



HAL
open science

Modal Logics with Composition on Finite Forests

Bartosz Bednarczyk, Stéphane Demri, Raul Fervari, Alessio Mansutti

► **To cite this version:**

Bartosz Bednarczyk, Stéphane Demri, Raul Fervari, Alessio Mansutti. Modal Logics with Composition on Finite Forests. 35th Annual ACM/IEEE Symposium on Logic in Computer Science, Jul 2020, Saarbruecken, Germany. 10.1145/3373718.3394787 . hal-03005865

HAL Id: hal-03005865

<https://hal.science/hal-03005865>

Submitted on 14 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Modal Logics with Composition on Finite Forests: Expressivity and Complexity

Bartosz Bednarczyk
TU Dresden & University of Wrocław

Raul Fervari
FAMAF, Universidad Nacional de Córdoba & CONICET

Stéphane Demri
LSV, CNRS, ENS Paris-Saclay, Université Paris-Saclay

Alessio Mansutti
LSV, CNRS, ENS Paris-Saclay, Université Paris-Saclay

Abstract

We study the expressivity and complexity of two modal logics interpreted on finite forests and equipped with standard modalities to reason on submodels. The logic $ML(\mathbb{I})$ extends the modal logic K with the composition operator \mathbb{I} from ambient logic, whereas $ML(*)$ features the separating conjunction $*$ from separation logic. Both operators are second-order in nature. We show that $ML(\mathbb{I})$ is as expressive as the graded modal logic GML (on trees) whereas $ML(*)$ is strictly less expressive than GML . Moreover, we establish that the satisfiability problem is TOWER-complete for $ML(*)$, whereas it is (only) $AEXP_{POL}$ -complete for $ML(\mathbb{I})$, a result which is surprising given their relative expressivity. As by-products, we solve open problems related to sister logics such as static ambient logic and modal separation logic.

CCS Concepts: • Theory of computation \rightarrow Modal and temporal logics.

Keywords: modal logic on trees, separation logic, static ambient logic, graded modal logic, expressive power, complexity

ACM Reference Format:

Bartosz Bednarczyk, Stéphane Demri, Raul Fervari, and Alessio Mansutti. 2020. Modal Logics with Composition on Finite Forests: Expressivity and Complexity. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '20)*, July 8–11, 2020, Saarbrücken, Germany. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3373718.3394787>

1 Introduction

The ability to quantify over substructures to express properties of a model is often instrumental to perform modular and local reasoning. Two well-known examples are provided

by separation logics [28, 35, 42], dedicated to reasoning on pointer programs, and ambient (or more generally, spatial) logics [9, 12, 14, 19], dedicated to reasoning on disjoint data structures. In the realm of modal logics dedicated to knowledge representation, submodel reasoning remains a key ingredient to express the dynamics of knowledge and belief, as done in the logics of public announcement [5, 33, 37], sabotage modal logics [4], refinement modal logics [11] and relation-changing logics [1–3]. Though the models may be of different nature (e.g. memory states for separation logics, epistemic models for logics of public announcement or finite edge-labelled trees for ambient logics), all those logics feature composition operators that enable to compose or decompose substructures in a very natural way.

From a technical point of view, reasoning about submodels requires a global analysis, unlike the local approach for classical modal and temporal logics (typically based on automata techniques [47, 48]). This makes the comparison between those formalisms quite challenging and often limited to a superficial analysis on the different classes of models and composition operators. For instance, the composition operator \mathbb{I} in ambient logics decomposes a tree into two disjoint pieces such that once a node has been assigned to one submodel, all its descendants belong to the same submodel. Instead, the separating conjunction $*$ from separation logic decomposes the memory states into two disjoint memory states. Obviously, these and other well-known operators are closely related but no uniform framework investigates exhaustively their relationships in terms of expressive power.

Most of these logics can be easily encoded in monadic second-order logic MSO (or in second-order modal logics [24, 30]). Complexity-wise, if models are tree-like structures, we can then infer decidability thanks to the celebrated Rabin's theorem [40]. However, most likely, this does not produce the best decision procedures when it comes to solving simple reasoning tasks (e.g. the satisfiability problem of MSO is TOWER-complete [43]). Thus, relying on MSO as a common umbrella to capture and understand the differences between those logical formalisms is often not satisfactory.

Our motivations. Our intention in this work is to provide an in-depth comparison between the composition operator \mathbb{I} from static ambient logic [12] and the separating conjunction $*$ from separation logics [42] by identifying a common

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

LICS '20, July 8–11, 2020, Saarbrücken, Germany

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7104-9/20/07...\$15.00

<https://doi.org/10.1145/3373718.3394787>

ground in terms of logical languages and models. As a consequence, we are able to study the effects of having these operators as far as expressivity and complexity are concerned. We aim at defining two logics whose only differences rest on their use of \mathbf{I} and $*$ syntactically and semantically (by considering the adequate composition operation). To do so, we pick as our common class of models, the Kripke-style finite trees (actually finite forests, so that the class is closed under taking submodels), which provides an ubiquitous class of structures, extremely well-studied in computer science. For the underlying logical language (i.e. apart from \mathbf{I} or $*$), we advocate the use of the standard modal logic K (i.e. to have Boolean connectives and the standard modality \diamond) so that the main operations on the models amount to quantify over submodels or to move along the edges. This framework is sufficiently fundamental to give us the possibility to take advantage of model theoretical tools from modal logics [6, 8, 20]. The benefits of settling a common ground for comparison may lead to further comparisons with other logics and new results.

Our contributions. We introduce $ML(\mathbf{I})$ and $ML(*)$, two logics interpreted on Kripke-style forest models, equipped with the standard modality \diamond , and respectively with the composition operator \mathbf{I} from static ambient logic [12] and with the separating conjunction $*$ from separation logic [42]. Both logical formalisms can state non-trivial properties about submodels, but the binary modalities \mathbf{I} and $*$ operate differently: whereas $*$ is able to decompose the models at any depth, \mathbf{I} is much less permissive as the decomposition is completely determined by what happens at the level of the children of the current node. We study their expressive power and complexity, obtaining surprising results. We show that $ML(\mathbf{I})$ is as expressive as the graded modal logic GML [6, 45] whereas $ML(*)$ is strictly less expressive than GML . Interestingly, this latter development partially reuses the result for $ML(\mathbf{I})$, hence showing how our framework allows us to transpose results between the two logics. To show that GML is strictly more expressive than $ML(*)$, we define Ehrenfeucht-Fraïssé games for $ML(*)$. In terms of complexity, the satisfiability problem for $ML(\mathbf{I})$ is shown $AEXP_{POL}$ -complete¹, interestingly the same complexity as for the refinement modal logic RML [11] handling a quantifier over refinements (generalising the submodel construction). The $AEXP_{POL}$ upper bound follows from an exponential-size model property, whereas the lower bound is by reducing the satisfiability problem for an $AEXP_{POL}$ -complete team logic [27]. Much more surprisingly, although $ML(*)$ is strictly less expressive than $ML(\mathbf{I})$, its complexity is much higher (not even elementary). Precisely, we show that the satisfiability problem for $ML(*)$ is $TOWER$ -complete. The $TOWER$ upper bound is a consequence of [40], whereas hardness is shown by reduction from a $TOWER$ -complete tiling problem, adapting substantially the

¹ Problems in $AEXP_{POL}$ are decidable by an alternating Turing machine working in exponential-time and using polynomially many alternations [10].

$TOWER$ -hardness proof from [7] for second-order modal logic K on finite trees. To conclude, we get the best of our results on $ML(\mathbf{I})$ and $ML(*)$ to solve several open problems. We relate $ML(\mathbf{I})$ with an intensional fragment of static ambient logic $SAL(\mathbf{I})$ from [12] by providing polynomial-time reductions between their satisfiability problems. Consequently, we establish $AEXP_{POL}$ -completeness of $SAL(\mathbf{I})$, refuting hints from [12, Section 6]. Similarly, we show that the modal separation logic $MSL(\diamond^{-1}, *)$ from [21] is $TOWER$ -complete. *Omitted proofs can be found in the technical report on ArXiv.*

2 Preliminaries

In this section, we introduce the logics $ML(\mathbf{I})$ and $ML(*)$ interpreted on tree-like structures equipped with operators to split the structure into disjoint pieces. Due to the presence of such operators, we are required to consider a class of models that is closed under submodels, which we call Kripke-style finite forests (or finite forests for short).

Let AP be a countably infinite set of *atomic propositions*. A (Kripke-style) *finite forest* is a triple $\mathfrak{M} = (W, R, V)$ where W is a non-empty finite set of *worlds*, $V : AP \rightarrow \mathcal{P}(W)$ is a *valuation* and $R \subseteq W \times W$ is a binary relation whose inverse R^{-1} is functional and acyclic. Then, in particular the graph described by (W, R) is a finite collection of disjoint finite trees (where R encodes the child relation).

We define $R(w) \stackrel{\text{def}}{=} \{w' \in W \mid (w, w') \in R\}$. Worlds in $R(w)$ are understood as *children* of w . We inductively define R^n : $R^0 \stackrel{\text{def}}{=} \{(w, w) \mid w \in W\}$; $R^{n+1} \stackrel{\text{def}}{=} \{(w, w'') \mid \exists w' (w, w') \in R^n \text{ and } (w', w'') \in R\}$. R^+ denotes the transitive closure of R .

We define operators that chop a finite forest. It should be noted that these operators, as well as the resulting logics, can be cast under the umbrella of the logic of bunched implications BI [25, 39], with the exception that we do not explicitly require them to have an identity element (as enforced on the multiplicative operators of BI , see [25]). Let $\mathfrak{M} = (W, R, V)$ and $\mathfrak{M}_i = (W_i, R_i, V_i)$ (for $i \in \{1, 2\}$) be three finite forests.

The separation logic composition. We introduce the binary operator $+$ that performs the disjoint union at the level of parent-child relation. Formally,

$$\mathfrak{M} = \mathfrak{M}_1 + \mathfrak{M}_2 \stackrel{\text{def}}{\Leftrightarrow} R_1 \uplus R_2 = R, W_1 = W_2 = W, V_1 = V_2 = V.$$

This is the composition used in separation logic [21, 42]. The figure below depicts possible instances for \mathfrak{M} , \mathfrak{M}_1 and \mathfrak{M}_2 .




The ambient logic composition. We introduce the operator $+_w$, where $w \in W$, that constrains further $+$:

$$\mathfrak{M} = \mathfrak{M}_1 +_w \mathfrak{M}_2 \stackrel{\text{def}}{\Leftrightarrow} \mathfrak{M} = \mathfrak{M}_1 + \mathfrak{M}_2 \text{ and } R_i^+(w') = R^+(w') \text{ holds for all } i \in \{1, 2\} \text{ and } w' \in R_i(w).$$

\mathfrak{M} is a disjoint union between \mathfrak{M}_1 and \mathfrak{M}_2 except that, as soon as $w' \in R_i(w)$, the whole subtree of w' in R belongs

to \mathfrak{M}_i , like the composition in ambient logic [12]. Below, we illustrate a model decomposed with $+_w$.



We say that \mathfrak{M}_1 is a *submodel* of \mathfrak{M} , written $\mathfrak{M}_1 \sqsubseteq \mathfrak{M}$ if there is \mathfrak{M}_2 such that $\mathfrak{M} = \mathfrak{M}_1 + \mathfrak{M}_2$.

Modal logics on trees. The logic $\text{ML}(\mathbf{I})$ enriches the modal logic K (a.k.a. ML) with a binary connective \mathbf{I} , called *composition operator*, that admits submodel reasoning via the operator $+_w$. Similarly, $\text{ML}(\ast)$ enriches ML with the connective \ast , called *separating conjunction* (or *star*) that admits submodel reasoning via the operator $+$. Both connectives \mathbf{I} and \ast are understood as binary modalities. As we show throughout the paper, $\text{ML}(\mathbf{I})$ and $\text{ML}(\ast)$ are strongly related to the graded modal logic GML [20]. For conciseness, let us define all these logics by considering formulae that contain all of their ingredients. These formulae are built from

$$\varphi := \top \mid p \mid \varphi \wedge \varphi \mid \neg\varphi \mid \diamond\varphi \mid \diamond_{\geq k}\varphi \mid \varphi \ast \varphi \mid \varphi \mathbf{I} \varphi,$$

where $p \in \text{AP}$ and $k \in \mathbb{N}$ (encoded in binary). A *pointed forest* (\mathfrak{M}, w) is a finite forest $\mathfrak{M} = (W, R, V)$ together with a world $w \in W$. The satisfaction relation \models is defined as follows (standard clauses for \wedge , \neg and \top are omitted):

$$\begin{aligned} \mathfrak{M}, w \models p &\iff w \in V(p); \\ \mathfrak{M}, w \models \diamond\varphi &\iff \text{there is } w' \in R(w) \text{ s.t. } \mathfrak{M}, w' \models \varphi; \\ \mathfrak{M}, w \models \diamond_{\geq k}\varphi &\iff |\{w' \in R(w) \mid \mathfrak{M}, w' \models \varphi\}| \geq k; \\ \mathfrak{M}, w \models \varphi_1 \ast \varphi_2 &\iff \text{there are } \mathfrak{M}_1, \mathfrak{M}_2 \text{ s.t. } \mathfrak{M} = \mathfrak{M}_1 + \mathfrak{M}_2, \\ &\quad \mathfrak{M}_1, w \models \varphi_1 \text{ and } \mathfrak{M}_2, w \models \varphi_2; \\ \mathfrak{M}, w \models \varphi_1 \mathbf{I} \varphi_2 &\iff \text{there are } \mathfrak{M}_1, \mathfrak{M}_2 \text{ s.t. } \mathfrak{M} = \mathfrak{M}_1 +_w \mathfrak{M}_2, \\ &\quad \mathfrak{M}_1, w \models \varphi_1 \text{ and } \mathfrak{M}_2, w \models \varphi_2. \end{aligned}$$

The formulae $\varphi \Rightarrow \psi$, $\varphi \vee \psi$ and \perp are defined as usual. We use the following standard abbreviations: $\square\varphi \stackrel{\text{def}}{=} \neg\diamond\neg\varphi$, $\diamond_{\leq k}\varphi \stackrel{\text{def}}{=} \neg\diamond_{\geq k+1}\varphi$ and $\diamond_{=k}\varphi \stackrel{\text{def}}{=} \diamond_{\geq k}\varphi \wedge \diamond_{\leq k}\varphi$. We write $\text{size}(\varphi)$ to denote the *size* of φ with a tree representation of formulae and with a reasonably succinct encoding of atomic formulae. Besides, we write $\text{md}(\varphi)$ to denote the *modal degree* of φ understood as the maximal number of nested unary modalities (i.e. \diamond or $\diamond_{\geq k}$) in φ . Similarly, the *graded rank* $\text{gr}(\varphi)$ of φ is defined as $\max(\{k \mid \diamond_{\geq k}\psi \in \text{subf}(\varphi)\} \cup \{0\})$, where $\text{subf}(\varphi)$ is the set of all the subformulae of φ .

Given the formulae φ and ψ , $\varphi \equiv \psi$ denotes that φ and ψ are *logically equivalent*; i.e., for every pointed forest (\mathfrak{M}, w) , $\mathfrak{M}, w \models \varphi$ iff $\mathfrak{M}, w \models \psi$. For instance ($k \geq 1$ and $p \in \text{AP}$):

1. $\diamond\varphi \equiv \diamond_{\geq 1}\varphi$;
2. $(\square\square\perp \mathbf{I} \square\square\perp) \not\equiv (\square\square\perp \ast \square\square\perp)$;
3. $\diamond_{\geq k} p \equiv \underbrace{\diamond p \ast \dots \ast \diamond p}_{k \text{ times}}$;
4. $\diamond_{\geq k} \varphi \equiv \underbrace{\diamond \varphi \mathbf{I} \dots \mathbf{I} \diamond \varphi}_{k \text{ times}}$.

The modal logic ML is the logic restricted to formulae with the unique modality \diamond [8]. Similarly, the graded modal logic GML is restricted to the *graded modalities* $\diamond_{\geq k}$ [20]. We introduce the modal logics $\text{ML}(\mathbf{I})$ and $\text{ML}(\ast)$, which are

restricted to the suites of modalities (\diamond, \mathbf{I}) and (\diamond, \ast) , respectively. The two equivalences (3) and (4) already shed some light on $\text{ML}(\mathbf{I})$ and $\text{ML}(\ast)$: the two logics are similar when it comes to their formulae of modal degree one.

Lemma 2.1. *Let φ be a formula in $\text{ML}(\mathbf{I})$ with $\text{md}(\varphi) \leq 1$. Then, $\varphi \equiv \varphi[\mathbf{I} \leftarrow \ast]$ where $\varphi[\mathbf{I} \leftarrow \ast]$ is the formula in $\text{ML}(\ast)$ obtained from φ by replacing every occurrence of \mathbf{I} by \ast .*

However, as shown by the non-equivalence (2), it is unclear how the two logics compare when it comes to formulae of modal degree greater than one. Indeed, since $\mathfrak{M} = \mathfrak{M}_1 +_w \mathfrak{M}_2$ implies $\mathfrak{M} = \mathfrak{M}_1 + \mathfrak{M}_2$, but not vice-versa, the separating conjunction \ast is more permissive than the operator \mathbf{I} . However, further connections between the two operators can be easily established. Let us introduce the auxiliary operator \blacklozenge defined as $\blacklozenge\varphi \stackrel{\text{def}}{=} \varphi \ast \square\perp$. Formally,

$$(W, R, V), w \models \blacklozenge\varphi \iff \text{there is } R' \subseteq R \text{ s.t. } R'(w) = R(w) \text{ and } (W, R', V), w \models \varphi.$$

Similar operators are studied in [2, 4, 11]. We show that \blacklozenge and \mathbf{I} are sufficient to capture \ast (essential property for Section 5).

Lemma 2.2. *Let $\varphi, \psi \in \text{GML}$. We have $\varphi \ast \psi \equiv \blacklozenge(\varphi \mathbf{I} \psi)$.*

Unlike \mathbf{I} , when \ast splits a finite forest \mathfrak{M} into \mathfrak{M}_1 and \mathfrak{M}_2 , it may disconnect in both submodels worlds that are otherwise reachable, from the current world, in \mathfrak{M} . Applying \blacklozenge before \mathbf{I} allows us to imitate this behaviour. Indeed, even though \mathbf{I} preserves reachability in either \mathfrak{M}_1 or \mathfrak{M}_2 , \blacklozenge deletes part of \mathfrak{M} , making some world inaccessible. This way of expressing the separating conjunction allows us to reuse some methods developed for $\text{ML}(\mathbf{I})$ in order to study $\text{ML}(\ast)$.

The logic QK^t . Both $\text{ML}(\mathbf{I})$ and $\text{ML}(\ast)$ can be seen as fragments of the logic QK^t , which in turn is known to be a fragment of monadic second-order logic on trees [7]. The logic QK^t extends ML with second-order quantification and is interpreted on finite trees. Its formulae are defined according to the following grammar:

$$\varphi := p \mid \diamond\varphi \mid \varphi \wedge \varphi \mid \neg\varphi \mid \exists p \varphi.$$

Given $\mathfrak{M} = (W, R, V)$ and $w \in W$, the satisfaction relation \models of ML is extended as follows:

$$\mathfrak{M}, w \models \exists p \varphi \text{ iff } \exists W' \subseteq W \text{ s.t. } (W, R, V[p \leftarrow W']), w \models \varphi.$$

One can show logspace reductions from $\text{ML}(\mathbf{I})$ and $\text{ML}(\ast)$ to QK^t , by simply reinterpreting the operators \ast and \mathbf{I} as restrictive forms of second-order quantification, and by relativising \diamond to appropriate propositional symbols in order to capture the notion of submodel (details are omitted).

Satisfiability problem. The *satisfiability problem* for a logic \mathcal{L} , written $\text{Sat}(\mathcal{L})$, takes as input a formula φ in \mathcal{L} and checks whether there is a pointed forest (\mathfrak{M}, w) such that $\mathfrak{M}, w \models \varphi$.

Note that any \mathcal{L} among ML, GML, $\text{ML}(\mathbf{I})$ or $\text{ML}(\ast)$ has the tree model property, i.e. any satisfiable formula is also satisfied in some tree structure. The problems $\text{Sat}(\text{ML})$ and $\text{Sat}(\text{GML})$ are known to be PSPACE-complete, see e.g. [8, 29,

[44, 45], and therefore $\text{Sat}(\text{ML}(\mathbf{I}))$ and $\text{Sat}(\text{ML}(*))$ are PSPACE-hard. As an upper bound, by Rabin's theorem [40], the satisfiability problem for QK^t is decidable in TOWER, which transfers directly to $\text{Sat}(\text{ML}(\mathbf{I}))$ and $\text{Sat}(\text{ML}(*))$.

Expressive power. Given two logics \mathfrak{L}_1 and \mathfrak{L}_2 , we say that \mathfrak{L}_2 is *at least as expressive as* \mathfrak{L}_1 (written $\mathfrak{L}_1 \leq \mathfrak{L}_2$) whenever for every formula φ of \mathfrak{L}_1 , there is a formula ψ of \mathfrak{L}_2 such that $\varphi \equiv \psi$. $\mathfrak{L}_1 \approx \mathfrak{L}_2$ denotes that \mathfrak{L}_1 and \mathfrak{L}_2 are *equally expressive*, i.e. $\mathfrak{L}_1 \leq \mathfrak{L}_2$ and $\mathfrak{L}_2 \leq \mathfrak{L}_1$. Lastly, $\mathfrak{L}_1 < \mathfrak{L}_2$ denotes that \mathfrak{L}_2 is *strictly more expressive* than \mathfrak{L}_1 , i.e. $\mathfrak{L}_1 \leq \mathfrak{L}_2$ and $\mathfrak{L}_1 \not\approx \mathfrak{L}_2$. The equivalence (1) recalls us that $\text{ML} < \text{GML}$ [20]. From the equivalence (4), we get $\text{GML} \leq \text{ML}(\mathbf{I})$.

3 $\text{ML}(\mathbf{I})$: Expressiveness and Complexity

In this section, we study the expressive power of $\text{ML}(\mathbf{I})$ and the complexity of $\text{Sat}(\text{ML}(\mathbf{I}))$. We show constructively that $\text{ML}(\mathbf{I}) \leq \text{GML}$, hence proving $\text{ML}(\mathbf{I}) \approx \text{GML}$. Next, we show that $\text{Sat}(\text{ML}(\mathbf{I}))$ is AExp_{POL} -complete. The upper bound is achieved by proving an exponential-size model property. The lower bound is by reduction from the satisfiability problem for propositional team logic [27, Thm. 4.9].

3.1 $\text{ML}(\mathbf{I})$ is not more expressive than GML

Establishing $\text{ML}(\mathbf{I}) \leq \text{GML}$ amounts to show that given φ_1, φ_2 in GML, one can construct ψ in GML such that $\varphi_1 \mathbf{I} \varphi_2 \equiv \psi$. For instance, a simple case analysis yields the equivalence $(p \vee \diamond_{\geq 3} r) \mathbf{I} (q \vee \diamond_{\leq 5} q) \equiv (p \vee \diamond_{\geq 3} r)$. With this property, the general algorithm consists in iteratively replacing innermost subformulae of the form $\varphi_1 \mathbf{I} \varphi_2$ by a counterpart in GML, allowing us to eliminate all the occurrences of \mathbf{I} and obtain an equivalent formula in GML. The base case involves subformulae φ_1 and φ_2 in ML (a fragment of GML).

Let us provide a few definitions. Let φ be a formula in GML. We write $\text{max}_{\text{PC}}(\varphi)$ to denote the set of subformulae ψ of φ that are maximal and modality-free, i.e.

1. ψ is modality-free: it does not contain modalities $\diamond_{\geq k}$ and one of its occurrences is not in the scope of $\diamond_{\geq k}$;
2. ψ is maximal: one of its occurrences does not belong to a larger modality-free subformula of φ .

For instance, $\text{max}_{\text{PC}}((p \vee \diamond_{\geq 3} r) \wedge (q \vee p)) = \{p, q \vee p\}$. Similarly, $\text{max}_{\text{GM}}(\varphi)$ denotes the set of subformulae ψ of φ such that ψ is of the form $\diamond_{\geq k} \psi'$ and one of its occurrences in φ is not in the scope of graded modalities $\diamond_{\geq k}$. For instance, $\text{max}_{\text{GM}}((p \vee \diamond_{\geq 3} r) \wedge (q \vee \diamond_{\geq 5} \diamond_{\geq 2} q)) = \{\diamond_{\geq 3} r, \diamond_{\geq 5} \diamond_{\geq 2} q\}$. Every formula φ in GML is a Boolean combination of formulae from $\text{max}_{\text{PC}}(\varphi) \cup \text{max}_{\text{GM}}(\varphi)$. Lastly, φ is in *good shape* if the properties (1) and (2) below hold:

1. $\text{max}_{\text{PC}}(\varphi) \subseteq \{\perp, \top\}$. Consequently, every propositional variable in φ occurs in the scope of a graded modality;
2. For all $\diamond_{\geq k} \psi, \diamond_{\geq k'} \psi'$ in $\text{max}_{\text{GM}}(\varphi)$ with $\psi \neq \psi'$, the conjunction $\psi \wedge \psi'$ is unsatisfiable.

Let φ_1 and φ_2 be GML formulae. First, we show that when $\varphi_1 \wedge \varphi_2$ is in good shape, there is a GML formula ψ such that $\varphi_1 \mathbf{I} \varphi_2 \equiv \psi$. To do so, we take a slight detour through Presburger arithmetic (PA), see e.g. [26, 38]. Given two formulae φ_1, φ_2 in GML, we will characterise the formula $\varphi_1 \mathbf{I} \varphi_2$ by using arithmetical constraints for the number of successors. Then, we will take advantage of basic properties of PA in order to eliminate quantifiers, and obtain a GML formula. Below, the variables x, y, z, \dots , possibly decorated and occurring in formulae, are from PA and therefore they are interpreted by natural numbers.

Let φ be in GML s.t. $\text{max}_{\text{PC}}(\varphi) \subseteq \{\top, \perp\}$ and $\{\psi_1, \dots, \psi_n\}$ contains the set $\{\psi \mid \diamond_{\geq k} \psi \in \text{max}_{\text{GM}}(\varphi)\}$. We define formulae in PA that state constraints about the number of children satisfying a formula ψ_j . The variable x_j is intended to be interpreted as the number of children satisfying ψ_j . We write $\varphi^{\text{PA}}(x_1, \dots, x_n)$ to denote the arithmetical formula obtained from φ by replacing with $x_j \geq k$ every occurrence of $\diamond_{\geq k} \psi_j$ that it is not in the scope of a graded modality. For instance, assuming that $\varphi = \diamond_{\geq 5} (p \wedge q) \vee \neg \diamond_{\geq 4} \neg p$, the expression $\varphi^{\text{PA}}(x_1, x_2)$ denotes the formula $x_1 \geq 5 \vee \neg(x_2 \geq 4)$.

Let φ_1, φ_2 be GML formulae such that $\varphi_1 \wedge \varphi_2$ is in good shape and $\{\psi_1, \dots, \psi_n\} = \{\psi \mid \diamond_{\geq k} \psi \in \text{max}_{\text{GM}}(\varphi_1 \wedge \varphi_2)\}$. We consider the formula $[\varphi_1, \varphi_2]^{\text{PA}}$ in PA defined below:

$$[\varphi_1, \varphi_2]^{\text{PA}} \stackrel{\text{def}}{=} \exists y_1^1, y_1^2, \dots, y_n^1, y_n^2 (\bigwedge_{j=1}^n x_j = y_j^1 + y_j^2) \wedge \varphi_1^{\text{PA}}(y_1^1, \dots, y_n^1) \wedge \varphi_2^{\text{PA}}(y_1^2, \dots, y_n^2).$$

The formula $[\varphi_1, \varphi_2]^{\text{PA}}$ states that there is a way to divide the children in two distinct sets and each set allows to satisfy φ_1^{PA} or φ_2^{PA} , respectively. As PA admits quantifier elimination [16, 38, 41], there is a quantifier-free formula χ equivalent to $[\varphi_1, \varphi_2]^{\text{PA}}$ and its free variables are among x_1, \dots, x_n . A priori, the atomic formulae of χ may not be of the simple form $x_j \geq k$ (e.g. 'modulo constraints' or constraints of the form $\sum a_i x_j \geq k$ may be involved). However, if the atomic formulae of χ are restricted to expressions of the form $x_j \geq k$, then we write χ^{GML} to denote the GML formula obtained from χ by replacing every occurrence of $x_j \geq k$ by $\diamond_{\geq k} \psi_j$.

Lemma 3.1. *Let φ_1, φ_2 be in GML such that $\varphi_1 \wedge \varphi_2$ is in good shape. $[\varphi_1, \varphi_2]^{\text{PA}}$ is equivalent to a quantifier-free PA formula χ whose atomic formulae are only of the form $x_j \geq k$. Moreover, $\varphi_1 \mathbf{I} \varphi_2 \equiv \chi^{\text{GML}}$ and $\text{gr}(\chi^{\text{GML}}) \leq \text{gr}(\varphi_1) + \text{gr}(\varphi_2)$.*

The bound on $\text{gr}(\chi^{\text{GML}})$ stated in this key lemma is essential to obtain an exponential bound on the smallest model satisfying a formula in $\text{ML}(\mathbf{I})$ (see Section 3.2). Thanks to Lemma 3.1, we can show that GML is closed under the operator \mathbf{I} by reducing the occurrences of this operator to formulae in good shape. In particular, we show that given two arbitrary formulae φ_1 and φ_2 in GML, $\varphi_1 \mathbf{I} \varphi_2$ is equivalent to a disjunction of formulae of the form $(\psi_1 \mathbf{I} \psi_2) \wedge \chi$, where χ is a Boolean combination of atomic propositions and $\psi_1 \wedge \psi_2$ is in good shape (hence $\psi_1 \mathbf{I} \psi_2$ is equivalent to a formula in GML

by Lemma 3.1). This is shown syntactically: atomic propositions are dealt with by propositional reasoning, whereas to produce ψ_1 and ψ_2 we use axioms from GML [6] and rely on the following equivalences:

- (**guess**) $\diamond_{\geq k} \varphi \equiv \diamond_{\geq k} ((\varphi \wedge \psi) \vee (\varphi \wedge \neg\psi))$;
 ($\diamond_{\geq k}$ **dist**) if $\varphi \wedge \psi$ unsat., $\diamond_{\geq k} (\varphi \vee \psi) \equiv \bigvee_{k=k_1+k_2} (\diamond_{\geq k_1} \varphi \wedge \diamond_{\geq k_2} \psi)$;
 (**dist**) $(\varphi \vee \psi) \upharpoonright \chi \equiv (\varphi \upharpoonright \chi) \vee (\psi \upharpoonright \chi)$.

Notice that the conjunction of $\varphi \wedge \psi$ and $\varphi \wedge \neg\psi$ from (**guess**) is trivially unsatisfiable, allowing us to use ($\diamond_{\geq k}$ **dist**). As GML is shown to be closed under the operator \upharpoonright , we conclude.

Theorem 3.2. $\text{ML}(\upharpoonright) \leq \text{GML}$. Therefore, $\text{ML}(\upharpoonright) \approx \text{GML}$.

To prove $\text{ML}(\upharpoonright) \leq \text{GML}$, we iteratively put subformulae in good shape and apply Lemma 3.1. This is done several times, potentially causing an exponential blow-up each time a formula is transformed. To provide an optimal complexity upper bound, we need to tame this combinatorial explosion.

3.2 AEXP_{POL}-completeness

In order to show that $\text{Sat}(\text{ML}(\upharpoonright))$ is in AEXP_{POL}, the main ingredient is to show that given φ in $\text{ML}(\upharpoonright)$, we build φ' in GML such that $\varphi' \equiv \varphi$ and the models for φ' (if any) do not require a number of children per node more than exponential in $\text{size}(\varphi)$. The proof of Theorem 3.2 needs to be refined to improve the way φ' is computed. In particular, this requires a strategy for the application of the equivalences used to put a formula in good shape.

We need to introduce a few more simple notions. Let φ be a GML formula with $\max_{\text{GML}}(\varphi) = \{\diamond_{\geq k_1} \psi_1, \dots, \diamond_{\geq k_n} \psi_n\}$. We define $\text{bd}(0, \varphi) \stackrel{\text{def}}{=} k_1 + \dots + k_n$. For all $m \geq 0$, we define $\text{bd}(m+1, \varphi) \stackrel{\text{def}}{=} \max\{\text{bd}(m, \psi) \mid \diamond_{\geq k} \psi \in \max_{\text{GML}}(\varphi)\}$. Hence, $\text{bd}(m, \varphi)$ can be understood as the maximal $\text{bd}(0, \psi)$ for some subformula ψ occurring at the modal depth m within φ . We write $\max_{\text{bd}}(\varphi)$ for the value $\max\{\text{bd}(m, \varphi) \mid m \in [0, \text{md}(\varphi)]\}$. If φ is satisfiable, we can use $\max_{\text{bd}}(\varphi)$ to obtain a bound on the smallest model satisfying it, as stated in Lemma 3.3 below.

Lemma 3.3. Every satisfiable φ in GML is satisfied by a pointed forest with at most $\max_{\text{bd}}(\varphi)^{\text{md}(\varphi)+1}$ worlds.

To show that $\text{ML}(\upharpoonright)$ has the exponential-size model property, we establish that given φ in $\text{ML}(\upharpoonright)$, there is φ' in GML such that $\varphi' \equiv \varphi$, $\text{md}(\varphi') \leq \text{md}(\varphi)$ and $\max_{\text{bd}}(\varphi')$ is exponential in $\text{size}(\varphi)$. First, we consider the fragment F of $\text{ML}(\upharpoonright)$: $\varphi ::= \diamond_{\geq k} \psi \mid p \mid \varphi \upharpoonright \varphi \mid \varphi \wedge \varphi \mid \neg\varphi$, where $p \in \text{AP}$ and $\diamond_{\geq k} \psi$ is a formula in GML (abusively assumed in $\text{ML}(\upharpoonright)$ but we know $\text{GML} \leq \text{ML}(\upharpoonright)$). Given φ in $\text{ML}(\upharpoonright)$ or in F , we write $\text{cd}(\varphi)$ to denote its *composition degree*, i.e. the maximal number of imbrications of \upharpoonright in φ . We extend the notion of bd to formulae in F , so that $\text{bd}(m, \varphi) = \text{bd}(m, \varphi[\upharpoonright \leftarrow \wedge])$, where $\varphi[\upharpoonright \leftarrow \wedge]$ is the formula obtained from φ by replacing every occurrence of \upharpoonright by \wedge . Similarly, $\max_{\text{GML}}(\varphi) \stackrel{\text{def}}{=} \max_{\text{GML}}(\varphi[\upharpoonright \leftarrow \wedge])$.

Let φ be in F such that $\max_{\text{GML}}(\varphi) = \{\diamond_{\geq k_1} \chi_1, \dots, \diamond_{\geq k_n} \chi_n\}$. The key step to show the exponential-size model property essentially manipulates the formulae in $\max_{\text{GML}}(\varphi)$ in order to

produce equivalent formulae ψ_1, \dots, ψ_n , so that for all distinct i and j , $\psi_i \wedge \psi_j$ is in good shape. Moreover, by replacing in φ every $\diamond_{\geq k_i} \chi_i$ with the equivalent formula ψ_i , we only witness an exponential blow-up on $\text{bd}(0, \varphi)$, whereas for every $m > 1$, $\text{bd}(m, \varphi)$ remains polynomially bounded by the bd of the original formula. With the bound on the graded rank found in Lemma 3.1, we derive Lemma 3.4.

Lemma 3.4. Let φ be a formula of the fragment F such that $\max_{\text{GML}}(\varphi) = \{\diamond_{\geq k_1} \chi_1, \dots, \diamond_{\geq k_n} \chi_n\}$ and $\widehat{k} = \max\{k_1, \dots, k_n\}$. There is a GML formula ψ such that $\varphi \equiv \psi$ and,

1. $\text{md}(\psi) \leq \text{md}(\varphi)$;
2. $\text{bd}(0, \psi) \leq \widehat{k} \times 2^{n+\text{cd}(\varphi)}$;
3. $\text{bd}(1, \psi) \leq n \times \text{bd}(1, \varphi)$;
4. $\forall m \geq 2, \text{bd}(m, \psi) = \text{bd}(m, \varphi)$.

In the proof of Lemma 3.4, a first step essentially consists in applying multiple times (**guess**) in order to derive, for every $i \in [1, n]$, an equivalence $\diamond_{\geq k_i} \chi_i \equiv \psi'_i$ where

$$\psi'_i \stackrel{\text{def}}{=} \diamond_{\geq k_i} \bigvee_{f: [1, n] \rightarrow \{\top, \perp\}} (\chi_i \wedge [\chi_1]^{f(1)} \wedge \dots \wedge [\chi_n]^{f(n)}).$$

Here, $[\chi_j]^\top \stackrel{\text{def}}{=} \chi_j$ and $[\chi_j]^\perp \stackrel{\text{def}}{=} \neg\chi_j$. Roughly speaking, in this step, we expand χ_i by considering all the possible truth values for the formulae χ_1, \dots, χ_n (the disjuncts where χ_i is negated can be simply discharged from the disjunction, as they are unsatisfiable). Substituting every $\diamond_{\geq k_i} \chi_i$ by ψ'_i in φ leads to a formula φ' such that $\text{bd}(1, \varphi') \leq n \times \text{bd}(1, \varphi)$ (as in Lemma 3.4) and for every $m \neq 1$, $\text{bd}(m, \varphi') = \text{bd}(m, \varphi)$. Afterwards, we repeatedly apply ($\diamond_{\geq k}$ **dist**) to ψ'_i and obtain the formula ψ_i satisfying the aforementioned property, i.e. for all distinct i and j , $\psi_i \wedge \psi_j$ is in good shape. With (**dist**), this allows us to apply Lemma 3.1 until all the operators \upharpoonright are removed. Besides, replacing every ψ'_i by ψ_i in φ' leads to a formula having the same bd as the formula ψ in Lemma 3.4.

Applying adequately the transformation from Lemma 3.4 to a formula in $\text{ML}(\upharpoonright)$, i.e. by considering maximal subformulae of the fragment F , allows us to get a logically equivalent GML formula having small models.

Lemma 3.5. Every satisfiable φ in $\text{ML}(\upharpoonright)$ is satisfied by a pointed forest of size at most exponential in $\text{size}(\varphi)$.

The proof of Lemma 3.5 (relying on Lemma 3.4) consists in showing that for all φ in $\text{ML}(\upharpoonright)$, there is φ' in GML such that $\varphi' \equiv \varphi$ and $\max_{\text{bd}}(\varphi')$ is exponential in $\text{size}(\varphi)$, which is sufficient by Lemma 3.3 to get the exponential-size model property, whence the upper bound AEXP_{POL}.

Theorem 3.6. $\text{Sat}(\text{ML}(\upharpoonright))$ is in AEXP_{POL}.

The (standard) proof consists in observing that to check the satisfiability status of φ in $\text{ML}(\upharpoonright)$, first guess a pointed forest of exponential-size (thanks to Lemma 3.5) and check whether it satisfies φ . This can be done in exponential-time using an alternating Turing machine with a linear amount of alternations (between universal states and existential states) by viewing $\text{ML}(\upharpoonright)$ as a fragment of MSO.

It remains to establish AEXP_{POL}-hardness. We provide a logspace reduction from the satisfiability problem for the team logic $\text{PL}[\sim]$ shown AEXP_{POL}-complete in [27, Thm. 4.9].

PL[\sim] formulae are defined by the following grammar:

$$\varphi := p \mid \dot{\neg}p \mid \varphi \wedge \varphi \mid \sim\varphi \mid \varphi \dot{\vee} \varphi,$$

where $p \in \text{AP}$ and the connectives $\dot{\neg}$ and $\dot{\vee}$ are dotted to avoid confusion with those of $\text{ML}(\mathbf{I})$. PL[\sim] is interpreted on sets of (Boolean) propositional valuations over a finite subset of AP. They are called *teams* and are denoted by $\mathfrak{T}, \mathfrak{T}_1, \dots$. A model for φ is a team \mathfrak{T} over a set of propositional variables including those occurring in φ and such that $\mathfrak{T} \models \varphi$ with:

$$\begin{aligned} \mathfrak{T} \models p &\iff \text{for all } v \in \mathfrak{T}, \text{ we have } v(p) = \top; \\ \mathfrak{T} \models \dot{\neg}p &\iff \text{for all } v \in \mathfrak{T}, \text{ we have } v(p) = \perp; \\ \mathfrak{T} \models \varphi_1 \dot{\vee} \varphi_2 &\iff \exists \mathfrak{T}_1, \mathfrak{T}_2 \text{ s.t. } \mathfrak{T} = \mathfrak{T}_1 \cup \mathfrak{T}_2, \mathfrak{T}_1 \models \varphi_1, \mathfrak{T}_2 \models \varphi_2. \end{aligned}$$

The connectives \sim and \wedge are interpreted as the classical negation and conjunction, respectively. Notice that, in the clause for $\dot{\vee}$, the teams \mathfrak{T}_1 and \mathfrak{T}_2 are not necessarily disjoint.

Let us discuss the reduction from $\text{Sat}(\text{PL}[\sim])$ to $\text{Sat}(\text{ML}(\mathbf{I}))$. A direct encoding of a team \mathfrak{T} into a pointed forest (\mathfrak{M}, w) consists in having a correspondence between the propositional valuations in \mathfrak{T} and the propositional valuations of the children of w . This would work fine if there were no mismatch between the semantics for \mathbf{I} (disjointness of the children) and the one for $\dot{\vee}$ (disjointness not required). To handle this, when checking the satisfaction of φ in PL[\sim] with n occurrences of $\dot{\vee}$, we impose that if a propositional valuation occurs among the children of w , then it occurs in at least $n + 1$ children. This property must be maintained after applying $\dot{\vee}$ several times, always with respect to the number of occurrences of $\dot{\vee}$ in the subformula of φ that is evaluated. Non-disjointness of the teams is encoded by carefully separating the children of w having identical valuations.

We now formalise the reduction. Assume that we wish to translate φ from PL[\sim], written with atomic propositions in $P = \{p_1, \dots, p_m\}$ and containing at most n occurrences of the operator $\dot{\vee}$. We introduce a set $Q = \{q_1, \dots, q_{n+1}\}$ of auxiliary propositions disjoint from P . The elements of Q are used to distinguish different copies of the same propositional valuation of a team. Thus, with respect to a pointed forest (\mathfrak{M}, w) , we require each child of w to satisfy exactly one element of Q . This can be done with the formula

$$\text{uni}(Q) \stackrel{\text{def}}{=} \Box(\bigwedge_{i \neq i' \in [1, n+1]} \neg(q_i \wedge q_{i'}) \wedge \bigvee_{i \in [1, n+1]} q_i).$$

We require that if a child of w satisfies a propositional valuation over (elements in) P , then there are $n + 1$ children satisfying that valuation over P , each of them satisfying a distinct symbol in Q . So, every valuation over P occurring in some child of w , occurs at least in $n + 1$ children of w . However, as the translation of the operator $\dot{\vee}$ modifies the set of copies of a propositional valuation, this property must be extended to arbitrary subsets of Q . Given $\emptyset \neq X \subseteq [1, n + 1]$, we require that for all $k \neq k' \in X$, if a children of w satisfies q_k , then there is a child satisfying $q_{k'}$ with the same valuation over P . The formula $\text{cp}(X)$ below does the job:

$$\bigwedge_{k \neq k' \in X} \neg(\Box q_k \llbracket \Diamond_{=1} q_k \wedge \neg(\top \llbracket \Diamond_{=1} q_k \wedge \Diamond_{=1} q_{k'} \wedge \bigwedge_{j \in [1, m]} \Diamond p_j \Rightarrow \Box p_j \rrbracket)).$$

Lastly, before defining the translation map τ , we describe how different copies of the same propositional valuation are split. We introduce two auxiliary choice functions c_1 and c_2 that take as arguments $X \subseteq [1, n + 1]$, and $n_1, n_2 \in \mathbb{N}$ with $|X| \geq n_1 + n_2$ such that for each $i \in \{1, 2\}$, we have $c_i(X, n_1, n_2) \subseteq X$, $|c_i(X, n_1, n_2)| \geq n_i$. Moreover $c_1(X, n_1, n_2) \dot{\cup} c_2(X, n_1, n_2) = X$. The maps c_1 and c_2 are instrumental to decide how to split X into two disjoint subsets respecting basic cardinality constraints. The translation map τ is designed as follows ($\emptyset \neq X \subseteq [1, n + 1]$):

$$\begin{aligned} \tau(p, X) &\stackrel{\text{def}}{=} \Box((\bigvee_{j \in X} q_j) \Rightarrow p); \\ \tau(\dot{\neg}p, X) &\stackrel{\text{def}}{=} \Box((\bigvee_{j \in X} q_j) \Rightarrow \neg p); \\ \tau(\varphi_1 \wedge \varphi_2, X) &\stackrel{\text{def}}{=} \tau(\varphi_1, X) \wedge \tau(\varphi_2, X); \quad \tau(\sim\varphi, X) \stackrel{\text{def}}{=} \neg\tau(\varphi, X); \\ \tau(\varphi_1 \dot{\vee} \varphi_2, X) &\stackrel{\text{def}}{=} (\tau(\varphi_1, X_1) \wedge \text{cp}(X_1)) \mathbf{I} (\tau(\varphi_2, X_2) \wedge \text{cp}(X_2)), \end{aligned}$$

where **(i)** $|X_i|$ is greater or equal to the number of occurrences of $\dot{\vee}$ in $\varphi_1 \dot{\vee} \varphi_2$ plus one; **(ii)** given n_1, n_2 such that n_1 (resp. n_2) is the number of occurrences of $\dot{\vee}$ in φ_1 (resp. φ_2) plus one, for each $i \in \{1, 2\}$ we have $c_i(X, n_1, n_2) = X_i$.

Lemma 3.7 below guarantees that starting with a linear number of children with the same propositional valuation is sufficient to encode $\dot{\vee}$ within $\text{ML}(\mathbf{I})$.

Lemma 3.7. *Let φ be in PL[\sim] with n occurrences of $\dot{\vee}$ and built upon p_1, \dots, p_m . Then, φ is satisfiable iff so is*

$$\text{uni}(q_1, \dots, q_{n+1}) \wedge \text{cp}([1, n + 1]) \wedge \tau(\varphi, [1, n + 1]).$$

The $\text{ML}(\mathbf{I})$ formula involved in Lemma 3.7 has modal depth one. By Theorem 3.6, $\text{Sat}(\text{ML}(\mathbf{I}))$ is AExp_{POL} -complete even restricted to formulae of modal depth at most one.

Corollary 3.8. *$\text{Sat}(\text{ML}(\mathbf{I}))$ is AExp_{POL} -complete.*

As we show in the next section, the complexity of $\text{ML}(\ast)$ does not collapse to modal depth one: $\text{Sat}(\text{ML}(\ast))$ restricted to formulae of modal depth k is exponentially easier than $\text{Sat}(\text{ML}(\ast))$ restricted to formulae of modal depth $k + 1$.

4 $\text{ML}(\ast)$ is TOWER-complete

We show that $\text{Sat}(\text{ML}(\ast))$ is TOWER-complete, i.e. complete for the class of all problems of time complexity bounded by a tower of exponentials whose height is an elementary function [43]. Given $k, n \geq 0$, we inductively define the tetration function t as $t(0, n) \stackrel{\text{def}}{=} n$ and $t(k + 1, n) = 2^{t(k, n)}$. Intuitively, $t(k, n)$ defines a tower of exponentials of height k . By k -NEXPTIME, we denote the class of all problems decidable with a nondeterministic Turing machine (NTM) of working time $O(t(k, p(n)))$ for some polynomial $p(\cdot)$, on each input of length n . To show TOWER-hardness, we design a uniform elementary reduction allowing us to get k -NEXPTIME-hardness for all k greater than a certain (fixed) integer. In our case, we achieve an exponential-space reduction from the k -NEXPTIME variant of the tiling problem, for all $k \geq 2$.

The tiling problem Tile_k takes as input a triple $\mathcal{T} = (\mathcal{T}, \mathcal{H}, \mathcal{V})$ where \mathcal{T} is a finite set of tile types, $\mathcal{H} \subseteq \mathcal{T} \times \mathcal{T}$

(resp. $\mathcal{V} \subseteq \mathcal{T} \times \mathcal{T}$) represents the horizontal (resp. vertical) matching relation, and an initial tile type $c \in \mathcal{T}$. A solution for the instance (\mathcal{T}, c) is a mapping $\tau : [0, t(k, n) - 1] \times [0, t(k, n) - 1] \rightarrow \mathcal{T}$ such that **(first)** $\tau(0, 0) = c$, and

(hor&vert) for all $i \in [0, t(k, n) - 1]$ and $j \in [0, t(k, n) - 2]$,
 $(\tau(j, i), \tau(j + 1, i)) \in \mathcal{H}$ and $(\tau(i, j), \tau(i, j + 1)) \in \mathcal{V}$.

The problem of checking whether an instance of Tile_k has a solution is known to be $k\text{-NEXPTime}$ -complete (see [36]).

The reduction below from Tile_k to $\text{Sat}(\text{ML}(*))$ recycles ideas from [7] to reduce Tile_k to $\text{Sat}(\text{QK}^t)$. To provide the adequate adaptation for $\text{ML}(*)$, we need to solve two major issues. First, QK^t admits second-order quantification, whereas in $\text{ML}(*)$, the second-order features are limited to the separating conjunction $*$. Second, the second-order quantification of QK^t essentially colours the nodes in Kripke-style structures without changing the frame (W, R) . By contrast, the operator $*$ modifies the accessibility relation, possibly making worlds that were reachable from the current world, unreachable in submodels. The TOWER -hardness proof for $\text{Sat}(\text{ML}(*))$ becomes then much more challenging: we would like to characterise the position on the grid encoded by a world w by exploiting properties of its descendants (as done for QK^t), but at the same time, we need to be careful and only consider submodels where w keeps encoding the same position. In a sense, our encoding is robust: when the operator $*$ is used to reason on submodels, we can enforce that no world changes the position of the grid that it encodes.


4.1 Enforcing $t(j, n)$ children.

Let $\mathfrak{M} = (W, R, V)$ be a finite forest. We consider two disjoint sets of atomic propositions $P = \{p_1, \dots, p_n, \text{val}\}$ and $\text{Aux} = \{x, y, 1, s, r\}$ (whose respective role is later defined). Elements from Aux are understood as *auxiliary* propositions. We call *ax-node* (resp. *Aux-node*) a world satisfying the proposition $\text{ax} \in \text{Aux}$ (resp. satisfying some proposition in Aux). We call *t-node* a world that satisfies the formula $t \stackrel{\text{def}}{=} \bigwedge_{\text{ax} \in \text{Aux}} \neg \text{ax}$. Every world of \mathfrak{M} is either a *t-node* or an *Aux-node*. We say that w' is a *t-child* of $w \in W$ if $w' \in R(w)$ and w' is a *t-node*. We define the concepts of *Aux-child* and *ax-child* similarly.

The key development of our reduction is given by the definition of a formula, of exponential size in $j \geq 1$ and polynomial size in $n \geq 1$, that when satisfied by (\mathfrak{M}, w) forces every *t-node* in $R^i(w)$, where $0 \leq i < j$, to have exactly $t(j-i, n)$ *t-children*, each of them encoding a different number in $[0, t(j-i, n) - 1]$. As we impose that w is a *t-node*, it must have $t(j, n)$ *t-children*. We assume n to be fixed throughout the section and denote this formula by $\text{type}(j)$. From the property above, if $\mathfrak{M}, w \models \text{type}(j)$ then for all $i \in [1, j-1]$ and all *t-nodes* $w' \in R^i(w)$ we have $\mathfrak{M}, w' \models \text{type}(j-i)$.

First, let us informally describe how numbers are encoded in the model (\mathfrak{M}, w) satisfying $\text{type}(j)$. Let $i \in [1, j]$. Given a *t-node* $w' \in R^i(w)$, $\mathbf{n}_i(w')$ denotes the number encoded by w' . We omit the subscript i when it is clear from the

context. When $i = j$, we represent $\mathbf{n}(w')$ by using the truth values of the atomic propositions p_1, \dots, p_n . The proposition p_b is responsible for the b -th bit of the number, with the least significant bit being encoded by p_1 . For example, for $n = 3$, we have $\mathfrak{M}, w' \models p_3 \wedge p_2 \wedge \neg p_1$ whenever $\mathbf{n}(w') = 6$. The formula $\text{type}(1)$ forces the parent of w' (i.e. is a t -node in $R^{j-1}(w)$) to have exactly 2^n *t-children* by requiring one *t-child* for each possible valuation upon p_1, \dots, p_n . Otherwise, for $i < j$ (and therefore $j \geq 2$), the number $\mathbf{n}_i(w')$ is represented by the binary encoding of the truth values of val on the *t-children* of w' which, since $(\mathfrak{M}, w') \models \text{type}(j-i)$, are $t(j-i, n)$ children implicitly ordered by the number they, in turn, encode. The essential property of $\text{type}(j)$ is therefore the following: the numbers encoded by the *t-children* of a *t-node* $w'' \in R^i(w)$, represent positions in the binary representation of the number $\mathbf{n}_i(w'')$. Thanks to this property, the formula $\text{type}(j)$ forces w to have exactly $t(j, n)$ children, all encoding different numbers in $[0, t(j, n) - 1]$. This is roughly represented in the picture below, where “1” stands for val being true whereas “0” stands for val being false.



To characterise these trees in $\text{ML}(*)$, we simulate second-order quantification by using *Aux-nodes*. Informally, we require a pointed forest (\mathfrak{M}, w) satisfying $\text{type}(j)$ to be such that **(i)** every *t-node* $w' \in R(w)$ has exactly one *x-child*, and one (different) *y-child*. These nodes do not satisfy any other auxiliary proposition; **(ii)** for every $i \geq 2$, every *t-node* $w' \in R^i(w)$ has exactly five *Aux-children*, one for each $\text{ax} \in \text{Aux}$. We can simulate second-order existential quantification on *t-nodes* with respect to the symbol $\text{ax} \in \text{Aux}$ by using the operator $*$ in order to remove edges leading to *ax-nodes*. Then, we evaluate whether a property holds on the resulting model where a *t-node* “satisfies” $\text{ax} \in \text{Aux}$ if it has a child satisfying ax . To better emphasise the need to move along *t-nodes*, given a formula φ , we write $\langle t \rangle \varphi$ for the formula $\diamond(t \wedge \varphi)$. Dually, $[t] \varphi \stackrel{\text{def}}{=} \square(t \Rightarrow \varphi)$. $\langle t \rangle^i$ and $[t]^i$ are also defined, as expected.

Let us start to formalise this encoding. Let $j \geq 1$. First, we restrict ourselves to models where every *t-node* reachable in at most j steps does not have two *Aux-children* satisfying the same proposition. Moreover, these *Aux-nodes* have no children and only satisfy exactly one $\text{ax} \in \text{Aux}$. We express this condition with the formula $\text{init}(j)$ below:

$$\boxplus_{\text{ax} \in \text{Aux}}^j \bigwedge \left((t \Rightarrow \neg(\diamond \text{ax} * \diamond \text{ax})) \wedge \square (\text{ax} \Rightarrow \square \perp \wedge \bigwedge_{\text{bx} \in \text{Aux} \setminus \{\text{ax}\}} \neg \text{bx}) \right),$$

where $\boxplus^0 \varphi \stackrel{\text{def}}{=} \varphi$ and $\boxplus^{m+1} \varphi \stackrel{\text{def}}{=} \varphi \wedge \square \boxplus^m (\varphi)$. Notice that if $\mathfrak{M}, w \models \text{init}(j)$ and $\mathfrak{M}' \sqsubseteq \mathfrak{M}$, then $\mathfrak{M}', w \models \text{init}(j)$.

Among the models $((W, R, V), w)$ satisfying $\text{init}(j)$, we define the ones satisfying $\text{type}(j)$ described below (see similar conditions in [7, Section IV]):

- (**sub**_{*j*}) every *t*-node in $R(w)$ satisfies $\text{type}(j-1)$;
- (**zero**_{*j*}) there is a *t*-node $\tilde{w} \in R(w)$ such that $\mathbf{n}(\tilde{w}) = 0$;
- (**uniq**_{*j*}) distinct *t*-nodes in $R(w)$ encode different numbers;
- (**compl**_{*j*}) for every *t*-node $w_1 \in R(w)$, if $\mathbf{n}(w_1) < t(j, n) - 1$ then $\mathbf{n}(w_2) = \mathbf{n}(w_1) + 1$ for some *t*-node $w_2 \in R(w)$;
- (**aux**) *w* is a *t*-node, every *t*-node in $R(w)$ has one *x*-child and one *y*-child, and every *t*-node in $R^2(w)$ has three children satisfying *l*, *r* and *s*, respectively.

We define $\text{type}(0) \stackrel{\text{def}}{=} \top$, and for $j \geq 1$, $\text{type}(j)$ is defined as

$$\text{type}(j) \stackrel{\text{def}}{=} \text{sub}(j) \wedge \text{zero}(j) \wedge \text{uniq}(j) \wedge \text{compl}(j) \wedge \text{aux},$$

where each conjunct expresses its homonymous property. The formulae for $\text{sub}(j)$, aux and $\text{zero}(j)$ can be defined as

$$\text{sub}(j) \stackrel{\text{def}}{=} [t] \text{type}(j-1);$$

$$\text{aux} \stackrel{\text{def}}{=} t \wedge [t](\diamond x * \diamond y) \wedge [t]^2(\diamond l * \diamond s * \diamond r);$$

$$\text{zero}(1) \stackrel{\text{def}}{=} \langle t \rangle \wedge \bigwedge_{b \in [1, n]} \neg p b;$$

$$\text{zero}(j+1) \stackrel{\text{def}}{=} \langle t \rangle [t] \neg \text{val}.$$

The challenge is therefore how to express $\text{uniq}(j)$ and $\text{compl}(j)$, to guarantee that the numbers of children of *w* span all over $[0, t(j, n) - 1]$. The structural properties expressed by $\text{type}(j)$ lead to strong constraints, which permits to control the effects of $*$ when submodels are constructed. This is a key point in designing $\text{type}(j)$ as it helps us to control which edges are lost when considering a submodel.

Nominals, forks and number comparisons. In order to define $\text{uniq}(j)$ and $\text{compl}(j)$ (completing the definition of $\text{type}(j)$), we introduce auxiliary formulae, characterising classes of models that emerge naturally when trying to capture the semantics of (**uniq**_{*j*}) and (**compl**_{*j*}).

Let us consider a finite forest $\mathfrak{M} = (W, R, V)$ and $w \in W$. A first ingredient is given by the concept of *local nominals*, borrowed from [7]. We say that $ax \in \text{Aux}$ is a (local) nominal for the depth $i \geq 1$ if there is exactly one *t*-node $w' \in R^i(w)$ having an *ax*-child. In this case, w' is said to be the world that corresponds to the local nominal *ax*. The following formula states that *ax* is a local nominal for the depth *i*:

$$\text{nom}_i(ax) \stackrel{\text{def}}{=} \langle t \rangle^i \diamond ax \wedge \bigwedge_{k \in [0, i-1]} [t]^k \neg (\langle t \rangle^{i-k} \diamond ax * \langle t \rangle^{i-k} \diamond ax).$$


We define the formula $\text{@}_{ax}^i \varphi \stackrel{\text{def}}{=} \langle t \rangle^i (\diamond ax \wedge \varphi)$ which, under the hypothesis that *ax* is a local nominal for the depth *i*, states that φ holds on the *t*-node that corresponds to *ax*. Moreover, we define $\text{nom}_i(ax \neq bx) \stackrel{\text{def}}{=} \text{nom}_i(ax) \wedge \text{nom}_i(bx) \wedge \neg \text{@}_{ax}^i \diamond bx$, which states that *ax* and *bx* are two nominals for the depth *i* with respect to two distinct *t*-nodes.

As a second ingredient, we introduce the notion of *fork* that is a specific type of models naturally emerging when trying to compare the numbers $\mathbf{n}(w_1)$ and $\mathbf{n}(w_2)$ of two worlds

$w_1, w_2 \in R^i(w)$ (e.g. when checking whether $\mathbf{n}(w_1) = \mathbf{n}(w_2)$ or $\mathbf{n}(w_2) = \mathbf{n}(w_1) + 1$ holds). Given $j \geq i \geq 1$ we introduce the formula $\text{fork}_j^i(ax, bx)$ that is satisfied by (\mathfrak{M}, w) iff:

- *ax* and *bx* are nominals for the depth *i*.
- *w* has exactly two *t*-children, say w_U and w_D .
- For every $k \in [1, i-1]$, both $R^k(w_U)$ and $R^k(w_D)$ contain exactly one *t*-child.
- The only *t*-node in $R^{i-1}(w_U)$, say w_{ax} , corresponds to the nominal *ax*. The only *t*-node in $R^{i-1}(w_D)$, say w_{bx} , corresponds to the nominal *bx*.
- If $i < j$, then (\mathfrak{M}, w_{ax}) and (\mathfrak{M}, w_{bx}) satisfy $\text{type}_{1sr}(j-i) \stackrel{\text{def}}{=} \text{type}(j-i) \wedge [t](\diamond l \wedge \diamond s \wedge \diamond r)$.

It should be noted that, whenever (\mathfrak{M}, w) satisfies the formula $\text{fork}_j^i(ax, bx)$, we witness two paths of length *i*, both starting at *w* and leading to w_{ax} and w_{bx} , respectively. Worlds in this path may have *Aux*-children. Below, we schematise a model satisfying $\text{fork}_j^i(ax, bx)$:



Since the definition of $\text{fork}_j^i(ax, bx)$ is recursive on *i* and *j* (due to $\text{type}(j-i)$), we postpone its formal definition to the next two sections where we treat the base cases for $i = j$ and the inductive case for $j > i$ separately.

The last auxiliary formulae are $[ax < bx]_j^i$ and $[bx = ax+1]_j$. Under the hypothesis that (\mathfrak{M}, w) satisfies $\text{fork}_j^i(ax, bx)$, the formula $[ax < bx]_j^i$ is satisfied whenever the two (distinct) worlds $w_{ax}, w_{bx} \in R^i(w)$ corresponding to the nominals *ax* and *bx* are such that $\mathbf{n}(w_{ax}) < \mathbf{n}(w_{bx})$. Similarly, under the hypothesis that (\mathfrak{M}, w) satisfies $\text{fork}_j^i(ax, bx)$, the formula $[bx = ax+1]_j$ is satisfied whenever $\mathbf{n}(w_{bx}) = \mathbf{n}(w_{ax}) + 1$ holds. Both formulae are recursively defined, with base cases for $i = j$ and $j = 1$, respectively.

For the base case, we define the formulae $\text{fork}_j^j(ax, bx)$ and $[ax < bx]_j^j$ (for arbitrary *j*), as well as $[bx = ax+1]_1$. From these formulae, we are then able to define $\text{uniq}(1)$ and $\text{compl}(1)$, which completes the characterisation of $\text{type}(1)$ and $\text{type}_{1sr}(1)$. Afterwards, we consider the case $1 \leq i < j$ and $j \geq 2$, and define $\text{fork}_j^i(ax, bx)$, $[ax < bx]_j^i$, $[bx = ax+1]_j$, as well as $\text{uniq}(j)$ and $\text{compl}(j)$, by only relying on formulae that are already defined (by inductive reasoning).

Base cases: $i = j$ or $j = 1$. In what follows, we consider a finite forest $\mathfrak{M} = (W, R, V)$ and a world *w*. Following its informal description, we have

$$\text{fork}_j^j(ax, bx) \stackrel{\text{def}}{=} \diamond_{=2} t \wedge [t] \boxplus^{j-2}(t \Rightarrow \diamond_{=1} t) \wedge \text{nom}_j(ax \neq bx),$$

where $\boxplus^j \varphi \stackrel{\text{def}}{=} \top$ for $j < 0$. As previously explained, in the base case, the number $\mathbf{n}(w')$ encoded by a *t*-node $w' \in R^j(w)$

is represented by the truth values of p_1, \dots, p_n . Then, the formula $[ax < bx]_j^j$ is defined as

$$[ax < bx]_j^j \stackrel{\text{def}}{=} \bigvee_{u \in [1, n]} (@_{ax}^j \neg p_u \wedge @_{bx}^j p_u \wedge \bigwedge_{v \in [u+1, n]} (@_{ax}^j p_v \Leftrightarrow @_{bx}^j p_v)).$$

The satisfaction of $(\mathfrak{M}, w) \models \text{fork}_j^j(ax, bx)$ enforces that the distinct t -nodes $w_{ax}, w_{bx} \in R^j(w)$ corresponding to ax and bx satisfy $\mathbf{n}(w_{ax}) < \mathbf{n}(w_{bx})$, which can be shown by using standard properties about bit vectors.

The formula $[bx = ax+1]_1$ is similarly defined:

$$\bigvee_{u \in [1, n]} (@_{ax}^1 (\neg p_u \wedge \bigwedge_{v \in [1, u-1]} p_v) \wedge @_{bx}^1 (p_u \wedge \bigwedge_{v \in [1, u-1]} \neg p_v) \wedge \bigwedge_{v \in [u+1, n]} (@_{ax}^1 p_v \Leftrightarrow @_{bx}^1 p_v)).$$

Assuming $(\mathfrak{M}, w) \models \text{fork}_1^1(ax, bx)$, this formula states that the two distinct t -nodes $w_{ax}, w_{bx} \in R(w)$ corresponding to ax and bx are such that $\mathbf{n}(w_{bx}) = \mathbf{n}(w_{ax}) + 1$. Again, correctness is guaranteed by standard analysis on bit vectors.

To define $\text{uniq}(1)$, we recall that a model satisfying $\text{type}(1)$ satisfies the formula aux and hence every t -node in $R(w)$ has two auxiliary children, one x -node and one y -node. The idea is to use these two Aux-children and rely on $*$ to state that it is not possible to find a submodel of \mathfrak{M} such that w has only two distinct children w_x and w_y corresponding to the nominals x and y , respectively, and such that $\mathbf{n}(w_x) = \mathbf{n}(w_y)$. In a sense, the operator $*$ simulates a second-order quantification on x and y . Let $[x = y]_1^1 \stackrel{\text{def}}{=} \neg([x < y]_1^1 \vee [y < x]_1^1)$. We define $\text{uniq}(1) \stackrel{\text{def}}{=} \neg(\top * (\text{fork}_1^1(x, y) \wedge [x = y]_1^1))$.

To capture $\text{compl}(1)$ we state that it is not possible to find a submodel of \mathfrak{M} that loses x -nodes from $R^2(w)$, keeps all y -nodes, and is such that (i) x is a local nominal for the depth 1, corresponding to a world w_x encoding $\mathbf{n}(w_x) < 2^n - 1$; (ii) there is no submodel where w has two t -children, w_x and a second world w_y , such that w_y corresponds to the nominal y and $\mathbf{n}(w_y) = \mathbf{n}(w_x) + 1$. Thus, $\text{compl}(1)$ is defined as:

$$\neg(\Box \perp * ([t] \diamond y \wedge @_x^1 \neg 1_1 \wedge \neg(\top * (\text{fork}_1^1(x, y) \wedge [y = x+1]_1^1)))).$$

The subscript “1” in the formula 1_1 refers to the fact that we are treating the base case of $\text{compl}(j)$ with $j = 1$. We have $1_1 \stackrel{\text{def}}{=} \bigwedge_{i \in [1, n]} p_i$, reflecting the encoding of $2^n - 1$.

This concludes the definition of $\text{type}(1)$ (and $\text{type}_{1sr}(1)$), which is established correct with respect to its specification.

Lemma 4.1. *Let $\mathfrak{M}, w \models \text{init}(1)$. We have $\mathfrak{M}, w \models \text{type}(1)$ iff (\mathfrak{M}, w) satisfies (sub_1) , (zero_1) , (uniq_1) , (compl_1) and (aux) .*

Inductive case: $1 \leq i < j$. As an implicit inductive hypothesis used to prove that the formulae are well-defined, we assume that $[bx = ax+1]_{j'}$ and $\text{type}(j')$ are already defined for every $j' < j$, whereas $\text{fork}_{j'}^{j'}(ax, bx)$, and $[ax < bx]_{j'}^{j'}$ are already defined for every $1 \leq i' \leq j'$ such that $j' - i' < j - i$. Therefore, we define:

$$\text{fork}_j^j(ax, bx) \stackrel{\text{def}}{=} \text{fork}_i^i(ax, bx) \wedge [t]^i \text{type}_{1sr}(j - i).$$

It is easy to see that this formula is well-defined: $\text{fork}_i^i(ax, bx)$ is from the base case, whereas $\text{type}_{1sr}(j - i)$ is defined by inductive hypothesis, since we have $j - i < j$.

Consider now $[ax < bx]_j^j$. Assuming $\mathfrak{M}, w \models \text{fork}_j^j(ax, bx)$, we wish to express $\mathbf{n}(w_{ax}) < \mathbf{n}(w_{bx})$ for the two distinct worlds $w_{ax}, w_{bx} \in R^j(w)$ corresponding to the nominals ax and bx , respectively. As $i < j$, $\mathbf{n}(w_{ax})$ (resp. $\mathbf{n}(w_{bx})$) is encoded using the truth value of val on the t -children of w_{ax} (resp. w_{bx}). To rely on arithmetical properties of binary numbers used to define $[ax < bx]_j^j$, we need to find two partitions $P_{ax} = \{L_{ax}, S_{ax}, R_{ax}\}$ and $P_{bx} = \{L_{bx}, S_{bx}, R_{bx}\}$, one for the t -children of w_{ax} and another one for those of w_{bx} s.t.:

(LSR): Given $b \in \{ax, bx\}$, P_b splits the t -children as follows:

- there is a t -child s_b of w_b such that $S_b = \{s_b\}$;
- $\mathbf{n}(r) < \mathbf{n}(s_b) < \mathbf{n}(l)$, for every $r \in R_b$ and $l \in L_b$.

(LESS): P_{ax} and P_{bx} have constraints to satisfy $<$:


- $\mathbf{n}(s_{ax}) = \mathbf{n}(s_{bx})$, $\mathfrak{M}, s_{ax} \models \neg \text{val}$ and $\mathfrak{M}, s_{bx} \models \text{val}$;
- for every $l_{ax} \in L_{ax}$ and $l_{bx} \in L_{bx}$, if $\mathbf{n}(l_{ax}) = \mathbf{n}(l_{bx})$ then $\mathfrak{M}, l_{ax} \models \text{val}$ iff $\mathfrak{M}, l_{bx} \models \text{val}$.

It is important to notice that these conditions essentially revolve around the numbers encoded by t -children, which will be compared using the already defined (by inductive reasoning) formulae $[ax < bx]_{j'}^{j'}$, where $j' - i' < j - i$. Since the semantics of $[ax < bx]_j^j$ is given under the hypothesis that $\mathfrak{M}, w \models \text{fork}_j^j(ax, bx)$, we can assume that every child of w_{ax} and w_{bx} has all the possible Aux-children. Then, we rely on the auxiliary propositions in $\{1, s, r\}$ in order to mimic the reasoning done in **(LSR)** and **(LESS)**.

We start by considering the constraints involved in **(LSR)** and express them with the formula $1sr(j)$, which is satisfied by a pointed forest $(\mathfrak{M} = (W, R, V), w)$ whenever:

- (\mathfrak{M}, w) satisfies $\text{type}(j)$.
- Every t -child of w has exactly one $\{1, s, r\}$ -child, and only one of these t -children (say w') has an s -child.
- Every t -child of w that has an l -child (resp. r -child) encodes a number greater (resp. smaller) than $\mathbf{n}(w')$.

Despite this formula being defined in terms of $\text{type}(j)$, we only rely on $1sr(j - i)$ (which is defined by inductive reasoning) in order to define $[ax < bx]_j^j$. The picture below schematises a model satisfying $1sr(j)$.



The definition of $1sr(j)$ follows closely its specification:

$$1sr(j) \stackrel{\text{def}}{=} \text{type}(j) \wedge \text{nom}_1(s) \wedge \neg(\top * (\text{fork}_j^j(s, 1) \wedge \neg[s < 1]_j^j)) \wedge \neg(\top * (\text{fork}_j^j(s, r) \wedge \neg[r < s]_j^j)) \wedge [t] \diamond_{=1} (1 \vee s \vee r).$$

We define the formula $[ax < bx]_j^i$ as follows:

$$\top * \left(\text{nom}_i(ax \neq bx) \wedge [t]^i \text{lsr}(j-i) \wedge S_j^i(ax, bx) \wedge L_j^i(ax, bx) \right),$$

where $S_j^i(ax, bx)$ and $L_j^i(ax, bx)$ check the first and second condition in (LESS), respectively. In particular, by defining $[ax = bx]_j^i \stackrel{\text{def}}{=} \neg([ax < bx]_j^i \vee [bx < ax]_j^i)$, we have

$$S_j^i(ax, bx) \stackrel{\text{def}}{=} \top * \left(\text{fork}_j^{i+1}(x, y) \wedge @_{ax}^i \langle t \rangle (\diamond s \wedge \diamond x) \wedge @_{bx}^i \langle t \rangle (\diamond s \wedge \diamond y) \wedge [x = y]_j^{i+1} \wedge @_x^{i+1} \neg \text{val} \wedge @_y^{i+1} \text{val} \right)$$

$$L_j^i(ax, bx) \stackrel{\text{def}}{=} \neg \left(\top * \left(\text{fork}_j^{i+1}(x, y) \wedge @_{ax}^i \langle t \rangle (\diamond l \wedge \diamond x) \wedge @_{bx}^i \langle t \rangle (\diamond l \wedge \diamond y) \wedge [x = y]_j^{i+1} \wedge \neg (@_x^{i+1} \text{val} \Leftrightarrow @_y^{i+1} \text{val}) \right) \right).$$

Both $\text{fork}_j^{i+1}(x, y)$ and $[x = y]_j^{i+1}$ used in these formulae are defined recursively. The formula $S_j^i(ax, bx)$ states that there is a submodel $\mathfrak{M}' \sqsubseteq \mathfrak{M}$ such that

- I. $\mathfrak{M}', w \models \text{fork}_j^{i+1}(x, y)$;
- II. s_{ax} corresponds to the nominal x at depth $i+1$;
- III. s_{bx} corresponds to the nominal y at depth $i+1$;
- IV-VI. $\mathbf{n}(s_{ax}) = \mathbf{n}(s_{bx})$, $\mathfrak{M}, s_{ax} \not\models \text{val}$ and $\mathfrak{M}, s_{bx} \models \text{val}$.

(The enumeration I-VI refers to the conjuncts in the formula)

$S_j^i(ax, bx)$ correctly models the first condition of (LESS).

Regarding $L_j^i(ax, bx)$ and (LESS), a similar analysis can be performed. We define $LS_j^i(ax, bx) \stackrel{\text{def}}{=} L_j^i(ax, bx) \wedge S_j^i(ax, bx)$.

Let us consider $[bx = ax+1]_j$. Under the hypothesis that $\mathfrak{M}, w \models \text{fork}_j^i(ax, bx)$, this formula must express $\mathbf{n}(w_{bx}) = \mathbf{n}(w_{ax}) + 1$ for the two (distinct) worlds $w_{ax}, w_{bx} \in R^i(w)$. Then, as done for defining $[ax < bx]_j^i$, we take advantage of arithmetical properties on binary numbers and we search for two partitions $P_{ax} = \{L_{ax}, S_{ax}, R_{ax}\}$ and $P_{bx} = \{L_{bx}, S_{bx}, R_{bx}\}$ of the t -children of w_{ax} and w_{bx} , respectively, such that P_{ax} and P_{bx} satisfy (LSR) as well as the condition below:

(PLUS): P_{ax} and P_{bx} have the arithmetical properties of $+1$:

- P_{ax} and P_{bx} satisfy (LESS);
- for every $r_{ax} \in R_{ax}$, we have $\mathfrak{M}, r_{ax} \models \text{val}$;
- for every $r_{bx} \in R_{bx}$, we have $\mathfrak{M}, r_{bx} \not\models \text{val}$,

where $S_{ax} = \{s_{ax}\}$ and $S_{bx} = \{s_{bx}\}$, as required by (LSR).

The definition of $[bx = ax+1]_j^i$ is similar to $[ax < bx]_j^i$:

$$\top * \left(\text{nom}_1(ax \neq bx) \wedge [t] \text{lsr}(j-1) \wedge LS_j^1(ax, bx) \wedge R(ax, bx) \right),$$

where $R(ax, bx) \stackrel{\text{def}}{=} @_{ax}^1 [t] (\diamond r \Rightarrow \text{val}) \wedge @_{bx}^1 [t] (\diamond r \Rightarrow \neg \text{val})$ captures the last two conditions of (PLUS).

To define $\text{uniq}(j)$ and $\text{compl}(j)$, we rely on $\text{fork}_j^i(ax, bx)$, $[ax < bx]_j^i$ and $[bx = ax+1]_j$.

$$\text{uniq}(j) \stackrel{\text{def}}{=} \neg \left(\top * \left(\text{fork}_j^1(x, y) \wedge [x = y]_j^1 \right) \right)$$

$$\text{compl}(j) \stackrel{\text{def}}{=} \neg \left(\square \perp * \left([t] (\text{type}_{\text{lsr}}(j-1) \wedge \diamond y) \wedge \text{nom}_1(x) \wedge @_x^1 \neg 1_j \wedge \neg \left(\top * \left(\text{fork}_j^1(x, y) \wedge [y = x+1]_j^1 \right) \right) \right) \right),$$

where $1_j \stackrel{\text{def}}{=} [t] \text{val}$ reflects the encoding of $t(j, n) - 1$ for $j > 1$. The main difference between $\text{compl}(1)$ and $\text{compl}(j)$ ($j > 1$) is that the conjunct $[t] \diamond y$ of $\text{compl}(1)$ is replaced

by $[t] (\text{type}_{\text{lsr}}(j-1) \wedge \diamond y)$ in $\text{compl}(j)$, as needed to correctly evaluate $\text{fork}_j^1(x, y)$. Indeed, the difference between $\text{fork}_1^1(x, y)$ and $\text{fork}_j^1(x, y)$ is precisely that the latter requires $[t] \text{type}_{\text{lsr}}(j-1)$. The definition of $\text{type}(j)$ is now complete. We can state its correctness.

Lemma 4.2. *Let $\mathfrak{M}, w \models \text{init}(j)$. We have $\mathfrak{M}, w \models \text{type}(j)$ iff (\mathfrak{M}, w) satisfies (sub_j), (zero_j), (uniq_j), (compl_j) and (aux).*

The size of $\text{type}(j)$ is exponential in $j > 1$ and polynomial in $n \geq 1$. As its size is elementary, we can use this formula as a starting point to reduce Tile_k .

4.2 Tiling a grid $[0, t(k, n) - 1] \times [0, t(k, n) - 1]$

Below, we briefly explain how to use previous developments to define a uniform reduction from Tile_k , for every $k \geq 2$. Several adaptations are needed to encode smoothly the grid but the hardest part was the design of $\text{type}(j)$. Let $k \geq 2$ and (\mathcal{T}, c) be an instance of Tile_k . We can construct a formula $\text{tiling}_{\mathcal{T}, c}(k)$ that is satisfiable if and only if (\mathcal{T}, c) as a solution. To represent $[0, t(k, n) - 1]^2$ in some pointed forest (\mathfrak{M}, w) , where $\mathfrak{M} = (W, R, V)$, we recycle the ideas for defining $\text{type}(k)$. From Lemma 4.2, we know that if $\mathfrak{M}, w \models \text{init}(k) \wedge \text{type}(k)$ then the t -children of w encode the interval $[0, t(k, n) - 1]$. A position in the grid is however a pair of numbers, hence the *crux of our encoding* rests on the fact that each $w' \in R(w)$ encodes *two* numbers $\mathbf{n}_{\mathcal{H}}(w')$ and $\mathbf{n}_{\mathcal{V}}(w')$. Similarly to $\text{type}(k)$, these numbers are represented by the truth values on the t -children of w' , with the help of *new propositions* $\text{val}_{\mathcal{H}}$ and $\text{val}_{\mathcal{V}}$. We are in luck: since both numbers are from $[0, t(k, n) - 1]$, w' just needs as many children as when encoding a single number, and therefore if $\mathfrak{M}, w \models \text{tiling}_{\mathcal{T}, c}(k)$ then $\mathfrak{M}, w' \models \text{type}(k-1)$. In fact, the portion of $\text{tiling}_{\mathcal{T}, c}(k)$ that encodes the grid can be described quite naturally by slightly updating the characterisation of $\text{type}(k)$. For example, (uniq_j) becomes

(uniq_{\mathcal{T}, k}) for all distinct t -nodes $w_1, w_2 \in R(w)$
 $\mathbf{n}_{\mathcal{H}}(w_1) \neq \mathbf{n}_{\mathcal{H}}(w_2)$ or $\mathbf{n}_{\mathcal{V}}(w_1) \neq \mathbf{n}_{\mathcal{V}}(w_2)$.

The formula $\text{uniq}(k)$ has to be updated accordingly, but without major differences or complications. Of course, more is required as $\text{tiling}_{\mathcal{T}, c}(k)$ must also encode the tiling conditions (first) and (hor&vert). Fortunately, the kit of formulae defined for $\text{type}(k)$ allows us to have access to $\mathbf{n}_{\mathcal{H}}(w')$ and $\mathbf{n}_{\mathcal{V}}(w')$ in such a way that both conditions can be expressed rather easily. For example, to express vertical constraints, we design a formula stating that for all t -nodes $w_1, w_2 \in R(w)$, if $\mathbf{n}_{\mathcal{V}}(w_2) = \mathbf{n}_{\mathcal{V}}(w_1) + 1$ and $\mathbf{n}_{\mathcal{H}}(w_2) = \mathbf{n}_{\mathcal{H}}(w_1)$ then there is $(c_1, c_2) \in \mathcal{V}$ such that $w_1 \in V(c_1)$ and $w_2 \in V(c_2)$. Further details are omitted by lack of space.

Theorem 4.3. *Sat(ML(*)) is TOWER-complete.*

5 ML(*) Strictly Less Expressive Than GML

Below, we focus on the expressivity of ML(*). We first show $\text{ML}(\ast) \leq \text{GML}$ and then we prove the strictness of the inclusion. The former result takes advantage of the notion of g-bisimulation, i.e. the underlying structural indistinguishability relation of GML, studied in [20]. To show $\text{ML}(\ast) < \text{GML}$, we define an ad hoc notion of Ehrenfeucht-Fraïssé games for ML(*), see e.g. classical definitions in [31] and similar approaches in [13, 18]. Then, we design a simple formula in GML that cannot be expressed in ML(*) .

5.1 ML(*) is not more expressive than GML

To establish that $\text{ML}(\ast) \leq \text{GML}$, we proceed as in Section 3.1. In fact, by Lemma 2.2, given φ_1, φ_2 in GML, the formula $\varphi_1 \ast \varphi_2$ is equivalent to $\blacklozenge(\varphi_1 \upharpoonright \varphi_2)$. Moreover, we know that given φ_1, φ_2 in GML, $\varphi_1 \upharpoonright \varphi_2$ is equivalent to some formula in GML, as shown in Section 3. So, to prove that $\text{ML}(\ast) \leq \text{GML}$ by applying the proof schema of Theorem 3.2, it is sufficient to show that given φ in GML, there is ψ in GML such that $\blacklozenge\varphi \equiv \psi$. To do so, we rely on the indistinguishability relation of GML, called g-bisimulation [20].

A g-bisimulation is a refinement of the classical back-and-forth conditions of a bisimulation (see e.g. [8]), tailored towards capturing graded modalities. It relates models with similar structural properties, but up to parameters $m, k \in \mathbb{N}$ responsible for the modal degree and the graded rank, respectively. The following invariance result holds: g-bisimilar models are modally equivalent in GML (up to formulae of modal degree m and graded rank at most k). For simplicity, we present the construction of the above-mentioned formula ψ by directly using the notion of model equivalence, without going explicitly through g-bisimulations.

Given $m, k \in \mathbb{N}$ and $P \subseteq_{\text{fin}} \text{AP}$, we write $\text{GML}[m, k, P]$ to denote the set of GML formulae ψ having $\text{md}(\psi) \leq m$, $\text{gr}(\psi) \leq k$ and propositional variables from P . $\text{GML}[m, k, P]$ is finite up to logical equivalence [20]. Given pointed forests (\mathfrak{M}, w) and (\mathfrak{M}', w') , we write $(\mathfrak{M}, w) \equiv_{m, k}^P (\mathfrak{M}', w')$ whenever (\mathfrak{M}, w) and (\mathfrak{M}', w') are $\text{GML}[m, k, P]$ -indistinguishable, i.e. for every ψ in $\text{GML}[m, k, P]$, $\mathfrak{M}, w \models \psi$ iff $\mathfrak{M}', w' \models \psi$. We write $\mathcal{T}^P(m, k)$ to denote the quotient set induced by the equivalence relation $\equiv_{m, k}^P$. As $\text{GML}[m, k, P]$ is finite up to logical equivalence, we get that $\mathcal{T}^P(m, k)$ is finite.

To establish that GML is closed under \blacklozenge , we show that there is a function $\mathfrak{f} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that for all $m, k \in \mathbb{N}$ and $P \subseteq_{\text{fin}} \text{AP}$, if two models are in the same equivalence class of $\equiv_{m, \mathfrak{f}(m, k)}^P$, then they satisfy the same formulae of the form $\blacklozenge\varphi$, where φ is in $\text{GML}[m, k, P]$. By standard arguments and using the fact that $\text{GML}[m, \mathfrak{f}(m, k), P]$ is finite up to logical equivalence, we then conclude that $\blacklozenge\varphi$ is equivalent to a formula in $\text{GML}[m, \mathfrak{f}(m, k), P]$. Similar approaches are followed in [22, 23, 34]. As we are not interested in the size of

the equivalent formula, we can simply use the cardinality of $\mathcal{T}^P(m, k)$ in order to inductively define a suitable function:

$$\mathfrak{f}(0, k) \stackrel{\text{def}}{=} k, \quad \mathfrak{f}(m+1, k) \stackrel{\text{def}}{=} k \times (|\mathcal{T}^P(m, \mathfrak{f}(m, k))| + 1).$$

In conformity with the results in Section 4, the map \mathfrak{f} can be shown to be a non-elementary function. To prove that \mathfrak{f} satisfies the required properties, we start by showing a technical lemma which essentially formalises a simulation argument on the relation $\equiv_{m, \mathfrak{f}(m, k)}^P$ with respect to the submodel relation. By taking submodels as with the \blacklozenge operator, equivalence in GML is preserved.

Lemma 5.1. *Let $(\mathfrak{M}, w) \equiv_{m, \mathfrak{f}(m, k)}^P (\mathfrak{M}', w')$ where $m, k \in \mathbb{N}$, $P \subseteq_{\text{fin}} \text{AP}$, $\mathfrak{M} = (W, R, V)$ and $\mathfrak{M}' = (W', R', V')$. Let $R_1 \subseteq R$. There is $R'_1 \subseteq R'$ s.t. $((W, R_1, V), w) \equiv_{m, k}^P ((W', R'_1, V'), w')$ and if $R_1(w) = R(w)$, then $R'_1(w') = R'(w')$.*

The proof of Lemma 5.1 is by induction on m . The last condition about $R_1(w) = R(w)$ will serve in the proof of Lemma 5.2, as it allows us to capture the semantics of \blacklozenge , by preserving the children of the world w' . In the proof, we rely on the properties of g-bisimulations [20] to define a binary relation \leftrightarrow between worlds of $R(w)$ and $R'(w')$. Every $w_1 \leftrightarrow w'_1$ is such that $(\mathfrak{M}, w_1) \equiv_{m-1, \mathfrak{f}(m-1, k)}^P (\mathfrak{M}', w'_1)$. The operator \blacklozenge does not necessarily preserve the children of w_1 and w'_1 , so that the induction hypothesis, naturally defined from the statement of Lemma 5.1, is applied on models where the condition $R_1(w_1) = R(w_1)$ may not hold. We show that for all $R_1 \subseteq R$, it is possible to construct $R'_1 \subseteq R'$ such that, for all $w_1 \leftrightarrow w'_1$, $((W, R_1, V), w_1) \equiv_{m-1, k}^P ((W', R'_1, V'), w'_1)$. The result is then lifted to $((W, R_1, V), w) \equiv_{m, k}^P ((W', R'_1, V'), w')$, again thanks to the properties of the g-bisimulation.

Intuitively, Lemma 5.1 states that given two models satisfying the same formulae up to the parameters m and $\mathfrak{f}(m, k)$, we can extract submodels satisfying the same formulae up to m and k (reduced graded rank). This allows us to conclude that if φ is in GML, there is some GML formula equivalent to $\blacklozenge\varphi$ (Lemma 5.2). In other words, the operator \blacklozenge can be eliminated to obtain a GML formula. This, together with Lemma 2.2 and Theorem 3.2 entail $\text{ML}(\ast) \leq \text{GML}$.

Lemma 5.2. *For every $\varphi \in \text{GML}[m, k, P]$ there is a formula $\psi \in \text{GML}[m, \mathfrak{f}(m, k), P]$ such that $\blacklozenge\varphi \equiv \psi$.*

5.2 Showing $\text{ML}(\ast) < \text{GML}$ with EF games for ML(*)

We tackle the problem of showing that $\text{ML}(\ast)$ is strictly less expressive than GML. To do so, we adapt the notion of Ehrenfeucht-Fraïssé games (EF games, in short) [31] to $\text{ML}(\ast)$, and use it to design a GML formula that is not expressible in $\text{ML}(\ast)$. We write $\text{ML}(\ast)[m, s, P]$ for the set of formulae φ of $\text{ML}(\ast)$ having $\text{md}(\varphi) \leq m$, at most s nested \ast , and atomic propositions from $P \subseteq_{\text{fin}} \text{AP}$. It is easy to see that $\text{ML}(\ast)[m, s, P]$ is finite up to logical equivalence.

We introduce the EF games for $\text{ML}(\ast)$. A game is played between two players: the *spoiler* and the *duplicator*. A game

Game on $[(\mathfrak{M}_1=(W_1, R_1, V_1), w_1), (\mathfrak{M}_2=(W_2, R_2, V_2), w_2), (m, s, P)]$.

if there is $p \in P$ s.t. $w_1 \in V_1(p)$ iff $w_2 \notin V_2(p)$ then the spoiler wins.
else the spoiler chooses $i \in \{1, 2\}$ and plays on \mathfrak{M}_i . The duplicator replies on \mathfrak{M}_j where $j \neq i$. The spoiler must choose one of the following moves, otherwise the duplicator wins:

modal move: if $m \geq 1$ and $R_i(w_i) \neq \emptyset$ then the spoiler **can** choose to play a modal move by selecting an element $w'_i \in R_i(w_i)$. Then,

- the duplicator must reply with a $w'_j \in R_j(w_j)$ (else, the spoiler wins);
- the game continues on $[(\mathfrak{M}_1, w'_1), (\mathfrak{M}_2, w'_2), (m-1, s, P)]$.

spatial move: if $s \geq 1$ then the spoiler **can** choose to play a spatial move by selecting two finite forests \mathfrak{M}_i^1 and \mathfrak{M}_i^2 s.t. $\mathfrak{M}_i^1 + \mathfrak{M}_i^2 = \mathfrak{M}_i$. Then,

- the duplicator replies with two forests \mathfrak{M}_j^1 and \mathfrak{M}_j^2 s.t. $\mathfrak{M}_j^1 + \mathfrak{M}_j^2 = \mathfrak{M}_j$;
- The game continues on $[(\mathfrak{M}_1^k, w_1), (\mathfrak{M}_2^k, w_2), (m, s-1, P)]$, where $k \in \{1, 2\}$ is chosen by the spoiler.


Figure 1. Ehrenfeucht-Fraïssé games for $ML(*)$

state is a triple made of two pointed forests (\mathfrak{M}, w) and (\mathfrak{M}', w') and a rank (m, s, P) , where $m, s \in \mathbb{N}$ and $P \subseteq_{\text{fin}} AP$. The goal of the spoiler is to show that the two models are different. The goal of the duplicator is to counter the spoiler and to show that the two models are similar. Two models are different whenever there is $\varphi \in ML(*)[m, s, P]$ that is satisfied by only one of the two models. The EF games for $ML(*)$ are formally defined in Figure 1. The exact correspondence between the game and the logic is formalised in Lemma 5.3.

Using the standard definitions in [31], the duplicator has a *winning strategy* for the game $((\mathfrak{M}, w), (\mathfrak{M}', w'), (m, s, P))$ if she can play in a way that guarantees her to win regardless how the spoiler plays. When this is the case, we write $(\mathfrak{M}, w) \approx_{m,s}^P (\mathfrak{M}', w')$. Similarly, the spoiler has a *winning strategy*, written $(\mathfrak{M}, w) \not\approx_{m,s}^P (\mathfrak{M}', w')$, if he can play in a way that guarantees him to win, regardless how the duplicator plays. Lemma 5.3 guarantees that the games are well-defined.

Lemma 5.3. $(\mathfrak{M}, w) \not\approx_{m,s}^P (\mathfrak{M}', w')$ iff there is a formula φ in $ML(*)[m, s, P]$ such that $\mathfrak{M}, w \models \varphi$ and $\mathfrak{M}', w' \not\models \varphi$.

Lemma 5.3 is proven with standard arguments from [31], for instance the left-to-right direction, i.e. the *completeness of the game*, is by induction on the rank (m, s, P) . Thanks to the EF games, we are able to find a GML formula φ that is not expressible in $ML(*)$. By Lemma 2.1 and as $ML(\mathbf{I}) \approx GML$, such a formula is necessarily of modal degree at least 2. Happily, $\varphi = \diamond_{=2} \diamond_{=1} \top$ does the job and cannot be expressed in $ML(*)$. For the proof, we show that for every rank (m, s, P) , there are two structures (\mathfrak{M}, w) and (\mathfrak{M}', w') such that $(\mathfrak{M}, w) \approx_{m,s}^P (\mathfrak{M}', w')$, $\mathfrak{M}, w \models \varphi$ and $\mathfrak{M}', w' \not\models \varphi$. The inexpressibility of φ then stems from Lemma 5.3. The two structures are represented below $((\mathfrak{M}, w)$ on the left).



In the following, we say that a world has *type* i if it has i children. As one can see in the figure above, children of the current worlds w and w' are of three types: 0, 1 or 2. When the spoiler performs a spatial move in the game, a world of type i can take, in the submodels, a type between 0 and i . That is, the number of children of a world weakly monotonically decreases when taking submodels. This monotonicity, together with the finiteness of the game, lead to bounds on the number of children of each type, over which the duplicator is guaranteed to win. For instance, the bound for worlds of type 2 is given by the value $2^s(s+1)(s+2)$, where s is the number of spatial moves in the game. In the two presented pointed forests, one child of type 0 and one of type 2 are added with respect to these bounds, so that the duplicator can make up for the different numbers of children of type 1.

Lemma 5.4. $ML(*)$ cannot characterise the class of models satisfying the GML formula $\diamond_{=2} \diamond_{=1} \top$.

Notice that $ML(*)$ is more expressive than ML . Indeed, the formula $\diamond \top * \diamond \top$ distinguishes the two models on the right, which are bisimilar and hence indistinguishable in ML [46].



By $ML(*) \leq GML$, Lemma 5.4 and Theorem 3.2, we conclude.

Theorem 5.5. $ML < ML(*) < GML \approx ML(\mathbf{I})$.

6 $ML(\mathbf{I})$, $ML(*)$ and Sister Logics

Below, we show how our new results on $ML(\mathbf{I})$ and $ML(*)$ allow us to make substantial contributions for sister logics.

6.1 Static ambient logic

Static ambient logic (SAL) is a formalism proposed to reason about spatial properties of concurrent processes specified in the ambient calculus [15]. In [12], the satisfiability and validity problems for a very expressive fragment of SAL are shown to be decidable and conjectured to be in PSPACE (see [12, Section 6]). We invalidate this conjecture by showing that the intensional fragment of SAL (see [32]), herein denoted $SAL(\mathbf{I})$, is already $AEXP_{\text{POL}}$ -complete. More precisely, we design semantically faithful reductions between $\text{Sat}(ML(\mathbf{I}))$ and $\text{Sat}(SAL(\mathbf{I}))$ (in both directions), leading to the above-mentioned result by Corollary 3.8. $SAL(\mathbf{I})$ formulae are from

$$\varphi := \top \mid \emptyset \mid n[\varphi] \mid \varphi \wedge \varphi \mid \neg \varphi \mid \varphi \mid \varphi,$$

where $n \in AP$ is an *ambient name*. Historically, the semantics of SAL is given on a class of syntactically defined finite trees. However, this class of models is isomorphic to the class of finite trees $\mathfrak{M} = (W, R, V)$, such that each world in W satisfies exactly one atomic proposition (its ambient name). Then, the satisfaction relation \models for $SAL(\mathbf{I})$ is standard for \top and Boolean connectives, $\varphi_1 \mid \varphi_2$ is as in $ML(\mathbf{I})$, and otherwise


$$\begin{aligned} \mathfrak{M}, w \models \emptyset &\Leftrightarrow R(w) = \emptyset; \\ \mathfrak{M}, w \models n[\varphi] &\Leftrightarrow \text{there is } w' \in W \text{ such that } R(w) = \{w'\}, \\ &\quad w' \in V(n) \text{ and } \mathfrak{M}, w' \models \varphi. \end{aligned}$$

With such a presentation, $\text{SAL}(\mathbf{I})$ is a fragment of $\text{ML}(\mathbf{I})$, where \emptyset and $n[\varphi]$ correspond to $\Box\perp$ and $\Diamond_{=1}\top \wedge \Diamond(n \wedge \varphi)$, respectively. However, to reduce $\text{Sat}(\text{SAL}(\mathbf{I}))$ to $\text{Sat}(\text{ML}(\mathbf{I}))$, we must deal with the constraint on V (uniqueness of the ambient name). Let φ be in $\text{SAL}(\mathbf{I})$ written with the ambient names in $N = \{n_1, \dots, n_m\}$. It is known (see [12, Lemma 8]) that if φ is satisfiable, then it can be satisfied by a tree having ambient names from $N \cup \{\bar{n}\}$, where \bar{n} is a fresh name. Thus, we can show that φ is satisfiable iff so is the $\text{ML}(\mathbf{I})$ formula

$$\varphi \wedge \boxplus^{\text{md}(\varphi)} (\bigvee_{n \in N \cup \{\bar{n}\}} (n \wedge \bigwedge_{n' \in (N \cup \{\bar{n}\}) \setminus \{n\}} \neg n')),$$

where the right conjunct states that V , restricted to the propositions in $N \cup \{\bar{n}\}$, forms a partition of the worlds reachable from the current one in at most $\text{md}(\varphi)$ steps.

Reducing $\text{Sat}(\text{ML}(\mathbf{I}))$ to $\text{Sat}(\text{SAL}(\mathbf{I}))$ requires a bit more work. Let $\mathfrak{M} = (W, R, V)$ be a finite forest and $w \in W$. Assume we want to check the satisfiability status of φ in $\text{ML}(\mathbf{I})$ having atomic propositions from $P = \{p_1, \dots, p_m\}$ and with n occurrences of \mathbf{I} . We encode (\mathfrak{M}, w) into a model $(\mathfrak{M}' = (W', R', V'), w)$ of $\text{SAL}(\mathbf{I})$ as follows. Let rel and ap be two ambient names not in P . The ambient name rel encodes the relation R whereas ap can be seen as a *container* for propositional variables holding on the current world. **(i)** We require $W \subseteq W'$, $R \subseteq R'$ and $\bigcup_{i \in [0, \text{md}(\varphi)]} R^i(w) \subseteq V'(\text{rel})$, i.e., every world reachable from w in at most $\text{md}(\varphi)$ steps has the ambient name rel . Let w' be one of these worlds and suppose that $\{p \mid w' \in V(p)\} \cap P = \{q_1, \dots, q_l\}$. **(ii)** We require W' to contain $n+1$ worlds $w'_1, \dots, w'_{n+1} \in R'(w') \setminus R(w')$, all having ambient name ap . These worlds encode copies of w' 's valuation, similarly to what is done in Section 3.2 to encode teams from $\text{PL}[\sim]$. **(iii)** For all $j \in [1, n+1]$, $R'(w'_j)$ contains l worlds, all satisfying \emptyset and a distinct ambient name from $\{q_1, \dots, q_l\}$. Below we schematise the encoding (w.r.t. w').



Let $n \in \text{AP}$. We define the modality $\langle n \rangle \varphi \stackrel{\text{def}}{=} n[\varphi] \top$ and its dual $[n] \varphi \stackrel{\text{def}}{=} \neg \langle n \rangle \neg \varphi$. We write $\forall[n]$ for $\neg((\neg \emptyset \wedge \neg n[\top]) \top)$, so that $(\mathfrak{M}, w) \models \forall[n]$ whenever every child of w has the ambient name n . Moreover, $[\# \geq 0] \stackrel{\text{def}}{=} \top$ and $[\# \geq \beta + 1] \stackrel{\text{def}}{=} \neg \emptyset \mid [\# \geq \beta]$, so that $(\mathfrak{M}, w) \models [\# \geq \beta]$ whenever w has at least β children. Lastly, $[\# = \beta] \stackrel{\text{def}}{=} [\# \geq \beta] \wedge \neg[\# \geq \beta + 1]$. The models of $\text{SAL}(\mathbf{I})$ encoding models of $\text{ML}(\mathbf{I})$ are characterised by

$$C_\varphi \stackrel{\text{def}}{=} \bigwedge_{j \in [0, \text{md}(\varphi)]} [\text{rel}]^j \left(\forall [\text{rel}] \left(\forall [\text{ap}] \wedge [\# = n+1] \wedge [\text{ap}] \left((p_1[\emptyset] \vee \emptyset) \dots \mid (p_m[\emptyset] \vee \emptyset) \right) \wedge \bigwedge_{i \in [1, m]} (\langle \text{ap} \rangle \langle p_i \rangle \top \Rightarrow [\text{ap}] \langle p_i \rangle \top) \right) \right).$$

Lastly, we define the translation of φ , written $\tau(\varphi)$, into $\text{SAL}(\mathbf{I})$. It is homomorphic for Boolean connectives and \top , $\tau(p) \stackrel{\text{def}}{=} \langle \text{ap} \rangle \langle p \rangle \top$ and otherwise it is inductively defined:

$$\tau(\Diamond \psi) \stackrel{\text{def}}{=} \langle \text{rel} \rangle \tau(\psi);$$

$$\tau(\psi_1 \mid \psi_2) \stackrel{\text{def}}{=} \left(\tau(\psi_1) \wedge \langle \text{ap} \rangle_{\geq j} \top \right) \mid \left(\tau(\psi_2) \wedge \langle \text{ap} \rangle_{\geq k} \top \right),$$

where in $\tau(\psi_1 \mid \psi_2)$, j (resp. k) is the number of occurrences of \mathbf{I} in ψ_1 (resp. ψ_2) plus one and $\langle \text{ap} \rangle_{\geq \alpha} \top \stackrel{\text{def}}{=}} (\forall [\text{ap}] \wedge [\# = \alpha]) \top$.

We show that φ is satisfiable in $\text{ML}(\mathbf{I})$ iff $C_\varphi \wedge \tau(\varphi)$ is satisfiable in $\text{SAL}(\mathbf{I})$, leading to the following results about the complexity of static ambient logics.

Corollary 6.1. *Sat(SAL(I)) is AExp_{POL} -complete. Sat(SAL) with SAL from [12] is AExp_{POL} -hard.*

6.2 Modal separation logic

The family of *modal separation logics* (MSL), combining separating and modal connectives, has been recently introduced in [21]. Its models, inspired from the memory states used in separation logic (see also [17]), are Kripke-style structures $\mathfrak{M} = (W, R, V)$, where $W = \mathbb{N}$ and $R \subseteq W \times W$ is finite and functional. Hence, unlike finite forests, \mathfrak{M} may have loops.

Among the fragments studied in [21], the modal separation logic $\text{MSL}(*, \Diamond^{-1})$ was left with a huge complexity gap (between PSPACE and TOWER). Its formulae are defined from

$$\varphi := p \mid \Diamond^{-1} \varphi \mid \varphi \wedge \varphi \mid \neg \varphi \mid \varphi * \varphi.$$

The satisfaction relation is as in $\text{ML}(*)$ for $p \in \text{AP}$, Boolean connectives and $\varphi_1 * \varphi_2$, otherwise

$$\mathfrak{M}, w \models \Diamond^{-1} \varphi \Leftrightarrow \exists w' \text{ s.t. } (w', w) \in R \text{ and } \mathfrak{M}, w' \models \varphi.$$

Since $\text{MSL}(*, \Diamond^{-1})$ is interpreted over a finite and functional relation, \Diamond^{-1} effectively works as the \Diamond modality of $\text{ML}(*)$. Then, assume we want to check the satisfiability of φ in $\text{ML}(*)$ by relying on an algorithm for $\text{Sat}(\text{MSL}(*, \Diamond^{-1}))$. We simply need to consider the formula $\varphi[\Diamond \leftarrow \Diamond^{-1}]$ obtained from φ by replacing every occurrence of \Diamond by \Diamond^{-1} , and check if it can be satisfied by a *locally acyclic* model (\mathfrak{M}, w) of MSL, i.e. one where w does not belong to a loop of length $\leq \text{md}(\varphi)$. Local acyclicity can be enforced by the formula

$$\text{locacycl} \stackrel{\text{def}}{=} r \wedge \bigwedge_{i \in [1, \text{md}(\varphi)]} (\Box^{-1})^i \neg r,$$

where $r \in \text{AP}$ is fresh. Then, φ in $\text{ML}(*)$ is satisfiable iff $\varphi[\Diamond \leftarrow \Diamond^{-1}] \wedge \text{locacycl}$ in $\text{MSL}(*, \Diamond^{-1})$ is satisfiable. Hence, the results in Section 4 allow us to close the complexity gap.

Corollary 6.2. *Sat(MSL(*, \Diamond^{-1})) is TOWER-complete.*

7 Conclusion

We have studied and compared $\text{ML}(\mathbf{I})$ and $\text{ML}(*)$, two modal logics interpreted on finite forests and featuring composition operators. We have not only characterised the expressive power and the complexity for both logics, but also identified remarkable differences and export our results to other logics. $\text{ML}(\mathbf{I})$ is shown as expressive as GML, and its satisfiability problem is found to be AExp_{POL} -complete. Besides the obvious similarities between $\text{ML}(\mathbf{I})$ and $\text{ML}(*)$, these results are counter-intuitive: though the logic $\text{ML}(*)$ is strictly less expressive than GML (and consequently, than $\text{ML}(\mathbf{I})$), $\text{Sat}(\text{ML}(*))$ is TOWER-complete. We also recalled that there

are logspace reductions from $ML(*)$ and $ML(\mathbf{I})$ to the second-order modal logic QK^t from [7].

Our proof techniques go beyond what is known in the literature. For instance, to design the TOWER-hardness proof we needed substantial modifications from the proof introduced in [7] for QK^t . On the other hand, to show the expressivity inclusion of $ML(*)$ within GML, we provided a novel definition of Ehrenfeucht-Fraïssé games for $ML(*)$.

Lastly, our framework led to the characterisation of the satisfiability problems for two sister logics. We proved that the satisfiability problem for the modal separation logic $MSL(*, \diamond^{-1})$ is TOWER-complete [21]. Moreover, the satisfiability problem for the static ambient logic $SAL(\mathbf{I})$ is $AEXP_{POL}$ -complete, solving open problems from [12, 21] and paving the way to study the complexity of the full SAL.

Acknowledgements

We would like to thank the anonymous reviewers for their comments and suggestions that helped us to improve the quality of the document. B. Bednarczyk is supported by the Polish Ministry of Science and Higher Education program “Diamantowy Grant” no. DI2017 006447. S. Demri and A. Mansutti are supported by the Centre National de la Recherche Scientifique (CNRS). R. Fervari is supported by ANPCyT-PICTs-2017-1130 and 2016-0215, and by the Laboratoire International Associé SINFIN.

References

- [1] C. Areces, R. Fervari, and G. Hoffmann. 2012. Moving Arrows and Four Model Checking Results. In *WoLLIC'12 (LNCS)*, Vol. 7456. Springer, 142–153.
- [2] C. Areces, R. Fervari, and G. Hoffmann. 2015. Relation-changing modal operators. *Logic Journal of the IGPL* 23, 4 (2015), 601–627.
- [3] G. Aucher, Ph. Balbiani, L. Fariñas del Cerro, and A. Herzig. 2009. Global and Local Graph Modifiers. *Electronic Notes in Theoretical Computer Science* 231 (2009), 293–307.
- [4] G. Aucher, J. van Benthem, and D. Grossi. 2018. Modal logics of sabotage revisited. *JLC* 28, 2 (2018), 269–303.
- [5] Ph. Balbiani, A. Baltag, H. van Ditmarsch, A. Herzig, T. Hoshi, and T. De Lima. 2008. ‘Knowable’ as ‘known after an announcement’. *Review on Symbolic Logic* 1, 3 (2008), 305–334.
- [6] M. Fattorosi Barnaba and F. De Caro. 1985. Graded Modalities. *Studia Logica* 44, 2 (1985), 197–221.
- [7] B. Bednarczyk and S. Demri. 2019. Why Propositional Quantification Makes Modal Logics on Trees Robustly Hard?. In *LiCS'19*. IEEE, 1–13.
- [8] P. Blackburn, M. de Rijke, and Y. Venema. 2001. *Modal Logic*. Cambridge University Press.
- [9] I. Boneva, J.-M. Talbot, and S. Tison. 2005. Expressiveness of a Spatial Logic for Trees. In *LiCS'05*. IEEE Computer Society, 280–289.
- [10] L. Bozzelli, A. Molinari, A. Montanari, and A. Peron. 2017. 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. 31–45.
- [11] L. Bozzelli, H. van Ditmarsch, and S. Pinchinat. 2015. The complexity of one-agent refinement modal logic. *TCS* 603 (2015), 58–83.
- [12] C. Calcagno, L. Cardelli, and A.D. Gordon. 2003. Deciding validity in a spatial logic for trees. In *TLDI'03*. ACM, 62–73.
- [13] C. Calcagno, T. Dinsdale-Young, and Ph. Gardner. 2010. Adjunct elimination in Context Logic for trees. *Inf. Comput.* 208, 5 (2010), 474–499.
- [14] C. Calcagno, Ph. Gardner, and U. Zarfaty. 2005. Context logic and tree update. In *POPL'05*. 271–282.
- [15] L. Cardelli and A.D. Gordon. 2001. *Formal Methods for Distributed Processing*. Cambridge University Press, New York, NY, USA, Chapter Mobile Ambients, 198–229.
- [16] D. Cooper. 1972. Theorem proving in arithmetic without multiplication. *Machine Learning* 7 (1972), 91–99.
- [17] J.-R. Courtault and D. Galmiche. 2018. A modal separation logic for resource dynamics. *JLC* 28, 4 (2018), 733–778.
- [18] A. Dawar, Ph. Gardner, and G. Ghelli. 2004. Adjunct Elimination Through Games in Static Ambient Logic. In *FST&TCS'04 (LNCS)*, Vol. 3328. Springer, 211–223.
- [19] A. Dawar, Ph. Gardner, and G. Ghelli. 2007. Expressiveness and complexity of graph logic. *I&C* 205, 3 (2007), 263–310.
- [20] M. de Rijke. 2000. A Note on Graded Modal Logic. *Studia Logica* 64, 2 (2000), 271–283.
- [21] S. Demri and R. Fervari. 2019. The power of modal separation logics. *JLC* 29, 8 (2019), 1139–1184.
- [22] S. Demri, R. Fervari, and A. Mansutti. 2019. Axiomatizing Logics with Separating Conjunction and Modalities. In *JELIA'19 (LNAI)*, Vol. 11468. Springer, 692–708.
- [23] M. Echenim, R. Iosif, and N. Peltier. 2019. The Bernays-Schönfinkel-Ramsey Class of Separation Logic on Arbitrary Domains. In *FoSSaCS'19 (LNCS)*, Vol. 11425. Springer, 242–259.
- [24] K. Fine. 1970. Propositional quantifiers in modal logic. *Theoria* 36 (1970), 336–346.
- [25] D. Galmiche, D. Méry, and D. Pym. 2005. The Semantics of BI and Resource Tableaux. *Mathematical Structures in Comp. Sci.* 15, 6 (Dec. 2005), 1033–1088.
- [26] Ch. Haase. 2018. A survival guide to Presburger arithmetic. *SIGLOG News* 5, 3 (2018), 67–82.
- [27] M. Hannula, J. Kontinen, J. Virtema, and H. Vollmer. 2018. Complexity of Propositional Logics in Team Semantics. *ACM ToCL* 19, 1 (2018), 2:1–2:14.
- [28] S. Ishtiaq and P. O’Hearn. 2001. BI as an assertion language for mutable data structures. In *POPL'01*. 14–26.
- [29] R. Ladner. 1977. The computational complexity of provability in systems of modal propositional logic. *SIAM Journal of Computing* 6, 3 (1977), 467–480.
- [30] F. Laroussinie and N. Markey. 2014. Quantified CTL: Expressiveness and Complexity. *LMCS* 10, 4:17 (2014).
- [31] L. Libkin. 2004. *Elements of Finite Model Theory*. Springer.
- [32] É. Lozes. 2004. Adjuncts elimination in the static ambient logic. *Electronic Notes in Theoretical Computer Science* 96 (2004), 51–72.
- [33] C. Lutz. 2006. Complexity and succinctness of public announcement logic. In *AAMAS'06*. ACM, 137–143.
- [34] A. Mansutti. 2018. Extending Propositional Separation Logic for Robustness Properties. In *FSTTCS (LIPIcs)*, Vol. 122. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 42:1–42:23.
- [35] P.W. O’Hearn, J.C. Reynolds, and H. Yang. 2001. Local Reasoning about Programs that Alter Data Structures. In *CSL'01 (LNCS)*, Vol. 2142. Springer, 1–19.
- [36] C. H. Papadimitriou. 1994. *Computational complexity*. Addison-Wesley, I–XV, 1–523 pages.
- [37] J. Plaza. 1989. Logics of public communication. In *ISMIS'89, Charlotte, North Carolina, USA*.
- [38] M. Presburger. 1929. Über die Vollständigkeit eines gewissen Systems der Arithmetik ganzer Zahlen, in welchem die Addition als einzige Operation hervortritt. In *Comptes Rendus du premier congrès de mathématiciens des Pays Slaves, Warszawa*. 92–101.
- [39] D. Pym. 2002. *The semantics and proof theory of the logic of bunched implications*. Applied Logic, Vol. 26. Kluwer Academic Publishers.
- [40] M. Rabin. 1969. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Math. Soc.* 41 (1969), 1–35.
- [41] C. Reddy and W. Loveland. 1978. Presburger arithmetic with bounded quantifier alternation. In *STOC'78*. ACM press, 320–325.
- [42] J.C. Reynolds. 2002. Separation logic: a logic for shared mutable data structures. In *LiCS'02*. IEEE, 55–74.
- [43] S. Schmitz. 2016. Complexity Hierarchies beyond Elementary. *TOCT* 8, 1 (2016), 3:1–3:36.
- [44] L. Schröder and D. Pattinson. 2006. PSPACE bounds for rank-1 modal logics. In *LiCS'06*. IEEE, 231–240.
- [45] S. Tobies. 2001. PSPACE Reasoning for Graded Modal Logics. *JLC* 11 (2001), 85–106.
- [46] J. van Benthem. 1976. *Modal Correspondence Theory*. Ph.D. Dissertation. University of Amsterdam.
- [47] M. Vardi and P. Wolper. 1986. Automata-theoretic techniques for modal logics of programs. *JCSS* 32 (1986), 183–221.
- [48] M. Vardi and P. Wolper. 1994. Reasoning about Infinite Computations. *I&C* 115 (1994), 1–37.