On the Universal and Existential Fragments of the μ-Calculus*

Thomas A. Henzinger, Orna Kupferman, and Rupak Majumdar

Department of EECS
University of California
Berkeley, CA 94720, USA
{tah,orna,rupak}@eecs.berkeley.edu

Abstract. One source of complexity in the μ-calculus is its ability to specify an unbounded number of switches between universal (AX) and existential (EX) branching modes. We therefore study the problems of satisfiability, validity, model checking, and implication for the universal and existential fragments of the μ-calculus, in which only one branching mode is allowed. The universal fragment is rich enough to express most specifications of interest, and therefore improved algorithms are of practical importance. We show that while the satisfiability and validity problems become indeed simpler for the existential and universal fragments, this is, unfortunately, not the case for model checking and implication. We also show the corresponding results for the alternation-free fragment of the μ-calculus, where no switches between least and greatest fixed points are allowed. Our results imply that efforts to find a polynomial time model-checking algorithm for the μ-calculus can be replaced by efforts to find such an algorithm for the universal or the existential fragments.

1 Introduction

In *formal verification*, we reason about systems and their properties by reasoning about formal models of systems and formal specifications of the properties [5]. The algorithmic nature of formal verification makes it fully automatic, convenient to use, and attractive to practitioners. At the same time, formal verification is very sensitive to the size of the formal model of the system and the formal specification. Commercial verification tools need to cope with the exceedingly large state spaces that are present in real-life designs. One of the most important developments in this area is the discovery of *symbolic* methods [2, 26]. Typically, symbolic model-checking tools proceed by computing fixed-point expressions over the model's set of states. For example, to find the set of states from which a state satisfying some predicate p is reachable, the model checker starts with the set y of states in which p holds, and repeatedly adds to y the set

* This work was supported in part by NSF grant CCR-9988172, AFOSR MURI grant F49620-00-1-0327, and a Microsoft Research Fellowship.

EXy of states that have a successor in y. Formally, the model checker calculates the least fixed point of the expression $y = p \vee EXy$.

The *μ-calculus* is a logic that contains the next modalities EX and AX, and the fixed-point operators μ and ν [20]. As such, it describes fixed-point computations in a natural form. The μ-calculus is an extremely general modal logic. It is as expressive as automata on infinite trees, and it subsumes most known specification formalisms, including dynamic logics such as PDL [13] and temporal logics such as LTL and CTL* [8, 7] (see [17] for a general result). The *alternation-free* fragment of μ-calculus (AFMC, for short) [12] has a restricted syntax that does not allow alternating nesting of fixed-point operators, making the evaluation of expressions very simple [6]. The alternation-free fragment subsumes the temporal logic CTL.

Four decision problems arise naturally for every specification formalism: the *satisfiability* problem (given a formula φ, is there a model that satisfies φ?) checks whether a specification can be implemented, and algorithms for deciding the satisfiability problems are the basis for program synthesis and control [3, 29, 28]; the *validity* problem (given φ, do all models satisfy φ?) checks whether the specification is trivially satisfied, and is used as a sanity check for requirements [24]; the *model-checking* problem (given a formula φ and a model M, does M satisfy φ?) is the basic verification problem; and the *implication* problem (given two formulas φ and ψ, is $\varphi \rightarrow \psi$ valid?) arises naturally in the context of modular verification, where it must be shown that a module satisfies a property under an assumption about the environment [27, 22].

The satisfiability, validity, and implication problems for the μ-calculus are all EXPTIME-complete [13, 1] (since the μ-calculus is closed under negation, it is easy to get EXPTIME completeness for the validity and implication problems by reductions to and from the satisfiability problem). The model-checking problem for the μ-calculus was first considered in [12], which describes an algorithm with complexity $O((mn)^{l+1})$, where m is the size of M, n is the size of φ, and l is the number of alternations between least and greatest fixed-point operators in φ. In [11], the problem was shown to be equivalent to the nonemptiness problem for parity tree automata, and thus to lie in NP $\cap$ co-NP. Today, it is known that the problem is in UP $\cap$ co-UP [18][1], and best known algorithm for μ-calculus model checking is of time complexity roughly $O(mn^{\frac{l}{2}})$ [25, 30, 19], which is still exponential in the number of alternations. The precise complexity of the problem, and in particular, the question whether a polynomial-time solution for it exists, is a long-standing open problem.

In this paper we study the complexity of the four decision problems for the *existential* and *universal* fragments of μ-calculus. The existential fragment consists of formulas where the only allowed next modality is the existential one (EX), and the universal fragment consists of formulas where the only allowed next modality is the universal one (AX). We consider μ-calculus in positive normal form, thus the strict syntactic fragments are also semantic fragments

[1] The class UP is a subset of NP, where each word accepted by the Turing machine has a unique accepting run.

— there is no way of specifying an existential next in the universal fragment by negating a universal one, and vice versa. Both sublogics induce the state equivalence *similarity* (mutual simulation), as opposed to bisimilarity, which is induced by the full μ-calculus [15]. The existential and universal fragments of the μ-calculus subsume the existential and universal fragments of the branching-time logics CTL and CTL*. For temporal logics, the universal and existential fragments have been well-studied (see, e.g., [22]). As we specify in the table in Figure 1, the satisfiability, validity, and implication problems for the universal and existential fragments of CTL and CTL* are all easier than the corresponding problems for the full logics [13, 32, 9, 22]. On the other hand, the model-checking complexities for the universal and existential fragments of CTL and CTL* coincide with the complexities of the full logics, and the same holds for the system complexities of model checking (i.e., the complexities in terms of the size of the model, assuming the specification is fixed) [4, 23].

Unlike CTL and CTL*, it is possible to express in μ-calculus *unbounded switching* of AX and EX modalities. Such an unbounded switching is an apparent source of complexity: for example, μ-calculus can express the reachability problem on And-Or graphs, which is PTIME-complete, while the reachability problem on plain graphs (existential reachability), and its universal counterpart, are NLOGSPACE-complete. Accordingly, the system complexity of the model-checking problem for μ-calculus is PTIME-complete, whereas the one for CTL and CTL* is only NLOGSPACE-complete [16, 12, 23]. By removing the switching of modalities from the μ-calculus, one may hope that the algorithms for the four decision problems, and model checking in particular, will become simpler. Since most specifications assert what a system must or must not do in *all* possible futures, the universal fragment of the μ-calculus is expressive enough to capture most specifications of interest. Also, the symbolic evaluation of LTL formula is reduced to model checking of an existential μ-calculus formula. Hence, our study is not only of a theoretical interest — an efficient algorithm for the universal or existential fragments of the μ-calculus is of practical interest.

We determine the complexities of the four decision problems for the universal and existential fragments of the μ-calculus, as well as their alternation-free fragments. Our results are summarized in Figure 1. It turns out that the hope to get easier algorithms for the universal and existential fragments is only partially fulfilled. We show that while the satisfiability and validity problems become easier for the existential and universal fragments, both the model-checking and implication problems stay as hard as for the full μ-calculus. In particular, our results imply that efforts to find a polynomial time model-checking algorithm for the μ-calculus can be replaced by efforts to find a polynomial time model-checking algorithms for the universal or the existential fragments. Note that the picture we obtain for μ-calculus does not coincide with the picture obtained in the study of the universal and existential fragments of CTL and CTL*, where the restriction to the fragments makes also the implication problem easier.

One key insight concerns the size of models for the existential and universal fragments of the μ-calculus. We prove the membership of the satisfiability

problem for the existential fragment of μ-calculus in NP via a *linear-size model property*. This is in contrast to the full μ-calculus, which has only an exponential-size model property [21]. This shows that extending propositional logic by the EX modality and fixed-point operators does not make the satisfiability problem harder. On the other hand, a similar extension with AX results in a logic for which the linear-size model property does not hold, and whose satisfiability problem is PSPACE-complete.

A second insight is that, in model-checking as well as implication problems, the switching of EX and AX modalities can be encoded by the boolean connectives $\vee$ and $\wedge$ in combination with either one of the two modalities and fixed-point operators. Let us be more precise. The model-checking problem for the μ-calculus is closely related to the problem of determining the winner in games on And-Or graphs. The system complexity of μ-calculus model checking is PTIME-hard, because a μ-calculus formula of a fixed size can specify an unbounded number of switches between universal and existential branching modes. In particular, the formula $\mu y.t \vee EX\,AX y$ specifies winning in a And-Or reachability game, and formulas with alternations between least and greatest fixed-point operators can specify the winner in a And-Or parity game. One would therefore expect that the universal and existential fragments of μ-calculus, in which no switching between branching modes is possible, would not be sufficiently strong for specifying And-Or reachability. Indeed, in [11] the authors define a fragment L_2 of μ-calculus that explicitly bounds the number of switches between both AX and EX modalities and $\wedge$ and $\vee$ boolean operators. This fragment is as expressive as extended CTL*, and it cannot specify reachability in And-Or graphs (the system complexity of model checking is NLOGSPACE-complete). However, in model checking as well as implication problems, we can consider models in which the successors of a state are labeled in a way that enables the specification to directly refer to them. Then, it is possible to replace the existential next modality by a disjunction over all the successors, and it is possible to replace the universal next modality by a conjunction that refers to each successor. More specifically, if we can guarantee that the successors of a state with branching degree two are labeled by l (left) and r (right), then the existential next formula EXy can be replaced by $AX(l \rightarrow y) \vee AX(r \rightarrow y)$, and the universal next formula AXy can be replaced by $EX(l \wedge y) \wedge EX(r \wedge y)$. While these observations are technically simple, they enable us to solve the open problems regarding the complexity of the universal and existential fragments of the μ-calculus.

2 Propositional μ-Calculus

The *propositional μ-calculus* (MC, for short) is a propositional modal logic augmented with least and greatest fixpoint operators [20]. Specifically, we consider a μ-calculus where formulas are constructed from Boolean propositions with Boolean connectives, the temporal operators EX and AX, as well as least (μ) and greatest (ν) fixpoint operators. We assume without loss of generality that μ-calculus formulas are written in positive normal form (negation only applied

	Satisfiability	Validity	Implication	Model checking	system complexity
CTL*	2EXPTIME	2EXPTIME	2EXPTIME	PSPACE	NLOGSPACE
$\forall$CTL*	PSPACE	PSPACE	EXPSPACE	PSPACE	NLOGSPACE
$\exists$CTL*	PSPACE	PSPACE	EXPSPACE	PSPACE	NLOGSPACE
CTL	EXPTIME	EXPTIME	EXPTIME	linear time	NLOGSPACE
$\forall$CTL	PSPACE	co-NP	PSPACE	linear time	NLOGSPACE
$\exists$CTL	NP	PSPACE	PSPACE	linear time	NLOGSPACE
MC	EXPTIME	EXPTIME	EXPTIME	NP $\cap$ co-NP	PTIME
$\forall MC$	PSPACE	co-NP	EXPTIME	NP $\cap$ co-NP	PTIME
$\exists MC$	NP	PSPACE	EXPTIME	NP $\cap$ co-NP	PTIME
$AFMC$	EXPTIME	EXPTIME	EXPTIME	linear time	PTIME
$\forall AFMC$	PSPACE	co-NP	EXPTIME	linear time	PTIME
$\exists AFMC$	NP	PSPACE	EXPTIME	linear time	PTIME

Fig. 1. Summary of known and new results

to atomic propositions). Formally, given a set AP of atomic propositions and a set V of variables, a μ-calculus formula is either:

- *true*, *false*, p or $\neg p$, for all $p \in AP$;
- y, for all $y \in V$;
- $\varphi_1 \wedge \varphi_2$ or $\varphi_1 \vee \varphi_2$, where φ_1 and φ_2 are μ-calculus formulas;
- $AX\varphi$ or $EX\varphi$, where φ is a μ-calculus formula;
- $\mu y.\varphi$ or $\nu y.\varphi$, where $y \in V$ and φ is a μ-calculus formula.

We say variable y is *bound* in $\mu y.\varphi$ and $\nu y.\varphi$. A variable is *free* if it is not bound. A *sentence* is a formula that contains no free variables. We refer to AX and EX as the *universal* and *existential* next modalities, respectively. For a μ-calculus formula φ, define the *size* $|\varphi|$ of φ as the size of the reduced DAG representation of φ. Note that this is always smaller than or equal to the usual definition of size in terms of the number of syntactic symbols occurring in the formula, our results hold for this definition as well.

The *universal μ-calculus*, ($\forall MC$, for short) is the fragment of μ-calculus in which the only next modality allowed is the universal one. Dually, the *existential μ-calculus* ($\exists MC$, for short) is the fragment of μ-calculus in which the only next modality allowed is the existential one. Note that since μ-calculus formulas are written in a positive normal form, there is no way to specify existential next in $\forall MC$ by negating a universal one.

A μ-calculus formula is *alternation free* if, for all $y \in V$, there are respectively no occurrences of ν (μ) on any syntactic path from an occurrence of μy (νy) to an occurrence of y. For example, the formula $\mu x.(p \vee \mu y.(x \vee EXy))$ is alternation free and the formula $\nu x.\mu y.((p \wedge x) \vee EXy)$ is not alternation free. The *alternation-free μ-calculus* ($AFMC$, for short) is a subset of μ-calculus containing only alternation-free formulas. We also refer to the universal and existential fragments of $AFMC$, denoted by $\forall AFMC$ and $\exists AFMC$, respectively.

A μ-calculus formula is *guarded* if for all $y \in V$, all the occurrences of y that are in a scope of a fixpoint modality $\lambda \in \{\mu, \nu\}$ are also in a scope of a next modality which is itself in the scope of λ. For example, the formula $\mu y.(p \vee EXy)$ is guarded and the formula $EX\mu y.(p \vee y)$ is not guarded. We assume that all μ-calculus formulas are guarded. As proved in [23], every μ-calculus formula can be linearly translated to an equivalent guarded one, thus we do not lose generality with our assumption.

The semantics of μ-calculus formulas is defined with respect to Kripke structures. A *Kripke structure* $\mathcal{K} = \langle AP, W, R, w_0, L\rangle$ consists of a set AP of atomic propositions, a set W of states, a total transition relation $R \subseteq W \times W$, an initial state $w_0 \in W$, and a labeling $L : W \to 2^{AP}$ that maps each state to the set of atomic propositions true in that state.

Given a Kripke structure $\mathcal{K} = \langle AP, W, R, w_0, L\rangle$ and a set $\{y_1, \ldots, y_n\}$ of free variables, a *valuation* $\mathcal{V} : \{y_1, \ldots, y_n\} \to 2^W$ is an assignment of subsets of W to the variables in $\{y_1, \ldots, y_n\}$. For a valuation $\mathcal{V}$, a variable y and a set $W' \subseteq W$, denote by $\mathcal{V}[y \leftarrow W']$ the valuation mapping y to W' and y' to $\mathcal{V}(y')$ for all $y' \neq y$. A formula φ with free variables $\{y_1, \ldots, y_n\}$ is interpreted over the structure $\mathcal{K}$ as a mapping $\varphi^{\mathcal{K}}$ from valuations to 2^W. Thus, $\varphi^{\mathcal{K}}(\mathcal{V})$ denotes the set of states that satisfy φ with the valuation $\mathcal{V}$. The mapping $\varphi^{\mathcal{K}}$ is defined inductively as follows:

- $true^{\mathcal{K}}(\mathcal{V}) = W$ and $false^{\mathcal{K}}(\mathcal{V}) = \emptyset$.
- For $p \in AP$, the set $p^{\mathcal{K}}(\mathcal{V}) = \{w \in W \mid p \in L(w)\}$ and $(\neg p)^{\mathcal{K}}(\mathcal{V}) = \{w \in W \mid p \notin L(w)\}$.
- $(\varphi_1 \wedge \varphi_2)^{\mathcal{K}}(\mathcal{V}) = \varphi_1^{\mathcal{K}}(\mathcal{V}) \cap \varphi_2^{\mathcal{K}}(\mathcal{V})$.
- $(\varphi_1 \vee \varphi_2)^{\mathcal{K}}(\mathcal{V}) = \varphi_1^{\mathcal{K}}(\mathcal{V}) \cup \varphi_2^{\mathcal{K}}(\mathcal{V})$.
- $(AX\varphi)^{\mathcal{K}}(\mathcal{V}) = \{w \in W \mid \forall w'.$ if $(w, w') \in R$ then $w' \in \varphi^{\mathcal{K}}(\mathcal{V})\}$.
- $(EX\varphi)^{\mathcal{K}}(\mathcal{V}) = \{w \in W \mid \exists w'.(w, w') \in R$ and $w' \in \varphi^{\mathcal{K}}(\mathcal{V})\}$.
- $(\mu x.\varphi)^{\mathcal{K}}(\mathcal{V}) = \bigcap \{W' \subseteq W \mid \varphi^{\mathcal{K}}(\mathcal{V}[x \leftarrow W']) \subseteq W'\}$.
- $(\nu x.\varphi)^{\mathcal{K}}(\mathcal{V}) = \bigcup \{W' \subseteq W \mid W' \subseteq \varphi^{\mathcal{K}}(\mathcal{V}[x \leftarrow W'])\}$.

For a sentence, no valuation is required. For a state $w \in W$ of the Kripke structure $\mathcal{K}$, and a sentence φ, we write $\mathcal{K}, w \models \varphi$ iff $w \in \varphi^{\mathcal{K}}$. By the Knaster-Tarski Theorem, the fixpoints always exist.

3 Satisfiability and Validity

The satisfiability and validity problems for μ-calculus and its alternation-free fragment are EXPTIME-complete [1, 13]. In this section we study the satisfiability and validity problems for the universal and existential fragments.

For a $\forall MC$ formula φ, let $[\varphi]$ denote the linear-time μ-calculus formula obtained from φ by omitting all its universal path quantifiers. It is easy to see that φ of $\forall MC$ is satisfiable iff $[\varphi]$ is satisfiable. Indeed, a model for $[\varphi]$ is also a model for φ, and each path in a model for φ is a model for $[\varphi]$. Since the satisfiability problem for linear-time μ-calculus and its alternation-free fragment is PSPACE-complete [31], so is the satisfiability problem for $\forall MC$ and $\forall AFMC$.

Theorem 1. *The satisfiability problem for $\forall MC$ and $\forall AFMC$ is PSPACE-complete.*

Since both $\exists MC$ and $\exists AFMC$ subsume propositional logic, the satisfiability problem for these logics is clearly hard for NP. We show that the satisfiability problem is in fact NP-complete. To show membership in NP, we prove a *linear-size model property* for $\exists MC$.

Lemma 1. *Let φ be a formula of $\exists MC$. If φ is satisfiable, then it has a model with at most $O(|\varphi|)$ states and $O(|\varphi|)$ transitions.*

Proof. The proof is similar to the one used in [22] to show a linear-size model property for $\exists$CTL. We proceed by induction on the structure of $\exists MC$ formulas. With each $\exists MC$ formula φ, we associate a set S_φ of models (Kripke structures) that satisfy φ. We define S_φ by structural induction. The states of the models in S_φ are labeled by both the atomic propositions and the variables free in φ. We use $S_{\varphi_1} \to S_{\varphi_2}$ to denote the set of models obtained by taking a model M_1 from S_{φ_1}, a model M_2 from S_{φ_2}, adding a transition from the initial state of M_1 to the initial state of M_2, and fixing the initial state to be the one of M_1. We use $S_{\varphi_1} \cap^* S_{\varphi_2}$ to denote the set of models obtained by taking a model M_1 from S_{φ_1} and a model M_2 from S_{φ_2}, such that M_1 and M_2 agree on the labeling of their initial states, fixing the initial state to be the initial state of M_1, redirecting transitions to the initial state of M_2 into the initial state of M_1, and adding transitions from the initial state of M_1 to all the successors of the initial state of M_2. Finally, we use $S_{\varphi(\#)}\downarrow$, where $\#$ is an atomic proposition not in AP, to denote the set of models obtained from a model in $S_{\varphi(\#)}$ by adding transitions from states labeled by $\#$ to all the successors of the initial state, and removing $\#$ from the labels of states. We can now define S_φ as follows. Note that we do not consider the case where $\varphi = x$, for $x \in V$, as we assume that φ is a sentence.

- $S_{\mathbf{true}}$ is the set of all one-state models over AP.
- $S_{\mathbf{false}} = \emptyset$.
- S_p, for $p \in AP$, is the set of all one-state models over AP in which p holds.
- $S_{\neg p}$, for $\neg p \in AP$, is the set of all one-state models over AP in which p does not hold.
- $S_{\varphi_1 \vee \varphi_2} = S_{\varphi_1} \cup S_{\varphi_2}$.
- $S_{\varphi_1 \wedge \varphi_2} = S_{\varphi_1} \cap^* S_{\varphi_2}$.
- $S_{EX\varphi_1} = S_{\mathbf{true}} \to S_{\varphi_1}$.
- $S_{\mu x.\varphi_1(x)} = S_{\varphi_1(\varphi_1(\mathit{false}))}$.
- $S_{\nu x.\varphi_1(x)} = S_{\varphi_1(\#\wedge(\varphi_1(\mathit{true})))}\downarrow$.

For example, if $AP = \{p\}$, and $\varphi = \nu x.\varphi_1(x)$ with $\varphi_1(x) = EX(p \wedge x) \wedge EX(\neg p \wedge x)$, then $\varphi_1(\# \wedge (\varphi_1(\mathit{true}))) = EX(p \wedge \# \wedge EXp \wedge EX\neg p) \wedge EX(\neg p \wedge \# \wedge EXp \wedge EX\neg p)$, and S_φ contains the two models obtained from the model M_1 described in Figure 2 by labeling the initial state by either p or $\neg p$. Also, if $AP = \{p, q\}$ and $\varphi = EXp \wedge (\mu y.q \vee (p \wedge EXy))$, then S_φ contains the models

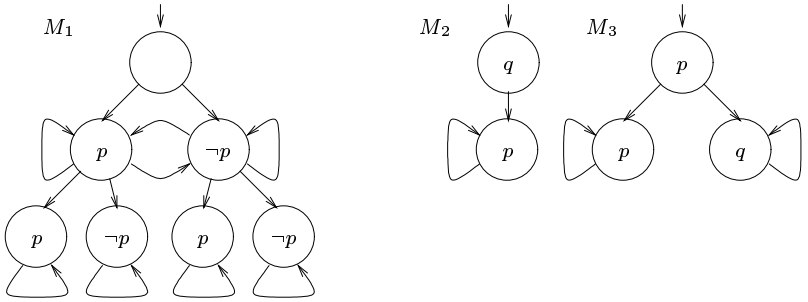

Fig. 2. Linear-size models.

obtained from the models M_2 and M_3 described there by completing labels of p or q the are left unspecified.

The models in S_φ are "economical" with respect to states that are required for satisfaction of formulas that refer to the strict future. For example, the initial state of models in $S_{EX\varphi_1}$ has a single successor that satisfies φ_1, and models in $S_{\mu y.\varphi(y)}$ that do not satisfy $\varphi(\mathit{false})$ in the initial state, are required to satisfy $\varphi(\mathit{false})$ in a successor state.

It is not hard to prove, by induction on the structure of φ, that each model in S_φ has $O(|\varphi|)$ states and $O(|\varphi|)$ transitions. We now prove, by an induction on the structure of φ, the following two claims.

1. For every model $M \in S_\varphi$, we have that M satisfies φ.
2. For every model M that satisfies φ, there is a model $M' \in S_\varphi$ such that M and M' agree on the labeling of their initial states.

Note that Claim (2) implies that if φ is satisfiable, then S_φ is not empty. Thus, the two claims together imply that if φ is satisfiable, then it has a satisfying model in S_φ, which is guaranteed to be of size linear in $|\varphi|$.

The proof for φ of the form *true*, *false*, p, $\neg p$, $\varphi_1 \vee \varphi_2$, and $\varphi_1 \wedge \varphi_2$ is easy. For the other cases, we proceed as follows.

Let $\varphi = EX\varphi_1$. By the induction hypothesis, all models in S_{φ_1} satisfy φ_1. Hence, (1) follows immediately from the definition of S_φ. To see (2), consider a model M that satisfies φ. Since $S_{\mathbf{true}}$ is the set of all one-state models over AP, it contains a model M' that agrees with M on the labeling of their initial states.

Let $\varphi = \mu x.\varphi_1(x)$. By the semantics of μ-calculus, a model that satisfies $\varphi_1(\varphi_1(\mathit{false}))$, satisfies φ as well. Hence, (1) follows immediately from the definition of S_φ. To see (2), consider a model M that satisfies φ. This means that for some $i > 0$, the model M satisfies $\varphi_1^i(\mathit{false})$. We construct a model M' that satisfies $\varphi_1(\varphi_1(\mathit{false}))$ and agrees with M on the label of the initial state. Let $\#$ be a proposition not in AP, and consider the formula $\varphi_1(\# \wedge \varphi_1^{i-1}(\mathit{false}))$. The model M can be attributed by $\#$ to satisfy $\varphi_1(\# \wedge \varphi_1^{i-1}(\mathit{false}))$. Moreover, since φ is guarded, the initial state of M is attributed by $\#$ only if there is a self loop in the initial state. Such a self loop can be unwound, so we can assume that the

initial state of M is not attributed by $\#$. Since φ is satisfiable, so is $\varphi_1(\mathit{false})$, and so there is a model N of $\varphi_1(\mathit{false})$. The structure M' is obtained from M by replacing all states attributed by $\#$ with N (i.e., all transitions leading into a state attributed by $\#$ are redirected to the initial state of N). Then, M' is a model of $\varphi_1(\varphi_1(\mathit{false}))$, and agrees with M in the labeling of the initial states.

Let $\varphi = \nu x.\varphi_1(x)$. By the semantics of μ-calculus, a model M satisfies φ iff M satisfies $\varphi_1^i(\mathit{true})$, for all $i \geq 0$. Consider a model $M \in S_\varphi$. By the definition of S_φ, the model M satisfies $\varphi_1(\mathit{true})$, and the states attributed $\#$ satisfy $\varphi_1(\mathit{true})$ as well. Since $\varphi_1(\mathit{true})$ is existential, the states attributed $\#$ continue to satisfy $\varphi_1(\mathit{true})$ after the new edges are added. In fact, it is not hard to see that after the new edges are added, the states attributed $\#$ also satisfy $\varphi_1(\#)$. Thus, for all $i \geq 1$, the model M can be unfolded $(i-1)$ times to show M satisfies $\varphi_1^i(\mathit{true})$, and we are done. To see (2), let M be a model of φ and let $\#$ be a proposition not in AP. Then, M satisfies $\varphi_1(\varphi_1(\mathit{true}))$, and it can be attributed by $\#$ to satisfy $\varphi_1(\# \wedge \varphi(\mathit{true}))$. As in the previous case, since φ is guarded, we can ensure that this leaves the labeling of the initial state unchanged, possibly after unwinding a self loop in the initial state. In addition, adding transitions from states attributed by $\#$ to all successors of the initial state, leaves the label of the initial state unchanged, and thus results in a model in S_φ that agrees with M on the labeling of their initial states. □

Note that μ-calculus with both universal and existential next modalities have only an exponential-size model property. Thus, the linear-size model property crucially depends on the fact that the only next modality that is allowed is the existential one. The linear-size model theorem shows that the satisfiability problem for $\exists MC$ and $\exists AFMC$ is in NP.

Theorem 2. *The satisfiability problem for $\exists MC$ and $\exists AFMC$ is NP-complete.*

Since a formula φ is satisfiable iff $\neg\varphi$ is valid, and since negating an $\exists MC$ formula results in a $\forall MC$ formula and vice versa, the following theorem is an immediate corollary of Theorems 1 and 2.

Theorem 3. *The validity problem is co-NP-complete for $\forall MC$ and $\forall AFMC$, and is PSPACE-complete for $\exists MC$ and $\exists AFMC$.*

4 Model Checking

In this section we study the model-checking problem for the universal and existential fragments of μ-calculus. We show that unlike the case of satisfiability, the model-checking problem for the restricted fragments is not easier than the model-checking problem for the μ-calculus.

The model-checking problem for μ-calculus is closely related to the problem of determining the winner in *turn-based games*. We first review here some definitions that will be used in the reduction of the model-checking problem for the full μ-calculus to the model-checking problem for the fragments. A *turn-based game*

graph is a directed graph $G = \langle V, E \rangle$, with a partition $V_e \cup V_u$ of V. The game is played between two players, player 1 and player 2. A position of the game is a vertex $v \in V$. At each step of the game, if the current position v is in V_e, player 1 chooses the next position among the vertices in $\{w \mid \langle v, w \rangle \in E\}$. Similarly, if $v \in V_u$, player 2 chooses the next position among the vertices in $\{w \mid \langle v, w \rangle \in E\}$. The game continues for an infinite number of steps, and induces an infinite path $\pi \in V^\omega$. The winner in the game depends on different conditions we can specify on words in V^ω. The simplest turn-based game is *reachability*. Then, the winning condition is some vertex $t \in V$, and player 1 wins the game if π eventually reaches the vertex t. Otherwise, player 2 wins. A richer turn-based game is *parity*. In parity games, there is a function $C : V \to \{0, \ldots, k-1\}$ that maps each vertex to a *color* in $\{0, \ldots, k-1\}$. Player 1 wins the parity game if the maximal color that repeats in π infinitely often is even.

A *strategy* for player 1 is a function $\xi_1 : V^* \times V_e \mapsto V$ such that for all $u \in V^*$ and $v \in V_e$, we have $\xi_1(u \cdot v) \in \{w \mid \langle v, w \rangle \in E\}$. A strategy for player 2 is defined similarly as $\xi_2 : V^* \times V_u \mapsto V$. For a vertex $s \in V$, and strategies ξ_1 and ξ_2 for player 1 and player 2 respectively, the *outcome* of ξ_1 and ξ_2 from s, denoted $\pi(\xi_1, \xi_2)(s)$, is the trace $v_0, v_1, \ldots \in V^\omega$ such that $v_0 = s$ and for all $i \geq 0$, we have $v_{i+1} \in \xi_1(v_0 \ldots v_{i-1}, v_i)$ if $v_i \in V_e$ and $v_{i+1} \in \xi_2(v_0 \ldots v_{i-1}, v_i)$ if $v_i \in V_u$. Finally, a vertex $s \in V$ is winning for player 1 in the turn-based game if there is a strategy ξ_1 of player 1 such that for all strategies ξ_2 of player 2, the outcome $\pi(\xi_1, \xi_2)(s)$ is winning for player 1. When G has an initial state s, we say that player 1 wins the game on G if s is winning for player 1 in G.

We start by considering the *system complexity* of the model-checking problem for the universal and existential fragments of μ-calculus; that is, the complexity of the problem in terms of the system, assuming the specification is fixed. As discussed in Section 1, the system complexity of $AFMC$ model checking is PTIME-complete, and hardness in PTIME [16] crucially depends on the fact that an $AFMC$ formula of a fixed size can specify an unbounded number of alternations between universal and existential branching mode. As we prove in Theorem 4 below, the setting of model checking enables us to trade an unbounded number of alternations between universal and existential branching modes by an unbounded alternation between disjunctions and conjunctions. The idea is that in model checking, unlike in satisfiability, we can consider models in which the successors of a state are labeled in a way that enables the specification to directly refer to them. Then, it is possible to replace the existential next modality by a disjunction over all the successors, and it is possible to replace the universal next modality by a conjunction that refers to each successor. Specifically, if the states of the Kripke structure are labeled by l and r so that a state labeled l is a left successor of its predecessor and a state labeled r is a right successor of its predecessor, then the existential next formula EXy can be replaced by $AX(l \to y) \vee AX(r \to y)$, and the universal next formula AXy can be replaced by $EX(l \wedge y) \wedge EX(r \wedge y)$.

Theorem 4. *The system complexity of $\forall AFMC$ and (so, also of $\exists AFMC$) model checking is PTIME-complete.*

Proof. Membership in PTIME follows from the linear time algorithm for $AFMC$ [6]. For hardness, we reduce the problem of deciding a winner in a turn-based reachability game to model checking of a $\forall AFMC$ formula of a fixed size. Since one can model check a specification φ by checking $\neg\varphi$ and negating the result, the same lower bound holds for $\exists AFMC$.

Deciding turn-based reachability games is known to be PTIME-hard already for acyclic graphs with branching degree two, where universal and existential vertices alternate, and both s and t are in V_e [14]. Given a bipartite, acyclic, turn-based game graph $G = \langle V, E\rangle$ with branching degree two, a partition of V to V_e and V_u, and two vertices s and t in V_e, we construct a Kripke structure $\mathcal{K} = \langle AP, W, R, w_0, L\rangle$ and a formula in $\forall AFMC$ such that $\mathcal{K}, w_0 \models \varphi$ iff player 1 wins the turn-based reachability game on G from state s and with target t.

We do the proof in two steps. First, we transform the graph G to another graph G', with some helpful properties, and then we construct the Kripke structure $\mathcal{K}$ from G'. Essentially, in G' each universal vertex is a left or right successor of exactly one existential vertex. Formally, $G' = \langle V', E'\rangle$, where $V' = V_e \cup V'_u$, and V'_u and E' are defined as follows. Let $E_e = E \cap (V_e \times V_u)$ and $E_u = E \cap (V_u \times V_e)$. Recall that each vertex in V_e has two successors. Let $E_e = E^l_e \cup E^r_e$ be a partition of E_e so that for each $v \in V_e$, one successor v_l of v is such that $\langle v, v_l\rangle \in E^l_e$ and the other successor v_r of v is such that $\langle v, v_r\rangle \in E^r_e$. Note that a vertex u may be the left successor of some vertex w_1 and the right successor of some other vertex w_2; thus $E^l_e(w_1, u)$ and $E^r_e(w_2, u)$. The goal of G' is to prevent such cases.

- $V'_u \subseteq V_u \times \{l, r\} \times V_e$ is such that $(v, l, w) \in V'_u$ iff $(w, v) \in E^l_e$ and $(v, r, w) \in V'_u$ iff $(w, v) \in E^r_e$. Thus, each edge $\langle w, v\rangle \in E_e$ contributes one vertex (v, l, w) or (v, r, w) to V'_u. Intuitively, visits to the vertex (v, l, w) correspond to visits to v in which it has been reached by following the left branch of w, and similarly for (v, r, w) and right.
- $E'_e = \{\langle v, (v_l, l, v)\rangle : \langle v, v_l\rangle \in E^l_e\} \cup \{\langle v, (v_r, r, v)\rangle : \langle v, v_r\rangle \in E^r_e\}$. Also, $E'_u = \{\langle (v, d, w), u\rangle : \langle v, u\rangle \in E_u\}$, and $E' = E'_e \cup E'_u$.

The size of G' is linear in the size of G. Indeed $|V'| = |V_e| + |E_e|$ and $|E'| = |E'_e| + |E'_u| \leq |E_e| + 2|E_u|$. It is not hard to see that player 1 can win the game in G iff he can win in G'. Note that the branching degree of G' remains two. The construction of G' ensures that the two successors of an existential vertex v can be referred to as the left or the right successor of v.

The graph $G' = \langle V', E'\rangle$, together with s and t, induces the Kripke structure $\mathcal{K} = \langle AP, V', E', s, L\rangle$ described below. The set of atomic propositions $AP = \{t, l\}$. For readability, we also introduce the shorthand r for $\neg l$.The proposition t holds in (and only in) the state t, and the propositions l and r hold in the left and right successor respectively for an existential node. Thus $l \in L(\langle v, l, w\rangle)$ and $r \in L(\langle v, r, w\rangle)$. Finally, let φ be the $\forall AFMC$ formula $\mu y.t \vee (AX(\neg l \vee AXy) \vee AX(\neg r \vee AXy))$. It is now easy to see that player 1 can win the turn-based reachability game for t from s in G' iff $\mathcal{K}, s \models \varphi$. □

Theorem 5. *The model-checking problem for $\forall MC$ (so, also for $\exists MC$) is as hard as the model-checking problem for μ-calculus.*

Proof. The idea is similar to the proof of Theorem 4, only that instead of talking about winning a turn-based reachability game, we talk about winning a *turn-based parity game* [10], to which and from which μ-calculus model checking can be reduced [11]. Without loss of generality, we assume that existential and universal vertices alternate (the turn-based game graph is bipartite), and each node has exactly two successors. We also assume that each vertex in V_u has the same color as the incoming existential nodes (otherwise, we can duplicate nodes and get an equivalent game with this property) We assume that each vertex of G is labeled by the color $C(v)$, thus we can refer to G as a Kripke structure with $AP = \{0, \ldots, k-1\}$: the proposition i holds at vertex v iff $C(v) = i$. From [10], player 1 wins the parity game G at an existential vertex $s \in V_e$ iff

$$G, s \models \lambda_{k-1} x_{k-1} \ldots \mu x_1 . \nu x_0 . \; (\bigvee_{i \in [0 \ldots (k-1)]} (i \wedge EXAXx_i)),$$

where $\lambda_n = \nu$ if n is even, and $\lambda_n = \mu$ if n is odd.

The formula above uses both universal and existential next modalities. By transforming G to a Kripke structure $\mathcal{K}$ as in the proof of Theorem 4, we can use left and right labels to vertices in the graph and use only one type of branching mode. Formally, let $\mathcal{K}$ be the Kripke structure induced by G. Then, player 1 wins the parity game in G at a node s iff

$$\mathcal{K}, s \models \lambda_{k-1} x_{k-1} \ldots \mu x_1 . \nu x_0 . \; (\bigvee_{i \in [0 \ldots (k-1)]} (i \wedge ((AX(\neg l \vee AXx_i)) \vee (AX(\neg r \vee AXx_i)))$$

□

If the syntax of the μ-calculus is equipped with actions on moves, then the above result follows immediately, because there is no distinction between existential and universal next modalities. Our proof shows that the result follows even if no such labeling is available. Our results in this section show that the efforts to find a polynomial-time model-checking algorithm for the μ-calculus can be replaced by efforts to find such an algorithm for the universal or existential fragments of μ-calculus.

5 Implication

The *implication problem* for a logic asks if one specification logically implies another specification; formally, given formulas φ and ψ of the logic, if the formula $\varphi \rightarrow \psi$ is valid. It arises naturally in *modular verification* [27, 22], where the antecedent of the implication is the assumption about the behavior of a component's environment, and the consequent is a guarantee about the behavior of the component. For logics closed under negation, the implication problem is equivalent to validity: a formula φ is valid iff $true \rightarrow \varphi$. Thus, the implication problem for the μ-calculus is EXPTIME-complete. However, for the existential and universal fragments of the μ-calculus, this is not the case: the implication problem combines both universal and existential formulas, and is more general than satisfiability or validity.

Theorem 6. *The implication problem for $\exists MC$ and $\exists AFMC$ (so, also for $\forall MC$ and $\forall AFMC$) is EXPTIME-complete.*

Proof. For formulas φ_1 and φ_2 of $\exists MC$, we have that $\varphi_1 \rightarrow \varphi_2$ iff the formula $\varphi_1 \wedge \neg\varphi_2$ is not satisfiable. Membership in EXPTIME follows from the complexity of the satisfiability problem for the μ-calculus. Note that $\neg\varphi_2$ is a formula of $\forall MC$, thus we cannot apply the results of Section 3.

To prove hardness in EXPTIME, we do a reduction from the satisfiability problem of $AFMC$, proved to be EXPTIME-hard in [13]. Given an $AFMC$ formula ψ, we construct a formula φ_A of $\forall AFMC$ and a formula φ_E of $\exists AFMC$ such that the conjunction $\varphi = \varphi_E \wedge \varphi_A$ is satisfiable iff ψ is satisfiable. For simplicity, we assume that ψ is satisfied iff it is satisfied in a tree of branching degree two. Note that while our assumption does not hold for all $AFMC$ formulas, the EXPTIME-hardness of the satisfiability problem for $AFMC$ holds already for such formulas, which is sufficiently good for our goal here.

Intuitively, the formula φ_E would require the states of models of φ to be attributed by directions so that at least one successor is labeled by l and at least one successor is labeled by r. In addition, φ_A would contain a conjunct that require each state to be labeled by at most one direction. Thus, states that are labeled by l cannot be labeled by r, and vice versa. Then, the other conjunct of φ_A is obtained from ψ by replacing an existential next modality by a disjunction over the successors of a state.

Formally, the formula $\varphi_E = \nu y.EX(l \wedge y) \wedge EX(r \wedge y)$ requires each state (except for the initial state) to have at least two successors, labeled by different directions, and the formula $\varphi_A^1 = \nu y.((\neg l) \vee (\neg r)) \wedge AXy$ requires each state to be labeled by at most one direction.

Then, the formula φ_A^2 is obtained from ψ by replacing a subformula of the form $EX\theta$ by the formula $AX(r \vee \theta) \vee AX(l \vee \theta)$. We show that for every ψ such that ψ is satisfiable iff it is satisfiable in a model of branching degree two, we have that ψ is satisfiable iff $\varphi_E \wedge \varphi_A^1 \wedge \varphi_A^2$ is satisfiable. First, if ψ is satisfiable, then there is a tree of branching degree two that satisfies it. This tree can be attributed with l and r so that it satisfies the formula $\varphi_E \wedge \varphi_A^1 \wedge \varphi_A^2$, by labeling the left successor of each node with $\{l, \neg r\}$ and the right successor of each node with $\{\neg l, r\}$. On the other hand, assume that the formula $\varphi_E \wedge \varphi_A^1 \wedge \varphi_A^2$ is satisfiable in a model M. The subformula $\varphi_E \wedge \varphi_A^1$ guarantees that each state of M has at least one successor that is not labeled l and at least one successor that is not labeled r. Accordingly, each subformula of the form $AX(r \vee \theta) \vee AX(l \vee \theta)$ is satisfied in a state w of M iff w has a successor that satisfies θ, thus w satisfies $EX\theta$. Hence, the model M also satisfies ψ. □

The above proof constructs, given a formula φ of μ-calculus, two formulas φ_E and φ_A such that φ_E is an existential formula, φ_A is a universal formula, and φ is satisfiable iff it is satisfiable in a model of branching degree two iff $\varphi_E \wedge \varphi_A$ is satisfiable. However, one cannot in general construct formulas φ_E and φ_A such that $\varphi_E \wedge \varphi_A$ is equivalent to φ. This can be proved considering two states of a Kripke structure that are similar, but not bisimilar, and the formula of μ-calculus that distinguishes them.

Note that the implication problem for $\forall$CTL* and $\exists$CTL* are EXPSPACE-complete [22], easier than the satisfiability problem for CTL*, which is 2EXPTIME-complete. The above construction does not work for $\forall$CTL*, as the formula φ_E used to label the states of a model by directions specifies an unbounded number of unfoldings of the structure. On the other hand, the number of unfoldings expressible by an $\exists$CTL* formula is bounded by the size of the formula; thus, the formula φ_E does not have an equivalent formula in $\exists$CTL*.

6 Discussion

We studied the complexity of the satisfiability, validity, model checking, and implication problems for the universal and existential fragments of the μ-calculus and its alternation-free fragment. We proved that the linear-size model property, which is known for $\exists$CTL, holds also for $\exists MC$. Interestingly, the property does not hold for $\exists$CTL*, which is less expressive than $\exists MC$. Thus, the picture we get here for $\exists MC$ and $\exists AFMC$ is different than the one known for $\exists$CTL* and $\exists$CTL. For the universal fragments $\forall MC$ and $\forall AFMC$, the picture does agree with the one known for $\forall$CTL* and $\forall$CTL, and the complexity of the satisfiability problem coincides with the complexity known for the linear versions of the logics.

We showed how labeling of states with directions makes the model-checking and the implication problems for the universal and existential problems as hard as for the full logics. While such a labeling is straightforward in the case of model checking, it is not always possible for implication. Indeed, in the case of CTL* and CTL, formulas cannot specify a legal labeling, making the implication problem for $\forall$CTL* and $\exists$CTL* strictly easier than the implication problem for CTL*, and similarly for CTL. In contrast, here we were able to label the directions legally using a $\forall AFMC$ formula, making the implication problems for $\forall MC$ and $\exists MC$ as hard as the one for MC, and similarly for the alternation free fragments. Another way to see the importance of the fixpoint modalities is to observe that the implication problem for Modal Logic (μ-calculus with no fixpoint modalities) and its universal and existential fragments is co-NP-complete.

Finally, the *equivalence problem* for a logic asks, given formulas φ and ψ, if the formula $\varphi \leftrightarrow \psi$ is valid. The equivalence problem for the μ-calculus is EXPTIME-complete, by easy reductions to and from satisfiability. This gives an EXPTIME upper bound for the equivalence problems for $\exists MC$ and $\forall MC$. By a reduction from satisfiability or validity (whichever is harder), we also get a PSPACE lower bound. However, the exact complexity for the equivalence problem remains open.

The gap above highlights the difficulty in studying the universal and existential fragments of the μ-calculus. It is easy to see that in all formalisms that are closed under complementation (in particular, full MC), equivalence is as hard as satisfiability. Indeed, φ and ψ are equivalent iff $(\varphi \wedge \neg\psi) \vee (\psi \wedge \neg\varphi)$ is not satisfiable. When a formalism is not closed under complementation, equivalence is not harder than implication, and is not easier than satisfiability or validity, whichever is harder. In the case of CTL, for example, it is easy to see that the

equivalence problem for $\forall$CTL and $\exists$CTL is PSPACE-complete, as implication has the same complexity as satisfiability or validity (whichever is harder). The same holds for word automata: if we relate the existential fragment with nondeterministic automata and the universal fragment with universal automata, then in both cases the language-containment problem (the automata-theoretic counterpart of implication) has the same complexity as the (harder between the) nonemptiness and the universality problems (the automata-theoretic counterparts of satisfiability and validity). Once we do not allow fixpoint modalities, the same holds for the μ-calculus: the equivalence problem for Modal Logic and its universal and existential fragments is co-NP-complete, as the co-NP-hardness of the implication problem applies already for the validity problem. So, in all the cases we know, except for the universal and existential fragments of CTL* and the μ-calculus, the above immediate upper and lower bounds do not induce a gap, and the tight complexity of the equivalence problem is known.

References

1. B. Banieqbal and H. Barringer. Temporal logic with fixed points. *Temporal Logic in Specification*, LNCS 398, pages 62–74. Springer-Verlag, 1987.
2. J. Burch, E. Clarke, K. McMillan, D. Dill, and L. Hwang. Symbolic model checking: 10^{20} states and beyond. *Information and Computation*, 98(2):142–170, June 1992.
3. E. Clarke and E. Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In *Proc. Workshop on Logic of Programs*, LNCS 131, pages 52–71. Springer-Verlag, 1981.
4. E. Clarke, E. Emerson, and A. Sistla. Automatic verification of finite-state concurrent systems using temporal-logic specifications. *ACM Transactions on Programming Languages and Systems*, 8(2):244–263, 1986.
5. E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 1999.
6. R. Cleaveland and B. Steffen. A linear-time model-checking algorithm for the alternation-free modal μ-calculus. *Formal Methods in System Design*, 2:121–147, 1993.
7. M. Dam. CTL* and ECTL* as fragments of the modal μ-calculus. *Theoretical Computer Science*, 126:77–96, 1994.
8. E. Emerson and J. Halpern. Sometimes and not never revisited: On branching versus linear time. *Journal of the ACM*, 33(1):151–178, 1986.
9. E. Emerson and C. Jutla. The complexity of tree automata and logics of programs. In *Proc. Foundations of Computer Science*, pages 328–337, 1988.
10. E. Emerson and C. Jutla. Tree automata, mu-calculus and determinacy. In *Proc. Foundations of Computer Science*, pages 368–377. IEEE Computer Society Press, 1991.
11. E. Emerson, C. Jutla, and A. Sistla. On model-checking for fragments of μ-calculus. In *Computer Aided Verification*, LNCS 697, pages 385–396. Springer-Verlag, 1993.
12. E. Emerson and C.-L. Lei. Efficient model checking in fragments of the propositional μ-calculus. In *Proc. Symp. on Logic in Computer Science*, pages 267–278. IEEE Press, 1986.
13. M. Fischer and R. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and Systems Sciences*, 18:194–211, 1979.

14. R. Greenlaw, H. Hoover, and W. Ruzzo. *Limits of parallel computation*. Oxford University Press, 1995.
15. M. Hennessy and R. Milner. Algebraic laws for nondeterminism and concurrency. *Journal of ACM*, 32:137–161, 1985.
16. N. Immerman. Number of quantifiers is better than number of tape cells. *Journal of Computer and System Sciences*, 22(3):384–406, 1981.
17. D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ-calculus with respect to the monadic second order logic. In *Concurrency Theory*, LNCS 1119, pages 263–277. Springer-Verlag, 1996.
18. M. Jurdzinski. Deciding the winner in parity games is in UP $\cap$ co-UP. *Information Processing Letters*, 68(3):119–124, 1998.
19. M. Jurdzinski. Small progress measures for solving parity games. In *Symposium on Theoretical Aspects of Computer Science*, LNCS 1770, pages 290–301. Springer-Varlag, 2000.
20. D. Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27:333–354, 1983.
21. D. Kozen. A finite model theorem for the propositional μ-calculus. *Studia Logica*, 47(3):333–354, 1988.
22. O. Kupferman and M. Vardi. An automata-theoretic approach to modular model checking. *ACM Transactions on Programming Languages and Systems*, 22:87–128, 2000.
23. O. Kupferman, M. Vardi, and P. Wolper. An automata-theoretic approach to branching-time model checking. *Journal of the ACM*, 47(2):312–360, March 2000.
24. R. Kurshan. *FormalCheck User's Manual*. Cadence Design, Inc., 1998.
25. D. Long, A. Brown, E. Clarke, S. Jha, and W. Marrero. An improved algorithm for the evaluation of fixpoint expressions. In *Computer Aided Verification*, LNCS 818, pages 338–350. Springer-Verlag, 1994.
26. K. McMillan. *Symbolic Model Checking*. Kluwer Academic Publishers, 1993.
27. A. Pnueli. In transition from global to modular temporal reasoning about programs. In *Logics and Models of Concurrent Systems*, volume F-13 of *NATO Advanced Summer Institutes*, pages 123–144. Springer-Verlag, 1985.
28. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *ACM Symp. on Principles of Programming Languages*, pages 179–190. ACM Press, 1989.
29. P. Ramadge and W. Wonham. The control of discrete event systems. *IEEE Transactions on Control Theory*, 77:81–98, 1989.
30. H. Seidl. Fast and simple nested fixpoints. *Information Processing Letters*, 59(6):303–308, 1996.
31. M. Vardi. A temporal fixpoint calculus. In *Proceedings of the 15th Annual Symposium on Principles of Programming Languages*, pages 250–259. ACM Press, 1988.
32. M. Vardi and L. Stockmeyer. Improved upper and lower bounds for modal logics of programs. In *Proc 17th ACM Symp. on Theory of Computing*, pages 240–251, 1985.