Branching-Time Temporal Logics with Minimal Model Quantifiers*

Fabio Mogavero and Aniello Murano

Universitá degli Studi di Napoli "Federico II", 80126 Napoli, Italy.
{mogavero, murano}@na.infn.it

Abstract. Temporal logics are a well investigated formalism for the specification and verification of reactive systems. Using formal verification techniques, we can ensure the correctness of a system with respect to its desired behavior (specification), by verifying whether a model of the system satisfies a temporal logic formula modeling the specification.
From a practical point of view, a very challenging issue in using temporal logic in formal verification is to come out with techniques that automatically allow to select small critical parts of the system to be successively verified. Another challenging issue is to extend the expressiveness of classical temporal logics, in order to model more complex specifications.
In this paper, we address both issues by extending the classical branching-time temporal logic CTL* with minimal model quantifiers (MCTL*). These quantifiers allow to extract, from a model, minimal submodels on which we check the specification (also given by an MCTL* formula). We show that MCTL* is strictly more expressive than CTL*. Nevertheless, we prove that the model checking problem for MCTL* remains decidable and in particular in PSPACE. Moreover, differently from CTL*, we show that MCTL* does not have the tree model property, is not bisimulation-invariant and is sensible to unwinding. As far as the satisfiability concerns, we prove that MCTL* is highly undecidable. We further investigate the model checking and satisfiability problems for MCTL* sublogics, such as MPML, MCTL, and MCTL$^+$, for which we obtain interesting results. Among the others, we show that MPML retains the finite model property and the decidability of the satisfiability problem.

1 Introduction

Temporal logics, which are a special kind of *modal logics* geared towards the description of the temporal ordering of events [Pnu77], have been adopted as a powerful tool for specifying and verifying correctness of concurrent systems [Pnu81], as they allow to express the temporal ongoing behavior of a system in a well-structured way.

Two possible views regarding the nature of time induce two different types of temporal logics: *linear* and *branching-time* [Lam80]. In linear-time temporal logics, such as LTL [Pnu77], time is treated as if each moment in time has a unique possible future. Thus, linear temporal logic formulas are interpreted over linear sequences. In branching-time temporal logics, such as CTL [CE81], CTL$^+$, and CTL* [EH85], each

* Work partially supported by MIUR PRIN Project no.2007-9E5KM8.

moment in time may split into various possible futures. Accordingly, the structures over which branching temporal logic formulas are interpreted are infinite trees. Many important parallel computer programs exhibit ongoing behavior that is characterized naturally in terms of infinite execution traces, possibly organized into tree-like structures that reflect the high degree of nondeterminism inherent in parallel computation.

In formal system design, one of the most significant developments has been the discovery of algorithmic methods for verifying temporal-logic properties of finite-state systems [CE81, QS82]. In temporal-logic model checking, we verify the correctness of a finite-state system with respect to a desired behavior by checking whether a labeled state-transition graph, called *Kripke structure*, that models the system satisfies a temporal logic formula that specifies this behavior. Hence, the name *model checking* for the verification method derived from this viewpoint. Since model checking has many practical applications (see [Eme90] for more motivations and background) it is important to classify temporal logics according to the computational complexity of their model checking problem. Indeed, the complexity for branching-time temporal logics is well understood: for CTL, CTL$^+$, and CTL* it is PTIME-COMPLETE, Δ_2^p-COMPLETE, and PSPACE-COMPLETE, respectively.

From a practical point of view, a very challenging issue in using temporal logics in formal specification and verification is to come out with automatic techniques that allow to select small critical parts of the system in order to restrict system verification to them. This necessity is mainly due to the fact that in a concurrent setting, the system under consideration is typically a parallel composition of many modules. Promising approaches to restrict the verification techniques to subsystems of interest are assume guarantee techniques [AL93], modular model checking [KV95, KV97], the exploitation of partial order information [Pel96], localization reduction [Kur94], and semantic minimization for eliminate unnecessary states from a system model [ECJB97]. Note that all these approaches have in common the fact that the modularity of the system is known in advance. Another important issue in system design and verification is to look for new temporal logics that are more expressive than the classical ones. In fact, although CTL* is a very powerful logic, there are several important but complex properties that require a more powerful framework. To overcome this limitation, several attempts have been carried out in literature in order to extend these logics by introducing appropriate semantics or operators usually guided by embedded contexts [AHK02, BLMV06, BMM09].

In this paper, we address both the above issues by introducing the branching-time temporal logic MCTL*. This logic is an extension of the classical branching-time temporal logic CTL* with minimal model quantifiers, which allow to extract, given a model, minimal and conservative submodels of it on which we successively check a given property. The goal is to check local properties of system components in order to deduce the global behavior of the entire one. Therefore, the introduced logic exploits the novel idea of checking a particular module of a whole composition system while its single modules are not known in advance. In more details, MCTL* extends CTL* by also allowing two special (*minimal model*) quantifiers: Λ and Ξ. These quantifiers allow to write state formulas such as $\varphi_1 \Lambda \varphi_2$ and $\varphi_1 \Xi \varphi_2$, which respectively read as *"all minimal and conservative models of* φ_2 *are models of* φ_1*"* and *"there exists a minimal model of* φ_2

that is model of φ_1*"*, for suitable and well-found concepts of minimality and conservativeness among Kripke structures. In accordance with this point of view, we call φ_2 the *submodel extractor*, φ_1 the *submodel verifier*, and our modular verification method an *extract-verify* paradigm. Our choice of considering only minimal and conservative submodels is justified by the fact that in this way we precisely select the parts of the system that are actually responsible for the particular behavior of interest. For an example of an application of the introduced logic see Example 1 in Section 3. It is worth recalling that logics having the ability to modify the model under evaluation (and then check the specification on the resulting part) have been also considered in other contexts. For example, we recall the arbitrary public announcement logic [FvD08] and the sabotage modal logic [LR03]. However, the first allow to extract, according to a submodel extractor formula, submodels that do not necessarily satisfy the formula itself, and the second does not extract submodels using a formula at all.

In this paper, we investigate MCTL* and its sublogics MCTL$^+$, MCTL and MPML (where M indicates the extension of the respective logics with minimal model quantifiers) from a theoretical point of view. As far as the expressivity regards, we show that all these logics are strictly more expressive than the corresponding classical ones. Unfortunately, this power comes at a price. Indeed, we show that the satisfiability for MCTL is highly undecidable. Moreover and differently from CTL, we have that introduced logics neither have the tree model property nor are bisimulation-invariant, while they all are sensible to unwinding. We also investigate succinctness and the model checking problem for the introduced logics, from which we got interesting results. Among the others, we show that MCTL is as succinct as MCTL$^+$(differently from the classical case of CTL and CTL$^+$). Moreover, as CTL$^+$ [LMS01], both MCTL and MCTL$^+$ have a Δ_2^p-COMPLETE (i.e., PTIME$^{\text{NPTIME}}$) model checking. As far as we know, our result provides the second example, after CTL$^+$, of Δ_2^p-COMPLETE problems in the field of formal verification. Since for this class very few complete problems are known, we believe that the obtained result is interesting as its own. Finally, we show that the propositional modal logic (PML) augmented with minimal model quantifiers (MPML) retains both the finite model property and the decidability of the satisfiability problem.

2 Preliminaries

Given a *set* X of *objects* (numbers, words, etc.), we denote by $|\mathrm{X}|$ its cardinality, called *size* of X, and by 2^{X} the *powerset* of X. As special sets, we consider $\mathbb{Z}$, $\mathbb{N}$, and $\mathbb{N}_+ = \mathbb{N} \setminus \{0\}$, as respectively, the sets of *relative*, *natural*, and *positive natural numbers*.

A *Kripke structure* $\mathcal{K} = \langle \mathrm{AP}, \mathrm{W}, \mathrm{R}, \mathrm{L} \rangle$ is an ordered tuple, where AP is a set of *atomic propositions*, $\mathrm{W} = \mathsf{dom}(\mathcal{K})$ is a non-empty set of *worlds*, $\mathrm{R} \subseteq \mathrm{W} \times \mathrm{W}$ is a *binary relation*, and $\mathrm{L} : \mathrm{W} \mapsto 2^{\mathrm{AP}}$ is the labeling function that maps each world to a set of atomic propositions true in that world. We denote the size $|\mathcal{K}|$ of $\mathcal{K}$ by $|\mathrm{W}| + |\mathrm{R}|$. An infinite Kripke structure is a structure of infinite size. Now, let $\mathcal{K}' = \langle \mathrm{AP}', \mathrm{W}', \mathrm{R}', \mathrm{L}' \rangle$ be another Kripke structure. We say that $\mathcal{K}'$ is a *substructure* of $\mathcal{K}$, in symbols $\mathcal{K}' \preccurlyeq \mathcal{K}$, iff *(i)* $\mathrm{AP}' \subseteq \mathrm{AP}$, *(ii)* $\mathrm{W}' \subseteq \mathrm{W}$, *(iii)* $\mathrm{R}' \subseteq \mathrm{R} \cap (\mathrm{W}' \times \mathrm{W}')$, and *(iv)* for all $w \in \mathrm{W}'$, it holds that $\mathrm{L}'(w) = \mathrm{L}(w) \cap \mathrm{AP}'$. Moreover, we say that $\mathcal{K}$ and $\mathcal{K}'$ are *comparable* iff *(i)* $\mathcal{K} \preccurlyeq \mathcal{K}'$ or *(ii)* $\mathcal{K}' \preccurlyeq \mathcal{K}$ holds, otherwise they are *incomparable*. For a set of structures $\mathfrak{S}$, we define the set of *minimal substructures* $\mathsf{minstructs}(\mathfrak{S})$ as the set consisting of the $\preccurlyeq$-

minimal elements of $\mathfrak{S}$. I.e., it is the set containing all and only the structures $\mathcal{K} \in \mathfrak{S}$ such that for all $\mathcal{K}' \in \mathfrak{S}$, it holds that *(i)* $\mathcal{K} \preccurlyeq \mathcal{K}'$, or *(ii)* $\mathcal{K}'$ is not comparable with $\mathcal{K}$. Note that all structures in $\mathsf{minstructs}(\mathfrak{S})$ are incomparable among them. A structure $\mathcal{K}$ is *minimal* w.r.t. a set $\mathfrak{S}$ (or simply minimal, when the context clarify the set $\mathfrak{S}$) iff $\mathcal{K} \in \mathsf{minstructs}(\mathfrak{S})$. A set of structures $\mathfrak{S}$ is minimal iff $\mathfrak{S} = \mathsf{minstructs}(\mathfrak{S})$.

For sake of space, all other classical concepts of *tree*, *path*, *set of maximal paths* $\mathsf{paths}(\mathcal{K}, w)$ of a structure $\mathcal{K}$ starting in a world $w \in \mathsf{dom}(\mathcal{K})$, and unwinding $\mathcal{U}_w^{\mathcal{K}}$ of $\mathcal{K}$ in w, are omitted (see [KVW00] for detailed definitions).

3 The Minimal Model Quantifiers temporal logic extensions

In this section, we introduce an extension of the classical branching-time temporal logic CTL* with minimal model quantifiers, which allow to extract minimal submodels on which we successively check a given property. To formally define the extended logic, we use the CTL* state and path formulas framework.

The *full computation tree logic with minimal model quantifiers* (MCTL*, for short) extends CTL* by further using two special quantifiers, the universal Λ and the existential Ξ ones. Informally, a model satisfies a state formula $\varphi_1 \Lambda \varphi_2$ iff all its minimal and conservative submodels satisfying φ_2 (φ_2 is the *submodel extractor*) are also models satisfying φ_1 (φ_1 is the *submodel verifier*). As in CTL*, in MCTL* the two path quantifiers A and E can prefix a linear time formula composed by an arbitrary combination and nesting of the four linear temporal operators X (*"effective next"*), $\tilde{\mathsf{X}}$ (*"hypothetical next"*), U (*"until"*), and R (*"release"*). The formal syntax of MCTL* follows.

Definition 1. (Syntax) MCTL* state *(φ) and* path *(ψ)* formulas *are built inductively from* AP *using the following context-free grammar, where* $p \in \mathrm{AP}$*:*

1. $\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \Lambda \varphi \mid \varphi \Xi \varphi \mid \mathsf{A}\psi \mid \mathsf{E}\psi$,
2. $\psi ::= \varphi \mid \neg\psi \mid \psi \wedge \psi \mid \psi \vee \psi \mid \mathsf{X}\psi \mid \tilde{\mathsf{X}}\psi \mid \psi \mathsf{U} \psi \mid \psi \mathsf{R} \psi$.

The class of MCTL* *formulas is the set of state formulas generated by the above grammar. In addition, the simpler classes of* MCTL$^+$, MCTL, *and* MPML *formulas are obtained, respectively, by avoiding nesting of temporal operators, by forcing each temporal operator occurring into a formula to be coupled with a path quantifier, and by excluding from* MCTL *path formulas the until and release operators.* □

The *length* $|\varphi|$ of a formula φ is defined inductively on the structure of φ in the classical way, and by also considering $|\varphi_1 \Lambda \varphi_2|$ and $|\varphi_1 \Xi \varphi_2|$ to be equal to $1 + |\varphi_1| + |\varphi_2|$.

We now define the semantics of MCTL* w.r.t. a Kripke structure $\mathcal{K}$. For a world $w \in \mathsf{dom}(\mathcal{K})$, we write $\mathcal{K}, w \models \varphi$ to indicate that a state formula φ holds at w, and, for a path $\pi \in \mathsf{paths}(\mathcal{K})$, we write $\mathcal{K}, \pi, k \models \psi$ to indicate that a path formula ψ holds on π at position $0 \leq k < |\pi|$. Note that, the relation $\mathcal{K}, \pi, k \models \psi$ does not hold for any point $k \in \mathbb{N}$, with $k \geq |\pi|$. The semantics of state and path formulas involving $\neg$, $\wedge$, and $\vee$, the classical path quantifiers E and A, and the classical temporal operators is defined as usual in CTL*. Here we only give the semantics of the remaining part.

Definition 2. (Semantics) *Given a Kripke structure* $\mathcal{K} = \langle \mathrm{AP}, \mathrm{W}, \mathrm{R}, \mathrm{L} \rangle$*, a world* $w \in \mathrm{W}$*, and two state formulas* φ_1 *and* φ_2 *it holds:*

1. $\mathcal{K}, w \models \varphi_1 \Lambda \varphi_2$ *iff for all* $\mathcal{K}' \in \mathsf{minstructs}(\mathfrak{S}(\mathcal{K}, w, \varphi_2))$ *it holds that* $\mathcal{K}', w \models \varphi_1$*;*
2. $\mathcal{K}, w \models \varphi_1 \Xi \varphi_2$ *iff there is a* $\mathcal{K}' \in \mathsf{minstructs}(\mathfrak{S}(\mathcal{K}, w, \varphi_2))$ *such that* $\mathcal{K}', w \models \varphi_1$*;*

where $\mathfrak{S}(\mathcal{K}, w, \varphi) = \{\mathcal{K}' \preccurlyeq \mathcal{K} \mid w \in \mathsf{dom}(\mathcal{K}') \wedge \forall \mathcal{K}'' \preccurlyeq \mathcal{K} : \mathcal{K}' \preccurlyeq \mathcal{K}'' \rightarrow \mathcal{K}'', w \models \varphi\}$

It is clear that, MCTL* (resp., MPML, MCTL, and MCTL$^+$) formulas without minimal model quantifiers are CTL* (resp., PML, CTL, and CTL$^+$) formulas.

Let $\mathcal{K}$ be a Kripke structure and φ a MCTL* formula. Then, $\mathcal{K}$ is a *model* for φ, denoting this by $\mathcal{K} \models \varphi$, iff there is $w \in \mathsf{dom}(\mathcal{K})$ such that $\mathcal{K}, w \models \varphi$. In this case, we also say that $\mathcal{K}$ is a model for φ on w. A MCTL* formula φ is said *satisfiable* iff there exists a model for it, moreover it is *invariant* on two Kripke structures $\mathcal{K}$ and $\mathcal{K}'$ iff either $\mathcal{K} \models \varphi$ and $\mathcal{K}' \models \varphi$ or $\mathcal{K} \not\models \varphi$ and $\mathcal{K}' \not\models \varphi$, i.e., $\mathcal{K}$ and $\mathcal{K}'$ agree on φ.

A Kripke structure $\mathcal{K}$ is conservative w.r.t. a formula φ iff, for all models $\mathcal{K}'$ extending $\mathcal{K}$, i.e., with $\mathcal{K} \preccurlyeq \mathcal{K}'$, it holds that $\mathcal{K}' \models \varphi$. Note that this concept of conservativeness is automatically embedded in the definition of $\mathfrak{S}(\mathcal{K}, w, \varphi)$, since we consider only models $\mathcal{K}' \in \mathfrak{S}(\mathcal{K}, w, \varphi)$ that, if extended, continue to satisfy the formula φ. To better understanding the meaning and the importance of the conservativeness, consider the Kripke structure $\mathcal{K}$ built by a chain of three states $w_0 \rightarrow w_1 \rightarrow w_2$, in which the final state w_2 is the only labeled with p. Moreover, consider the two submodels $\mathcal{K}'$ and $\mathcal{K}''$ of $\mathcal{K}$ built, respectively, by w_0 and $w_0 \rightarrow w_1$. Clearly $\mathcal{K}' \preccurlyeq \mathcal{K}'' \preccurlyeq \mathcal{K}$ and, for $\varphi = \mathsf{EX}\,\mathsf{t} \rightarrow \mathsf{EF}\,p$, we have that $\mathcal{K}' \models \varphi$, $\mathcal{K}'' \not\models \varphi$, and $\mathcal{K} \models \varphi$. Hence, we have that $\mathcal{K}'$ satisfies φ, but it is not conservative, since $\mathcal{K}''$ (that extend $\mathcal{K}'$) does not satisfy φ.

For all state formulas φ_1 and φ_2 (resp., path formulas ψ_1 and ψ_2), we say that φ_1 is *equivalent* to φ_2, formally $\varphi_1 \equiv \varphi_2$, (resp., ψ_1 is *equivalent* to ψ_2, formally $\psi_1 \equiv \psi_2$) iff for all Kripke structures $\mathcal{K}$ and worlds $w \in \mathsf{dom}(\mathcal{K})$, it holds that $\mathcal{K}, w \models \varphi_1$ iff $\mathcal{K}, w \models \varphi_2$ (resp., for all paths $\pi \in \mathsf{paths}(\mathcal{K}, w)$, it holds that $\mathcal{K}, \pi, 0 \models \psi_1$ iff $\mathcal{K}, \pi, 0 \models \psi_2$).

In the rest of the paper, we mainly consider formulas in *existential normal form* or in *positive normal form*, i.e., formulas in which only existential (minimal model and path) quantifiers occur or negation is applied only to atomic propositions, respectively. In fact, it is to this aim that we have considered in the syntax of MCTL* both the connectives $\wedge$ and $\vee$, the quantifiers Λ, Ξ, A and E, and the dual operators $\tilde{\mathsf{X}}$ and R. Indeed, all formulas can be converted in existential or positive normal form by using De Morgan's laws and the following equivalences, which directly follow from the semantics of the logic. Let φ_1 and φ_2 be state formulas and ψ, ψ_1, and ψ_2 be path formulas, then it holds that $\neg(\varphi_1 \Lambda \varphi_2) \equiv \neg\varphi_1 \Xi \varphi_2$, $\neg \mathsf{A}\psi \equiv \mathsf{E}\neg\psi$, $\neg \mathsf{X}\psi \equiv \tilde{\mathsf{X}}\neg\psi$, and $\neg(\psi_1 \mathsf{U} \psi_2) \equiv \neg\psi_1 \mathsf{R} \neg\psi_2$. In order to abbreviate writing formulas, we also use the boolean values t (*"true"*) and f (*"false"*) and the path temporal operators $\mathsf{F}\,\psi \equiv \mathsf{t}\,\mathsf{U}\,\psi$ (*"future"*) and $\mathsf{G}\,\psi \equiv \mathsf{f}\,\mathsf{R}\,\psi$ (*"globally"*). Moreover, note that the following equivalences also hold: $\mathsf{E}(\psi_1 \vee \psi_2) \equiv \mathsf{E}\psi_1 \vee \mathsf{E}\psi_2$, $\tilde{\mathsf{X}}\,\psi \equiv \tilde{\mathsf{X}}\,\mathsf{f} \vee \mathsf{X}\,\psi$, $\psi_1 \mathsf{R}\,\psi_2 \equiv (\psi_2 \mathsf{U}\,(\psi_1 \wedge \psi_2)) \vee \mathsf{G}\,\psi_2$, $\mathsf{X}\,(\psi_1 \wedge \psi_2) \equiv \mathsf{X}\,\psi_1 \wedge \mathsf{X}\,\psi_2$, and $\mathsf{G}\,(\psi_1 \wedge \psi_2) \equiv \mathsf{G}\,\psi_1 \wedge \mathsf{G}\,\psi_2$.

Example 1. ***(Arbiter system)*** Consider an arbiter system used to control a two-users access to a shared memory location (see Figure 1 for a model $\mathcal{K}$ of it), where only the resquest (`r`) and the acknowledge (`a`) signals are known. Suppose now that we want to verify that the idle state i and the common request state $\langle r_1, r_2 \rangle$ are unique w.r.t. the order of user request or arbiter acknowledge. We can perform this check by applying MCTL* model checking in the state i using a formula $\varphi = \varphi_1 \wedge \varphi_2$, where $\varphi_1 = \mathsf{AG}\,(i \rightarrow$

$\mathsf{X}\,\mathfrak{t}) \wedge \mathsf{E}(\mathsf{F}\,(a_1 \wedge \mathsf{X}\,\mathsf{F}\,i) \wedge \mathsf{F}\,(a_2 \wedge \mathsf{X}\,\mathsf{F}\,i))$ checks if the "acknowledge subsystem" reaches the same idle state and $\varphi_2 = \mathsf{AG}\,(r_1 \wedge r_2 \rightarrow \mathsf{X}\,\mathfrak{t}) \wedge (\mathsf{EX}\,(r_1 \wedge \mathsf{X}\,(r_2 \wedge \mathsf{X}\,\mathfrak{t})) \wedge \mathsf{EX}\,(r_2 \wedge \mathsf{X}\,r_1))$ checks if the common request state reached by the "request subsystem" is unique. For two minimal and conservative submodels of φ_1 and φ_2 in $\mathcal{K}$ see $\mathcal{K}_1$ and $\mathcal{K}_2$ in Figure 1. Note that also their "mirror images" are submodels of φ_1 and φ_2.

One may note that the above check can not be achieved using a classical logic such as CTL*. Indeed, we may have a bisimilar model of $\mathcal{K}$, with more idle or common request states, in which no CTL* formula can check that these states are not unique. □

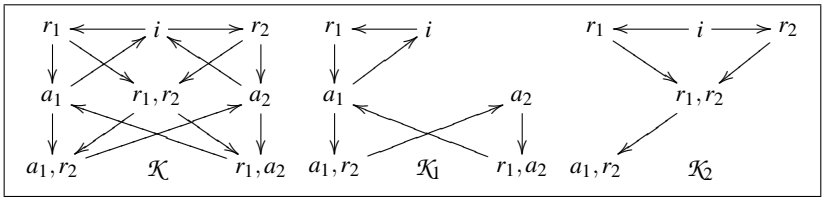

Fig. 1. A model of an arbiter system for shared memory locations and its submodels.

By means of counterexamples, we show that the introduced extended logics are more expressive than the corresponding classical ones. Indeed, since they can distinguish among models that are invariant for the classical logics, as described in the previous example. The result is reported in the following theorem.

Theorem 1. *For* MPML, MCTL, MCTL$^+$, *and* MCTL* *it holds that they* (i) *do not have the tree model property;* (ii) *are neither invariant under unwinding nor under partial unwinding;* (iii) *are not invariant under bisimulation; and* (iv) *are more expressive than* PML, CTL, CTL$^+$, *and* CTL*, *respectively.*

Proof. Item (i) To prove this item, we consider a formula with an existential minimal model quantifier such that it requires to extract a graph submodel, which can not be a tree, in order to be satisfied. Consider the MPML formula $\varphi = \varphi_1 \,\Xi\, \varphi_2$, where $\varphi_1 = \mathsf{EX}(\beta \wedge \mathsf{EX}\,\mathsf{EX}\,\gamma)$, $\varphi_2 = \alpha \wedge \mathsf{EX}(\beta \wedge \mathsf{EX}\,\delta) \wedge \mathsf{EX}(\gamma \wedge \mathsf{EX}(\delta \wedge \mathsf{EX}\,\gamma))$, $\alpha = a \wedge b$, $\beta = \neg a \wedge b$, $\gamma = a \wedge \neg b$, and $\delta = \neg a \wedge \neg b$. This formula is satisfiable. In Figure 2 we show $\mathcal{K}_1$, $\mathcal{K}_2$, $\mathcal{K}_3$, and $\mathcal{K}_4$ as the only minimal models of φ_2, where only $\mathcal{K}_1$ is a tree and $\mathcal{K}_3$ and $\mathcal{K}_4$ are the only models of φ. Indeed, only $\mathcal{K}_3$ and $\mathcal{K}_4$ satisfy φ_1. Since any model of φ must include $\mathcal{K}_3$ or $\mathcal{K}_4$ as submodel, it follows that no tree model can satisfy φ. Since MPML is a sublogic of MCTL, MCTL$^+$, and MCTL* the thesis easily follows.

Item (ii) Suppose by contradiction that MPML is invariant under unwinding. Then, for each satisfiable formula φ, since there exists a model $\mathcal{K}$ such that $\mathcal{K}, w \models \varphi$, it holds that $\mathcal{U}^{\mathcal{K}}_w, w \models \varphi$. But $\mathcal{U}^{\mathcal{K}}_w$ is a tree, so each satisfiable formula has a tree model, but this contradicts the previous item. To prove that this logic is also not invariant under partial unwinding, consider the model $\mathcal{K}$, built by a single world w with a loop relation on it, and its "one step" unwinding $\mathcal{K}'$, formed by two worlds, w and v, linked together by a relation and with a relation loop on the second one (see Figure 2). Moreover, consider, the MPML formula $\varphi = (\mathsf{EX}\,\mathsf{EX}\,\mathfrak{t})\,\Xi\,(\mathsf{EX}\,\mathfrak{t})$. It holds that $\mathfrak{S}(\mathcal{K}, w, \mathsf{EX}\,\mathfrak{t}) = \{\mathcal{K}\}$, while $\mathfrak{S}(\mathcal{K}', w, \mathsf{EX}\,\mathfrak{t}) = \{\mathcal{K}', \mathcal{K}''\}$, with $\mathcal{K}'' \preccurlyeq \mathcal{K}'$ and where $\mathcal{K}''$ is equal to $\mathcal{K}'$ once the loop on the last node is removed. It is easy to verify that $\mathcal{K} \models \mathsf{EX}\,\mathsf{EX}\,\mathfrak{t}$, but $\mathcal{K}'' \not\models \mathsf{EX}\,\mathsf{EX}\,\mathfrak{t}$,

so we have that only $\mathcal{K}$ is a model of φ. This shows that MPML is able to distinguish between a model and one of its partial unwindings, thus the thesis follows.

Item (iii) Since an unwinding is a particular case of a bisimilarity relation, we have also that MPML is not bisimilar, i.e., it is possible to express a MPML property satisfied on a model $\mathcal{K}$, but not on a bisimilar model $\mathcal{K}'$ of $\mathcal{K}$.

Item (iv) This item follows from the fact that PML, CTL, CTL$^+$, and CTL* are invariant under bisimulation, while their extensions with minimal model quantifiers are not. Therefore, the extended logics can characterize more models than the classical ones and thus they are more expressive. □

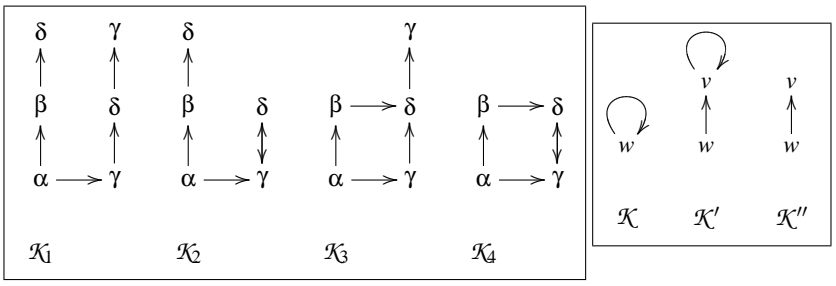

Fig. 2. Minimal models of φ_2 in item *i* and models of $\mathsf{EX}\,\mathsf{t}$ in item *ii* of Theorem 1.

We now show that MPML has the strong finite model property. To this aim, we first introduce some extra notations. For a Kripke structure $\mathcal{K}$, by $\mathsf{dep}(\mathcal{K})$ we denote the maximal length of a path in the unwinding $\mathcal{U}^w_{\mathcal{K}}$, for all worlds $w \in \mathsf{dom}(\mathcal{K})$. Moreover, given a MPML formula φ, we denote by $\mathsf{dep}(\varphi)$ the *depth* of φ, i.e., the maximal number of nested occurrences of path quantifiers in φ, but those appearing in its submodel verifiers. Formally, the depth function is inductively defined as follows: $\mathsf{dep}(p) = 0$, for $p \in \mathrm{AP}$; $\mathsf{dep}(\neg\varphi) = \mathsf{dep}(\varphi)$; $\mathsf{dep}(\varphi_1 \wedge \varphi_2) = \mathsf{dep}(\varphi_1 \vee \varphi_2) = \max\{\mathsf{dep}(\varphi_1), \mathsf{dep}(\varphi_2)\}$; $\mathsf{dep}(\varphi_1 \,\Lambda\, \varphi_2) = \mathsf{dep}(\varphi_1 \,\Xi\, \varphi_2) = \mathsf{dep}(\varphi_2)$; $\mathsf{dep}(\mathsf{A}\psi) = \mathsf{dep}(\mathsf{E}\psi) = 1 + \mathsf{dep}(\varphi)$, where $\psi = \mathsf{X}\varphi$ or $\psi = \tilde{\mathsf{X}}\varphi$. It is easy to see that $\mathsf{dep}(\varphi) = O(|\varphi|)$.

Theorem 2. MPML *has the strong finite model property, i.e., each* MPML *satisfiable formula* φ *has a finite model* $\mathcal{K}$ *with size* $|\mathcal{K}| \leq g(|\varphi|)$*, where* g *is a recursive function, and depth* $\mathsf{dep}(\mathcal{K}) \leq \mathsf{dep}(\varphi)$.

In [EH85] it is shown that CTL$^+$ is equivalent to CTL by using an exponential blow-up translation. Also, in [Wil99] it is shown that this blow-up is unavoidable. In the next theorem, we show that MCTL and MCTL$^+$ are polynomially equivalent and then, as an immediate corollary, we obtain that MCTL is exponentially more succinct than CTL.

Theorem 3. MCTL *is polynomially equivalent to* MCTL$^+$.

Proof. (Sketch.) Given a MCTL$^+$ formula φ we show that there exists a MCTL formula φ' equivalent to φ such that $|\varphi'| = O(|\varphi|^3)$.

W.l.o.g we assume that φ is in existential normal form (we recall that any MCTL$^+$ formula can be linearly translated in this form). Moreover, due to classical formula equivalences [EH85], we can also assume that φ has one E quantifier by recursively

applying the translation algorithm to nested subformulas containing an E. So we can assume that φ is of the form $\mathsf{E}\psi$ and ψ is a Boolean combination of subformulas of the form $\varphi'_i \mathsf{U} \varphi''_i$, $\mathsf{G}\varphi_1$, $\mathsf{X}\varphi_2$, and $\tilde{\mathsf{X}}\mathfrak{f}$, where each φ'_i, φ''_i, φ_1 and φ_2 are MCTL formulas (found by recursive applications of the translation algorithm). In practice, this turns out to use, as base case of the translation idea, the four equivalences listed below:

i) $\mathsf{E}(\bigwedge_{i=1}^{n} \varphi'_i \mathsf{U} \varphi''_i \wedge \mathsf{G}\varphi_1 \wedge \tilde{\mathsf{X}}\mathfrak{f}) \equiv \bigwedge_{i=1}^{n} \varphi''_i \wedge \varphi_1 \wedge \mathsf{E}\tilde{\mathsf{X}}\mathfrak{f}$;
ii) $\mathsf{E}(\mathsf{G}\varphi_1 \wedge \mathsf{X}\varphi_2) \equiv \varphi_1 \wedge \mathsf{EX}(\varphi_2 \wedge \mathsf{EG}\varphi_1)$;
iii) $\mathsf{E}(\bigwedge_{i=1}^{n} \varphi'_i \mathsf{U} \varphi''_i \wedge \mathsf{G}\varphi_1) \equiv \bigvee_{i=1}^{n} (f'_i \,\Xi\, f''_i)$;
iv) $\mathsf{E}(\bigwedge_{i=1}^{n} \varphi'_i \mathsf{U} \varphi''_i \wedge \mathsf{G}\varphi_1 \wedge \mathsf{X}\varphi_2) \equiv \bigwedge_{i=1}^{n} \varphi''_i \wedge \varphi_1 \wedge \mathsf{EX}(\varphi_2 \wedge \mathsf{EG}\varphi_1) \vee \bigvee_{i=1}^{n} (f'_i \,\Xi\, f''_i)$;

where $f'_i = \bigwedge_{1 \le h < k \le n}^{h,k \neq i} (\mathsf{EF}(\varphi''_h \wedge \mathsf{EF}\varphi''_k) \vee \mathsf{EF}(\varphi''_k \wedge \mathsf{EF}\varphi''_h))$ and $f''_i = \mathsf{E}((\varphi'_i \wedge \varphi_1) \mathsf{U} (\varphi''_i \wedge \mathsf{EG}\varphi_1)) \wedge \bigwedge_{j=1; j\neq i}^{n} \overline{\mathsf{E}}(\varphi'_j \mathsf{U} (\varphi''_j \wedge \mathsf{EF}\varphi''_i))$.

The first two equivalences, which do not contain the minimal model quantifier Ξ, are derivable by simply applying classical transformations. The proof of the last two, instead, can be obtained by formally showing (by induction) that each model satisfying the first member of an equivalence must satisfy also the second one and vice versa. Here, we omit this part for the sake of space, while we give an intuition of the third equivalence, which shows, as in the fourth one, how to avoid the exponential blow-up incurred by the classical translation in CTL for the corresponding case.

The key step in the translation is the selection of the right submodel of the extractor formula f''_i, through the verifier formula f'_i, which must satisfy $\varphi = \mathsf{E}(\bigwedge_{i=1}^{n} \varphi'_i \mathsf{U} \varphi''_i \wedge \mathsf{G}\varphi_1)$. If a model $\mathcal{K}$ satisfies φ in a world $w \in \mathsf{dom}(\mathcal{K})$, for all $\mathcal{K}' \in \mathsf{minstructs}(\mathfrak{S}(\mathcal{K}, w, \varphi))$ it holds that $\mathcal{K}' \in \mathsf{minstructs}(\mathfrak{S}(\mathcal{K}, w, f''_i))$, for a given index i. Moreover, for all paths $\pi \in \mathsf{paths}(\mathcal{K}', w)$ such that $\mathcal{K}', \pi, 0 \models \bigwedge_{i=1}^{n} \varphi'_i \mathsf{U} \varphi''_i$, we have that $\mathcal{K}', \pi, 0 \models \mathsf{F}(\varphi''_h \wedge \mathsf{F}\varphi''_k)$ or $\mathcal{K}', \pi, 0 \models \mathsf{F}(\varphi''_k \wedge \mathsf{F}\varphi''_h)$, for all indexes h and k, with $h < k$ and $h, k \neq i$. Hence, it holds that $\mathcal{K}', w \models f'_i$ and then $\mathcal{K}, w \models f'_i \,\Xi\, f''_i$.

Vice versa, consider a model $\mathcal{K}$ such that $\mathcal{K}, w \models f'_i \,\Xi\, f''_i$, for a given index i. Then, it holds that there exists a minimal model $\mathcal{K}' \in \mathsf{minstructs}(\mathfrak{S}(\mathcal{K}, w, f''_i))$ such that $\mathcal{K}', w \models f'_i$. Now, suppose by contradiction that $\mathcal{K}', w \not\models \varphi$. Then there exist at least three different and not directly connected parts $\mathcal{K}'_1$, $\mathcal{K}'_2$, and $\mathcal{K}'_3$, with $\mathcal{K}'_1, \mathcal{K}'_2, \mathcal{K}'_3 \preccurlyeq \mathcal{K}'$, and three paths $\pi_1 \in \mathsf{paths}(\mathcal{K}'_1, w)$, $\pi_2 \in \mathsf{paths}(\mathcal{K}'_2, w)$, and $\pi_3 \in \mathsf{paths}(\mathcal{K}'_3, w)$, such that each path formula $\varphi'_j \mathsf{U} \varphi''_j$, with $j \neq i$, is satisfied on just two of these paths. Then, each formula $\mathsf{E}(\varphi'_j \mathsf{U} (\varphi''_j \wedge \mathsf{EF}\varphi''_i))$ is satisfied in at least two ways in two different submodels of $\mathcal{K}'$ and then there exists a submodel $\mathcal{K}'' \preccurlyeq \mathcal{K}'$, $\mathcal{K}'' \neq \mathcal{K}'$ such that $\mathcal{K}'', w \models f''_i$. So, $\mathcal{K}'$ is not minimal, but this contradicts the assumption. Hence, $\mathcal{K}', w \models \varphi$.

Finally, note that it is fundamental that minimal model quantifiers are conservative. Otherwise, we could have a model $\mathcal{K}$ such that $\mathcal{K}, w \models \mathfrak{t} \,\Xi\, \mathsf{EG}\varphi_1$ even if $\mathcal{K}, w \not\models \mathsf{EG}\varphi_1$. This means that in the above discussion, we could have $\mathcal{K}, w \not\models \varphi$, since there are no paths satisfying $\mathsf{G}\varphi_1$, but $\mathcal{K}, w \models f'_i \,\Xi\, f''_i$, for some i. □

Corollary 1. MCTL *is exponentially more succinct than* CTL.

4 Model checking

In this section, we solve the model checking for the introduced logics, showing that the considered extract-verify paradigm retains the decidability of this problem.

We start with a lemma that shows how to calculate a polynomial certificate for particular MCTL and MCTL* formulas. This result will be then useful to show the corresponding upper bound results for the addressed model checking problems.

Lemma 1. *Let $\mathcal{K}$ be a Kripke structure, $w \in \mathsf{dom}(\mathcal{K})$ be a world and $\varphi = \varphi_1 \,\Xi\, \varphi_2$ be a* MCTL *(resp.,* MCTL**) formula, with φ_1 and φ_2* CTL *(resp.,* CTL**) formulas. Then, there exists a polynomial certificate $\mathcal{K}'$ of the testing $\mathcal{K}, w \models \varphi$, which is verifiable in* PTIME *(resp.,* PSPACE*).*

Proof. To check that the test $\mathcal{K}, w \models \varphi$ is in NPTIME (resp., in PSPACE), we verify that there exists a minimal and conservative submodel $\mathcal{K}'$ of $\mathcal{K}$ (the certificate of the test) of polynomial size (since $|\mathcal{K}'| \leq |\mathcal{K}|$), satisfying φ_2 in w, such that $\mathcal{K}', w \models \varphi_1$. To this aim we split the verification procedure into the following four phases: *(i)* testing of $\mathcal{K}', w \models \varphi_2$, *(ii)* checking the minimality of $\mathcal{K}'$, *(iii)* checking the conservativeness for $\mathcal{K}'$, and *(iv)* testing of $\mathcal{K}', w \models \varphi_1$. The first and last items are easily achievable in PTIME (resp., in PSPACE) by applying a classical CTL (resp., CTL*) model checking algorithm. Instead, to verify that $\mathcal{K}'$ is minimal w.r.t. the formula φ_2, we check that for all maximal and proper submodels $\mathcal{K}''$ of $\mathcal{K}'$, it holds that $\mathcal{K}'', w \not\models \varphi_2$. Now, note that all models $\mathcal{K}''$ are in number $O(|\mathcal{K}'|)$ since each of them is obtained by removing only one component from $\mathcal{K}'$. So, we deduce that also the check for minimality can be done in PTIME (resp., in PSPACE). Finally, it remains to verify whether $\mathcal{K}'$ is conservative, i.e., for all models $\mathcal{K}''$, with $\mathcal{K}' \preccurlyeq \mathcal{K}''$, it holds $\mathcal{K}'', w \models \varphi_2$. To do this, we can check that, for all subformula φ' of φ_2 and for all worlds $w \in \mathsf{dom}(\mathcal{K}')$, it holds that $\mathcal{K}', w \models \varphi'$ iff $\mathcal{K}, w \models \varphi'$. Since the number of all subformulas φ' is polynomial in the size of φ_2, and thus in the size of φ, it follows that also the check for conservativeness is in PTIME (resp., in PSPACE). To sum up, we have that to verify a certificate for the test $\mathcal{K}, w \models \varphi$ is in PTIME (resp., PSPACE), and therefore we have done with the proof. □

Using the above result, we are now able to prove the following two theorems.

Theorem 4. MCTL* *has a* PSPACE-COMPLETE *model checking problem.*

Proof. For the lower bound, we recall that for CTL*, which is a sublogic of MCTL*, the model checking problem is already PSPACE-HARD. We now proceed with the upper bound. To this aim, let $\mathcal{K}$ be a Kripke structure and φ an MCTL* in existential normal form, we construct a recursive algorithm that checks in PSPACE whether $\mathcal{K}, w \models \varphi$.

First of all, we enumerate all subformulas $\overline{\varphi} = \varphi' \,\Xi\, \varphi''$ of φ and we associate to each of them a fresh different atomic proposition. More formally, suppose that we have a sequence $(\overline{\varphi_1}, \ldots, \overline{\varphi_n})$ of such subformulas, then we associate to each $\overline{\varphi_i}$ the proposition ep_i. Now, consider Θ_m as the set of all formulas $\overline{\varphi} = \varphi' \,\Xi\, \varphi''$ subformulas of φ such that φ' contains just m and φ'' contains at most m nested occurrences of the Ξ quantifier, or vice versa. Also, consider Θ'_m as the set of formulas $\tilde{\varphi}$ obtained from each $\overline{\varphi} \in \Theta_m$ by replacing every occurrence of a minimal model quantifier, but the most external one, with the relative atomic proposition. Note now that, for all m, each $\tilde{\varphi} \in \Theta'_m$ is a MCTL* formula of the type $\tilde{\varphi} = \varphi' \,\Xi\, \varphi''$, with φ' and φ'' CTL* formulas and that $|\Theta'_m| = O(|\varphi|)$.

Now, set $\mathcal{K}_0 = \mathcal{K}$, we construct a sequence of Kripke structures $\mathcal{K}_m$ such that, for all $\tilde{\varphi} \in \Theta_m$ and $w \in \mathsf{dom}(\mathcal{K})$, $ep_i \in \mathrm{L}_m(w)$ iff $\mathcal{K}_{m-1}, w \models \tilde{\varphi}$, where ep_i is the atomic proposition relative to $\overline{\varphi}$. The latter can be checked by applying the PSPACE procedure of Lemma 1. Then, the result follows by recursively applying the above procedure. □

Theorem 5. MCTL *and* MCTL$^+$ *have a* Δ_2^p-COMPLETE *model checking problem.*

Proof. First note that, since by Theorem 3 MCTL and MCTL$^+$ are one into the other polynomial reducible, MCTL$^+$ simply extends CTL$^+$with minimal model quantifiers, and the latter has a Δ_2^p-COMPLETE model checking problem [LMS01], we have that MCTL has a Δ_2^p-HARD model checking. What remains to show is that it is in Δ_2^p. To prove this, we use a variation of the deterministic algorithm of Theorem 4, which, instead to call a PSPACE procedure to know if $\mathcal{K}, w \models \varphi_1 \,\Xi\, \varphi_2$ or not, call a NPTIME oracle (in accordance with Lemma 1), which solve the check in a single step. Now, since all other instructions of the algorithm are based on a classical CTL model checking procedure that can be executed in PTIME, we easily obtain a Δ_2^p model checking procedure for MCTL. Moreover, by Theorem 3, it holds that a MCTL$^+$ formula can be polynomially translated into a MCTL one, so the thesis follows also for this logic. □

Corollary 2. MPML *has a* Δ_2^p *model checking problem.*

5 Satisfiability

In this section we study the satisfiability for the introduced logics. We show that for all of them, but MPML, the question is undecidable. For MPML, we show decidability by using a brute force procedure via strong finite model property [BdRV04]. The result is reported in the following theorem.

Theorem 6. *The satisfiability problem for* MPML *is decidable.*

Proof. By Theorem 2, MPML has the strong finite model property w.r.t. a precise recursive function g. So for a given MPML formula φ we can construct a non deterministic Turing machine that, once the value of the function $g(|\varphi|)$ is computed, it guesses a model $\mathcal{K}$ of size at most equal to this value and then checks if it satisfies the formula by applying the decidable model checking procedure given by Corollary 2. Since φ is satisfiable iff it is satisfied on a model $\mathcal{K}$ of size at most $g(|\varphi|)$ and the built machine systematically examines all these kinds of models, the thesis easily follows. □

In rest of this section, we show undecidability of the satisfiability problem for MCTL, MCTL$^+$, and MCTL*through a reduction of a *domino problem* to it. This approach, often used in undecidability proofs in logic (see for example [BS99]), is classically known as *"undecidability via tiling"* [BdRV04].

The well-known domino problem, proposed for the first time by Wang [Wan61], consists of placing a given number of tile types on an infinite grid, satisfying a predetermined set of constraints on adjacent tiles. Its standard version asks for a compatible tiling of the whole plane $\mathbb{Z} \times \mathbb{Z}$. However, as stated by Knuth [Knu68], a compatible tiling of the first quadrant yields compatible tilings of arbitrary large finite rectangles, which in turn yields a compatible tiling of the whole plane. Since the existence of a solution for the original problem is known to be Π_0^1-COMPLETE [Ber66, Rob71], we have undecidable results (Π_0^1-HARD) also for the above variants of the classical domino problem. A formal definition of the $\mathbb{N} \times \mathbb{N}$ tiling problem follows.

Definition 3. (Tiling System) *A* $\mathbb{N}\times\mathbb{N}$ tiling system $\mathcal{D}=(\mathcal{T},\mathcal{H},\mathcal{V})$ *is a structure built on a non-empty set* $\mathcal{T}$ *of* domino types *and two* horizontal *and* vertical matching pairs $\mathcal{H},\mathcal{V}\subseteq\mathcal{T}^2$. *The domino problem asks for a compatible tiling of the first quadrant* $(\mathbb{N}\times\mathbb{N})$ *of the plane, which is a* solution mapping $\tau:\mathbb{N}^2\mapsto\mathcal{T}$ *such that, for all* $x,y\in\mathbb{N}$ *with* $\tau(x,y)=t$, *it holds that if* $\tau(x+1,y)=t'$, *then* $(t,t')\in\mathcal{H}$, *and if* $\tau(x,y+1)=t'$, *then* $(t,t')\in\mathcal{V}$. □

In the literature, an extension of the above problem has been also introduced as the *recurrent domino problem*. This problem, in addition to the tiling of the semiplane $\mathbb{N}\times\mathbb{N}$, asks whether there exists a distinguished tile type that occurs infinitely often in the first row of the grid. This problem is known to be more complex of the classical one. Indeed, it turns to be Σ^1_1-COMPLETE [Har84]. The formal definition follows.

Definition 4. (Recurrent Tiling System) *A* $\mathbb{N}\times\mathbb{N}$ recurrent tiling system $\mathcal{RD}=(\mathcal{T},\mathcal{H},\mathcal{V},t^*)$ *is a structure in which* $\mathcal{D}=(\mathcal{T},\mathcal{H},\mathcal{V})$ *is a* $\mathbb{N}\times\mathbb{N}$ *tiling system and* $t^*\in\mathcal{T}$ *is a distinguished tile type. The recurrent domino problem asks for a solution mapping* τ *such that the set of horizontal index* $\{m\mid\tau(m,0)=t^*\}$ *has an infinite cardinality.* □

By showing a reduction from the recurrent domino problem, we can prove in particular that the satisfiability for the MCTL logic is Σ^1_1-HARD. We achieve this reduction by showing that a given recurrent tiling system $\mathcal{RD}$ can be *"embedded"* into a model $\mathcal{K}_{\mathcal{RD}}$ of a particular formula $\varphi_{\mathcal{RD}}$ in such a way that $\varphi_{\mathcal{RD}}$ is satisfiable (i.e., it has a model) if and only if $\mathcal{RD}$ allows for a compatible tiling. To this aim we extend the proof structure used by Baader and Sattler [BS99]. For the sake of clarity, we split the reduction into four tasks, as described as follows:

Task 1 - (*Grid Specification*): It is possible to represent a "square structure" of $\mathbb{N}\times\mathbb{N}$, which consists of points (x,y), $(x+1,y)$, $(x,y+1)$, and $(x+1,y+1)$, in order to yield a complete covering of the semi-plane via a repeating regular grid structure. The basic idea is to use the minimal model quantifiers to force the horizontal successor of $(x,y+1)$ and the vertical successor of $(x+1,y)$ to correspond to the unique point $(x+1,y+1)$, with the aim to represent a square structure model on which to place the domino types. Formally, this can be expressed by using the following formula φ_{GS}, with $\alpha=a\wedge b$, $\beta=\neg a\wedge b$, $\gamma=a\wedge\neg b$, and $\delta=\neg a\wedge\neg b$:

$$\begin{aligned}
\varphi_H(\varphi') &= (\alpha\rightarrow \mathsf{EX}(\gamma\wedge\varphi'))\wedge(\beta\rightarrow \mathsf{EX}(\delta\wedge\varphi'))\wedge(\gamma\rightarrow \mathsf{EX}(\alpha\wedge\varphi'))\wedge(\delta\rightarrow \mathsf{EX}(\beta\wedge\varphi'));\\
\varphi_V(\varphi') &= (\alpha\rightarrow \mathsf{EX}(\beta\wedge\varphi'))\wedge(\beta\rightarrow \mathsf{EX}(\alpha\wedge\varphi'))\wedge(\gamma\rightarrow \mathsf{EX}(\delta\wedge\varphi'))\wedge(\delta\rightarrow \mathsf{EX}(\gamma\wedge\varphi'));\\
\varphi_S &= \varphi_V(\varphi_H(\varphi_V(t)))\,\Xi\,(\varphi_V(\varphi_H(t))\wedge\varphi_H(\varphi_V(\varphi_V(t))));\\
\varphi_{U_H} &= \varphi_H(\varphi_H(t)\wedge\varphi_V(t))\,\Lambda\,(\varphi_H(\varphi_H(t))\wedge\varphi_H(\varphi_V(t)));\\
\varphi_{U_V} &= \varphi_V(\varphi_H(t)\wedge\varphi_V(t))\,\Lambda\,(\varphi_V(\varphi_H(t))\wedge\varphi_V(\varphi_V(t)));\\
\varphi_A &= ((\alpha\vee\delta)\rightarrow \mathsf{AX}(\beta\vee\gamma))\wedge((\beta\vee\gamma)\rightarrow \mathsf{AX}(\alpha\vee\delta));\\
\varphi_{GS} &= \varphi_S\wedge\varphi_{U_H}\wedge\varphi_{U_V}\wedge\varphi_A.
\end{aligned}$$

Task 2 - (*Compatible Tiling*): It is possible to express that a tiling is locally compatible, i.e., the two horizontal $(x+1,y)$ and vertical $(x,y+1)$ points have admissible domino types with respect to the (x,y) point. The idea here is to associate to each domino type $t\in\mathcal{T}$ an atomic proposition T_t and express the horizontal and vertical matching conditions via suitable object labeling. Note that these constraints are

very easy to express. Indeed, they can be expressed in PML. Formally, we have:
$\varphi_{CT} = \bigvee_{i \in \mathcal{T}} (T_i \wedge \bigwedge_{j \neq i, j \in \mathcal{T}} \neg T_j \wedge \varphi_H(\bigvee_{(i,j) \in \mathcal{H}} T_j) \wedge \varphi_V(\bigvee_{(i,j) \in \mathcal{V}} T_j))$.

Task 3 - (*Recurrent Tile*): It is possible to assert that the distinguished tile type t^* occurs infinitely often on the first row of the semi-plane. This task can be easily achieved by using the kind of recursion available in the basic logic CTL. By means of this recursion, we can impose that the relative atomic proposition T_{t^*} is satisfied in an infinite number of worlds $v \in \mathsf{dom}(\mathcal{K}_{\mathcal{RD}})$, linearly reachable from the origin $w \in \mathsf{dom}(\mathcal{K}_{\mathcal{RD}})$ of the grid. Formally, we have:
$\varphi_{RT} = \varphi_V(\neg\varepsilon) \wedge (\varepsilon \rightarrow \varphi_H(\mathsf{EF}(\varepsilon \wedge T_{t^*})))$.

Task 4 - (*Global Reachability*): Finally, it is possible to impose that the above three conditions hold on all points in $\mathbb{N} \times \mathbb{N}$. As for the recurrent tile condition, also this task can be achieved by the simple recursion given by CTL. Formally, we have:
$\varphi_{GR} = \mathsf{AG}(\varphi_{GS} \wedge \varphi_{CT} \wedge \varphi_{RT})$.

We now give a formal proof of the undecidability, introducing the formula $\varphi_{\mathcal{RD}}$ who assemble all the above concepts.

Theorem 7. *The satisfiability problem for* MCTL, MCTL$^+$, *and* MCTL* *is highly undecidable. In particular, it is* Σ^1_1-HARD.

Proof. To prove the undecidability of the logic, we show the equivalence between find the solution of the recurrent tiling problem, with the distinguished tile type t^*, and the satisfiability of the formula $\varphi_{\mathcal{RD}} = \alpha \wedge \varepsilon \wedge \varphi_{GR}$.

Assume, for the direct reduction, that there exists a solution mapping τ. Then, we can build a satisfying model $\mathcal{K} = \langle \mathrm{AP}, \mathrm{W}, \mathrm{R}, \mathrm{L} \rangle$ that satisfies $\varphi_{\mathcal{RD}}$ as follows:

- $\mathrm{AP} = \{a, b, \varepsilon\} \cup \{T_t \mid t \in \mathcal{T}\}$;
- $\mathrm{W} = \mathbb{N} \times \mathbb{N}$;
- $\mathrm{R} = \{((m,n),(m+1,n)) \mid m,n \in \mathbb{N}\} \cup \{((m,n),(m,n+1)) \mid m,n \in \mathbb{N}\}$;
- for all $m,n \in \mathbb{N}$, it holds that: $a \in \mathrm{L}((m,n))$ if $n \equiv 0 \pmod 2$, $b \in \mathrm{L}((m,n))$ if $m \equiv 0 \pmod 2$, $\varepsilon \in \mathrm{L}((0,0))$ and $\varepsilon \in \mathrm{L}((m,0))$ if $\tau(m,0) = t^*$, and $T_t \in \mathrm{L}((m,n))$ if $\tau(m,n) = t$.

It is easy to see that $\mathcal{K} \models \varphi_{\mathcal{RD}}$, since $\mathcal{K}, (0,0) \models \varphi_{\mathcal{RD}}$.

Conversely, let $\mathcal{K}$ be a model such that there exists a world $w \in \mathsf{dom}(\mathcal{K})$ such that $\mathcal{K}, w \models \varphi_{\mathcal{RD}}$. First, we show that $\mathcal{K}$ is a *grid-like model* and then that is possible to construct a solution mapping τ from it. Indeed, since $\mathcal{K}, w \models \varphi_{\mathcal{RD}}$, we have that for all worlds $v \in \mathsf{dom}(\mathcal{K})$ reachable from w, (i.e., $(w,v) \in \mathrm{R}^n$, for some $n \in \mathbb{N}$) it holds that $\mathcal{K}, v \models \varphi_{GS}$ and thus $\mathcal{K}, v \models \varphi_S$. Now, it is not difficult to see that $\mathcal{K}$ must contain a square submodel in v (see proof of item *(i)* of Theorem 1 and structures $\mathcal{K}_3$ and $\mathcal{K}_4$ in Figure 1 for an example of models of φ_{GS}, where also holds that $\mathcal{K}_i, v \models \alpha$, for $i \in \{3,4\}$). Moreover, $\mathcal{K}, v \models \varphi_A$, so there are only two kinds of successors for v, i.e., if $\mathcal{K}, v \models \alpha$ or $\mathcal{K}, v \models \delta$ then v has successor worlds u, with $\mathcal{K}, u \models \beta$ or $\mathcal{K}, u \models \gamma$ and vice versa. Finally, since $\mathcal{K}, v \models \varphi_{U_H} \wedge \varphi_{U_V}$, if $\mathcal{K}, v \models \alpha$ or $\mathcal{K}, v \models \delta$, v has only one successor u' with $\mathcal{K}, u' \models \beta$ and only one successor u'' with $\mathcal{K}, u'' \models \gamma$ and viceversa. Now, it is clear that each world v reachable from w (including w itself) has only two successors u' and u'', which have a common successor o. Hence, $\mathcal{K}$ is a *grid-like model*. To extract a solution mapping τ from $\mathcal{K}$ is a routine task, so left to the reader. □

6 Conclusions

In this paper, we have introduced the branching-time temporal logic MCTL* as an extension of the classical branching-time temporal logic CTL* with minimal model quantifiers. These quantifiers allow to extract minimal submodels of a system model (even when the modularity of the system is not known in advance) on which we successively check a given property of the introduced logic.

We have deeply investigated, from a theoretical point of view, MCTL* and some of its sublogics. As far as the expressivity regards, we have showed that MCTL* is strictly more expressive than CTL*. Unfortunately, this power comes at a price. Indeed, the satisfiability problem for MCTL*, as well as for its sublogic MCTL, has been proved to be highly undecidable. Moreover, MCTL* does not have the tree model property, it is not bisimulation-invariant, and it is sensible to unwinding, opposed to CTL*.

As good news, we have showed that the sublogic MPML of MCTL* retains both the finite model property and the decidability of the satisfiability problem. Moreover, we have showed that the model checking problem for MCTL* remains decidable and in PSPACE. In more details and differently from CTL*, the PSPACE upper bound we provide is both in the size of the system and in the size of specification. Since for CTL it is only PSPACE in the size of the formula, it is left as an open question whether this extra complexity can be avoided. Anyway, although practical applications of MCTL* are not in the target of this paper, we argue that the extra blow-up for MCTL*should not have any consequence in practical applications as it can be absorbed in classical symbolic model checking algorithms, which are already exponential. Last but not least, we have investigated succinctness and the model checking problems for MCTL$^+$ and MCTL. We have shown that, differently from the classical case of CTL and CTL$^+$, MCTL is as succinct as MCTL$^+$. Moreover, as for CTL$^+$, the model checking problem for MCTL and MCTL$^+$ is Δ_2^p-COMPLETE (i.e., PTIME$^{\text{NPTIME}}$).

As future work, it would be worth investigating if the bisimulation-invariant fragment of MCTL* (i.e., the set of formulae that agree on bisimilar Kripke structures) is equally expressive as CTL*. In other words, we would like to check whether there exists an MCTL* formula which does not distinguish bisimilar structures, but it is still not expressible in CTL*. Then, if such a fragment is not equivalent to CTL*, it would be also relevant to investigate the related decidability problems.

Acknowledgement. We wish to thank Moshe Y. Vardi for useful discussions and the DLT 2009 referees for many helpful comments and suggestions.

References

[AHK02] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-Time Temporal Logic. *JACM*, 49(5):672–713, 2002.

[AL93] M. Abadi and L. Lamport. Composing Specifications. *TOPLAS*, 15(1):73–132, 1993.

[BdRV04] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic.* Cambridge University Press, 2004.

[Ber66] R. Berger. The Undecidability of the Domino Problem. *MAMS*, 66:1–72, 1966.

[BLMV06] P.A. Bonatti, C. Lutz, A. Murano, and M.Y. Vardi. The Complexity of Enriched μ-Calculi. In *ICALP'06*, LNCS 4052, pages 540–551. Springer-Verlag, 2006.

[BMM09] Alessandro Bianco, Fabio Mogavero, and Aniello Murano. Graded Computation Tree Logic. In *LICS'09*, To appear., 2009.

[BS99] F. Baader and U. Sattler. Expressive Number Restrictions in Description Logics. *JLC*, 9(3):319–350, 1999.

[CE81] E.M. Clarke and E.A. Emerson. Design and Synthesis of Synchronization Skeletons Using Branching-Time Temporal Logic. In *LP'81*, LNCS 131, pages 52–71. Springer-Verlag, 1981.

[ECJB97] W.M. Elseaidy, R. Cleaveland, and Jr. J.W. Baugh. Modeling and Verifying Active Structural Control Systems. *SCP*, 29(1-2):99–122, 1997.

[EH85] E.A. Emerson and J.Y. Halpern. Decision Procedures and Expressiveness in the Temporal Logic of Branching Time. *JCSS*, 30(1):1–24, 1985.

[Eme90] E. A. Emerson. Temporal and Modal Logic. In *Handbook of Theoretical Computer Science, Volume B: Formal Models and Sematics (B)*, pages 995–1072. 1990.

[FvD08] T. French and H.P. van Ditmarsch. Undecidability for Arbitrary Public Announcement Logic. In *AIML*, pages 23–42, 2008.

[Har84] D. Harel. A Simple Highly Undecidable Domino Problem. In *CLC'84*, 1984.

[Knu68] D.E. Knuth. *The Art of Computer Programming, Volume I: Fundamental Algorithms.* Addison-Wesley, 1968.

[Kur94] R.P. Kurshan. The Complexity of Verification. In *STOC'94*, pages 365–371, 1994.

[KV95] O. Kupferman and M.Y. Vardi. On the Complexity of Branching Modular Model Checking. In *CONCUR'95*, LNCS 962, pages 408–422. Springer-Verlag, 1995.

[KV97] O. Kupferman and M.Y. Vardi. Modular Model Checking. In *COMPOS'97*, LNCS 1536, pages 381–401. Springer-Verlag, 1997.

[KVW00] O. Kupferman, M.Y. Vardi, and P. Wolper. An Automata-Theoretic Approach to Branching-Time Model Checking. *JACM*, 47(2):312–360, 2000.

[Lam80] L. Lamport. "Sometime" is Sometimes "Not Never": On the Temporal Logic of Programs. In *POPL'80*, pages 174–185, 1980.

[LMS01] F. Laroussinie, N. Markey, and P. Schnoebelen. Model Checking CTL+ and FCTL is Hard. In *FOSSACS'01*, LNCS 2030, pages 318–331. Springer-Verlag, 2001.

[LR03] C. Löding and P. Rohde. Model Checking and Satisfiability for Sabotage Modal Logic. In *FSTTCS'03*, LNCS 2914, pages 302–313. Springer-Verlag, 2003.

[Pel96] D. Peled. Combining Partial Order Reductions with On-the-Fly Model Checking. *FMSD*, 8(1):39–64, 1996.

[Pnu77] A. Pnueli. The Temporal Logic of Programs. In *FOCS'77*, pages 46–57, 1977.

[Pnu81] A. Pnueli. The Temporal Semantics of Concurrent Programs. *TCS*, 13:45–60, 1981.

[QS82] J.-P. Queille and J. Sifakis. Specification and Verification of Concurrent Systems in CESAR. In *CISP'82*, pages 337–351. Springer-Verlag, 1982.

[Rob71] R.M. Robinson. Undecidability and Nonperiodicity for Tilings of the Plane. *IM*, 12:177–209, 1971.

[Wan61] H. Wang. Proving Theorems by Pattern Recognition II. *BSTJ*, 40:1–41, 1961.

[Wil99] T. Wilke. CTL+ is Exponentially More Succinct than CTL. In *FSTTCS'99*, LNCS 1738, pages 110–121. Springer-Verlag, 1999.

A Full definition of the MCTL* semantics

Given a Kripke structure $\mathcal{K} = \langle \mathrm{AP}, \mathrm{W}, \mathrm{R}, \mathrm{L} \rangle$ and $w \in \mathrm{W}$, for all MCTL* state formulas φ, the relation $\mathcal{K}, w \models \varphi$, is inductively defined as follows.

1. $\mathcal{K}, w \models p$, with $p \in \mathrm{AP}$, iff $p \in \mathrm{L}(w)$.
2. For all state formulas φ, φ_1, and φ_2, it holds:
   (a) $\mathcal{K}, w \models \neg\varphi$ iff not $\mathcal{K}, w \models \varphi$, that is $\mathcal{K}, w \not\models \varphi$;
   (b) $\mathcal{K}, w \models \varphi_1 \wedge \varphi_2$ iff $\mathcal{K}, w \models \varphi_1$ and $\mathcal{K}, w \models \varphi_2$;
   (c) $\mathcal{K}, w \models \varphi_1 \vee \varphi_2$ iff $\mathcal{K}, w \models \varphi_1$ or $\mathcal{K}, w \models \varphi_2$.
3. For all state formulas φ_1 and φ_2, it holds:
   (a) $\mathcal{K}, w \models \varphi_1 \,\Lambda\, \varphi_2$ iff for all Kripke structures $\mathcal{K}' \in \mathsf{minstructs}(\mathfrak{S}(\mathcal{K}, w, \varphi_2))$ it holds that $\mathcal{K}', w \models \varphi_1$;
   (b) $\mathcal{K}, w \models \varphi_1 \,\Xi\, \varphi_2$ iff there exists a Kripke structure $\mathcal{K}' \in \mathsf{minstructs}(\mathfrak{S}(\mathcal{K}, w, \varphi_2))$ such that $\mathcal{K}', w \models \varphi_1$;
   where $\mathfrak{S}(\mathcal{K}, w, \varphi) = \{\mathcal{K}' \preccurlyeq \mathcal{K} \mid \forall \mathcal{K}'' \preccurlyeq \mathcal{K} : \mathcal{K}' \preccurlyeq \mathcal{K}'' \text{ implies } \mathcal{K}'', w \models \varphi_2\}$
4. For a path formula ψ, it holds:
   (a) $\mathcal{K}, w \models \mathsf{A}\psi$ iff for all $\pi \in \mathsf{paths}(\mathcal{K}, w)$ it holds that $\mathcal{K}, \pi, 0 \models \psi$;
   (b) $\mathcal{K}, w \models \mathsf{E}\psi$ iff there exists $\pi \in \mathsf{paths}(\mathcal{K}, w)$ such that $\mathcal{K}, \pi, 0 \models \psi$.

For all MCTL* path formulas ψ, paths $\pi \in \mathsf{paths}(\mathcal{K})$, and natural numbers $k < |\pi|$, the relation $\mathcal{K}, \pi, k \models \psi$ is inductively defined as follows.

4. $\mathcal{K}, \pi, k \models \varphi$, with φ state formula, iff $\mathcal{K}, \pi(k) \models \varphi$.
5. Where ψ, ψ_1, and ψ_2 are path formulas, we have:
   (a) $\mathcal{K}, \pi, k \models \neg\psi$ iff not $\mathcal{K}, \pi, k \models \psi$, that is $\mathcal{K}, \pi, k \not\models \psi$;
   (b) $\mathcal{K}, \pi, k \models \psi_1 \wedge \psi_2$ iff $\mathcal{K}, \pi, k \models \psi_1$ and $\mathcal{K}, \pi, k \models \psi_2$;
   (c) $\mathcal{K}, \pi, k \models \psi_1 \vee \psi_2$ iff $\mathcal{K}, \pi, k \models \psi_1$ or $\mathcal{K}, \pi, k \models \psi_2$.
6. Where ψ, ψ_1, and ψ_2 path formulas, we have:
   (a) $\mathcal{K}, \pi, k \models \mathsf{X}\psi$ iff $k < |\pi| - 1$ and $\mathcal{K}, \pi, (k+1) \models \psi$;
   (b) $\mathcal{K}, \pi, k \models \tilde{\mathsf{X}}\psi$ iff $k = |\pi| - 1$ or $k < |\pi| - 1$ and $\mathcal{K}, \pi, (k+1) \models \psi$;
   (c) $\mathcal{K}, \pi, k \models \psi_1 \mathsf{U} \psi_2$ iff there exists an index i, with $k \leq i < |\pi|$, such that $\mathcal{K}, \pi, i \models \psi_2$ and, for all indexes j with $k \leq j < i$, it holds $\mathcal{K}, \pi, j \models \psi_1$;
   (d) $\mathcal{K}, \pi, k \models \psi_1 \mathsf{R} \psi_2$ iff for all indexes i, with $k \leq i < |\pi|$, it holds $\mathcal{K}, \pi, i \models \psi_2$ or there exists an index j with $k \leq j < i$, such that $\mathcal{K}, \pi, j \models \psi_1$. □

B Proof of Theorem 2

We show that if there is a model $\mathcal{K}$ for φ then there exists a model $\mathcal{K}' \preccurlyeq \mathcal{K}$, with $|\mathcal{K}'| \leq g(|\varphi|) = 2^{f(|\varphi|)}$, where f is recursive and monotone, and $\mathsf{dep}(\mathcal{K}) \leq \mathsf{dep}(\varphi)$, such that for all models $\mathcal{K}''$, with $\mathcal{K}' \preccurlyeq \mathcal{K}'' \preccurlyeq \mathcal{K}$, it holds that $\mathcal{K}'' \models \varphi$, i.e., $\mathcal{K}'$ is conservative w.r.t. the model $\mathcal{K}$ and the formula φ. We procede by mutual induction on the number of nested occurrences of minimal model quantifiers in φ (external induction) and on the structure of the formula itself (internal induction). W.l.o.g we assume φ is in positive normal form.

The base step for the external induction follows directly by applying to φ (since it is a PML formula) the well-known selection procedure used to prove the finite model property for PML [BdRV04].

We now proceed with the base step for the internal induction in the external inductive case, where φ is of the form $\varphi_1 \,\Xi\, \varphi_2$ (resp., $\varphi_1 \,\Lambda\, \varphi_2$). Since φ is satisfiable, there is a model $\mathcal{K}$ and a world $w \in \mathsf{dom}(\mathcal{K})$ for which there exists a minimal model $\mathcal{K}' \in \mathsf{minstructs}(\mathfrak{S}(\mathcal{K}, w, \varphi_2))$ such that $\mathcal{K}', w \models \varphi_1$ (resp. for all minimal models $\mathcal{K}' \in \mathsf{minstructs}(\mathfrak{S}(\mathcal{K}, w, \varphi_2))$ it holds that $\mathcal{K}', w \models \varphi_1$). Now, by inductive hypothesis, for each $\mathcal{K}'' \in \mathfrak{S}(\mathcal{K}, w, \varphi_2)$ there exists a conservative model $\mathcal{K}''' \preccurlyeq \mathcal{K}''$, such that the properties $|\mathcal{K}'''| \leq 2^{f(|\varphi_2|)}$, $\mathsf{dep}(\mathcal{K}''', w) \leq \mathsf{dep}(\varphi_2)$, and $\mathcal{K}''', w \models \varphi_2$ hold. Hence, $\mathcal{K}'$ is a conservative model satisfying the same properties, such that $\mathcal{K}', w \models \varphi$. Moreover, $|\mathcal{K}'| \leq 2^{f(|\varphi_2|)} \leq 2^{f(|\varphi|)}$ and $\mathsf{dep}(\mathcal{K}', w) \leq \mathsf{dep}(\varphi_2) = \mathsf{dep}(\varphi)$, so the thesis follows for this base case.

Consider now the inductive steps for the internal induction in the external inductive case.

If $\varphi = \varphi_1 \vee \varphi_2$ then there exists a $\mathcal{K}$ such that $\mathcal{K} \models \varphi$, so it holds that $\mathcal{K} \models \varphi_1$ or $\mathcal{K} \models \varphi_2$. Now, by inductive hypothesis, there exist a conservative model $\mathcal{K}_1 \preccurlyeq \mathcal{K}$ of φ_1, with $|\mathcal{K}_1| \leq 2^{f(|\varphi_1|)}$ and $\mathsf{dep}(\mathcal{K}_1) \leq \mathsf{dep}(\varphi_1)$, or a conservative model $\mathcal{K}'' \preccurlyeq \mathcal{K}$ of φ_2, with $|\mathcal{K}_2| \leq 2^{f(|\varphi_2|)}$ and $\mathsf{dep}(\mathcal{K}_2) \leq \mathsf{dep}(\varphi_2)$. Consider now the model $\mathcal{K}' = \mathcal{K}_i$, for some i, if $\mathcal{K}_i$ exists. It is obvious that, $\mathcal{K}' \models \varphi$. Moreover, $|\mathcal{K}'| = |\mathcal{K}_i| \leq 2^{f(|\varphi_i|)} \leq 2^{f(|\varphi|)}$ and $\mathsf{dep}(\mathcal{K}') = \mathsf{dep}(\mathcal{K}_i) \leq \mathsf{dep}(\varphi_i) \leq \mathsf{dep}(\varphi)$, so the thesis follows.

If $\varphi = \varphi_1 \wedge \varphi_2$ then there exists a $\mathcal{K}$ such that $\mathcal{K} \models \varphi$, so it holds that $\mathcal{K} \models \varphi_1$ and $\mathcal{K} \models \varphi_2$. Now, by inductive hypothesis, that there exist a conservative model $\mathcal{K}_1 \preccurlyeq \mathcal{K}$ of φ_1, with $|\mathcal{K}_1| \leq 2^{f(|\varphi_1|)}$ and $\mathsf{dep}(\mathcal{K}_1) \leq \mathsf{dep}(\varphi_1)$, and a conservative model $\mathcal{K}'' \preccurlyeq \mathcal{K}$ of φ_2, with $|\mathcal{K}_2| \leq 2^{f(|\varphi_2|)}$ and $\mathsf{dep}(\mathcal{K}_2) \leq \mathsf{dep}(\varphi_2)$. Consider now the model $\mathcal{K}' \preccurlyeq \mathcal{K}$ formed by the union[1] of $\mathcal{K}_1$ and $\mathcal{K}_2$. By conservativeness of $\mathcal{K}_1$ and $\mathcal{K}_2$, for all models $\mathcal{K}''$, with $\mathcal{K}' \preccurlyeq \mathcal{K}'' \preccurlyeq \mathcal{K}$, it holds that $\mathcal{K}'' \models \varphi_1$ and $\mathcal{K}'' \models \varphi_2$ and so $\mathcal{K}'' \models \varphi$. Moreover, $|\mathcal{K}'| \leq |\mathcal{K}_1| + |\mathcal{K}_2| \leq 2^{f(|\varphi_1|)} + 2^{f(|\varphi_2|)} \leq 2^{f(|\varphi|)}$ and $\mathsf{dep}(\mathcal{K}') = \max\{\mathsf{dep}(\mathcal{K}_1), \mathsf{dep}(\mathcal{K}_2)\} \leq \max\{\mathsf{dep}(\varphi_1), \mathsf{dep}(\varphi_2)\} = \mathsf{dep}(\varphi)$, so the thesis follows.[2]

If $\varphi = \mathsf{EX}\,\varphi'$ (resp., $\varphi = \mathsf{AX}\,\varphi'$) then there exists a $\mathcal{K}$ and a $w \in \mathsf{dom}(\mathcal{K})$ such that $\mathcal{K}, w \models \varphi$. So there exists a world $v \in \mathsf{dom}(\mathcal{K})$, with $(w, v) \in \mathrm{R}$, such that $\mathcal{K}, v \models \varphi'$ (resp., for all world $v \in \mathsf{dom}(\mathcal{K})$, with $(w, v) \in \mathrm{R}$, it holds that $\mathcal{K}, v \models \varphi'$). Now, by inductive hypothesis, there exists a conservative model $\mathcal{K}' \preccurlyeq \mathcal{K}$ of φ', with $|\mathcal{K}'| \leq 2^{f(|\varphi'|)}$ and $\mathsf{dep}(\mathcal{K}', v) \leq \mathsf{dep}(\varphi')$. If there is $u \in \mathsf{dom}(\mathcal{K})$, with $(u, v) \in \mathrm{R}'$, then it is obvious that $\mathcal{K}', u \models \varphi$, moreover, $|\mathcal{K}'| \leq 2^{f(|\varphi'|)} \leq 2^{f(|\varphi|)}$ and $\mathsf{dep}(\mathcal{K}', u) \leq \mathsf{dep}(\mathcal{K}', v) + 1 \leq \mathsf{dep}(\varphi') + 1 = \mathsf{dep}(\varphi)$, so the thesis easily follows. If such an u does not exists, consider the model $\mathcal{K}''$ constructed adding to $\mathcal{K}'$ the world w of $\mathcal{K}$ such that $(w, v) \in \mathrm{R}$. Then, it is obvious that $\mathcal{K}'', w \models \varphi$, moreover, $|\mathcal{K}''| = |\mathcal{K}'| + 2 \leq 2^{f(|\varphi'|)} + 2 \leq 2^{f(|\varphi|)}$ and $\mathsf{dep}(\mathcal{K}'', w) = \mathsf{dep}(\mathcal{K}'', v) + 1 = \mathsf{dep}(\mathcal{K}', v) + 1 \leq \mathsf{dep}(\varphi') + 1 = \mathsf{dep}(\varphi)$, so the thesis holds also in this case.

---

[1] Union of models is defined in the classical way: union of sets of worlds, union of relations, union of sets of atomic propositions, etc.

[2] Observe that this step deeply makes use of the conservativeness, indeed if $\mathcal{K}_1$ or $\mathcal{K}_2$ are not conservative we can not use the fact that their union still satisfy the formula φ. As an example, consider the formula $\varphi = \mathsf{EX}\,a \wedge (\mathsf{EX}\,a \rightarrow \mathsf{EX}\,b)$. It is not hard to show a model with two nodes for $\mathsf{EX}\,a$ and a non conservative model with one world for $\mathsf{EX}\,a \rightarrow \mathsf{EX}\,b$ whose union does not satisfy φ.

For the cases $\varphi = \mathsf{E}\tilde{\mathsf{X}}\varphi'$ and $\varphi = \mathsf{A}\tilde{\mathsf{X}}\varphi'$ the proof procedes in a similar way, so we omit the details.