-ups. For simplicity, let us consider the case of φ being a well-quantified p.r.b. formula of modal depth at most $2k$. The arguments used for this case can be adapted for formulae of arbitrary modal depth. First, we look at the formulae of $\mathsf{msf}(k,\varphi)$, whose modal depth is at most k, and eliminate all local quantifiers from each of these formulae, as described above. In doing so, $|\mathsf{gsf}(k,\varphi)|$ witnesses a k-exponential blow-up, but the size of $\mathsf{msf}(k,\varphi)$ is unchanged. We consider the quantification prefix of φ, and eliminate all its quantifiers over P to produce an equivalent formula from GML. The first step is to make the set $\mathsf{msf}(k,\varphi)$ a disjoint set over P. Because of the k-exponential blow-up on $\mathsf{gsf}(k,\varphi)$, the first strategy of Step (I) is of no use. We appeal to the second one, which modifies $\mathsf{msf}(k,\varphi)$ into a disjoint set of size only doubly-exponential in the size of the original formula φ. By relying on the inductive reasoning discussed above, we produce the equivalent GML formula in DisjNF. Because of the doubly-exponential bound on $\mathsf{msf}(k,\varphi)$, this elimination is exponentially worse than the one done for formulae of modal depth at most k. Then, appealing to Prop. 2 yields Prop. 1.

5 Further connections

In introducing $\mathsf{ML}(\exists^k_{FO})$ and $\mathsf{ML}(\exists^k_{SO})$, one of our goals is to provide a common framework to relate several modal logics featuring propositional quantification in disguise. Apart from the relations stated in Sec. 2, in an extended version of this work we aim at establishing connections between $\mathsf{ML}(\exists^1_{SO})$ and *propositional team logics* [21], *propositional logic of dependence* [32] and *ambient logics* [13]; as well as connections bwteen $\mathsf{ML}(\exists^{\infty}_{FO})$ and *sabotage logics* [8,4].

Acknowledgments. R. Fervari is supported by CONICET project PIP 11220200100812CO, and by the LIA SINFIN. A. Mansutti is supported by the ERC project ARiAT (Grant agreement No. 852769).

References

1. Andréka, H., Németi, I., van Benthem, J.: Modal languages and bounded fragments of predicate logic. Journal of Philosophical Logic **27**(3), 217–274 (1998)
2. Areces, C., Blackburn, P., Marx, M.: Hybrid logics: characterization, interpolation and complexity. The Journal of Symbolic Logic **66**(3), 977–1010 (2001)
3. Areces, C., ten Cate, B.: Hybrid logics. In: Handbook of Modal Logic, Studies in logic and practical reasoning, vol. 3, pp. 821–868. North-Holland (2007)
4. Areces, C., Fervari, R., Hoffmann, G.: Relation-changing modal operators. Logic Journal of the IGPL **23**(4), 601–627 (2015)
5. Barnaba, M.F., Caro, F.D.: Graded modalities. Studia Logica **44**(2), 197–221 (1985)
6. Bednarczyk, B., Demri, S.: Why propositional quantification makes modal logics on trees robustly hard? In: Logic in Computer Science. pp. 1–13. IEEE (2019)
7. Bednarczyk, B., Demri, S., Fervari, R., Mansutti, A.: Modal logics with composition on finite forests: Expressivity and complexity. In: Logic in Computer Science. pp. 167–180. ACM (2020)
8. van Benthem, J.: An essay on sabotage and obstruction. In: Mechanizing Mathematical Reasoning. LNCS, vol. 2605, pp. 268–276 (2005)
9. Blackburn, P., Braüner, T., Kofod, J.: Remarks on Hybrid Modal Logic with Propositional Quantifiers, pp. 401–426. No. 4 in Logic and Philosophy of Time (2020)
10. Blackburn, P., Wolter, F., van Benthem, J. (eds.): Handbook of Modal Logics, Studies in logic and practical reasoning, vol. 3. Elsevier (2006)
11. Bozzelli, L., Molinari, A., Montanari, A., Peron, A.: On the complexity of model checking for syntactically maximal fragments of the interval temporal logic HS with regular expressions. In: GandALF'17. EPTCS, vol. 256, pp. 31–45 (2017)
12. Bull, R.A.: On modal logic with propositional quantifiers. The Journal of Symbolic Logic **34**(2), 257–263 (1969)
13. Calcagno, C., Cardelli, L., Gordon, A.: Deciding validity in a spatial logic for trees. In: International Workshop on Types in Languages Design and Implementation. pp. 62–73. ACM (2003)
14. ten Cate, B., Franceschet, M.: On the complexity of hybrid logics with binders. In: Ong, L. (ed.) Computer Science Logic. pp. 339–354 (2005)
15. Chandra, A.K., Kozen, D.C., Stockmeyer, L.J.: Alternation. Journal of the ACM **28**(1), 114–133 (1981)
16. Demri, S., Fervari, R.: The power of modal separation logics. Journal of Logic and Computation **29**(8), 1139–1184 (2019)
17. Ding, Y.: On the logics with propositional quantifiers extending s5Π. In: Advances in Modal Logic. pp. 219–235. College Publications (2018)
18. Fine, K.: Propositional quantifiers in modal logic. Theoria **36**, 336–346 (1970)
19. Fischer, M.J., Ladner, R.E.: Propositional modal logic of programs. In: ACM Symposium on Theory of Computing. p. 286–294 (1977)
20. Fischer, M.J., Rabin, M.O.: Super-exponential complexity of presburger arithmetic. In: Complexity of Computation, SIAM–AMS Proceedings. pp. 27–41 (1974)
21. Hannula, M., Kontinen, J., Virtema, J., Vollmer, H.: Complexity of propositional logics in team semantic. ACM Transactions on Computational Logic **19**(1), 2:1–2:14 (2018)
22. Kaplan, D.: S5 with quantifiable propositional variables. The Journal of Symbolic Logic **35**(2), 355 (1970)

23. Mansutti, A.: Reasoning with Separation Logics: Complexity, Expressive Power, Proof Systems. Ph.D. thesis, Université Paris-Saclay (December 2020)
24. Mansutti, A.: Notes on kAExp(pol) problems for deterministic machines (2021)
25. Meier, A., Mundhenk, M., Thomas, M., Vollmer, H.: The complexity of satisfiability for fragments of CTL and CTL*. Electronic Notes in Theoretical Computer Science **223**, 201–213 (2008)
26. Prior, A.: Past, Present and Future. Oxford Books (1967)
27. Rabin, M.: Decidability of second-order theories and automata on infinite trees. Transactions of the American Mathematical Society **41**, 1–35 (1969)
28. de Rijke, M.: A note on graded modal logic. Studia Logica **64**(2), 271–283 (2000)
29. Schmitz, S.: Complexity hierarchies beyond elementary. ACM Transactions on Computation Theory **8**(1), 3:1–3:36 (2016)
30. Schneider, T.: The complexity of hybrid logics over restricted frame classes. Ph.D. thesis, Friedrich Schiller University of Jena (2007)
31. Sistla, A.P., Clarke, E.M.: The complexity of propositional linear temporal logics. Journal of the ACM **32**(3), 733–749 (1985)
32. Yang, F., Väänänen, J.: Propositional logics of dependence. Annals of Pure and Applied Logic **167**(7), 557–589 (2016)