10

Controlled Model Exploration

GABRIEL G. INFANTE-LOPEZ, CARLOS ARECES, AND
MAARTEN DE RIJKE

ABSTRACT. We provide a detailed analysis of very weak fragments of modal logic. Our fragments lack connectives that introduce non-determinism and they feature restrictions on the modal operators, which may lead to substantial reductions in complexity. Our main result is a general game-based characterization of the expressive power of our fragments over the class of finite structures.

1 Introduction

The search for computationally well-behaved fragments of languages such as first-order and second-order logic has a long history. For instance, early in the twentieth century, Löwenheim already gave a decision procedure for the satisfiability of first-order sentences with only unary predicates. Some familiar fragments of first-order logic are defined by means of restrictions of the quantifier prefix of formulas in prenex normal forms. Finite-variable fragments of first-order logic are yet another family of fragments whose computational properties have been studied extensively, with decidability results going back to the early 1960s [20], while the late 1990s saw detailed complexity analyses of the two-variable fragment [10, 11, 16]. Despite the fact that the computational properties of prenex normal form and finite variable fragments have been (almost) completely investigated, these fragments leave something to be desired: their meta-logical properties are often poor, and, in particular, they usually do not enjoy a decent model theory that helps us to understand their computational properties. To overcome these drawbacks, there are ongoing research efforts to identify fragments of first-order or second-order logic that manage to combine good computational behavior with good logical properties.

One such effort takes modal logic as its starting point. Through the *standard* or *relational translation*, modal languages may be viewed as fragments of first-order languages [4]. Modal fragments are computationally very well-behaved; their satisfiability and model checking problems are of reasonably low complexity, and they are so in a robust way [21, 9]. The *guarded fragment* [1] was introduced as a generalization of the modal fragment, one that

Advances in Modal Logic, Volume 4, 205–220.

retains the good computational properties of modal fragments as much as possible. The good computational behavior of modal and guarded fragments has been explained in terms of the tree model property, and generalizations thereof.

In this paper we also search for well-behaved fragments of first-order logic by considering modal and modal-like languages, but we aim at a more fine-grained analysis. We start by taking a computationally well-behaved logic that can be translated into first-order logic, and try to generalize what we believe to be the main features responsible for the good computational behavior. Instead of modal logic, however, our starting point is taken from description logic. The description logic $\mathcal{FL}^-$ may be viewed as a restriction of the traditional modal language, where disjunctions are disallowed and the diamond operator is severely constrained. The restrictions built into $\mathcal{FL}^-$ yield significant reductions in computational complexity.

The aim of the paper is to provide a systematic exploration of the logical aspects of the restrictions built into $\mathcal{FL}^-$. We define a family of modal fragments inspired by $\mathcal{FL}^-$, briefly survey the computational complexity of their satisfiability problems, and spend most of the paper on providing a game-based characterization of their expressive power.

2 Description logics and $\mathcal{FL}^-$

Description logics have been proposed in the area of knowledge representation to specify systems in which structured knowledge can be expressed and reasoned with in a principled way [2]. They provide a logical basis to the well-known traditions of frame-based systems, semantic networks and KL-ONE-like languages, and now also for the semantic web. The main building blocks of languages of description logic are *concepts* and *roles*. The former are interpreted as subsets of a given domain, and the later as binary relations on the domain. Description logics differ in the constructions they admit for building complex concepts and roles.

Our starting point here is the logic $\mathcal{FL}^-$ [5]; its language has universal quantification, conjunction and unqualified existential quantification. That is, the legal concepts are generated by the following rule: $C ::= A \mid C \sqcap C \mid \forall R.C \mid \exists R.\top$, where A is an atomic concept, and R is an atomic role. In traditional modal logic notation, this production rule would be written as $\phi ::= p \mid \phi \wedge \phi \mid [R]\phi \mid \langle R \rangle \top$, or as $\phi ::= p \mid \phi \wedge \phi \mid \Box\phi \mid \Diamond\top$ when considering only one role.

Interpretations for description logics such as $\mathcal{FL}^-$ are pairs $\mathcal{I} = (\Delta, I)$, where Δ is a non-empty set, and I is a mapping that takes concepts to subsets of Δ and roles to subsets of $\Delta \times \Delta$. In (uni-)modal notation, a model is a tuple $\mathcal{A} = (W, R, V)$ where W is a non-empty set, R is a binary relation on W, and V is a function assigning subsets of W to proposition letters.

3 Taking a cue from $\mathcal{FL}^-$

The logic $\mathcal{FL}^-$ was carefully designed to control two important sources of computational complexity: non-determinism and deep model exploration. This aim shows up clearly in the syntactic constraints imposed on the language. The elimination of negation and disjunction restricts non-determinism (partial information cannot be expressed), while the restriction to unqualified existential quantification reduces model exploration to the bare minimum. As we will see in detail in Section 4, these design decisions have a significant impact on the computational complexity, making satisfiability checking trivial and subsumption checking polynomially tractable.

In contrast, standard modal logics (allowing full Boolean expressivity and qualified existential quantification) have PSPACE-complete satisfiability problems, as they allow one to code up models that are exponential in the size of the input formula [4]. The fact that restrictions on modal operators (the modal counterparts of description logic's quantifiers) produce computationally well behaved languages has also been studied in the modal logic community. Specifically, bounding the depth of nesting of modal operators may bring the complexity of the satisfiability problem down in dramatic ways, especially if one restricts the language even further by allowing only finitely many proposition letters (see [12]).

Despite the considerable computational impact of restricting non-determinism and existential quantification, a thorough analysis of its logical aspects, and especially of the expressive power, has been missing so far. The definitions below allow us to capture not just $\mathcal{FL}^-$ but a wide variety of additional fragments as well. In what follows we take the Boolean restrictions as they occur in $\mathcal{FL}^-$ mostly for granted (but we do include $\top$, $\bot$), and focus instead on its modal restrictions in a systematic way.

First, as we saw above, description and modal languages encode two kinds of information: *local* information depending only on the current node of evaluation, and *non-local* or *relational* information requiring model exploration (in controlled ways).

DEFINITION 1 (Local Formulas). A formula ϕ is a *local formula* if it is in the set of formulas LF generated by: $\phi ::= \top \mid \bot \mid p \mid \phi \wedge \phi \mid \Diamond\top$, where p is a proposition letter.

Second, we generalize the notion of unqualified existential quantification, by allowing complete control on which quantifiers are permitted at each level of nesting.

DEFINITION 2 (Fragment of $\mathcal{ML}$ modulo f). Let X be either $\mathbb{N}$ or an initial segment $\{1, \ldots, k\}$ of $\mathbb{N}$. Let $f : X \to \{\Diamond, \Box, \boxdiamond\}$. The *fragment of* $\mathcal{ML}$ *modulo* f (notation: $\mathcal{ML}^f$) is defined inductively as

$$
\begin{aligned}
\mathcal{ML}^f_0 &= \text{LF (the set of local formulas)}\\
\mathcal{ML}^f_{n+1} &= \text{the closure under taking conjunctions of (LF}\\
&\qquad \cup \{\Diamond\phi \mid \phi \in \mathcal{ML}^f_n \text{ and } f(n+1) = \Diamond\} \ \cup\\
&\qquad \cup \{\Box\phi \mid \phi \in \mathcal{ML}^f_n \text{ and } f(n+1) = \Box\} \ \cup\\
&\qquad \cup \{\Diamond\phi, \Box\phi \mid \phi \in \mathcal{ML}^f_n \text{ and } f(n+1) = \boxdiamond\}).
\end{aligned}
$$

The *language* $\mathcal{ML}^f$ is defined as $\mathcal{ML}^f = \bigcup_{n \in X} \mathcal{ML}^f_n$.

A few comments are in order. First, the definition of our $\mathcal{ML}^f$-fragments depends on the choice of LF, the set of local formulas; in Section 7 we will vary this set.

Second, the function f used in the definition allows us to precisely control the legal arguments of the modalities at each node in the construction tree of a formula in $\mathcal{ML}^f$. In this manner we are able to cut up the full modal language in novel ways. However, the present definition does not yet allow us to define all of the standard modal language $\mathcal{ML}$; see Section 7 for more on this.

Third, let $f^\Box : \mathbb{N} \to \{\Diamond, \Box, \boxdiamond\}$ be such that $f(n) = \Box$ for all n. Obviously, if we were to allow $\top$ and $\bot$ in $\mathcal{FL}^-$, we would have $\mathcal{ML}^{f^\Box} = \mathcal{FL}^-$. Our definition captures $\mathcal{FL}^-$ in a very natural way: the function $f^\Box$ dictates that the modal box (and only the modal box) can have arguments of arbitrary complexity.

4 Computational aspects

In this section we provide a brief overview of some computational aspects of our $\mathcal{ML}^f$-fragments. First of all, recall that the satisfiability problem for the standard modal logic **K** is PSPACE-complete. By going down to $\mathcal{FL}^-$, that is, by disallowing disjunction (as well as negation, $\top$ and $\bot$) and by restricting ourselves to unqualified existential quantification, the satisfiability problem becomes trivial as all formulas in $\mathcal{FL}^-$ are satisfiable. More interesting is the fact that deciding subsumption (given two formulas $\phi, \psi \in \mathcal{FL}^-$ decide whether $\phi \to \psi$ is a theorem) is solvable in polynomial time [6]. We refer the reader to [8] for further discussion on the computational aspects of $\mathcal{FL}^-$ and its extensions.

In [12], Halpern shows that finiteness restrictions (both on the number of propositional symbols and on the nesting of operators) also lowers the complexity of the inference tasks. Satisfiability of the basic modal logic **K** becomes NP-complete when we only allow finite nesting of modalities, and it drops to linear time when we furthermore restrict the language to only a finite number of propositional symbols.

These results can immediately be extended to the appropriate $\mathcal{ML}^f$ fragments. For example, the results for $\mathcal{FL}^-$ directly imply similar results for the fragment $\mathcal{ML}^{f^\Box}$ defined above. The following two results are more general, but also straightforward.

THEOREM 3. *Let $f : X \to \{\Diamond, \Box, \Box\!\!\!\!\!\Diamond\}$, where X is an initial segment or $X = \mathbb{N}$. The problem of deciding whether a formula in $\mathcal{ML}^f$ is satisfiable is in co-NP.*

Proof. For each $\mathcal{ML}^f$-fragment, we can reduce its satisfiability problem to the satisfiability problem for the description logic $\mathcal{ALE}$ [19]. $\mathcal{ALE}$ extends $\mathcal{FL}^-$ by allowing atomic negation, $\top$, $\bot$, and qualified existential quantification. That is, its set of legal concepts is given by $C ::= \top \mid \bot \mid \neg A \mid C \sqcap C \mid \forall R.C \mid \exists RC$. The satisfiability problem for $\mathcal{ALE}$ is known to be co-NP-complete [8]. ■

THEOREM 4. *Let X be $\mathbb{N}$ or an initial segment of $\mathbb{N}$, and let $f : X \to \{\Diamond, \Box, \Box\!\!\!\!\!\Diamond\}$ be such that $|\{n \mid f(n) = \Diamond$ or $f(n) = \Box\!\!\!\!\!\Diamond\}|$ is finite. Assume that the set of local formulas LF is built using only finitely many proposition letters. Then deciding if a formula in $\mathcal{ML}^f$ is satisfiable can be done in linear time.*

Proof. The proof follows the lines of the similar proof in [12]. Let k be the maximal n such that $f(n) = \Diamond$ or $f(n) = \Box\!\!\!\!\!\Diamond$. Given a formula ϕ in $\mathcal{ML}^f$, define ϕ^* by replacing every $\Box$-subformula of ϕ that occurs at depth $k+1$ or deeper by $\Box\bot$. It is easy to see that ϕ is satisfiable iff ϕ^* is. Hence, we only have to consider the fragment with formulas of modal depth at most $k+1$. A straightforward induction shows that there are only finitely many non-equivalent formulas in such fragments. Using this, one can find a fixed number of finite models such that a formula is satisfiable iff it is satisfiable on one of these models. This can be checked in time linear in the size of the formula being checked. ■

5 A game-based characterization

Our next aim is to obtain an exact semantic characterization of the $\mathcal{ML}^f$-fragments. Games are a flexible and popular tool for obtaining results of this kind; see e.g., [7] for an introduction at the textbook level. Given an appropriate function f we define a game G^f that precisely characterizes $\mathcal{ML}^f$; in Section 6 below we build on this to capture the expressive power of our $\mathcal{ML}^f$-fragments.

DEFINITION 5. Let $\mathcal{A}$ be a model, and let X and X' be two subsets of its universe. We use $\mathcal{A}, X \models \phi$ to denote that $\mathcal{A}, w \models \phi$, for all $w \in X$.

The *children* of w in $\mathcal{A}$ are all v such that wRv. We say that $XR^{\uparrow}X'$ if for every x in X there exists x' in X' with xRx'. We say that $XR^{\downarrow}X'$ if for all x' of X' there exists x in X with xRx' (i.e., X' is a subset of the children of X).

Let $\mathcal{A}$, $\mathcal{B}$ be models with domains $W_{\mathcal{A}}$ and $W_{\mathcal{B}}$, respectively, and let $X_0 \subseteq W_{\mathcal{A}}$, $Y_0 \subseteq W_{\mathcal{B}}$. Let f be such that $\operatorname{dom} f = \{1, \ldots, k\}$ or $\operatorname{dom} f = \mathbb{N}$, and assume n in the domain of f. We write $G^f(\mathcal{A}, X_0, \mathcal{B}, Y_0, n)$ to denote the following game. The game is played by two players, called `Di` and `Si`, on relational structures $\mathcal{A}$ and $\mathcal{B}$ (intuitively, `Di` is trying to proof that $\mathcal{A}$ and $\mathcal{B}$ are *different*, while `Si` wants to show they are *similar*). A *position* in the game $G^f(\mathcal{A}, X_0, \mathcal{B}, Y_0, n)$ is given by a pair $\langle X, Y \rangle$ such that X is a set of elements in $\mathcal{A}$ and Y a set of elements in $\mathcal{B}$; $\langle X_0, Y_0 \rangle$ is the initial position. During the $(i+1)$-th round, the current position $\langle X_i, Y_i \rangle$ will change to the new position $\langle X_{i+1}, Y_{i+1} \rangle$ according to the following rules.

Rule 1 If $f(n-i) = \Box$ then `Di` has to choose a set $Y_{i+1} \subseteq W_{\mathcal{B}}$ such that $Y_i R^{\uparrow} Y_{i+1}$, a counter-move of `Si` consists of choosing a set $X_{i+1} \subseteq W_{\mathcal{A}}$ such that $X_i R^{\uparrow} X_{i+1}$.

Rule 2 If $f(n-i) = \Diamond$ then `Di` has to choose a set $X_{i+1} \subseteq W_{\mathcal{A}}$ such that $X_i R^{\downarrow} X_{i+1}$. `Si` has to answer by choosing a set $Y_{i+1} \subseteq W_{\mathcal{B}}$ such that $Y_i R^{\downarrow} Y_{i+1}$.

Rule 3 If $f(n-i) = \mathrlap{\Box}\Diamond$ then `Di` can choose any of the previous rules to play by during this round.

The game *ends* on position $\langle X_i, Y_i \rangle$ when one of the following conditions fires:

Condition 1 There is a formula $\phi \in \mathrm{LF}$ st. $\mathcal{A}, X_i \models \phi$ but $\mathcal{B}, Y_i \not\models \phi$.

Condition 2 $i < n$ and `Si` cannot move.

Condition 3 $i < n$ and `Di` cannot move.

Condition 4 Both players have made n moves ($i = n$) and none of the conditions above holds.

We say that `Si` *wins* the game if the game finishes because of conditions 3 or 4, otherwise `Di` wins. By $G^f(\mathcal{A}, X_0, \mathcal{B}, Y_0, \infty)$ (for f such that $\operatorname{dom} f = \mathbb{N}$) we denote the game with an unbounded number of moves. In this case the game ends if Conditions 1, 2 or 3 are fired for some n. `Di` wins if the game ends under Condition 1 or 2. `Si` wins if Condition 3 holds, or if the game can continue indefinitely.

Two important characteristics of the definition of G^f are its directedness (in the rules and in Condition 1), and its use of sets, instead of elements, to represent positions. We will see that they are crucial in the following example.

EXAMPLE 6. To illustrate the definitions given so far we play $G^f(\mathcal{A}, \{a_1\}, \mathcal{B}, \{b_1\}, 1)$ with $\mathcal{A}$ and $\mathcal{B}$ as shown in Figure 1 for different values of f.

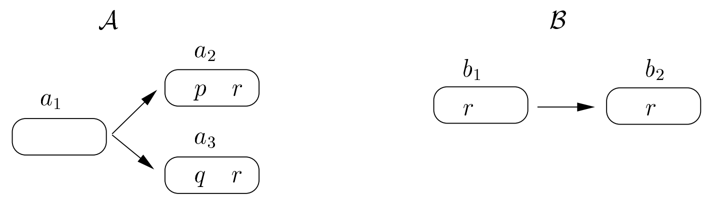

Figure 1. Playing a game

First take $f(1) = \Box$, then `Si` has a winning strategy: `Di` has to move in $\mathcal{B}$ with $Y' = \{b_2\}$. `Si` can choose either $\{a_1\}$, $\{a_2\}$ or $\{a_1, a_2\}$, and win the game. Note that all formulas $\Box\phi$, with ϕ local, that are satisfied in $(\mathcal{A}, a_1)$ are satisfied in $(\mathcal{B}, b_1)$. Now take $f(1) = \Diamond$; this time `Di` has a winning strategy: choose $\{a_2\}$ or $\{a_3\}$. In any of the two possibilities `Si` can only choose $Y' = \{b_2\}$ and in both cases there is a local formula, namely p or q such that $\mathcal{B}, \{b_2\} \not\models p$ or $\mathcal{B}, \{b_2\} \not\models q$, respectively.

The definition of G^f has been tailored to the restricted expressivity of the language $\mathcal{ML}^f$. We will now show that making the definition less tight would produce a mismatch in expressive power.

Suppose we weaken Condition 1 to make it symmetric, requiring that there is a formula $\phi \in$ LF such that either $\mathcal{A}, X_i \models \phi$ but $\mathcal{B}, Y_i \not\models \phi$; or $\mathcal{B}, Y_i \models \phi$ but $\mathcal{A}, X_i \not\models \phi$. Under this definition `Di` has a winning strategy for $G^f(\mathcal{A}, \{a_1\}, \mathcal{B}, \{b_1\}, k)$ for any f and k. This would be the case, in general, whenever two states disagree on the formulas in LF they satisfy. But this would be equivalent to allow atomic negation in LF! Under the new definition, the game G^f would be too discriminating for the expressive power of $\mathcal{ML}^f$. More generally, making Rules 1 to 3 symmetric would correspond to allowing full negation.

In a similar way, if we restrict positions to singleton sets (or equivalently, elements in the domain) the game would 'preserve' disjunctions, allowing `Di` to define a winning strategy on two models that only differ on disjunctive statements.

Games provide a mechanism for identifying differences between two models. Such differences may also be captured by logical formulas (in some language) that are true in one model but not in the other. The following theorem relates these two ideas for G^f-games and $\mathcal{ML}^f$-equivalence.

THEOREM 7. *Fix k and n such that $k \geq n > 0$, and let $f : \{1, \dots, k\} \rightarrow \{\Diamond, \Box, \boxtimes\}$ be given.*

1. *`Si` has a winning strategy for the game $G^f(\mathcal{A}, X, \mathcal{B}, Y, n)$ iff for every formula $\phi \in \mathcal{ML}^f_n$, if $\mathcal{A}, X \models \phi$ then $\mathcal{B}, Y \models \phi$.*

2. *`Di` has a winning strategy for the game $G^f(\mathcal{A}, X, \mathcal{B}, Y, n)$ iff there is a formula $\phi \in \mathcal{ML}^f_n$, such that $\mathcal{A}, X \models \phi$ and $\mathcal{B}, Y \not\models \phi$.*

Proof. *1.* ($\Rightarrow$) We will prove this direction using induction on $n \leq k$.

Assume that `Si` has a winning strategy for the game $G^f(\mathcal{A}, X, \mathcal{B}, Y, 0)$. The theorem says that all the formulas in LF that are satisfied in all the elements of X have to be satisfied in all the elements of Y. As `Si` has a winning strategy, Condition 3 or 4 should hold. Condition 3 does not apply and hence Condition 4 ensures the needed condition. Assume that the result holds for $G^f(\mathcal{A}, X, \mathcal{B}, Y, n)$. Let ϕ in $\mathcal{ML}^f_{n+1}$ such that $\mathcal{A}, X \models \phi$. We first consider the case $f(n+1) = \Diamond$. If ϕ is a formula in LF, the truth of ϕ in $\mathcal{B}$ is given by Rule 1. If ϕ is a conjunction $\phi_1 \wedge \phi_2$, we can use a second inductive argument (on the number of $\wedge$-signs) to establish the claim. Next, ϕ may be of the form $\Diamond\phi_1$ with $\phi_1 \in \mathcal{ML}^f_n$. Since $\mathcal{A}, X \models \Diamond\phi_1$, we have that every element $x \in X$ has an R-child x' such that $\mathcal{A}, x' \models \phi_1$. Let us play with `Di` choosing $X' = \{x' \mid \mathcal{A}, x' \models \phi_1\}$. Since `Si` has a winning strategy for $G^f(\mathcal{A}, X, \mathcal{B}, Y, n+1)$, she will counter-play with a set Y' such that $YR^{\uparrow}Y'$ and will still have a winning strategy for $G^f(\mathcal{A}, X, \mathcal{B}, Y, n)$. Using the induction hypothesis we have that $\mathcal{B}, Y' \models \phi_1$ since $\mathcal{A}, X' \models \phi_1$. Hence, $\mathcal{B}, Y \models \phi$.

Suppose $f(n+1) = \Box$. `Di` has to move in $\mathcal{B}$. Suppose that $\phi = \Box\phi_1$ and $\mathcal{B}, Y \not\models \phi$. Then there is a set Y' such that $YR^{\uparrow}Y'$ and $\mathcal{B}, Y' \not\models \phi_1$. Let this Y' be the set chosen by `Di`. Since `Si` has a wining strategy she will choose a set X' such that $XR^{\uparrow}X'$. X' will be such that $\mathcal{B}, X' \models \phi_1$. As `Si` has a winning strategy for the game $G^f(\mathcal{A}, X', \mathcal{B}, Y', n)$, from this and the inductive hypothesis we can conclude that X' and Y' do indeed satisfy the same set of formulas, a contradiction.

The case where $f(n+1) = \boxtimes$ reduces to one of the two cases above.

($\Leftarrow$) Assume that for every $\phi \in \mathcal{ML}^f_n$ we have that $\mathcal{A}, X \models \phi$ implies $\mathcal{B}, Y, \models \phi$. We need to define a winning strategy for `Si` when she plays the game $G^f(\mathcal{A},\ X,\ \mathcal{B},\ Y,\ n)$. The proof is by induction on n. For $n = 0$ we need to check that $G^f(\mathcal{A}, X, \mathcal{B}, Y, 0)$ starts with a winning position for `Si`. This is true because by hypothesis we have that formulas in LF valid in X are also valid in Y.

Assume the result holds for n. Suppose that $f(n+1) = \Diamond$ and that `Di` has chosen a set X' such that $XR^{\downarrow}X'$. Let $\Phi = \{\phi \in \mathcal{ML}^f_n \mid \mathcal{A}, X' \models \phi\}$. Let `Si` choose a set Y' such that $YR^{\downarrow}Y'$ and $\mathcal{B}, Y' \models \Phi$. The existence of such a set is given by hypothesis. After this move all the formulas in $\mathcal{ML}^f_n$ satisfied in X' are also satisfied in Y', and, by induction, `Si` can complete the strategy. For $f(n+1) = \Box$, `Di` has to move in $\mathcal{B}$. Let us suppose that `Di` has chosen a set Y'. Define Φ as before. Let X' be a set of elements in $\mathcal{A}$ such that $\mathcal{A}, X' \models \phi$. X' will be the move of `Si`— by using the induction hypothesis again we have the complete strategy.

The $f(n+1) = \Box\!\!\!\!\Diamond$ case reduces to one of the two cases above.

2. ($\Rightarrow$) The left-to-right implication is similar to item 1, left-to-right. To prove the right-to-left implication one can build the required strategy for `Di` by induction on the size of the formula ϕ.

If the formula is an atomic proposition letter then `Si` wins immediately because of Rule 1. For the inductive case, we should decide the move for `Di` in the current position and the induction hypothesis will provide the rest of the winning strategy. Since $\mathcal{B}, Y \not\models \phi$, there is an element $w' \in Y$ such that $\mathcal{B}, w' \not\models \phi$. Suppose that $\phi = \Box\phi_1$, since $\mathcal{B}, w' \not\models \Box\phi_1$ implies that there is w'_1 such that $w'Rw'_1$ and $\mathcal{B}, w'_1 \not\models \phi_1$. Following the rules of the game, `Di` has to move in $\mathcal{B}$: `Di` has to choose any subset of the children of Y such that w_1 is included. Suppose that $\phi = \Diamond\phi_1$, then there exists a set $YR^{\downarrow}Y'$ such that $\mathcal{B}, Y' \not\models \phi_1$. By inductive hypothesis, `Di` has a winning strategy for the game $G^f(\mathcal{A}, X', \mathcal{B}, Y', n)$. Y' as a first move together with the previous strategy, gives `Di` the complete strategy. ■

COROLLARY 8. *For every game $G^f(\mathcal{A}, X, \mathcal{B}, Y, n)$ either `Di` or `Si` has a winning strategy, i.e., the game $G^f(\mathcal{A}, X, \mathcal{B}, Y, n)$ is deterministic.*

6 The expressive power of $\mathcal{ML}^f$

Van Benthem [3] proved the following preservation result: a class of models defined by a first-order sentence is closed under bisimulations iff it can be defined by a modal formula. Rosen [18] proved that this result remains true over the class of finite structures. Kurtonina and de Rijke [13] extended Van Benthem's result in a different direction, by proving analogous preservation results for broad classes of description logics, including both restrictions and extension of the basic modal language such as $\mathcal{FL}^-$; see also [14].

Below we prove a general preservation result, for each of the fragments defined in Definition 2, over the class of finite structures. Our proof, which is based on the games introduced in the previous section, follows the structure of Rosen's proof.

For the formulation of our results it is convenient to work with so-called *pointed models* $(\mathcal{A}, c^{\mathcal{A}})$; these are models with a distinguished element.

DEFINITION 9. Let $\mathcal{A}$ and $\mathcal{B}$ two models with distinguished elements $c^{\mathcal{A}}$ and $c^{\mathcal{B}}$ respectively. We write $\mathcal{A} \sim^n_{Gf} \mathcal{B}$ to denote that `Si` has winning strategies for both of the following games: $G^f(\mathcal{A}, \{c^{\mathcal{A}}\}, \mathcal{B}, \{c^{\mathcal{B}}\}, n)$ and $G^f(\mathcal{B}, \{c^{\mathcal{B}}\}, \mathcal{A}, \{c^{\mathcal{A}}\}, n)$. We write $\mathcal{A} \sim^\infty_{Gf} \mathcal{B}$ to denote that `Si` has winning strategies for the corresponding infinite games.

The first key theorem in Rosen's paper is the following.

THEOREM 10 (Rosen [18]). *Let $\mathcal{C}$ be any class of models (each model $\mathcal{A}$ with a distinguished node $c^{\mathcal{A}}$), closed under isomorphism. Let $\mathcal{C}'$ be any subclass of $\mathcal{C}$, also closed under isomorphism. Then for all n, the following conditions are equivalent:*

1. *For all $\mathcal{A} \in \mathcal{C}'$, $\mathcal{B} \in \mathcal{C} - \mathcal{C}'$, $\mathcal{A} \not\sim^n \mathcal{B}$ (where $\sim^n$ stands for n-bisimulation).*

2. *There is a modal formula of quantifier rank less or equal than n that defines $\mathcal{C}'$ over $\mathcal{C}$.*

We extend this result to be able to cope with our restricted fragments.

THEOREM 11. *Let $\mathcal{C}$ be any class of models (each model $\mathcal{A}$ with a distinguished node $c^{\mathcal{A}}$), closed under isomorphism. Let $\mathcal{C}'$ be a subclass of $\mathcal{C}$ also closed under isomorphism. For every n, let M_n be the biggest subset of $\mathcal{ML}^f_n$ such that all the formulas in M_n are satisfied in $c^{\mathcal{A}}$ by all the models $\mathcal{A}$ in $\mathcal{C}'$. If $M_n \neq \emptyset$ then the following conditions are equivalent:*

1. *For all $\mathcal{A} \in \mathcal{C}'$, $\mathcal{B} \in \mathcal{C} - \mathcal{C}'$, $\mathcal{A} \not\sim^n_{Gf} \mathcal{B}$.*

2. *There is a formula ϕ in M_n such that ϕ defines $\mathcal{C}'$ over $\mathcal{C}$.*

Observe the following relations between Theorems 10 and 11. The set $\mathcal{ML}^f_n$ in Rosen's theorem is given by all modal formulas in which the maximal number of nested modal operators is at most n. The set M_n (of formulas satisfied in all elements of $\mathcal{C}'$) is never empty: the set of modal formulas containing at most n nested modal operators is logically finite, and every model in $\mathcal{C}'$ will satisfy ϕ or $\neg\phi$ for ϕ with modal depth less than or equal to n. Form the disjunction of one such formula ϕ per model in $\mathcal{C}$, and this formula will be true in all the elements of $\mathcal{C}'$.

Proof of Theorem 11. ($1 \Rightarrow 2$) Suppose that $M_n \neq \emptyset$ and suppose that for all $\mathcal{A} \in \mathcal{C}'$, $\mathcal{B} \in \mathcal{C} - \mathcal{C}'$, $\mathcal{A} \not\sim^n_{Gf} \mathcal{B}$. By Theorem 7 this implies that for all $\mathcal{A} \in \mathcal{C}'$, $\mathcal{B} \in \mathcal{C} - \mathcal{C}'$, there is a formula ϕ in M_n such that $\mathcal{A}, c^{\mathcal{A}} \models \phi$ but $\mathcal{B}, c^{\mathcal{B}} \not\models \phi$. Let $\mathcal{A}$ be any model in $\mathcal{C}'$, we define $\Phi_{\mathcal{A}}$ by putting

$$\Phi_{\mathcal{A}} = \bigwedge\{\phi \in \mathcal{ML}^f_n \mid \mathcal{A} \models \phi \text{ and } \mathcal{B} \not\models \phi \text{ with } \mathcal{B} \in \mathcal{C} - \mathcal{C}'\}.$$

Note that since $\mathcal{ML}^f_n$ is finite, $\Phi_{\mathcal{A}}$ is a finite conjunction. Note also that $\Phi_{\mathcal{A}}$ belongs to $\mathcal{ML}^f_n$ and that it is satisfied by all the models in $\mathcal{C}$ (it belongs to M_n) but not by any model in $\mathcal{C}'$. Hence $\Phi_{\mathcal{A}}$ is the needed definition.

($2 \Rightarrow 1$) Suppose that $M_n \neq \emptyset$ and that there is a formula ϕ in M_n such that ϕ defines $\mathcal{C}'$ over $\mathcal{C}$. By Theorem 7, `Di` will have a winning strategy for the appropriate games and for all $\mathcal{A} \in \mathcal{C}'$, $\mathcal{B} \in \mathcal{C} - \mathcal{C}'$, it will be true that $\mathcal{A} \not\sim^n_{Gf} \mathcal{B}$. ∎

We need some further terminology. Given a model $\mathcal{A}$ and a node w in $\mathcal{A}$, we say that w is a *descendant* of v if wR^*v, where R^* is the transitive closure of R. The *family* of w in $\mathcal{A}$, written $\mathcal{F}^w_{\mathcal{A}}$, is the submodel of $\mathcal{A}$ with universe $\{w\} \cup \{v \mid v \text{ is a descendant of } w\}$. We say that w and v are *disjoint* iff $\mathcal{F}^w_{\mathcal{A}} \cap \mathcal{F}^v_{\mathcal{A}} = \emptyset$. The *$r$-neighborhood* of a node w, denoted $\mathcal{N}_r(w)$, is defined inductively. $\mathcal{N}_0(w)$, is the submodel of $\mathcal{A}$ with universe $\{w\}$, and for all $r+1$, $v \in \mathcal{N}_{r+1}(w)$ iff $v \in \mathcal{N}_r(w)$ or there is a $w' \in \mathcal{N}_r(w)$ such that $\mathcal{A} \models Rw'v \vee Rvw'$. An *$r$-tree* is a directed tree rooted at u of height $\leq r$. An *r-pseudotree* is a model such that $\mathcal{N}_r(u)$ is a tree with the property that all distinct pairs of its leaves are disjoint, as defined above. As is standard, $\cong$ denotes isomorphism.

PROPOSITION 12. *Assume that $(\mathcal{A}, c^{\mathcal{A}})$ and $(\mathcal{B}, c^{\mathcal{B}})$ are two models with $(\mathcal{A}, c^{\mathcal{A}}) \sim^n_{Gf} (\mathcal{B}, c^{\mathcal{B}})$. Then there are n-pseudotrees $(\mathcal{A}', c^{\mathcal{A}'})$ and $(\mathcal{B}', c^{\mathcal{B}'})$ with $(\mathcal{A}, c^{\mathcal{A}}) \sim^\infty_{Gf} (\mathcal{A}', c^{\mathcal{A}'})$, $(\mathcal{B}, c^{\mathcal{B}}) \sim^\infty_{Gf} (\mathcal{B}', c^{\mathcal{B}'})$ and $\mathcal{N}_n(c^{\mathcal{A}'}) \cong \mathcal{N}_n(c^{\mathcal{B}'})$.*

Proof. We specify an algorithm that transforms the two pointed models into models with isomorphic n-neighborhoods. After each step s ($s \leq n$) we have models $(\mathcal{A}_s, c^{\mathcal{A}}_s)$ and $(\mathcal{B}_s, c^{\mathcal{B}}_s)$ such that $(\mathcal{A}, c^{\mathcal{A}}) \sim^\infty_{Gf} (\mathcal{A}_s, c^{\mathcal{A}}_s)$ and $(\mathcal{B}, c^{\mathcal{B}}) \sim^\infty_{Gf} (\mathcal{B}_s, c^{\mathcal{B}}_s)$ while $c^{\mathcal{A}_s}$ and $c^{\mathcal{B}_s}$ have isomorphic s neighborhoods. At each step $s+1$, $\mathcal{A}_{s+1}$ (respectively $\mathcal{B}_{s+1}$) is obtained from $\mathcal{A}_s$ ($\mathcal{B}_s$) by adding or removing copies of families of nodes at distance $s+1$ from the root.

Let $\{a_1, \ldots, a_l, b_1, \ldots, b_m\}$ be the set of children of $c^{\mathcal{A}}$ and $c^{\mathcal{B}}$. We will build the models using the following two rules: If $f(n) = \Box$ then for constructing $\mathcal{A}_1$ and $\mathcal{B}_1$ we just choose one a_i and b_j and drop all the remaining children. We will redefine the set of local formulas satisfied in a_i and b_j as the local formulas that are common to all states $\{a_1, \dots, a_l, b_1, \dots, b_m\}$.

All formulas in $\mathcal{ML}_n^f$ will either start with a box or will be local formulas. If a formula $\phi = \Box\phi_1$ is satisfied in $\mathcal{A}$ then ϕ_1 will be satisfied in all children of $c^{\mathcal{A}}$, and, in particular, in a_i, hence ϕ will be satisfied in $c^{\mathcal{A}_1}$.

If $f(n) = \Diamond$, the relation $\sim_{Gf}^{n-1}$ induces an equivalence classes on the set $\{a_1, \dots, a_l, b_1, \dots, b_m\}$. Note that not every equivalence class necessarily has a member in each $\mathcal{A}$ and $\mathcal{B}$. An example of such a configuration is as in (a) below:

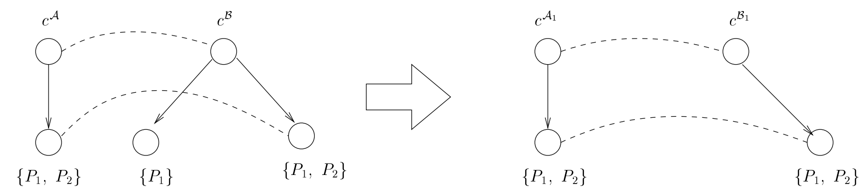

(a) (b)

To obtain $\mathcal{A}_1$ and $\mathcal{B}_1$ with isomorphic 1-neighborhoods of c such that $\mathcal{A} \sim_{Gf}^1 \mathcal{A}_1$, we have to do two things. First we should add enough copies of families of the children a_i and b_j such that each equivalence class has an equal number of members in $\mathcal{A}_1$ and in $\mathcal{B}_1$. Second, if an element d is not related by $\sim_{Gf}^{n-1}$ to an element in the opposite model, we just drop its family.

We should now verify that we are not throwing away any states that provide the only way to satisfy a certain formula. Suppose for contradiction that this is the case: there is a state b_i and a formula ϕ such that $\mathcal{B}_1, b_j \models \phi$ (hence $\mathcal{B}, c^{\mathcal{B}} \models \Diamond\phi$) and b_j is the only child of $c^{\mathcal{B}}$ that satisfies this formula. Since $\mathcal{A}, c^{\mathcal{A}} \models \Diamond\phi$ there is a child a_i of $c^{\mathcal{A}}$ such that $\mathcal{A}, a_i \models \phi$. By hypothesis a_i is not $\sim_{Gf}^{n-1}$-related to b_j, meaning that there is a formula ϕ_1 such that $\mathcal{A}, a_i \models \phi_1$ but $\mathcal{B}, b_j \not\models \phi_1$. As $\phi \wedge \phi_1 \in \mathcal{ML}^f$ there is another child of $c^{\mathcal{B}}$ that satisfies ϕ, a contradiction. The case where $f(n+1) = \Box\!\!\!\!\Diamond$ reduces to one of the two cases above. The next step on the algorithm is to move to each of the elements in the isomorphic neighborhoods and apply the same schema for each pair of nodes related by the isomorphism.

The models $\mathcal{A}_s$ and $\mathcal{B}_s$ constructed during the proof will be both isomorphic and $\sim_{Gf}^{\infty}$-related to $\mathcal{A}$ and $\mathcal{B}$, respectively, as needed. ■

Before we can formulate our main expressiveness result, we need one more auxiliary result, due to Rosen. In formulating it, we write $\mathcal{A} \equiv^n \mathcal{B}$ to denote that $\mathcal{A}$ and $\mathcal{B}$ satisfy the same first-order sentences with at most n nested quantifiers.

THEOREM 13 (Rosen [18]). *Let $\mathcal{A}$ and $\mathcal{B}$ be two $(2H(n))$-pseudotrees for which $\mathcal{N}_{H(n)}(c^{\mathcal{A}}) \cong \mathcal{N}_{H(n)}(c^{\mathcal{B}})$ holds, where $H(x)$ is the Hanf function. Then there are $\mathcal{A}'$ and $\mathcal{B}'$ such that $\mathcal{A} \sim^{\infty}_{G^f} \mathcal{A}'$, $\mathcal{B} \sim^{\infty}_{G^f} \mathcal{B}'$, and $\mathcal{A}' \equiv^n \mathcal{B}'$.*

Actually, Rosen used $\sim$ (bisimulation) in his theorem instead of $\sim^{\infty}_{G^f}$, but if two models are related by $\sim$ they will be related by $\sim^{\infty}_{G^f}$.

THEOREM 14. *Let $\mathcal{C}$ be a class of finite pointed models and $\mathcal{C}'$ a subclass of $\mathcal{C}$ such that the set of formulas $\Phi \subseteq \mathcal{ML}^f$ that are satisfied in all the models in $\mathcal{C}'$ is non-empty. Let $\mathcal{C}'$ be defined by a first-order formula α. If $\mathcal{C}'$ is closed under $\sim^{\infty}_{G^f}$, then $\mathcal{C}'$ is definable by a formula in $\mathcal{ML}^f$.*

Proof. Suppose that $\mathcal{C}'$ is defined by a first-order sentence and closed under $\sim^{\infty}_{G^f}$ but not definable by a formula in $\mathcal{ML}^f$. We want to prove that for all n there are pointed models $(\mathcal{A}, c^{\mathcal{A}}) \in \mathcal{C}'$ and $(\mathcal{B}, c^{\mathcal{B}}) \in \mathcal{C} - \mathcal{C}'$ such that $(\mathcal{A}, c^{\mathcal{A}}) \equiv^n (\mathcal{B}, c^{\mathcal{B}})$, which would contradict the hypothesis. For any $n \in \mathbb{N}$, Theorem 11 says that there are models $(\mathcal{A}', c^{\mathcal{A}'})$ and $(\mathcal{B}', c^{\mathcal{A}'})$ such that $(\mathcal{A}', c^{\mathcal{A}'}) \sim^{H(n)}_{G^f} (\mathcal{B}', c^{\mathcal{A}'})$. Proposition 12 then lets us construct models such that $(\mathcal{A}'', c^{\mathcal{A}''}) \sim^{\infty}_{G^f} (\mathcal{B}'', c^{\mathcal{A}''})$, and $\mathcal{N}_{H(n)}(c^{\mathcal{A}''}) \cong \mathcal{N}_{H(n)}(c^{\mathcal{B}''})$. Finally, we apply Theorem 13 to obtain the needed $(\mathcal{A}, c^{\mathcal{A}})$ and $(\mathcal{B}, c^{\mathcal{B}})$ and the contradiction. ∎

7 Extensions

In this section we discuss some possible extensions of $\mathcal{ML}^f$ for which Theorem 14 still holds. Such extensions involve two main issues: modifying $\mathcal{ML}^f$ and finding the corresponding game.

First of all, we can easily cater for atomic negations, simply by expanding the definition of local formulas to also include negations of proposition letters. In this case the game definition is not affected.

Second, adding disjunctions is also straightforward. The new $\mathcal{ML}^f(\vee)$-fragments are closely related to the description logic $\mathcal{FLU}^-$ [5], just like the original $\mathcal{ML}^f$-fragments are closely related to $\mathcal{FL}^-$. Their definitions are given by:

$$\begin{array}{rcl}
\mathcal{ML}^f_0(\vee) & = & \text{LF} \\
\mathcal{ML}^f_{n+1}(\vee) & = & \text{the closure under } \wedge \text{ and } \vee \text{ of (LF} \\
 & & \cup \{\Diamond\phi \mid \phi \in \mathcal{ML}^f_n(\vee) \text{ and } f(n+1) = \Diamond\} \\
 & & \cup \{\Box\phi \mid \phi \in \mathcal{ML}^f_n(\vee) \text{ and } f(n+1) = \Box\} \\
 & & \cup \{\Diamond\phi, \Box\phi \mid \phi \in \mathcal{ML}^f_n(\vee) \text{ and } f(n+1) = \Diamond\!\!\!\Box\}).
\end{array}$$

`Di` and `Si` will have to play using singletons. In other words, the first designated position in each model will be a singleton and both `Di` and

`Si` have to choose singletons in following moves. Note that once we have added the usual connectives and modal operators, $\sim_{Gf}^{\infty}$ is equivalent to bisimulation and Theorem 10 is actually equivalent to Theorem 11.

Another natural extension is to go multi-modal; this can be done in many different ways. The obvious one is to replace $\Box$ with $[R_i]$ in each of the production rules in Definition 2. This method will not control the modal depth at which a particular relation is used. Alternatively, we can let our functions f choose which subset of modal operators is to be considered legal at each level in the definition of $\mathcal{ML}^f$. Our main complexity, characterization, and expressiveness results hold for both ways of going multi-modal.

Finally, we can go a step further and allow unqualified number restrictions, thus moving to modal counterparts of fragments of the description logic $\mathcal{FLN}^-$ [5]. Recall that unqualified number restrictions are formulas of the form $\Diamond^{\leq n}\top$ that are true in a state w iff there are $w_1, \ldots, w_k$ with $k \leq n$ such that $wRw_1, \ldots, wRw_k$. Unqualified number restrictions are still very local, and because of that it is easy to extend our setup to deal with them: we can simply add them to the set of local formulas, and our characterization and expressiveness results will continue to hold.

8 Conclusions and future work

We have introduced a novel mechanism for decomposing modal logic into fragments. Each of these fragments can be specified in a very fine-grained manner, and for each of them we have defined a notion of game that allows us to characterize the fragment's expressive power. We have also provided uniform upper bounds for the complexity of the satisfiability problem for each fragment. Our games provide a natural tool to understand the fragments and the constructors they admit.

The first natural next step is to extend our fragment so as to also capture more expressive modal logics, especially ones with qualified number restrictions. We also aim to further explore how these fragments behave computationally: not only by providing better upper bounds for the complexity of the satisfiability problem, but also by considering different reasoning tasks, for example model checking.

Finally, in our characterization of the expressive power of the $\mathcal{ML}^f$-fragments we adapted a proof due to Rosen. Otto [15] has recently given a alternative proof of Rosen's result, of a far less combinatorial nature than Rosen's. It would be instructive to see what additional insights this proof yields when adapted to our fragments.

Acknowledgmentss

We thank the anonymous referees for their comments. Gabriel Infante-Lopez is supported by the Netherlands Organization for Scientific Research (NWO) under project number 220-80-001. Carlos Areces is supported by a grant from NWO under project number 612.069.006. Maarten de Rijke is supported by NWO under project numbers 612.013.001, 612.069.006, 365-20-005, 220-80-001, 612.000.106, and 612.000.207.

BIBLIOGRAPHY

[1] H. Andréka, J. van Benthem, and I. Németi. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217–274, 1998.

[2] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications.* Cambridge University Press, 2003.

[3] J. van Benthem. Correspondence theory. In D. Gabbay and F. Guenthner, editors, *Handbook of philosophical logic. Vol. II*, volume 165 of *Synthese Library*, pages 167–247. D. Reidel Publishing Co., Dordrecht, 1984. Extensions of classical logic.

[4] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic.* Cambridge University Press, 2001.

[5] R. Brachman and H. Levesque. The tractability of subsumption in frame-based description languages. In *Proc. of the 4th Nat. Conf. on Artificial Intelligence (AAAI-84)*, pages 34–37, 1984.

[6] R. Brachman and H. Levesque. Expressiveness and tractability in knowledge representation and reasoning. *Computational Intelligence*, 3:78–93, 1987.

[7] K. Doets. *Basic Model Theory.* CSLI Publications, 1997.

[8] F. Donini. *Complexity of Reasoning*, chapter 3, pages 101–141. In Baader et al. [2], 2003.

[9] E. Grädel. Why are modal logics so robustly decidable? *Bulletin EATCS*, 68:90–103, 1999.

[10] E. Grädel, P. Kolaitis, and M. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolc Logic*, 3:53–69, 1997.

[11] E. Grädel, M. Otto, and E. Rosen. Two-variable logic with counting is decidable. In *Proceedings of LICS 1997*, pages 306–317, 1997.

[12] J. Halpern. The effect of bounding the number of primitive propositions and the depth of nesting on the complexity of modal logic. *Artificial Intelligence*, 75:361–372, 1995.

[13] N. Kurtonina and M. de Rijke. Expressiveness of concept expressions in first-order description logics. *Artificial Intelligence*, 107:303–333, 1999.

[14] A. Nonnengart, H. Ohlbach, M. de Rijke, and D. Gabbay. *Encoding two-valued non-classical logics in classical logic*, pages 1403–1486. In Robinson and Voronkov [17], 2001.

[15] M. Otto. Modal and guarded characterisation theorems over finite transition systems. In *Proceedings LICS 2002*, 2002.

[16] L. Pacholsky, W. Szwast, and L. Tendera. Complexity of two-variable logic with counting. In *Proceedings of LICS 1997*, pages 318–327, 1997.

[17] A. Robinson and A. Voronkov, editors. *Handbook of Automated Reasoning.* Elsevier Science Publishers, 2001.

[18] E. Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6(4):427–439, 1997.

[19] M. Schmidt-Schauss and G. Smolka. Attributive concept descriptions with complements. *Artificial Intelligence*, 48(1):1–26, 1991.

[20] D. Scott. A decision method for validity of sentences in two variables. *Journal of Symbolic Logic*, 27:377, 1962.
[21] M. Vardi. Why is modal logic so robustly decidable? In *DIMACS Ser. Disc. Math. Theoret. Comp. Sci. 31*, pages 149–184. AMS, 1997.

Carlos Areces
Langue et Dialogue Group, INRIA Lorraine
615, rue du Jardin Botanique
54602 Villers lès Nancy Cedex, France
E-mail: carlos.areces@loria.fr

Gabriel G. Infante-Lopez, Maarten de Rijke
Language & Inference Technology Group
ILLC, U. of Amsterdam
Nieuwe Achtergracht 166
1018 WV Amsterdam, The Netherlands
E-mail: {infante,mdr}@science.uva.nl