Inflationary Fixed Points in Modal Logic

ANUJ DAWAR
University of Cambridge Computer Laboratory
ERICH GRÄDEL and STEPHAN KREUTZER
RWTH Aachen

We consider an extension of modal logic with an operator for constructing inflationary fixed points, just as the modal μ-calculus extends basic modal logic with an operator for least fixed points. Least and inflationary fixed point operators have been studied and compared in other contexts, particularly in finite model theory, where it is known that the logics IFP and LFP that result from adding such fixed point operators to first order logic have equal expressive power. As we show, the situation in modal logic is quite different, as the modal iteration calculus (MIC) we introduce has much greater expressive power than the μ-calculus. Greater expressive power comes at a cost: the calculus is algorithmically much less manageable.

Categories and Subject Descriptors: F.4.1 [**Mathematical Logic**]: Mathematical Logic

General Terms: Theory

Additional Key Words and Phrases: Fixed-point logics, modal logic, expressive power, decidability, complexity

1. INTRODUCTION

The modal μ-calculus L_μ is an extension of multi-modal logic with an operator for forming least fixed points. This logic has been extensively studied, having acquired importance for a number of reasons. In terms of expressive power, it subsumes a variety of modal and temporal logics used in verification, in particular LTL, CTL, CTL*, PDL and also many logics used in other areas of computer science, for instance description logics. On the other hand, L_μ has a rich theory, and is well-behaved in model-theoretic and algorithmic terms.

The logic L_μ is only one instance of a logic with an explicit operator for forming least fixed points. Indeed, in recent years, a number of fixed point extensions of first order logic have been studied in the context of finite model theory. It may be argued that fixed point logics play a central role in finite model theory, more important than first order logic itself. The best known of these fixed point logics is LFP, which extends first order logic with an operator for forming the least fixed points of positive formulae, defining monotone operators. In this sense, it relates to

Author's Addresses: A. Dawar, University of Cambridge Computer Laboratory, Cambridge CB3 0FD, UK, anuj.dawar@cl.cam.ac.uk. Research supported by EPSRC grant GR/N23028.
E. Grädel and S. Kreutzer, Aachen University of Technology, Mathematical Foundations of Computer Science, D-52065 Aachen, {graedel,kreutzer}@informatik.rwth-aachen.de.

first order logic in much the same way as L_μ relates to propositional modal logic. However, a number of other fixed point operators have been extensively studied in finite model theory, including inflationary, partial, nondeterministic and alternating fixed points. All of these have in common that they allow the construction of fixed points of operators that are not necessarily monotone.

Furthermore, a variety of fragments of the fixed point logics formed have been studied, such as existential and stratified fragments, bounded fixed point logics, transitive closure logic and varieties of Datalog. Thus, there is a rich theory of the structure and expressive power of fixed point logics on finite relational structures and, to a lesser extent, on infinite structures.

In the present paper, we take a first step in the study of extensions of propositional modal logic by operators that allow us to form fixed points of non-monotone formulae. We focus on the simplest of these, that is the inflationary fixed point (also sometimes called the iterative fixed point). Though the inflationary fixed point extension of first order logic (IFP) is often used interchangeably with LFP, as the two have the same expressive power on finite structures, we show that in the context of modal logic, the inflationary fixed point behaves quite differently from the least fixed point.

Least and Inflationary Inductions. We begin by reviewing the known results on the logics LFP and IFP.

(1) On finite structures, LFP and IFP have the same expressive power [Gurevich and Shelah 1986]. A recent result shows that this equivalence of LFP and IFP also extends to infinite structures (see [Kreutzer 2002]).
(2) On ordered finite structures, LFP and IFP express precisely the properties that are decidable in polynomial time.
(3) Simultaneous least or inflationary inductions do not provide more expressive power than simple inductions.
(4) The complexity of evaluating a formula ψ in LFP or IFP on a given finite structure $\mathfrak{A}$ is polynomial in the size of the structure, but exponential in the length of the formula. For formulae with a bounded number k of variables, the evaluation problem is PSPACE-complete [Dziembowski 1996], even for $k = 2$ and on fixed (and very small) structures. If, in addition to bounding the number of variables one also forbids parameters in fixed point formulae, the evaluation problem for LFP is computationally equivalent to the model checking problem for L_μ [Grädel and Otto 1999; Vardi 1995] which is known to be in NP $\cap$ Co-NP, in fact in UP $\cap$ Co-UP [Jurdzinski 1998], and hard for PTIME. It is an open problem whether this problem can be solved in polynomial time. The model checking problem for bounded variable IFP does not appear to have been studied previously.

We also note that even though IFP does not provide more expressive power than LFP on finite structures, it is often more convenient to use inflationary inductions in explicit constructions. The advantage of using IFP is that one is not restricted to inductions over positive formulae. A non-trivial case in point is the formula defining an order on the k-variable types in a finite structure, an essential ingredient of the proof of the Abiteboul-Vianu Theorem, saying that least and partial fixed point logics coincide if and only if PTIME = PSPACE (see [Abiteboul and Vianu 1995; Dawar 1993; Dawar et al. 1995; Ebbinghaus and Flum 1999]). Furthermore, IFP is more robust, in the sense that inflationary fixed points are well-defined, even when other, non-monotone, operators are added to the language (see, for instance, [Dawar and Hella 1995]).

Inflationary Inductions in Modal Logic. Given the close relationship between LFP and IFP on finite structures, and the importance of the μ-calculus, it is natural to study also the properties and expressive power of inflationary fixed points in modal logic. In this paper, we undertake a study of an analogue of IFP for modal logic. We define a modal iteration calculus, MIC, by extending basic multi-modal logic with simultaneous inflationary inductions. While deferring formal definitions until Section 2, we begin with an informal explanation.

In L_μ, we can write formulae $\mu X.\varphi$, which are true in state s of a transition system $\mathcal{K}$ if, and only if, s is in the least set X satisfying $X \leftrightarrow \varphi$ in $\mathcal{K}$. We can do this, provided that the variable X appears only positively in φ. This guarantees that φ defines a monotone operator and has a least fixed point. Moreover, the fixed point can be obtained by an iterative process. Starting with the empty set, if we repeatedly apply the operator defined by φ (possibly through a transfinite series of stages), we obtain an increasing sequence of sets, which converges to the desired least fixed point. If, on the other hand, φ is not positive in X, we can still define an increasing sequence of sets, by starting with the empty set, and iteratively taking the union of the current set X with the set of states satisfying $\varphi(X)$, and this sequence must eventually converge to a fixed point (not necessarily of φ, but of the operator that maps X to $X \vee \varphi(X)$). More generally, we allow formulae $\mathbf{ifp}\ X_i : [X_1 \leftarrow \varphi_1, \dots, X_k \leftarrow \varphi_k]$ that construct sets by a simultaneous inflationary induction. At each stage α, we have a tuple of sets $X_1^\alpha, \dots, X_k^\alpha$. Substituting these into the formulae $\varphi_1, \dots, \varphi_k$ we obtain a new tuple of sets, which we *add* to the existing sets $X_1^\alpha, \dots, X_k^\alpha$, to obtain the next stage.

It is clear that MIC is a modal logic in the sense that it is invariant under bisimulation. In fact, on every class of bounded cardinality, inflationary fixed points can be unwound to obtain equivalent infinitary modal formulae. As a consequence, MIC has the tree model property. It is also clear that MIC is at least as expressive as L_μ. The following natural questions now arise.

(1) Is MIC more expressive than L_μ?
(2) Does MIC have the finite model property?
(3) What are the algorithmic properties of MIC? Is the satisfiability problem decidable? Can model checking be performed efficiently (as efficiently as for L_μ)?
(4) Can we eliminate, as in the μ-calculus and as in IFP, simultaneous inductions without losing expressive power?
(5) What is the relationship of MIC with monadic second-order logic (MSO) and with finite automata? Or more generally, what are the 'right' automata for MIC?
(6) Is MIC the bisimulation-invariant fragment of any natural logic (as L_μ is the bisimulation-invariant fragment of MSO [Janin and Walukiewicz 1996])? It

has recently been proved that the bisimulation invariant fragment of the most natural candidate for this, the monadic fragment of inflationary fixed-point logic, has much greater expressive power than MIC [Dawar and Kreutzer 2002]. Thus, if such a logic exists, its expressive power must be somewhere between MSO and monadic IFP.

We provide answers to most of these questions. From an algorithmic point of view, most of the answers are negative. From the point of view of expressiveness, we can say that in the context of modal logic, inflationary fixed points provide much more expressive power than least fixed points, and MIC has very different structural properties to L_μ. In particular, we establish the following results:

(1) There exist MIC-definable languages that are not regular. Hence MIC is more expressive than the μ-calculus, and does not translate to monadic second-order logic.
(2) MIC does not have the finite model property.
(3) The satisfiability problem for MIC is undecidable. In fact, it is not even in the arithmetic hierarchy.
(4) The model checking problem for MIC is PSPACE-complete.
(5) Simultaneous inflationary inductions do provide more expressive power than simple inflationary inductions. Nevertheless the algorithmic intractability results for MIC apply also to MIC without simultaneous inductions.
(6) There are bisimulation-invariant polynomial time properties that are not expressible in MIC.
(7) All languages in DTIME$(O(n))$ are MIC-definable.

No doubt, these properties exclude MIC as a candidate logic for hardware verification. On the other hand, the present study is an investigation into the structure of the inflationary fixed point operator and may suggest tractable fragments of the logic MIC, which involve crucial use of an inflationary operator, just as logics like CTL and alternation-free L_μ carve out efficiently tractable fragments of L_μ. In any case, it delineates the differences between inflationary and least fixed point constructs in the context of modal logic.

In the rest of this paper, we begin in Section 2 by giving the necessary background on modal logic and fixed points, and giving the definition of MIC, along with an example that illustrates how this calculus has higher expressive power than L_μ. Section 3 establishes that MIC fails to have the finite model property and that the satisfiability problem is highly undecidable. In Section 4 we investigate questions of the computational complexity of MIC in the context of finite transition systems. We show that the model checking problem is PSPACE-complete and that the class of models of any MIC formula is decidable in both polynomial time and linear space. In Section 5 we investigate the expressive power of MIC on finite words, establishing that there are languages definable in MIC that are not context-free, and that every linear time decidable language is expressible in MIC. In Section 6 we prove that there are polynomial time bisimulation-invariant properties that are not expressible in MIC. Further, we discuss the automaticity of MIC-definable sets of finite words and finite trees. Finally, in Section 7 we prove that 1MIC, the fragment of MIC that uses only simple induction, has less expressive power than full MIC.

2. THE MODAL ITERATION CALCULUS

Before we define the modal iteration calculus, we briefly recall the definitions of propositional modal logic ML and the μ-calculus L_μ.

2.1 Propositional Modal Logic.

Transition Systems. Modal logics are interpreted on transition systems (also called Kripke structures). Fix a set A of actions and a set $\mathcal{W}$ of atomic propositions. A transition system for A and $\mathcal{W}$ is a structure $\mathcal{K}$ with universe V (whose elements are called states) binary relations $E_a \subseteq V \times V$ for each $a \in A$ and monadic relations $p \subseteq V$ for each atomic proposition $p \in \mathcal{W}$ (we do not distinguish notationally between atomic propositions and their interpretations.)

Syntax of ML. For a set A of actions and a set $\mathcal{W}$ of proposition variables, the formulae of ML are built from *false*, *true* and the variables $p \in \mathcal{W}$ by means of Boolean connectives $\wedge$, $\vee$, $\neg$ and modal operators $\langle a \rangle$ and $[a]$. That is, if ψ is a formula of ML and $a \in A$ is an action, then $\langle a \rangle\psi$ and $[a]\psi$ are also formulae of ML. If there is only one action in A, one simply writes $\Box$ and $\Diamond$ for $[a]$ and $\langle a \rangle$, respectively.

Semantics of ML. The formulae of ML are evaluated on transition systems at a particular state. Given a formula ψ and a transition system $\mathcal{K}$ with state v, we write $\mathcal{K}, v \models \psi$ to denote that the formula ψ holds in $\mathcal{K}$ at state v. We also write $[\![\psi]\!]^{\mathcal{K}}$ to denote the set of states v, such that $\mathcal{K}, v \models \psi$. In the case of atomic propositions, $\psi = p$, we have $[\![p]\!]^{\mathcal{K}} = p$. Boolean connectives are treated in the natural way. Finally for the semantics of the modal operators we put

$$[\![\langle a \rangle\psi]\!]^{\mathcal{K}} := \{v : \text{ there exists a state } w \text{ such that } (v,w) \in E_a \text{ and } w \in [\![\psi]\!]^{\mathcal{K}}\}$$
$$[\![[a]\psi]\!]^{\mathcal{K}} := \{v : \text{ for all } w \text{ such that } (v,w) \in E_a, \text{ we have } w \in [\![\psi]\!]^{\mathcal{K}}\}.$$

Hence $\langle a \rangle$ and $[a]$ can be viewed as existential and universal quantifiers 'along a-transitions'.

For background on propositional modal logic, we recommend [Blackburn et al. 2001].

2.2 The μ-calculus L_μ.

Syntax of L_μ. The μ-calculus extends propositional modal logic ML by the following rule for building fixed point formulae: if ψ is a formula in L_μ and X is a propositional variable that occurs only positively in ψ, then $\mu X.\psi$ and $\nu X.\psi$ are L_μ formulae.

Semantics of L_μ. A formula $\psi(X)$ with a propositional variable X defines on every transition system $\mathcal{K}$ (with state set V, and with interpretations for free variables other than X occurring in ψ) an operator $\psi^{\mathcal{K}} : \mathcal{P}(V) \to \mathcal{P}(V)$ assigning to every set $X \subseteq V$ the set $\psi^{\mathcal{K}}(X) := [\![\psi]\!]^{\mathcal{K},X} = \{v \in V : (\mathcal{K}, X), v \models \psi\}$.

As X occurs only positively in ψ, the operator $\psi^{\mathcal{K}}$ is *monotone* for every $\mathcal{K}$, and therefore, by a well-known theorem due to Knaster and Tarski, has a least fixed point $\mathbf{lfp}(\psi^{\mathcal{K}})$ and a greatest fixed point $\mathbf{gfp}(\psi^{\mathcal{K}})$. Now we put $[\![\mu X.\psi]\!]^{\mathcal{K}} := \mathbf{lfp}(\psi^{\mathcal{K}})$ and $[\![\nu X.\psi]\!]^{\mathcal{K}} := \mathbf{gfp}(\psi^{\mathcal{K}})$.

Least (and greatest) fixed points can also be constructed inductively. Given a formula $\mu X.\psi(X)$, we define for each ordinal α, the stage X^α of the **lfp**-induction of $\psi^{\mathcal{K}}$ by $X^0 := \emptyset$, $X^{\alpha+1} := [\![\psi]\!]^{(\mathcal{K},X^\alpha)}$, and $X^\alpha := \bigcup_{\beta<\alpha} X^\beta$ if α is a limit ordinal.

By monotonicity, the stages of the **lfp**-induction increase until a fixed point is reached. The first ordinal at which this happens is called the *closure ordinal* of the induction. By ordinal induction, one easily proves that this inductively constructed fixed point coincides with the least fixed point. The cardinality of a closure ordinal cannot be larger than the cardinality of $\mathcal{K}$.

For any formula φ, the formula $\nu X.\varphi$ is equivalent to $\neg\mu X.\neg\varphi(\neg X)$, where $\varphi(\neg X)$ denotes the formula obtained from φ by replacing all occurrences of X with $\neg X$.

Simultaneous Fixed Points. There is a variant of L_μ that admits systems of simultaneous fixed points. These do not increase the expressive power but sometimes allow for more straightforward formalisations. Here one associates with any tuple $\overline{\psi} = (\psi_1, \dots, \psi_k)$ of formulae $\psi_i(\overline{X}) = \psi_i(X_1, \dots, X_k)$, in which all occurrences of all X_i are positive, a new formula $\varphi = \mu\overline{X}.\overline{\psi}$. The semantics of φ is induced by the least fixed point of the monotone operator $\psi^{\mathcal{K}}$ mapping $\overline{X}$ to $\overline{X}'$ where $X'_i = \{v \in V : (\mathcal{K}, \overline{X}), v \models \psi_i\}$. More precisely, $\mathcal{K}, v \models \varphi$ iff v is an element of the first component of the least fixed point of the above operator. It is known that simultaneous least fixed points can be eliminated in favour of nested individual fixed points (see e.g. [Arnold and Niwiński 2001, page 27]). Indeed, $\mu XY.[\psi(X,Y), \varphi(X,Y)]$ is equivalent to $\mu X.\psi(X, \mu Y.\varphi(X,Y))$, and this equivalence generalises to larger systems in the obvious way.

Bisimulations and Tree Model Property. Bisimulation is a notion of behavioural equivalence for transition systems. Modal logics, like ML, CTL, the μ-calculus etc. do not distinguish between transition systems that are bisimulation equivalent. Formally, given two transition systems $\mathcal{K}$ and $\mathcal{K}'$, with distinguished states v and v' respectively, we say that $\mathcal{K}, v$ is bisimulation equivalent to $\mathcal{K}', v'$, written $\mathcal{K}, v \sim \mathcal{K}', v'$, if there is a relation $R \subseteq V \times V'$ between the states of $\mathcal{K}$ and the states of $\mathcal{K}'$ such that: (1) $(v, v') \in R$; (2) for each atomic proposition $p \in \mathcal{W}$ and each $(u, u') \in R$, $u \in [\![p]\!]^{\mathcal{K}}$ if, and only if, $u' \in [\![p]\!]^{\mathcal{K}'}$; (3) for each $(u, u') \in R$, and each $t \in V$ such that $(u, t) \in E_a$, there is a $t' \in V'$ with $(u', t') \in E'_a$ and $(t, t') \in R$; and (4) for each $(u, u') \in R$, and each $t' \in V'$ such that $(u', t') \in E'_a$, there is a $t \in V$ with $(u, t) \in E_a$ and $(t, t') \in R$.

Bisimulation equivalence corresponds to equivalence in an infinitary modal logic ML^∞ [van Benthem 1983]. This logic is the extension of ML by disjunctions and conjunctions taken over arbitrary sets of formulae. Thus, if S is any set (possibly infinite) of formulae of ML^∞, then $\bigwedge S$ and $\bigvee S$ are also formulae of ML^∞. It can be shown that for any transition systems $\mathcal{K}$ and $\mathcal{K}'$, $\mathcal{K}, v \sim \mathcal{K}', v'$ if, and only if, $\mathcal{K}, v$ makes true exactly the same formulae of ML^∞ as $\mathcal{K}', v'$.

A transition system is called a *tree*, if for every state v, there is at most one state u, and at most one action a such that $(u, v) \in E_a$ and there is exactly one state r, called the *root* of the tree, for which there is no state having a transition to r, and if every state is reachable from the root. It is known that for every transition system $\mathcal{K}$, and any state v, there is a tree $\mathcal{T}$ with root r such that $\mathcal{K}, v \sim \mathcal{T}, r$. One consequence of this is that any logic that respects bisimulation has the tree model property. For instance, for any formula φ of L_μ, if φ is satisfiable, then there is a tree $\mathcal{T}$ such that $\mathcal{T}, r \models \varphi$.

2.3 The Modal Iteration Calculus.

We are now ready to introduce MIC. Informally, MIC is propositional modal logic ML, augmented with simultaneous inflationary fixed points.

DEFINITION 2.1. The modal iteration calculus MIC extends propositional modal logic by the following rule: if $\varphi_1, \dots, \varphi_k$ are formulae of MIC, and $X_1, \dots, X_k$ are propositional variables, then

$$S := \begin{cases} X_1 \leftarrow \varphi_1 \\ \quad\vdots \\ X_k \leftarrow \varphi_k \end{cases}$$

is a *system* of rules, and $(\textbf{ifp}\ X_i : S)$ is a formula of MIC. If S consists of a single rule $X \leftarrow \varphi$ we simplify the notation and write $(\textbf{ifp}\ X \leftarrow \varphi)$ instead of $(\textbf{ifp}\ X : X \leftarrow \varphi)$.

Semantics: On every Kripke structure $\mathcal{K}$, the system S defines, for each ordinal α, a tuple $\overline{X}^\alpha = (X_1^\alpha, \dots, X_k^\alpha)$ of sets of states, via the following inflationary induction (for $i = 1, \dots, k$).

$$\begin{aligned} X_i^0 &:= \emptyset, \\ X_i^{\alpha+1} &:= X_i^\alpha \cup [\![\varphi_i]\!]^{(\mathcal{K}, \overline{X}^\alpha)}, \\ X_i^\alpha &:= \bigcup_{\beta<\alpha} X_i^\beta \text{ if } \alpha \text{ is a limit ordinal.} \end{aligned}$$

We call $(X_1^\alpha, \dots, X_k^\alpha)$ the stage α of the inflationary induction of S on $\mathcal{K}$. As the stages are increasing (i.e. $X_i^\alpha \subseteq X_i^\beta$ for any $\alpha < \beta$), this induction reaches a fixed point $(X_1^\infty, \dots, X_k^\infty)$. Now we put $[\![(\textbf{ifp}\ X_i : S)]\!]^{\mathcal{K}} := X_i^\infty$.

Before introducing examples in Section 2.4 and 3, we establish some simple properties of MIC.

LEMMA 2.2. $L_\mu \subseteq$ MIC. *Further, on every class of structures of bounded cardinality* MIC $\subseteq$ ML$^\infty$.

Proof. Clearly, if X occurs only positively in ψ, then $\mu X.\psi \equiv \textbf{ifp}\ X \leftarrow \psi$. Hence $L_\mu \subseteq$ MIC.

Now, let S be a system of rules $X_i \leftarrow \varphi_i(X_1, \dots, X_k)$. It is clear that for each ordinal α there exist formulae $\varphi_1^\alpha, \dots, \varphi_k^\alpha \in$ ML$^\infty$ defining, over any Kripke structure, the stage α of the induction by S. As closure ordinals are bounded on structures of bounded cardinality, the second claim follows. ☐

COROLLARY 2.3. MIC *is invariant under bisimulation and has the tree model property.*

Note that on classes of structures with *unbounded* cardinality, L_μ and MIC are not contained in ML$^\infty$. For instance, well-foundedness is expressed by the L_μ-formula $\mu X.\Box X$, but is known not to be expressible in ML$^\infty$.

2.4 Non-Regular Languages

We now demonstrate that MIC is strictly more expressive than L_μ. Recall that every formula of L_μ can be translated into a formula of monadic second order logic (MSO). Moreover, it is a well-known classical result [Büchi 1960] that the only sets of finite words that are expressible in MSO are the regular languages.

DEFINITION 2.4. *For our purposes, a word w of length n, in an alphabet Σ is a transition system with n states $v_1, \ldots, v_n$, a single action E such that $(v_i, v_j) \in E$ if, and only if, $j = i+1$ and an atomic proposition s for each $s \in \Sigma$, such that for each v_i, there is a unique s with $v_i \in s$.*

PROPOSITION 2.5. *There is a language that is expressible in* MIC *but not in* MSO.

Proof. The language $L := \{a^n b^m : n \leq m\}$ is not regular, hence not definable in monadic second-order logic, but it is definable in MIC. To see this, we consider first the formula $\pi(X) = (\textbf{ifp}\ Y \leftarrow \Diamond(b \wedge \neg X) \vee \Diamond(a \wedge X \wedge Y))$ which (since the rule is positive in Y) is in fact equivalent to a L_μ-formula. On every word $w = w_0 \cdots w_{n-1} \in \{a,b\}^*$ and $X \subseteq \{0, \ldots, n-1\}$, the formula is true if the word w starts with a (possibly empty) a-sequence inside X followed by a b outside X. Now the formula $(\textbf{ifp}\ X \leftarrow (a \wedge \pi(X)) \vee (b \wedge \Box X))$ defines (inside a^*b^*) the language L. Note that the language a^*b^* is definable in L_μ, so we can conjoin this definition to the above formula to obtain a definition of L which works on all words in $\{a,b\}^*$. □

The observation in Proposition 2.5 was pointed out to us in discussion by Martin Otto, and was the starting point of the investigation reported here.

3. INTERPRETING ARITHMETIC IN MIC

In this section we prove that the satisfiability problem of MIC is undecidable in a very strong sense. Given that MIC is invariant under bisimulation, we can restrict attention to trees. In fact we will only consider well-founded trees (i.e. trees satisfying the formula $\textbf{ifp}\ X \leftarrow \Box X$.) The *height* $h(v)$ of a node v in a well-founded tree $\mathcal{T}$ is an ordinal, namely the least strict upper bound of the heights of its children. For any node v in a tree $\mathcal{T}$, we write $\mathcal{T}(v)$ for the subtree of $\mathcal{T}$ with root v. We first show that the nodes of finite height and the nodes of height ω are definable in MIC.

LEMMA 3.1. *Let S be the system*

$$X \leftarrow \Box \mathit{false} \vee (\Box X \wedge \Diamond \neg Y)$$
$$Y \leftarrow X.$$

Then, on every tree $\mathcal{T}$, $[\![\textbf{ifp}\ X : S]\!]^{\mathcal{T}} = [\![\textbf{ifp}\ Y : S]\!]^{\mathcal{T}} = \{v : h(v) < \omega\}$.

Proof. By induction we see that for each $i < \omega$, $X^i = \{v : h(v) < i\}$ and $Y^i = X^{i-1} = \{v : h(v) < i-1\}$. As a consequence $X^\omega = Y^\omega = \{v : h(v) < \omega\}$. One further iteration shows that $X^{\omega+1} = Y^{\omega+1} = X^\omega$. □

With the system S exhibited in Lemma 3.1 we obtain the formulae

$$\text{finite-height} := (\mathbf{ifp}\ X : S)$$

and

$$\omega\text{-height} := \neg\text{finite-height} \land \Box\text{finite-height}$$

which define, respectively, the nodes of finite height and the nodes of height ω. Note that ω-height is a satisfiable formula all of whose models are infinite.

PROPOSITION 3.2. MIC *does not have the finite model property.*

We next show that the satisfiability problem of MIC is undecidable. In fact MIC interprets full arithmetic on the heights of nodes. To prove this we first define some auxiliary formulae that will be used frequently throughout the paper. We always assume that the underlying structure is a well-founded tree.

—The formula $\mathit{somewhere}(\varphi) := (\mathbf{ifp}\ X_s \leftarrow \varphi \lor \Diamond X_s)$ expresses that φ holds somewhere in the subtree of the current node: $\mathcal{T}, v \models \mathit{somewhere}(\varphi)$ iff $[\![\varphi]\!]^{\mathcal{T}} \cap \mathcal{T}(v) \neq \emptyset$.

—Dually $\mathit{everywhere}(\varphi) := (\mathbf{ifp}\ X_e \leftarrow \varphi \land \Box X_e)$ says that φ holds at all nodes of the subtree $\mathcal{T}(v)$.

—We say that a set X (in a tree $\mathcal{T}$) *encodes* the ordinal α if $X = \{v : h(v) < \alpha\}$. Let ordinal$(X)$ be the conjunction of the formula $\mathit{everywhere}(X \to \Box X)$ with

$$\begin{aligned}\neg(\mathbf{ifp}\ Z : Y &\leftarrow \Box Y \\ Z &\leftarrow \mathit{somewhere}(\neg Y \land \Box Y \land X) \land \mathit{somewhere}(\neg Y \land \Box Y \land \neg X)).\end{aligned}$$

It expresses that X encodes some ordinal. Indeed $\mathit{everywhere}(X \to \Box X)$ says that with each node $v \in X$, the entire subtree rooted at v is contained in X. The second conjunct is intended to ensure that two nodes of the same height are either both in X or both outside it. To see that it does this, note that at stage $\beta + 1$ of the induction nodes of height β are added to Y. At this stage, these are exactly the nodes satisfying $\neg Y \land \Box Y$. Thus, a node r is added to Z at this stage if, and only if, somewhere in the subtree rooted at r there are two nodes of height $\beta + 1$, one of which is in X and the other is not. Hence, at the end of the induction the root of the tree will *not* be contained in Z if, and only if, X does not distinguish between nodes of the same height. Together the two conjuncts imply that X contains all nodes up to some height.

—The formula number(X) = ordinal$(X) \land \mathit{somewhere}(\text{finite-height} \land \neg X)$ says that X encodes a natural number n (inside a tree of height $> n$).

LEMMA 3.3. *Let $\mathcal{T}$ be a well-founded tree of height ω. There exist formulae* plus(S,T) *and* times(S,T) *of* MIC *such that, whenever the sets S and T encode, in the tree $\mathcal{T}$, the natural numbers s and t, then $[\![\text{plus}(S,T)]\!]^{\mathcal{T}}$ encodes $s + t$, and $[\![\text{times}(S,T)]\!]^{\mathcal{T}}$ encodes st.*

Proof. Let

$$\begin{aligned}\text{plus}(S,T) := \mathbf{ifp}\ Y :\ X &\leftarrow \Box X \\ Y &\leftarrow S \lor (\Box Y \land \mathit{somewhere}(X) \land \mathit{everywhere}(X \to T)).\end{aligned}$$

Obviously at each stage n, we have $X^n = \{v : h(v) < n\}$. We claim that for each n, $Y^{n+1} = \{v : h(v) < s + \min(n,t)\}$. For $n = 0$ this is clear (note that for the case $s = 0$ this is true because the conjunct *somewhere*(X) prevents the Y-rule from being active at stage 1). For $n > 0$ the inclusion $X^n \subseteq T$ is true iff $n \leq t$. Hence we have $Y^{n+1} = \{v : h(v) < s + n\}$ in the case that $n \leq t$ and $Y^{n+1} = Y^n = \cdots = Y^t$ otherwise. To express multiplication we define

$$\begin{aligned}\text{times}(S,T) := \textbf{ifp } Y : \; & X \leftarrow \Box X \\ & Y \leftarrow \text{plus}(Y,S) \wedge \textit{everywhere}(\Box X \rightarrow T).\end{aligned}$$

We claim that $Y^n = \{v : h(v) < s \cdot \min(n,t)\}$. This is trivially true for $n = 0$. If it is true for $n < t$, then $Y^{n+1} = \{v : h(v) < sn + s\} = \{v : h(v) < s(n+1)\}$. Finally for $n \geq t$, the extension of $\Box X^n$ is $\{v : h(v) < n+1\}$ which is not contained in $T = \{v : h(v) < t\}$, hence $Y^{n+1} = Y^n = \cdots = Y^t$. □

Corollary 3.4. *For every polynomial $f(x_1,\ldots,x_r)$ with coefficients in the natural numbers there exists a formula $\psi_f(X_1,\ldots,X_r) \in$ MIC such that for every tree $\mathcal{T}$ of height ω and all sets $S_1,\ldots,S_r$ encoding numbers $s_1,\ldots,s_r \in \omega$*

$$[\![\psi_f(S_1,\ldots,S_r)]\!]^{\mathcal{T}} = \{v : h(v) < f(s_1,\ldots,s_r)\}.$$

Proof. By induction on f.

—$\psi_0 := \textit{false}$.

—$\psi_1 := \Box\textit{false}$.

—$\psi_{x_i} := X_i$.

—$\psi_{f+g} := \text{plus}[S/\psi_f, T/\psi_g]$, i.e. the formula obtained by replacing in plus(S,T) the variables S and T by, respectively, ψ_f and ψ_g.

—$\psi_{f\cdot g} := \text{times}[S/\psi_f, T/\psi_g]$.

□

Theorem 3.5. *For every first order sentence ψ in the vocabulary $\{+,\cdot,0,1\}$ of arithmetic, there exists a formula $\psi^* \in$ MIC such that ψ is true in the standard model $(\mathbb{N},+,\cdot,0,1)$ of arithmetic if, and only if, ψ^* is satisfiable.*

Proof. We have already seen that there exists a MIC-axiom ω-height axiomatising the models that are bisimilar to a tree of height ω. Further, we can express set equalities $X = Y$ by *everywhere*$(X \leftrightarrow Y)$ and we know how to represent polynomials by MIC-formulae. What remains is to translate quantifiers.

More precisely, we need to show that for each first order formula $\psi(y_1,\ldots,y_r)$ in the language of arithmetic there exists a MIC-formula $\psi^*(Y_1,\ldots,Y_r)$ such that on rooted trees $\mathcal{T},w$ of height ω and for all sets $S_1,\ldots,S_r$ that encode numbers $s_1,\ldots,s_r$ on $\mathcal{T}$ we have that $(\mathbb{N},+,\cdot,0,1) \models \psi(s_1,\ldots,s_r)$ iff $\mathcal{T},w \models \psi^*(S_1,\ldots,S_r)$.

Only the case of formulae of the form $\psi(\overline{y}) := \exists x\varphi(x,\overline{y})$ remains to be considered. By induction hypothesis, we assume that for $\varphi(x,\overline{y})$ the corresponding MIC-formula $\varphi^*(X,\overline{Y})$ has already been constructed. Now let

$$\begin{aligned}\psi^*(\overline{Y}) := \textbf{ifp } Z : \; & X \leftarrow \Box X \\ & Z \leftarrow \varphi^*(X,\overline{Y}) \wedge \text{number}(X).\end{aligned}$$

□

Corollary 3.6. *The satisfiability problem for* MIC *is undecidable. In fact, it is not even in the arithmetical hierarchy.*

The proof given above appears to rely crucially on the use of simultaneous inductions. Indeed, one can show that formulae of MIC involving simultaneous inductions, in particular the formula constructed in the proof of Lemma 3.1, cannot be expressed without simultaneous inductions (see Theorem 7.2).

However, we will prove in Section 7 that first order arithmetic can actually also be reduced to the satisfiability problem for 1MIC (MIC without simultaneous inductions).

Countable models. While MIC does not have the finite model property, it does have the Löwenheim-Skolem property, i.e. every satisfiable MIC-formula has a countable model. Indeed the proof of the Löwenheim-Skolem property for LFP (see e.g. [Flum 1999; Grädel 2003]), readily extends to the inflationary fixed point logic IFP which contains MIC. As a consequence, the satisfiability problem for MIC is in Σ^1_2, the second level of the analytical hierarchy.

4. THE MODEL CHECKING PROBLEM FOR MIC

Recall that the model checking problem for the μ-calculus is in UP ∩ Co-UP, and is conjectured by some to be solvable in polynomial time. We now show that MIC is algorithmically more complicated (unless PSPACE = NP).

We first observe that the naive bottom-up evaluation algorithm for MIC-formulae uses polynomial time with respect to the size of the input structure, and polynomial space (and exponential time) with respect to the length of the formula. Let $\mathcal{K}$ be a transition system with n nodes and m edges. The size $||\mathcal{K}||$ of appropriate encodings of $\mathcal{K}$ as an input for a model checking algorithm is $O(n+m)$. It is well known that the extension $[\![\varphi]\!]^{\mathcal{K}}$ of a basic modal formula φ (without fixed points) on a finite transition system $\mathcal{K}$ can be computed in time $O(|\varphi|\cdot||\mathcal{K}||)$. Further, any inflationary induction **ifp** $X_i : [X_1 \leftarrow \varphi_1, \ldots, X_k \leftarrow \varphi_k]$ reaches a fixed point on $\mathcal{K}$ after at most kn iterations. Hence, the bottom-up evaluation of a MIC-formula ψ with d nested simultaneous inflationary fixed points, each of width k, on $\mathcal{K}$ needs at most $O((kn)^d)$ basic evaluation steps. For each fixed point variable occurring in the formula, $2n$ bits of workspace are needed to record the current value and the last value of the induction. This gives the following complexity results.

Proposition 4.1. *Any* MIC *formula ψ of nesting depth d and simultaneous inductions of width at most k on a transition system $\mathcal{K}$ with n nodes can be evaluated in time $O((kn)^d|\psi|\cdot||\mathcal{K}||)$ and space $O(|\psi|\cdot n)$.*

In terms of common complexity classes the results can be stated as follows.

Theorem 4.2. *(1) The combined complexity of the model checking problem for* MIC *on finite structures is in* PSPACE.

(2) For any fixed formula $\psi \in$ MIC, the model checking problem for ψ on finite structures is solvable in polynomial time and linear space.

We now show that, contrary to the case of the μ-calculus, the complexity results obtained by this naive algorithm cannot be essentially improved.

THEOREM 4.3. *There exist transition systems $\mathcal{K}$, such that the model checking problem for* MIC *on $\mathcal{K}$ is* PSPACE*-complete (even for* MIC*-formulae without simultaneous inductions).*

Proof. The proof is by reduction from QBF (the evaluation problem for quantified Boolean formulae). For MIC with simultaneous inductions, this is trivial. Let $\mathcal{K}$ be the "structure" consisting just of a single point v. We associate with every quantified Boolean formula ψ a MIC-formula ψ^*. First, we eliminate the universal quantifiers in ψ in favour of negations and existential quantifiers. Then we replace inductively all subformulae $\varphi(\overline{Y}) := \exists X\vartheta(X, \overline{Y})$ by inflationary fixed points

$$\begin{aligned}\varphi^*(\overline{Y}) := \textbf{ifp } Z : X &\leftarrow \mathit{true}\\ Z &\leftarrow \vartheta^*(X, \overline{Y}).\end{aligned}$$

We have to show that, for any interpretation of $\overline{Y}$, we have $(\mathcal{K}, \overline{Y}), v \models \varphi^*(\overline{Y})$ iff $\varphi(\overline{Y})$ is true. (We identify truth values for $\overline{Y}$ with their truth values at v.) Note that $X^0 = \emptyset$ and $X^1 = \{v\}$. Hence v is included into Z^1 if $\varphi(0, \overline{Y})$ holds, and into Z^2 if $\varphi(1, \overline{Y})$ is true.

Note that we have translated quantified Boolean formulae into purely *propositional* formulae with inflationary fixed points (no modal operators are used).

For MIC without simultaneous fixed points, the construction is somewhat more complicated. Let $\mathcal{K}$ be the Kripke-structure consisting of two points 0,1, the atomic proposition $p = \{1\}$ and the complete transition relation $\{0, 1\} \times \{0, 1\}$.

Let $\alpha(X) := \neg X \wedge (p \rightarrow \Diamond X)$. Further, let $\varphi[X/\alpha(X)]$ denote the formula obtained from φ by replacing every free occurrence of X by $\alpha(X)$. The transformation from QBF-formulae ψ to MIC-formulae ψ^* without simultaneous fixed points is defined inductively as follows.

(1) For $\psi := X$ we set $\psi^* := (p \wedge X) \vee (\neg p \wedge \Diamond X)$,

(2) $(\neg\psi)^* := \neg\psi^*$ and $(\psi \circ \varphi)^* := \psi^* \circ \varphi^*$ for $\circ \in \{\wedge, \vee\}$,

(3) for $\psi := \forall X \varphi$ we put $\psi^* := \Box(\textbf{ifp } X \leftarrow \alpha(X) \wedge \varphi^*[X/\alpha(X)])$.

We say that a set $Y^* \subseteq \{0, 1\}$ represents $Y = \mathit{true}$ if $1 \in Y^*$, and that it represents $Y = \mathit{false}$ otherwise. We claim that for any QBF-formula ψ, any interpretation $\overline{Y}$ for the free variables of ψ and any representation of $\overline{Y}$ by $\overline{Y}^*$, we have

$$[\![\psi^*]\!]^{(\mathcal{K}, \overline{Y}^*)} = \begin{cases} \{0, 1\} & \text{if } \psi(\overline{Y}) \text{ is } \mathit{true} \\ \emptyset & \text{if } \psi(\overline{Y}) \text{ is } \mathit{false}. \end{cases}$$

(1) For $\psi := X$ and for Boolean combinations of formulae the verification of the claim is straightforward.

(2) Let $\psi = \forall X \varphi(X, \overline{Y})$. Note that, for $X = \emptyset$, $[\![\alpha(X)]\!]^{(\mathcal{K}, X)} = \{0\}$ and, for $X = \{0\}$, $[\![\alpha(X)]\!]^{(\mathcal{K}, X)} = \{1\}$. Thus, the first induction step of $(\textbf{ifp } X \leftarrow \alpha(X) \wedge \varphi^*[X/\alpha(X)])$ produces $X^1 = \{0\} \cap [\![\varphi^*(0)]\!]^{(\mathcal{K}, Y^*)}$, which, by induction, is $\{0\}$ if $\varphi(0, \overline{Y})$ evaluates to true and $X^1 = \emptyset$ otherwise. If the value $X^1 =$

$\{0\}$ was produced, then the second iteration step produces $X^2 = X^1 \cup (\{1\} \cap [\![\varphi^*(1)]\!]^{(\mathcal{K},\overline{Y}^*)})$. Thus,

$$[\![(\textbf{ifp } X \leftarrow \alpha(X) \wedge \varphi^*[X/\alpha(X)])]\!]^{(\mathcal{K},\overline{Y}^*)} = \begin{cases} \emptyset & \text{if } \neg\varphi(0,\overline{Y}) \\ \{0\} & \text{if } \varphi(0,\overline{Y}) \wedge \neg\varphi(1,\overline{Y}) \\ \{0,1\} & \text{if } \varphi(0,\overline{Y}) \wedge \varphi(1,\overline{Y}). \end{cases}$$

As required, in the first two cases the formula $\Box(\textbf{ifp } X \leftarrow \alpha(X) \wedge \varphi^*[X/\alpha(X)])$ is false at both states, and in the last case it is true at both states.

Finally note that the formula ψ^* can be computed in linear time from ψ. □

5. LANGUAGES

In this section we investigate the expressive power of MIC on finite strings. In other words we attempt to determine what languages are definable by formulae of MIC. As in Definition 2.4, we represent a word by a transition system consisting of a single path.

We have already seen in Proposition 2.5, that there are non-regular languages that are definable in MIC. We begin this section by strengthening this result and showing that there are languages definable in MIC that are not even context-free.

5.1 Non-CFLs in MIC

THEOREM 5.1. *There is a language definable in* MIC *that is not context-free.*

Proof. We show that the language $L = \{ww \mid w \in \{a,b\}^*\}$ is definable in MIC. To be precise, it is given by the formula

$$\begin{aligned} \varphi := \textbf{ifp } Y : X &\leftarrow \Box X \\ Y &\leftarrow \neg\psi, \end{aligned}$$

where ψ is given by:

$$\begin{aligned} \psi := \textbf{ifp } W : U &\leftarrow X \wedge \Box U \\ U' &\leftarrow U \\ V &\leftarrow (\neg X \wedge \Box X) \vee \Box V \\ V' &\leftarrow V \\ W &\leftarrow (\mathit{diff} \vee \mathit{len}_< \vee \mathit{len}_>) \end{aligned}$$

where, further,

$$\begin{aligned} \mathit{diff} :=& (\text{somewhere}(a \wedge U \wedge \neg U') \wedge \text{somewhere}(b \wedge V \wedge \neg V')) \vee \\ & (\text{somewhere}(b \wedge U \wedge \neg U') \wedge \text{somewhere}(a \wedge V \wedge \neg V')), \end{aligned}$$

$$\begin{aligned} \mathit{len}_< :=& V \wedge \text{somewhere}(X \wedge \neg U) \\ \mathit{len}_> :=& \neg V \wedge \text{everywhere}(X \rightarrow U). \end{aligned}$$

To understand the construction of the formula, note that in the outermost induction defined by φ interpreted on a word, at stage i, X contains the last i symbols in the word. We claim that at this stage ψ is true at all positions in the word except the one at distance $2i$ from the end and it is false at this position if, and only if, the word of length $2i$ starting from here is of the form ww. Thus, in the induction defined by φ, Y eventually contains all positions from which the rest of the word is of the form ww.

To verify the claim regarding ψ, note that at stage j of its induction ($j \leq i$), U contains the last j positions in the word while V contains the j positions immediately preceding those in X, while U' and V' lag one stage behind. A position is included in W if one of three conditions holds: all of X is included in U, but V has not yet reached the current position ($len_>$) which indicates that the current position is at distance greater than $2i$ from the end; V has reached the current position but not all of X is in U ($len_<$) which indicates that the positions is at distance less than $2i$ from the end; or the last symbol added to U is different to the one added to V (*diff*). Thus, as intended, the only position not included in W is the one at distance $2i$ from the end and then only if it is the beginning of a word of the form ww. □

While the above construction relies crucially on simultaneous inductions, it is not difficult to construct examples of non-context-free languages definable even in 1MIC. For example, in [Dawar et al. 2001], we exhibit a 1MIC formula defining the language $\{cwdw \mid w \in \{a, b\}^*\}$.

Finally, to place the expressive power of MIC in the Chomsky hierarchy, we note that every language definable in MIC can be defined by a context-sensitive grammar. This follows from the observation made in Section 4 that any class of finite structures defined by a formula of MIC is decidable in linear space, and the result that all languages decidable by nondeterministic linear space machines are definable by context-sensitive grammars.

5.2 Capturing Linear Time Languages

We have seen in Section 4 that the data complexity of evaluating MIC-formulae is in polynomial time and linear space. It is also clear that MIC can express PTIME-complete properties, as this is already the case for the μ-calculus.

On words the situation is somewhat different. The μ-calculus defines precisely the regular languages. On the other hand we have already seen that there exist MIC-definable languages that are not even context-free. We will now show that MIC can in fact define all languages that are decidable in linear time (by a Turing machine).

An observation that we will use in the proof, but which may well be of independent interest, is that cardinality comparisons and addition of cardinalities are expressible in MIC on words (recall that none of these are MSO-definable).

LEMMA 5.2. *There exists a formula $\varphi(X, Y)$ of* MIC *such that on every word w, we have $w, X, Y \models \varphi$ if and only if $|X| = |Y|$. Similarly for $|X| < |Y|$ and $|X| + |Y| = |Z|$.*

Proof. $|X| \neq |Y|$ is equivalent to

$$\begin{aligned} \textbf{ifp}\ Z : X^+ &\leftarrow X \wedge \Box \mathit{everywhere}(X \rightarrow X^+) \\ Y^+ &\leftarrow Y \wedge \Box \mathit{everywhere}(Y \rightarrow Y^+) \\ Z &\leftarrow \mathit{everywhere}(X \rightarrow X^+) \oplus \mathit{everywhere}(Y \rightarrow Y^+) \end{aligned}$$

where $\oplus$ means exclusive or. Each iteration of this system includes into X^+ the rightmost node of X that is not yet included in X^+. Hence after $|X|$ iterations X^+ will contain all of X. Analogously for Y. Hence the root will be included into Z if and only if $|X| \neq |Y|$. For $|X| < |Y|$ replace the last rule by $Z \leftarrow \mathit{everywhere}(X \rightarrow X^+) \wedge \mathit{somewhere}(Y \wedge \neg Y^+)$.

For expressing $|X| + |Y| = |Z|$ a similar technique is used. We count $|X| + |Y|$ by the number of stages of an inflationary induction. We hence have to make sure that nodes contained in X or Y, but not in $X \cap Y$ trigger one iteration, and nodes in $X \cap Y$ trigger two iterations. To ensure this we use two variables U and V where V will always be a subset of U and there will always be at most one node in $U - V$. At each iteration we put into U the rightmost node in $X \cup Y$ that is not yet in V and we put into V the rightmost node in $X \cup Y$ not yet in V provided it is already in U, or it is not contained in $X \cap Y$. Hence, as desired, nodes in $X \cap Y$ take two iterations to be included in V whereas nodes in $(X \cup Y) - (X \cap Y)$ only take one iteration. Otherwise the construction is the same as above. Thus, $|X| + |Y| = |Z|$ is expressed by the formula

$$\begin{aligned} \neg\textbf{ifp}\ W : U &\leftarrow (X \vee Y) \wedge \Box \mathit{everywhere}((X \vee Y) \rightarrow V) \\ V &\leftarrow (X \vee Y) \wedge \Box \mathit{everywhere}((X \vee Y) \rightarrow V)) \wedge (U \vee \neg X \vee \neg Y) \\ Z^+ &\leftarrow Z \wedge \Box \mathit{everywhere}(Z \rightarrow Z^+) \\ W &\leftarrow \mathit{everywhere}((X \vee Y) \rightarrow V) \oplus \mathit{everywhere}(Z \rightarrow Z^+). \end{aligned}$$

□

THEOREM 5.3. *Every language $L \in \text{DTIME}(O(n))$ is* MIC*-definable.*

Proof. Let M be a deterministic Turing machine, that works on strings $w = w_0 \dots w_{n-1} \in A^*$ and decides after at most kn moves, whether $w \in L$ (for some fixed constant k).

To show that L is MIC-definable we view input words w as (Kripke) structures with universe $\{n-1, \dots, 0\}$, the predecessor relation $E = \{(i+1, i) : i < n-1\}$, and atomic propositions (unary predicates) P_a, for each symbol $a \in A$ encoding the positions of the letters. Hence $w, i+1 \models \Diamond\varphi$ and $w, i+1 \models \Box\varphi$ both mean that $w, i \models \varphi$. (The use of $\Box$ or $\Diamond$ makes a difference only at point 0, where $\Box\varphi$ is true and $\Diamond\varphi$ is false (no matter what φ is).

For simplicity of notation we will first treat the case where M has only one tape and where $k = 1$, i.e. the machine makes at most n steps on words of length n. This case is not very interesting as the machine can either not move backwards or not read the entire input, but as we will see, the proof easily extends to the general case.

Encoding the computation. We describe a system Γ of rules of the form $X_i \leftarrow \varphi(\overline{X})$ (with MIC-formulae $\varphi(\overline{X})$) such that, for each $m \leq n$, the stage $\overline{X}^m$ of the inflationary induction of Γ on w describes the computation of M on w up to time m. The description of the computation is slightly non-standard (compared to the usual proofs of this kind) in the sense that only the current state, the moves and the symbols written at each time are explicitly recorded by the second-order variables, but not the storage content and the current position of the head on the tape. Rather, the currently read symbol is at every iteration reconstructed by a MIC-formula, based on the history of moves and written symbols up to that time.

The variables of Γ are T (for time), S_q (for every state q of M), L and R (for left and right moves on the tape), and W_b (for every symbol used by M). The intended interpretation, after m iterations, are:

—$T^m = \{i : i < m\}$.
—$S_q^m = \{i < m : M \text{ is in state } q \text{ at time } i+1\}$.
—$L^m = \{i < m : M \text{ makes a move to the left on the work tape at time } i\}$.
Analogously for R^m.
—$W_b^m = \{i < m : M \text{ wrote symbol } b \text{ on the work tape at time } i\}$.

Let us assume for the moment, that we have defined, for each symbol c, a MIC-formula CURRENT$[c]$, which on any word w expanded by the relations $T^m, S_q^m, L^m, R^m, W_b^m$ is true at position m if, and only if, the currently read symbol on the tape is c. Under this assumption (which will be proved below) the construction of Γ is a relatively straightforward adaption of the usual techniques in descriptive complexity theory to our situation.

Construction of Γ. Let $\delta : Q \times \Sigma \to Q \times \Sigma \times \{L, R\}$ be the transition function of M, and let q_0 be the initial state. For every element $s \in Q \cup \Sigma \cup \{L, R\}$ we denote by $\delta^{-1}(s)$ the set of all pairs $(q, a) \in Q \times \Sigma$ such that s is one of the components of $\delta(q, a)$. The system Γ consists of the rule $T \leftarrow \Box T$ that updates the time variables and rules for each S_q, W_b, L and R. We describe the rules for a state variable S_q. At each iteration step, we want to include the current time m (satisfying $\neg T \wedge \Box T$) into S_q if M goes into state q at time m. This happens if one of the following two conditions hold.

(1) We are at time 0, the input starts with some a such that $(q_0, a) \in \delta^{-1}(q)$.
(2) There exist $(q', a) \in \delta^{-1}(q)$ such that M went into state q' in the last step and currently reads the symbol a.

Hence the rule for S_q has the form

$$S_q \leftarrow (\Box \mathit{false} \wedge \bigvee_{(q_0,a)\in\delta^{-1}(q)} P_a) \vee$$
$$(\neg T \wedge \Diamond T \wedge \bigvee_{(q',a)\in\delta^{-1}(q)} (\Diamond S_{q'} \wedge \text{CURRENT}[a])).$$

The rules for L, R and W_b are analogous. Further, we can assume that the machine has a unique accepting state *acc*. Then M accepts a word w if, and only if, $w \models \mathbf{ifp}\ S_{acc} : \Gamma$.

Construction of CURRENT$[c]$. It remains to show that the formula CURRENT$[c]$ can indeed by constructed. What the formula needs to assert is that one of the following two cases holds:

(1) At the last time where M was at the current position the symbol c was written. In other words, there is some node $i \in W_c$ (i.e. c has been written at time i) such that between time i and now (i.e. time m), there have been equally many left and right moves, and that i was the last time with this property.

(2) M has never before been at the current position, and the input symbol at this position is c.

Here is where we use cardinality comparisons.

To express condition (1) we go systematically to all previous times (by iterating a rule $T' \leftarrow \Box T'$), and put the current node m (which is the unique node satisfying $\neg T \wedge \Box T$) into Z if after the time i described by the current value of T' there have been the same number of left and right moves (i.e. $|L \wedge \neg T'| = |R \wedge \neg T'|$), if at the time described by T', symbol c was written, and if since then this position has not been visited again. All this is expressed by

$$\begin{aligned}&\textbf{ifp } Z : T' \leftarrow \Box T'\\ &\qquad Z \leftarrow \neg T \wedge \Box T \wedge |L \wedge \neg T'| = |R \wedge \neg T'| \wedge\\ &\qquad\quad \textit{somewhere}(\neg T' \wedge \Box T' \wedge W_c) \wedge\\ &\qquad\quad \textit{nowhere}(T \wedge \neg \Box T' \wedge |L \wedge \neg T'| = |R \wedge \neg T'|),\end{aligned}$$

where $\textit{nowhere}(\varphi) := \neg \textit{somewhere}(\varphi)$. Note that the formula in the last line is false if, and only if, there exists a node $j \in T - \Box T'$ (i.e. a time j with $i < j < m$) such that between i and j there have been the same number of left and right moves. Hence the last line indeed says that the position reached at times i and m has not been visited in between.

To express condition (2) note that the current position at time m is $|L| - |R|$. We go systematically through all positions (by iterating a rule $P \leftarrow \Box P$), and put the current node m into Z if P describes the current position (i.e. if $|P| + |L| = |R|$) and if at the node defined by P (the unique node satisfying $\neg P \wedge \Box P$) the proposition P_c is true. Putting this into a formula, we obtain

$$\begin{aligned}&\textbf{ifp } Z : P \leftarrow \Box P\\ &\qquad Z \leftarrow \neg T \wedge \Box T \wedge (|P| + |L| = |R|) \wedge \textit{somewhere}(\neg P \wedge \Box P \wedge P_c)\end{aligned}$$

which is true at the current node m described by T if, and only if, the input symbol at the current position is c. We have to take the conjunction of this with a formula expressing that the current position has not been visited before. This is done by the formula

$$\begin{aligned}&\neg\textbf{ifp } Z : T' \leftarrow \Box T'\\ &\qquad Z \leftarrow (|L \wedge \neg T'| = |R \wedge \neg T'|).\end{aligned}$$

Z becomes true at the current time if there is some earlier time since which we have seen the same number of left and right moves. Hence Z does not become true if, and only if, we have never visited before the current cell.

Arbitrary Values of k and Multi-tape Machines. The construction easily extends to linear time computations of length kn for arbitrary values of k. We replace each of the variables T, S_q, W_b, L and R by a k-tuple of variables. The iteration of the time parameter is described by rules

$$\begin{aligned} T_1 &\leftarrow \Box T_k \\ T_2 &\leftarrow T_1 \\ &\vdots \\ T_k &\leftarrow T_{k-1} \end{aligned}$$

The iteration of these rules on a word of length n has kn stages. At iteration $ik + j + 1$ (where $i < n, j < k$) the node i is put into T_j . Also the rules for S_q are replaced by k-rules for variables $S_{q,1}, \dots, S_{q_k}$ constructed in a similar way. That is the rule for $S_{q,i+1}$ refers to variables $S_{q',i}$ for the possible predecessor states and the rule for $S_{q,0}$ refers to $\Diamond S_{q',k}$. Similarly for the other variables. Also the formula CURRENT[c] has to be reformulated, but this poses no essential difficulties.

Finally the extension to multi-tape Turing machines is completely straightforward by taking additional variables for the symbols written on each tape and the moves of each head. □

Note that we cannot expect to extend the result for linear time to quadratic time or higher. This is because, as we have seen, every language definable in MIC is decidable in linear space, and it is not expected that quadratic time is included in linear space.

6. BISIMULATION-INVARIANT PTIME AND AUTOMATICITY

In [Otto 1999], Otto introduced a higher-dimensional μ-calculus, denoted L^ω_μ, which extends ML with an operator for forming least fixed points of arbitrary arity, rather than just sets. He showed that L^ω_μ can express every bisimulation-invariant, polynomial-time decidable property of finite structures. He further showed that every formula of LFP that expresses a bisimulation-invariant property is equivalent to a formula of L^ω_μ. Combining this with the result of Kreutzer [Kreutzer 2002] that every formula of IFP is equivalent to one of LFP and the fact that MIC is a fragment of IFP that only expresses bisimulation-invariant properties immediately yields the following proposition.

PROPOSITION 6.1. *Every formula of* MIC *is equivalent to a formula of* L^ω_μ.

We now show that the converse of Proposition 6.1 fails. In particular, there are properties of finite trees that are bisimulation-invariant and polynomial time decidable that cannot be expressed in MIC.

THEOREM 6.2. *There is a collection $\mathcal{F}$ of finite trees in* PTIME, *closed under bisimulation, which is not expressible in* MIC.

Proof. We take $\mathcal{F}$ to be the collection of all trees $\mathcal{T}$ of finite height such that if t_1 and t_2 are any two children of the root of $\mathcal{T}$, then $\mathcal{T}, t_1 \sim \mathcal{T}, t_2$. Since it is known that bisimulation equivalence is polynomial time decidable, it follows that $\mathcal{F}$ is in PTIME. It is also obvious that $\mathcal{F}$ is closed under bisimulations.

To show that $\mathcal{F}$ is not definable in MIC, assume, towards a contradiction, that there is a formula φ defining it. Assume, without loss of generality, that all bound variables in φ are distinct. Let these variables be $X_1, \ldots, X_k$. We note that if $\mathcal{T}$ is a tree of height n, then for any formula **ifp** $X : S$, the closure ordinal is at most $l(n+1)$, where l is the number of variables involved in the system S. For any k-tuple $(i_1, \ldots, i_k)$, where each i_j is in the range $[0, k(n+1)]$, and any subformula ψ of φ, we write $\mathcal{T} \models \psi[i_1, \ldots, i_k]$ if ψ holds at the root of $\mathcal{T}$, when any free variable X_j in ψ is interpreted by $X_j^{i_j}$.

If Φ denotes the collection of subformulae of φ, we define the φ-type of $\mathcal{T}$ to be the function $\varphi^{\mathcal{T}} : (k(n+1)+1)^k \to 2^{\Phi}$ such that $\psi \in \varphi^{\mathcal{T}}(\bar{\imath})$ if, and only if, $\mathcal{T} \models \psi[\bar{\imath}]$. Now, if we denote by $[\mathcal{T}_1, \mathcal{T}_2]$ the tree whose root has two children, one being the root of a subtree $\mathcal{T}_1$, and the other being the root of a subtree $\mathcal{T}_2$, we can establish the following facts:

LEMMA 6.3. *If* $\varphi^{\mathcal{T}_1} = \varphi^{\mathcal{T}_2}$, *then* $\varphi^{[\mathcal{T}_1, \mathcal{T}_2]} = \varphi^{[\mathcal{T}_1, \mathcal{T}_1]}$.

Proof of Lemma 6.3. We assume, without loss of generality, that the variables $X_1, \ldots, X_k$ are ordered such that, if X_j is bound inside the scope of the **ifp** operator that binds $X_{j'}$, then $j' < j$.

Let $\mathcal{P}$ denote the collection of pairs $(\psi, \bar{\imath})$, where ψ is a subformula of φ, and $\bar{\imath}$ is a tuple in $(k(n+1)+1)^k$ such that, if X_j is bound in ψ, then $i_j = k(n+1)$.

We prove Lemma 6.3 by showing, by an induction on the lexicographical ordering of k-tuples in $(k(n+1)+1)^k$ that, for every pair in $(\psi, \bar{\imath}) \in \mathcal{P}$,

$$[\mathcal{T}_1, \mathcal{T}_1] \models \psi[\bar{\imath}] \quad \text{if, and only if,} \quad [\mathcal{T}_1, \mathcal{T}_2] \models \psi[\bar{\imath}]. \tag{1}$$

Since, if X_j is not free in ψ, the value of i_j is not relevant to whether or not $\mathcal{T} \models \psi[\bar{\imath}]$, this suffices to establish the lemma.

Basis: ($\bar{\imath} = \overline{0}$). We proceed by induction on the structure of ψ. If $\psi \equiv X_j$, since $i_j = 0$, ψ is false in both $[\mathcal{T}_1, \mathcal{T}_1]$ and $[\mathcal{T}_1, \mathcal{T}_2]$. The cases of the Boolean connectives are trivial. If ψ is $\Diamond\vartheta$, we know by hypothesis that $\mathcal{T}_1 \models \vartheta[\bar{\imath}]$ if, and only if, $\mathcal{T}_2 \models \vartheta[\bar{\imath}]$, and therefore, $[\mathcal{T}_1, \mathcal{T}_1] \models \Diamond\vartheta[\bar{\imath}]$ if, and only if, $[\mathcal{T}_1, \mathcal{T}_2] \models \Diamond\vartheta[\bar{\imath}]$. Finally, ψ cannot be of the form **ifp** $X_j : S$, as X_j would be a bound variable in ψ, and therefore the pair $(\psi, \overline{0})$ is not in $\mathcal{P}$.

Induction Step: Assume that for all $\bar{\imath}'$ lexicographically below $\bar{\imath}$, and all pairs $(\psi, \bar{\imath}') \in \mathcal{P}$, (1) holds. We prove it for all pairs $(\psi, \bar{\imath}) \in \mathcal{P}$, again by induction on the structure of ψ.

Suppose that ψ is atomic, i.e. of the form X_j, and let **ifp** $X_j : [\cdots, X_j \leftarrow \vartheta, \cdots]$ be the minimal subformula of φ in which X_j is bound. In other words, ϑ is the formula defining X_j. Then, if $[\mathcal{T}_1, \mathcal{T}_1] \models X_j[\bar{\imath}]$, by the inflationary semantics, there must be some $\bar{\imath}'$ lexicographically earlier than $\bar{\imath}$, such that $[\mathcal{T}_1, \mathcal{T}_1] \models \vartheta[\bar{\imath}']$. By the induction hypothesis, we then have $[\mathcal{T}_1, \mathcal{T}_2] \models \vartheta[\bar{\imath}']$, and therefore $[\mathcal{T}_1, \mathcal{T}_2] \models X_j[\bar{\imath}]$. A similar argument establishes $[\mathcal{T}_1, \mathcal{T}_1] \models X_j[\bar{\imath}]$, if we assume $[\mathcal{T}_1, \mathcal{T}_2] \models X_j[\bar{\imath}]$.

The case of the Boolean connectives is trivial, and if ψ is of the form $\Diamond\vartheta$, we again rely on the hypothesis that the φ-types of $\mathcal{T}_1$ and $\mathcal{T}_2$ are the same.

Finally, if ψ is of the form **ifp** $X_j : S$, where S is $[\cdots, X_j \leftarrow \vartheta, \cdots]$, then $i_j = k(n+1)$, by the definition of $\mathcal{P}$. Therefore, if $[\mathcal{T}_1, \mathcal{T}_1] \models \psi[\bar{\imath}]$, there must be an $i' < i_j$ such that $[\mathcal{T}_1, \mathcal{T}_1] \models \vartheta[\bar{\imath}']$, where $\bar{\imath}'$ is the tuple obtained from $\bar{\imath}$ by replacing i_l by i', where l ranges over the indices of all variables bound by the system S, including j. The tuple $\bar{\imath}'$ is lexicographically prior to $\bar{\imath}$, and therefore, by induction hypothesis, $[\mathcal{T}_1, \mathcal{T}_2] \models \vartheta[\bar{\imath}']$, and then by the inflationary semantics, we have $[\mathcal{T}_1, \mathcal{T}_2] \models \psi[\bar{\imath}]$.

This completes the proof of Lemma 6.3. □

We next observe that there is a constant c (depending only on φ) and a polynomial p such that the number of distinct φ-types of trees of height n is at most $c^{p(n)}$. Indeed, we can take c to be the cardinality of 2^Φ, i.e. the number of sets of subformulae of φ, and we can take p to be $(k(n+1)+1)^k$.

However, the number of distinct $\sim$-equivalence classes of trees of height n grows as a non-elementary function of n. This follows from the fact that, if there are e distinct equivalence classes of trees of height n, there are $2^e - 1$ distinct classes of trees of height $n+1$. For, if $\mathcal{T}_1, \ldots, \mathcal{T}_e$ is an enumeration of representatives of the equivalence classes of trees of height n, we can form, for any non-empty set $S \subseteq \{1, \ldots, e\}$, a tree of height $n+1$ whose root has a child which is the root of a subtree $\mathcal{T}_j$ if, and only if, $j \in S$. It is easily seen that for distinct sets S, the resulting trees are not bisimilar.

It now follows that we can find two trees $\mathcal{T}_1$ and $\mathcal{T}_2$ which have the same φ-type, but $\mathcal{T}_1 \not\sim \mathcal{T}_2$. Thus, $[\mathcal{T}_1, \mathcal{T}_2] \notin \mathcal{F}$, while $[\mathcal{T}_1, \mathcal{T}_1] \in \mathcal{F}$, however, $[\mathcal{T}_1, \mathcal{T}_1] \models \varphi$ if, and only if, $[\mathcal{T}_1, \mathcal{T}_2] \models \varphi$, contradicting the assumption that φ defines the class $\mathcal{F}$. □

Automaticity. Another way of understanding the proof above is in terms of *automaticity.* This is a measure of the complexity of a formal language, considered, for instance, in [Shallit and Breitbart 1996].

The automaticity of a language $L \subseteq \Sigma^*$ is the function A_L which gives for each n the number of states in the smallest deterministic automaton which accepts a language that agrees with L on all strings of length at most n. Clearly, a language L is regular if, and only if, A_L is bounded by a constant. Moreover, for any language L whatsoever, including undecidable ones, A_L is, at worst, exponential.

Here, we observe that the language $L = \{ww \mid w \in \{a,b\}^*\}$ used in the proof of Theorem 5.1 has exponential automaticity, which is worst possible. To be precise, we can show that $A_L(n) \geq 2^{n/2}$. For, suppose to the contrary that, for some n, $A_L(n) < 2^{n/2}$ and let A be a minimal deterministic finite automaton accepting $L \cap \{a,b\}^n$. Then, there must be two distinct strings w_1 and w_2 of length $n/2$ such that A, starting in its initial state ends up in the same state on inputs w_1 and w_2. It then follows that A accepts the string $w_1 w_1$ if, and only if, it accepts $w_2 w_1$. However, since both of these strings are of length n, and the former is in L but the latter is not, we have a contradiction.

The observation that MIC can express languages of exponential automaticity, which is worst possible, would suggest that this is not a good measure of complexity for studying the expressive power of MIC. However, automaticity may be useful if we generalise it from strings to finite trees as MIC can still express only tree languages of at most exponential automaticity but this is no longer maximal.

Automata on trees have been extensively studied (see, for instance, [Gécseg and Steinby 1997]). For finite trees, we define a finite automaton as a tuple $\mathcal{A} =$

(Q, A, δ, F, S), where Q is a set of states, A is the alphabet, $S \subseteq Q$ is a collection of start states, $F \subseteq Q$ is the set of final states, and $\delta : Q \times A \to 2^Q$ is the transition function. If $\mathcal{T}$ is a finite tree, where edges are labelled by elements of A, we say that $\mathcal{A}$ *accepts* $\mathcal{T}$, if there is a labelling $l : \mathcal{T} \to Q$ of the nodes of $\mathcal{T}$ such that

—the root of $\mathcal{T}$ is labelled $q_0 \in S$;

—if $l(v) = q$, then $\{l(w) : v \xrightarrow{a} w\} = \delta(q, a)$, for each $a \in A$; and

—for every leaf v, $l(v) \in F$.

Definition 6.4. For a class $\mathcal{C}$ of finite trees, closed under bisimulation, the *automaticity* $A_{\mathcal{C}}$ of $\mathcal{C}$ is the function mapping n to the number of states in the smallest automaton that accepts a tree of *height* at most n if, and only if, it is in $\mathcal{C}$.

With this definition, we can extract the observations contained in the following two theorems from the proof of Theorem 6.2.

Theorem 6.5. *There is a polynomial time decidable, bisimulation-invariant, collection of finite trees with non-elementary automaticity.*

Proof. Let $\mathcal{F}$ be the class of trees used in the proof of Theorem 6.2, i.e. it is the collection of all trees $\mathcal{T}$ of finite height such that if t_1 and t_2 are any two children of the root, then $\mathcal{T}, t_1 \sim \mathcal{T}, t_2$. Suppose, towards a contradiction, that there is a k such that the automaticity $\mathcal{A}_{\mathcal{F}}$ of this class is bounded by a tower of 2s of height k. We can then choose $n > k$ so that the number of bisimulation equivalence classes of trees of height n is greater than $\mathcal{A}_{\mathcal{F}}(n+1)$, and let A_{n+1} be the minimal automaton accepting trees of height $n+1$ if, and only if, they are in $\mathcal{F}$. Now, there must be two trees $\mathcal{T}_1$ and $\mathcal{T}_2$ of height n such that the same set of states occurs at children of the root in an accepting run of A_{n+1} on $[\mathcal{T}_1, \mathcal{T}_1]$ as it does on $[\mathcal{T}_2, \mathcal{T}_2]$. It then follows that these two runs can be combined into an accepting run of A_{n+1} on $[\mathcal{T}_1, \mathcal{T}_2]$, which is a contradiction, as this tree is not in $\mathcal{F}$. □

Theorem 6.6. *Every class of finite trees that is definable in* MIC *has at most exponential automaticity.*

Proof. Let φ be a formula of MIC. For each n, we can construct an automaton $\mathcal{A} = (Q, A, \delta, S, F)$ that accepts the trees of height n satisfying φ. Here:

—Q is the collection of all φ-types $\varphi^{\mathcal{T}}$ of trees whose height is at most n;

—A is the set of actions occurring in the formula φ;

—$q' \in \delta(q, a)$ if, and only if, q' is the φ-type of an a-child of the root of a tree whose φ-type is q;

—S is the set of all φ-types q such that $\varphi \in q(\overline{k(n+1)})$; and

—F is the set of all φ-types of trees with only one node.

It is easily checked that a tree $\mathcal{T}$ is accepted by this automaton if, and only if, $\mathcal{T} \models \varphi$. Moreover, the number of states is exponential in n, by the argument given in the proof of Theorem 6.2. □

7. SIMPLE INDUCTIONS

In the previous sections we have seen a couple of expressibility results for MIC. Many of these made use of simultaneous inductions and we remarked at several places that the formulae used there could not be defined in 1MIC. In this section we investigate the relationship between simple and simultaneous inductions in the modal iteration calculus.

It is easy to see that the equivalence $\mu XY.(\psi, \varphi) \equiv \mu X.\psi(X, \mu Y.\varphi(X, Y))$ (sometimes called the Bekic-principle [Arnold and Niwiński 2001]) fails in both directions when we take inflationary instead of least fixed points. However, it still is conceivable that simultaneous inductions could be eliminated by more complicated techniques. We show here that this is not the case, i.e. simultaneous inflationary inductions provide more expressive power than simple ones. However, we have seen in Section 4 that the model checking problem was PSPACE-complete even for formulae without simultaneous inductions. In the second part of this section we show that, similar to MIC, the satisfiability problem for formulae without simultaneous inductions is still undecidable and not in the arithmetical hierarchy.

Let 1MIC denote the fragment of MIC that does not involve simultaneous inductions.

7.1 Simple vs. Simultaneous Inductions

For any ordinal α, let $\mathcal{O}_\alpha$ denote the structure $(\{\beta : \beta \leq \alpha\}, >)$. That is, the elements of the structure are all ordinals less than or equal to α, and γ is accessible from β if $\beta > \gamma$. Note that α is the maximal ordinal in the set, and we refer to it as the *root* of the structure $\mathcal{O}_\alpha$. For a formula φ, we write $\mathcal{O}_\alpha \models \varphi$ as shorthand for $\mathcal{O}_\alpha, \alpha \models \varphi$. It is easily seen that for any ordinal $\beta \leq \alpha$, $\mathcal{O}_\alpha, \beta \models \varphi$ if, and only if, $\mathcal{O}_\beta \models \varphi$. This is because the elements reachable by $>$-paths from β are exactly the ordinals below β, since $>$ is transitive. As a final bit of notation, if X is any atomic proposition on $\mathcal{O}_\alpha$, we write $\mathcal{O}_\alpha, \beta \models \varphi(X^-)$ to denote that $\mathcal{O}_\alpha, \beta \models \varphi(X - \{\beta\})$.

We begin with a few observations about the evaluation of inductive formulae on the structures $\mathcal{O}_\alpha$, which will be useful in the proof of the following lemma. First, we note that in $\mathcal{O}_\beta$, the maximum closure ordinal of any induction is $\beta + 1$. This can be proved by a straightforward induction on the ordinals. One consequence is that if $\mathcal{O}_\alpha, \beta \models (\textbf{ifp}\ X \leftarrow \psi)$, then $\beta \in X^{\beta+1}$. Furthermore, since the truth of a formula φ at β can only depend on elements $\gamma \leq \beta$, we have that for any sets X and Y, if $X \cap (\beta + 1) = Y \cap (\beta + 1)$, then $\mathcal{O}_\alpha, \beta \models \varphi(X)$ if, and only if, $\mathcal{O}_\alpha, \beta \models \varphi(Y)$.

LEMMA 7.1. *Let φ be a formula of* 1MIC. *If $X_1, \ldots, X_k \subseteq \omega$ are atomic propositions such that $\mathcal{O}_\omega \models \varphi(X_1, \ldots, X_k)$, then there is a finite N such that for all $n > N$, $\mathcal{O}_\omega, n \models \varphi(X_1^-, \ldots, X_k^-)$.*

Proof. Note that, by the hypothesis of the lemma $\omega \notin X_i$ for any i. As we are now working in a fixed structure $\mathcal{O}_\omega$, we will say a formula φ holds (or is true) at α, to mean that $\mathcal{O}_\omega, \alpha \models \varphi$, which is the case if, and only if, $\mathcal{O}_\alpha \models \varphi$.

The lemma is proved by induction on the depth d of nesting of **ifp**-operators.
Basis: If $d = 0$, the formula contains no occurrences of the **ifp** operator, and is therefore equivalent to a formula of ML, where all negations are at the atoms. A simple induction on the structure of the formula then establishes the result. Any atomic formula X_i is false at ω by hypothesis, and, by definition, for all n,

$\mathcal{O}_\omega, n \not\models X_i^-$, and therefore the claim holds for all atoms. For negated atoms the dual argument holds, i.e. any formula $\neg X_i$ is true at ω and for all n, $\mathcal{O}_\omega, n \models \neg X_i^-$. The case of the Boolean connectives $\wedge$ and $\vee$ is trivial. If φ is $\Diamond\psi$, then it is clear that $\mathcal{O}_\omega \models \varphi$ if, and only if, there is an $N < \omega$ such that $\mathcal{O}_N \models \psi$ if, and only if, for all $n > N$, $\mathcal{O}_\omega, n \models \Diamond\psi$ and therefore $\mathcal{O}_\omega, n \models \varphi(\overline{X}^-)$. Similarly, if φ is $\Box\psi$, $\mathcal{O}_\omega \models \varphi$ if, and only if, for all n, $\mathcal{O}_\omega, n \models \psi$ if, and only if, for all $n < \omega$, $\mathcal{O}_\omega, n \models \Box\psi(\overline{X}^-)$.
Induction step: If φ is a formula with depth $d+1$ of nesting of **ifp** operators, then it is a Boolean combination of formulae ϑ_i, each of which either has depth at most d, is of the form **ifp** $X \leftarrow \psi$, where ψ is a formula of depth at most d, or is of the form $\Diamond\psi$ or $\Box\psi$, where ψ has depth at most $d+1$. We assume that negations are always pushed inside modalities.

Clearly, if the claim holds for formulae ϑ_i and ϑ_j, then it also holds for $\vartheta_i \wedge \vartheta_j$ and $\vartheta_i \vee \vartheta_j$, just by taking N to be the maximum of N_i and N_j, which witness the claim for the two formulae. Thus, it suffices to prove the claim for the following four cases.

(1) $\vartheta \equiv \Diamond\psi$. If $\mathcal{O}_\omega \models \vartheta$, then there is an $N < \omega$ such that $\mathcal{O}_\omega, N \models \psi$, and therefore, for all $n > N$, $\mathcal{O}_\omega, n \models \vartheta(\overline{X}^-)$.

(2) $\vartheta \equiv \Box\psi$. If $\mathcal{O}_\omega \models \Box\psi$, then for all n, $\mathcal{O}_\omega, n \models \psi$ and therefore $\mathcal{O}_\omega, n \models \vartheta(\overline{X}^-)$.

(3) $\vartheta \equiv$ (**ifp** $X \leftarrow \psi$). Suppose $\mathcal{O}_\omega \models \vartheta$. Then, there is a stage α such that $\omega \notin X^\alpha$, and $\mathcal{O}_\omega \models \psi(X_1, \ldots, X_k, X^\alpha)$. By induction hypothesis, there is an N such that for all $n > N$, $\mathcal{O}_\omega, n \models \psi(X_1^-, \ldots, X_k^-, X^{\alpha-})$. Now, for each such n, if n is in stage α of the induction of $\psi(X_1^-, \ldots, X_k^-)$, then, by the inflationary semantics $\mathcal{O}_\omega, n \models$ (**ifp** $X \leftarrow \psi)(X_1^-, \ldots, X_k^-)$, and we are done. So, suppose n is not in stage α of the induction of $\psi(X_1^-, \ldots, X_k^-)$. However, since $X_i^- \cap n = X_i \cap n$, it follows that stage α of the induction of $\psi(X_1^-, \ldots, X_k^-)$ on $\mathcal{O}_n$ is exactly $X^{\alpha-} \cap (n+1)$. Thus, $\mathcal{O}_\omega, n \models \psi(X_1^-, \ldots, X_k^-, X^{\alpha-})$ implies $\mathcal{O}_\omega, n \models$ (**ifp** $X \leftarrow \psi)(X_1^-, \ldots, X_k^-)$.

(4) $\vartheta := \neg$(**ifp** $X \leftarrow \psi$). Suppose $\mathcal{O}_\omega \models \vartheta$. We have to show that there is some $N < \omega$ such that for all $n > N$, $\mathcal{O}_\omega, n \models \vartheta$. Towards a contradiction, assume that there are infinitely many n such that $\mathcal{O}_\omega, n \models$ (**ifp** $X \leftarrow \psi)(X_1^-, \ldots, X_k^-)$. Since $\mathcal{O}_\omega \models \neg$(**ifp** $X \leftarrow \psi$), it is the case that $\mathcal{O}_\omega \models \neg\psi(X_1, \ldots, X_k, X^\omega)$, and therefore, by induction hypothesis, there is an $N < \omega$ such that

$$\text{for all } n > N, \mathcal{O}_\omega, n \models \neg\psi(X_1^-, \ldots, X_k^-, X^{\omega-}). \qquad (*)$$

As we have noted, in any $\mathcal{O}_\beta$, the closure ordinal of any induction is at most $\beta + 1$. Hence, for any $\beta < N$, $\mathcal{O}_\omega, \beta \models \psi(X_1, \ldots, X_k, X^\omega)$ if, and only if, $\mathcal{O}_\omega, \beta \models \psi(X_1, \ldots, X_k, X^N)$. As, by assumption, there are infinitely many n such that $\mathcal{O}_\omega, n \models$ (**ifp** $X \leftarrow \psi)(X_1^-, \ldots, X_k^-)$, there are infinitely many such greater than N. For each of these, there must be a least finite ordinal a_n such that $\mathcal{O}_\omega, n \models \psi(X_1^-, \ldots, X_k^-, X^{a_n})$. We distinguish two cases:

  (a) There is a finite ordinal α that is the upper bound of all such a_n. In this case, we can show that there is a finite bound on the elements satisfying (**ifp** $X \leftarrow \psi)(X_1^-, \ldots, X_k^-)$, from which the claim follows. To establish the finite bound, we show by induction on the stages, that each stage X^β contains only nodes up to some finite height. Clearly, this is the case for $\beta = 0$, as X^β is empty. Inductively, if X^β is bounded in height, then there is a formula of ML defining the set. Substituting this formula for X in ψ, we obtain a formula ψ' equivalent to $\psi(X^\beta)$, with **ifp** nesting depth d. Therefore, by the main induction hypothesis, since $\mathcal{O}_\omega \models \neg\psi'(X_1, \dots, X_k)$, there is an M such that $\mathcal{O}_\omega, n \models \neg\psi'(X_1^-, \dots, X_k^-)$ for all $n > M$. It follows that for all $n > M$, $n \notin X^{\beta+1}$. This implies that there is a finite bound on the height of the elements in X^α, contradicting the assumption that there are infinitely many n such that $\mathcal{O}_\omega, n \models (\textbf{ifp}\ \ X \leftarrow \psi)(X_1^-, \dots, X_k^-)$.

(b) There is no finite bound on the stages a_n. Thus, for any finite stage α there is some $a_n > \alpha$ with $n \notin X^{a_n}$ and $\mathcal{O}_\omega, n \models \psi(X_1^-, \dots, X_k^-, X^{a_n})$. Choose α and a_n such that $a_n > N$. Let c be the minimal node such that $c \notin X^{a_n}$ but $c \in X^\beta$ for some $\beta > a_n$. Clearly, for every node $m < a_n$, $m \in X^{a_n}$ if, and only if, $m \in X^\omega$. It follows that $c \geq a_n$ and thus $c > N$. Further, as c was chosen minimal,

$$\{m : m < c\} \cap X^{a_n} = \{m : m < c\} \cap X^\omega,$$

i.e. on the set of nodes reachable but different from c the fixed point of X is reached at stage a_n. Therefore,

$$\begin{aligned} &\mathcal{O}_\omega, c \models \psi(X_1^-, \dots, X_k^-, X^{a_n}) \quad \text{if, and only if,} \\ &\mathcal{O}_\omega, c \models \psi(X_1^-, \dots, X_k^-, X^{\omega-}). \end{aligned}$$

As $c \geq a_n > N$ and, by $(*)$, $\mathcal{O}_\omega, c \not\models \psi(X_1^-, \dots, X_k^-, X^{\omega-})$ we get that $c \notin X^\beta$ for some $\beta > a_n$, contradicting the assumption.

□

It is a straightforward consequence of this lemma that the formula *finite-height* defined in Sect. 3 is not equivalent to any formula of 1MIC. This is because $\mathcal{O}_\alpha$ is bisimulation equivalent to a well-founded tree of height α. Suppose there was a formula $\varphi \in$ 1MIC equivalent to *finite-height*. By definition, $\mathcal{O}_\omega \not\models \varphi$ and therefore there is some $N < \omega$ such that $\mathcal{O}_n \models \neg\varphi$ for all $n > N$. As these $\mathcal{O}_n$ are of finite height, we get a contradiction to φ being equivalent to *finite-height*. We hence have established the following separation result.

THEOREM 7.2. MIC *is strictly more powerful than* 1MIC.

While the separation given in Theorem 7.2 is proved on an infinite structure, the proof of Lemma 7.1 actually shows that the separation holds even when we restrict ourselves to finite structures. For, consider the collection of finite structures $\mathcal{O}_n$, for all finite ordinals n. The construction in Lemma 7.1 shows that for any formula φ of 1MIC, the set $\{n : \mathcal{O}_n \models \varphi\}$ is either finite or co-finite. Now, consider the formula η

$$\begin{aligned} \textbf{ifp}\ Y :\ & X \leftarrow \Box\Box X \\ & Y \leftarrow \neg X \wedge \Box X. \end{aligned}$$

It can be verified that $\mathcal{O}_n \models \eta$ if, and only if, n is even. Hence, η is not equivalent to any formula of 1MIC.

The notation 1MIC was chosen, naturally, to suggest that we can define, for any natural number k, the fragment kMIC consisting of those formulae of MIC in which no **ifp** operation is defined over a system with more than k simultaneous formulae. The natural question that arises is whether increasing k gives rise to a hierarchy of increasing expressive power. We do not, at the moment, know of a method to settle this question one way or the other.

7.2 Infinity Axioms and the Satisfiability Problem for 1MIC

We have seen above that simple inductions in MIC provide less expressive power than simultaneous inductions. We show now that even without simultaneous inductions the finite model property fails and the satisfiability problem remains undecidable, in fact not even arithmetical.

Essentially, we use the same method as in Section 3 to reduce the decision problem for the first-order theory of arithmetic to Sat(1MIC). The proofs there relied crucially on simultaneous inductions on two variables X and Y and we have seen above that the formulae used there are not equivalent to any formulae of 1MIC. Instead, we use the following trick to simulate the simultaneous induction. Let $\mathcal{T}$ be a tree of height ω as used in Section 3. To simulate an induction on two variables we recursively make two copies of the successors of each node in $\mathcal{T}$ and label the root of one of the copies by the proposition a and the root of the other by b. Let $\mathcal{T}'$ be this new tree. Thus, every node in $\mathcal{T}'$ of finite height is labelled by a or b and the height of each of these nodes u equals the length of the longest path entirely labelled by a's from a successor of u to a leaf which, again, is the same as the length of the longest path entirely labelled by b's from a successor of u to a leaf. On this model, we can simulate the simultaneous induction on two variables X and Y by an induction on one variable Z by letting Z contain all copies of nodes contained in X which are labelled by a and all copies of nodes contained in Y labelled by b.

We now turn to the axiomatisation of this tree model. Besides the proposition symbols a and b already mentioned there are two other propositions, namely w, with which only the root is labelled, and s, labelling only the direct successors of the root. As in Section 3, we restrict attention to well-founded tree models, i.e. tree models of the formula $\mu X.\Box X$.

We first define a formula *label* ensuring that its models are labelled as described above. Let the formula *label* be defined as

$$\begin{aligned}\text{label} := \; & w \wedge \neg a \wedge \neg b \wedge \neg s \wedge \Box s \wedge \Box\Box \mathit{everywhere}(\neg s) \wedge \\ & \Box \mathit{everywhere}(\neg w \wedge (a \vee b) \wedge \neg(a \wedge b)),\end{aligned}$$

where *everywhere* is defined as in Section 3 as $\mathit{everywhere}(\varphi) := \textbf{ifp}\ \ X \leftarrow \varphi \wedge \Box X$.

Proposition 7.3. *In any model $\mathcal{T}, v$ of* label *the root and only the root v is labelled by w, all direct successors of v, and only those, are labelled by s, and all nodes reachable but different from v are either labelled by a or by b but not by both.*

The next step is to ensure that the models we are going to describe are of height ω. Towards this end, we first introduce some notation.

Definition 7.4. *A node labelled by a is called an* a-node. *An a-path between two nodes is a path between them consisting only of a-nodes. Finally, we inductively define the* a-height α *of an a-node u as follows.*

—*The a-height of leaves labelled by a is defined to be* 0.

—*For all other a-nodes, their a-height is defined as the least strict upper bound of all β such that β is the a-height of a successor of u labelled by a.*

The notion of b-nodes, b-paths, and the b-height of a node is defined analogously.

Let the formula ϑ be defined as

$$\vartheta := (b \wedge \Box(b \rightarrow X)) \vee (a \wedge \Box(a \rightarrow X)).$$

A simple induction on the stages establishes the following lemma.

LEMMA 7.5. *Let $\mathcal{T}$ be a tree. Then, for all α, the α-th stage X^α of the induction on ϑ contains exactly the nodes of a- or b-height less than α.* □

Consider the formula

$$\text{infinity} := \text{label} \wedge \textit{everywhere}(\Diamond a \leftrightarrow \Diamond b) \wedge \text{inf} \wedge \neg \textbf{ifp} \quad X \leftarrow \varphi,$$

where

$$\text{inf} := \neg \textbf{ifp} \quad X \leftarrow (b \wedge \Diamond b \wedge \Box(b \rightarrow \Box \textit{false})) \vee \varphi.$$

and

$$\begin{aligned}\varphi := \vartheta \vee (w \wedge \Box(b \rightarrow X) \wedge \Diamond(a \wedge \neg X)) \vee \\ (w \wedge \Box(a \rightarrow X) \wedge \Diamond(b \wedge \neg X)).\end{aligned}$$

LEMMA 7.6. *Let $\mathcal{T}, v \models \textit{label} \wedge \textit{everywhere}(\Diamond a \leftrightarrow \Diamond b)$ and let $\gamma_a := \sup\{\alpha : \alpha$ is the a-height of a direct successor of $v\}$ and $\gamma_b := \sup\{\alpha : \alpha$ is the b-height of direct successor of $v\}$.*

(i) $\mathcal{T}, v \models \neg \textbf{ifp} \quad X \leftarrow \varphi$ *if, and only if,* $\gamma_a = \gamma_b$.

(ii) $\mathcal{T}, v \models \text{inf}$ *if, and only if, γ_a and γ_b are finite and $\gamma_a + 1 = \gamma_b$ or both are infinite and $\gamma_a = \gamma_b$.*

Proof.

(i) For Part (i), we show that the root v satisfies $\textbf{ifp} \quad X \leftarrow \varphi$ if, and only if, there is a b-successor of v whose b-height is greater than the a-height of any a-successor of v or vice versa.
Towards the forth direction, suppose $\mathcal{T}, v \models \textbf{ifp} \quad X \leftarrow \varphi$. Thus, there is a stage α such that $(\mathcal{T}, X^\alpha), v \models \varphi$. As $\mathcal{T}$ is a model of the formula *label* above, this implies

$$\begin{aligned}(\mathcal{T}, X^\alpha), v \models \ & (w \wedge \Box(b \rightarrow X) \wedge \Diamond(a \wedge \neg X)) \vee \\ & (w \wedge \Box(a \rightarrow X) \wedge \Diamond(b \wedge \neg X)).\end{aligned}$$

Suppose $(\mathcal{T}, X^\alpha), v \models (w \wedge \Box(b \rightarrow X) \wedge \Diamond(a \wedge \neg X))$. Thus, using Lemma 7.5, all b-successors of v are of height less than α, whereas there is at least one a-successor whose height is greater than or equal to α. The case where $(\mathcal{T}, X^\alpha), v \models (w \wedge \Box(a \rightarrow X) \wedge \Diamond(b \wedge \neg X))$ is analogous.
For the converse, if v has a b-successor whose b-height α is greater than the a-height of any a-successor, then $X^{<\alpha}$ contains all a-successors but not all

b-successors of v. Thus, v would be in X^α. The case where there is an a-successor of v whose a-height is greater than the b-height of any b-successor of v is analogous.

Thus we have shown that the root v satisfies **ifp** $X \leftarrow \varphi$ if, and only if, there is a b-successor of v whose b-height is greater than the a-height of any a-successor of v or vice versa. This proves the first part of the lemma.

(ii) To establish Part (ii), we first prove by induction on the finite stages, that for all $0 < n < \omega$, X^n contains exactly the a-nodes of height less than n and the b-nodes of height less than or equal to n. For the case of nodes labelled by a this follows immediately from Lemma 7.5 as the additional disjuncts only affect b-nodes (and the root v.)

Let $n = 1$. Obviously, X^1 contains all b-nodes of height 0 and all b-nodes satisfying $\Diamond \mathit{true} \wedge \Box\Box \mathit{false}$. Clearly, such a node u must be of height 1, as it only has leaves as successors. As $\mathcal{T}$ is a model of the formula *label*, this means that u must have a b-successor and thus is of b-height 1.

For the induction step assume the claim has already been proved for all $n' < n$. An argument as in the proof of Part (i) above shows that X^n contains all b-nodes of b-height less than or equal to n. This finishes the induction.

Now consider the stage ω. We have seen that X^ω contains all nodes of finite a- or b-height and the same argument as in Lemma 7.5 shows that for all $\alpha \geq \omega$, X^α contains all nodes of a- or b-height less than α.

Now suppose that the root v of the tree occurs in X^∞. There must be a stage $\alpha + 1$ such that v occurs in $X^{\alpha+1}$ but not in X^α, i.e. $(\mathcal{T}, X^\alpha), v \models (w \wedge \Box(b \rightarrow X) \wedge \Diamond(a \wedge \neg X)) \vee (w \wedge \Box(a \rightarrow X) \wedge \Diamond(b \wedge \neg X))$. Suppose that $(\mathcal{T}, X^\alpha), v \models (w \wedge \Box(b \rightarrow X) \wedge \Diamond(a \wedge \neg X))$. Thus, there is an a-successor of v whose a-height is greater than the b-height of any b-successor. If α is infinite, this means that $\gamma_a > \gamma_b$. If α is finite, this means that all b-successors of v are of height less than or equal to α whereas there is an a-successor of v of height greater than or equal to α. This implies that $\gamma_a + 1 > \gamma_b$.

On the other hand, if $(\mathcal{T}, X^\alpha), v \models (w \wedge \Box(a \rightarrow X) \wedge \Diamond(b \wedge \neg X))$, this implies that $\gamma_b > \gamma_a$ if α is infinite and $\gamma_b > \gamma_a + 1$ otherwise.

This proves that if $\mathcal{T}, v \models \mathit{inf}$ then γ_a and γ_b are finite and $\gamma_a + 1 = \gamma_b$ or both are infinite and $\gamma_a = \gamma_b$. The converse direction follows immediately.

□

A simple consequence of the lemma is the following corollary.

Corollary 7.7. *For every tree $\mathcal{T}$ and node v, if $\mathcal{T}, v \models$ infinity then the height of v is infinite. Thus,* 1MIC *does not have the finite model property.*

To finish the axiomatisation of the intended model, we have to ensure that the models of our formula not only have infinite height but height exactly ω and further, that for all nodes of finite height, all of their successors are of the same height.

To formalise that the tree has height exactly ω we say that the root is of infinite height but none of its successors is. For this let $\mathrm{inf}_s := \mathrm{inf}[w/s, b/(b \wedge \neg s), a/(a \wedge \neg s)]$ be the result of replacing in the formula *inf* above each w by s, each b by $(b \wedge \neg s)$, and each a by $(a \wedge \neg s)$. Analogously, define $\varphi_s := \varphi[w/s, b/(b \wedge \neg s), a/(a \wedge \neg s)]$. Then a similar argument as in Lemma 7.6 shows the following.

LEMMA 7.8. *Let the formula* ω-height *be defined as*

$$\omega\text{-}height := infinity \wedge \Box\neg(inf_s \wedge \neg \mathbf{ifp} \quad X \leftarrow \varphi_s).$$

Then for any tree $\mathcal{T}$ with root v, if $\mathcal{T}, v \models \omega$-height then the height of v is infinite but the height of all successors of v is finite. Thus, the height of v is exactly ω.

Finally, we formalise that for all nodes of finite height all of their successors are of the same height.

LEMMA 7.9. *Consider the formula ψ defined as*

$$\psi := \neg \mathbf{ifp} \quad X \leftarrow (\neg w \wedge \Box X) \vee (w \wedge \mathbf{ifp} \quad Y \leftarrow \Diamond Y \vee (\neg w \wedge \Diamond X \wedge \Diamond \neg X)).$$

Let $\mathcal{T}, v \models \omega$-height $\wedge\, \psi$ be a tree. Then, for all nodes u of finite height all direct successors of u are of the same height.

Proof. As $\mathcal{T}, v \models \omega$-height, no node of finite height is labelled by w. A simple induction on the stages proves that at stage $n \in \omega$, X^n contains all nodes of height less than n. Now assume that at some stage n the root v is included into X, i.e. v satisfies $\mathbf{ifp} \quad Y \leftarrow \Diamond Y \vee (\neg w \wedge \Diamond X \wedge \Diamond \neg X)$. Thus there is a node u whose height is greater than n but has a successor in X^n, i.e. one whose height is less than n, and a successor not in X^n, i.e. one whose height is greater than or equal to n.

Conversely, if there is a node u of some finite height n which has a successor of height less than $n-1$, then $\Diamond X \wedge \Diamond \neg X$ becomes true at stage $n-1$ and the root v is included into X^n.

Together, we get that $\mathcal{T}, v \models \psi$ if, and only if, for all nodes u of finite height, all of its successors are of the same height. □

We now turn to the reduction of the decision problem for the first-order theory of arithmetic to the satisfiability problem for 1MIC. In the remainder of this section we only consider models of the formula ω-height $\wedge\, \psi$. As in Section 3, we code natural numbers by sets of nodes of certain height. The difference is, that here a number $n \in \omega$ is coded by the set of nodes u of height less than or *equal to* n. In particular, the number 0 is not coded by the empty set but by the set of leaves. To simplify the presentation, we do not reduce the first-order theory of arithmetic to Sat(1MIC) directly but first to the decision problem for the first-order theory of the following structure.

DEFINITION 7.10. *Let* $\mathfrak{N}' := (\mathbb{N}, +, \cdot, 0, 1)$ *be the structure over the universe* $\mathbb{N}$, *where the constants* $0^{\mathfrak{N}'}, 1^{\mathfrak{N}'}$ *and the function* $+^{\mathfrak{N}'}$ *are interpreted as in the standard model* $\mathfrak{N}$ *of arithmetic, and* $\cdot^{\mathfrak{N}'}$ *is interpreted as follows:* $s \cdot^{\mathfrak{N}'} t := 0$ *if* $s = 0$ *or* $t = 0$, *and otherwise* $s \cdot^{\mathfrak{N}'} t := s \cdot^{\mathfrak{N}} (t+1)$.

Clearly, the first order theories of $\mathfrak{N}$ and $\mathfrak{N}'$ can be reduced to each other. Thus reducing the decision problem for the theory of $\mathfrak{N}'$ to Sat(1MIC) shows this problem to be undecidable.

Given the encoding of natural numbers explained above, consider the formula *plus*(S, T) defined as

$$\begin{aligned}\text{plus}(S,T) \;:=\; & (\Box(b \to \psi_+) \wedge \neg \mathit{everywhere}(T \leftrightarrow \Box \mathit{false}) \wedge \\ & \qquad\qquad \neg \mathit{everywhere}(S \leftrightarrow \Box \mathit{false})) \vee \\ & \Box \mathit{false} \vee (\mathit{everywhere}(T \leftrightarrow \Box \mathit{false}) \wedge S) \vee \\ & \qquad (\mathit{everywhere}(S \leftrightarrow \Box \mathit{false}) \wedge T)\end{aligned}$$

where ψ_+ is defined as

$$\begin{aligned}\psi_+ := \mathbf{ifp} \quad X \leftarrow \; & (a \wedge \Box(a \to X)) \vee (b \wedge S) \vee \\ & (b \wedge \Box(b \to X) \wedge \mathit{everywhere}(\Box\Diamond(a \wedge X) \to T)).\end{aligned}$$

LEMMA 7.11. *If S and T encode natural numbers s and t as described above, then* plus(S, T) *encodes the sum $s + t$.*

Proof. Let s and t be the natural numbers coded by the sets S and T respectively. Assume $s, t > 0$ and consider the sub-formula ψ_+. We claim that at stage $i > 0$, X^i contains all a-nodes of height less than i and all b-nodes of height at most $\min\{s + i - 1, s + t - 1\}$. For a-nodes, this is trivial. The claim for the b-nodes is proved by induction on the stages $i > 0$.

—Clearly, X^1 contains exactly the b-nodes contained in S and thus, as $t \geq 1$, only nodes of height at most $s + t - 1$.

—Now assume that for $0 < i < t$ the claim has been proved, i.e. X^i contains all a-nodes of height less than i and all b-nodes of height at most $s + (i - 1)$. Therefore, the sub-formula *everywhere*$(\Box\Diamond(a \wedge X) \to T)$ is globally true and all b-nodes of height at most $s + i$ are included into X^{i+1}, as they satisfy $b \wedge \Box(b \to X)$.

—Now suppose $i = t$. Thus, X^i contains all a-nodes of height less than t. Clearly, the sub-formula $b \wedge \Box(b \to X)$ is true only for b-nodes of height at most $s + (t - 1) + 1$. Let u be a node of height exactly $s + t$ and thus, as $s \geq 1$, greater than i. Then there is a node of height $t + 1$ in the subtree rooted at u which satisfies $\Box\Diamond(a \wedge X)$ but not T. Therefore, *everywhere*$(\Box\Diamond(a \wedge X) \to T)$ is false at u and u is not included into X^{i+1}.

—The same argument shows that at no higher stage, a b-node of height greater than $s + t - 1$ can be added to the fixed-point. This proves the claim.

Now consider the formula *plus*(S, T). By assumption, $s, t > 0$ and therefore the formulae $\neg$*everywhere*$(T \leftrightarrow \Box$*false*$)$ and $\neg$*everywhere*$(S \leftrightarrow \Box$*false*$)$ are true for all nodes except for the leaves. Further, we have seen that the sub-formula $\Box(b \to \psi_+)$ becomes true for all nodes whose b-successors are of height less than or equal to $s + t - 1$ and, as we are working in models of the formula ω-height $\wedge\ \psi$, for all nodes of height less than or equal to $s + t$.

Now suppose $s = 0$. Then $\neg$*everywhere*$(S \leftrightarrow \Box$*false*$)$ is false for all nodes except for the leaves. Thus, $\Box$*false*$\vee$*everywhere*$(S \leftrightarrow \Box$*false*$)\wedge T$ defines all nodes contained in T and *plus*(S, T) defines the set of nodes representing the sum $0 + t = t$. The case where $t = 0$ is analogous. This finishes the proof. $\Box$

We now turn to the formalisation of multiplication in 1MIC. For this, consider the formula

$$\text{times}(S,T) := \Box \mathit{false} \;\vee\; (\neg \mathit{everywhere}(S \leftrightarrow \Box \mathit{false}) \wedge \neg \mathit{everywhere}(T \leftrightarrow \Box \mathit{false}) \wedge \Box(b \rightarrow \psi_*)),$$

where

$$\psi_* := \textbf{ifp}\;\; X \leftarrow (a \wedge \Box(a \rightarrow X) \vee (b \wedge \text{plus}(\Box(b \rightarrow X), S) \wedge \mathit{everywhere}(\Box\Diamond(a \wedge X)) \rightarrow T)).$$

LEMMA 7.12. *If S and T encode natural numbers s and t as described above, then* $\text{times}(S,T)$ *encodes the product* $s \cdot (t+1)$ *if* $s,t > 0$ *and* 0 *otherwise.*

Proof. First, suppose that s or t equals 0. In this case, one of the sub-formulae $\neg \mathit{everywhere}(S \leftrightarrow \Box \mathit{false})$ or $\neg \mathit{everywhere}(T \leftrightarrow \Box \mathit{false})$ becomes false on all nodes except the leaves and therefore the formula $\mathit{times}(S,T)$ is true only on the leaves and encodes the result 0.

Now suppose $s,t > 0$ and consider the induction on X in the sub-formula ψ_*. We claim that at stage $i > 0$, X^i contains all a-nodes of height less than i and all b-nodes of height at most $\min\{i \cdot s + i - 1, t \cdot s + t - 1\}$. For the a-nodes, the claim is obvious. Now consider a node labelled by b.

—In the first stage, X^1 contains all b-nodes satisfying $\text{plus}(\Box(b \rightarrow X), S)$. As X^0 is empty, $\Box(b \rightarrow X)$ is true only at the leaves and thus the formula *plus* evaluates to S.

—Now assume the claim has been proved for stage i. If $i < t$, then X^i contains a-nodes of height less than $i < t$. Thus, the sub-formula $\mathit{everywhere}(\Box\Diamond(a \wedge X)) \rightarrow T)$ is true for all nodes. By induction hypothesis, X^i contains exactly the b-nodes of height less than or equal to $i \cdot s + i - 1$ and therefore $\Box(b \rightarrow X)$ becomes true at all nodes of height at most $i \cdot s + i$. By Lemma 7.11, the formula $\text{plus}(\Box(b \rightarrow X), S)$ is then true for all nodes of height $(i \cdot s + i + s) = (i+1) \cdot s + (i+1) - 1$.

—Now suppose $i = t$. X^i contains exactly all a-nodes of height less than t and all b-nodes of height at most $t \cdot s + t - 1$. Let u be any b-node not in X^i, i.e. of height greater than $t \cdot s + t - 1$. Then there is a node of height $t + 1$ in the subtree rooted at u satisfying $\Box\Diamond(a \wedge X)$ but not T. Thus, the formula $\mathit{everywhere}(\Box\Diamond(a \wedge X) \rightarrow T)$ is false at u and u is not added to the fixed-point. The same argument, of course, holds for all stages $i > t$ and thus, in restriction to the b-nodes, the fixed point of X has been reached. This finishes the proof of the claim.

Now consider the formula *times*. By assumption, $s,t > 0$ and therefore the sub-formulae $\neg \mathit{everywhere}(S \leftrightarrow \Box \mathit{false})$ and $\neg \mathit{everywhere}(T \leftrightarrow \Box \mathit{false})$ are both true for all nodes except the leaves. Thus, $\text{times}(S,T)$ becomes true for all leaves – as they satisfy $\Box \mathit{false}$ – and for all nodes whose b-successors satisfy ψ_*, i.e. are of height $((t+1) \cdot s - 1)$. Together, $\text{times}(S,T)$ represents $(t+1) \cdot s$. □

The following corollary shows that the evaluation of polynomials in the structure $\mathfrak{N}'$ can be reduced to the evaluation of 1MIC formulae.

COROLLARY 7.13. *For every polynomial $f(x_1, \ldots, x_r)$ over $\mathfrak{N}'$ with coefficients in the natural numbers there exists a formula $\psi_f(X_1, \ldots, X_r) \in$ 1MIC such that for* *every tree* $\mathcal{T}, v \models \omega$*-height*$\wedge \psi$ *and all sets* $S_1, \dots, S_r$ *encoding numbers* $s_1, \dots, s_r \in \omega$

$$[\![\psi_f(S_1, \dots, S_r)]\!]^{\mathcal{T}} = \{v : h(v) \leq f^{\mathfrak{N}'}(s_1, \dots, s_r)\},$$

where $f^{\mathfrak{N}'}(s_1, \dots, s_r)$ *denotes the result of* f *in* $\mathfrak{N}'$.

Proof. The proof is by induction on f.

—$\psi_0 := \Box false$.

—$\psi_1 := \Box\Box false$.

—$\psi_x := X$.

—$\psi_{f+g} :=$ plus$[S/\psi_f, T/\psi_g]$, i.e. the formula obtained by replacing in plus(S,T) the variables S and T by ψ_f and ψ_g, respectively.

—$\psi_{f \cdot g} :=$ times$[S/\psi_f, T/\psi_g]$.

□

THEOREM 7.14. *For every* FO*-sentence* ψ *in the vocabulary* $\{+, \cdot, 0, 1\}$ *of arithmetic, there exists a formula* $\psi^* \in$ 1MIC *such that* ψ *is true in* $\mathfrak{N}'$ *if, and only if,* ψ^* *is satisfiable.*

Proof. By induction on formulae $\psi(\overline{x}) \in$ FO$[+, \cdot, 0, 1]$ we construct a formula $\psi^*(\overline{x}) \in$ 1MIC such that for all $\overline{n} \in \omega$ with encodings $\overline{N}$, $\psi(\overline{n})$ is true in $\mathfrak{N}'$ if, and only if, $\psi^*(\overline{N})$ is true at the root v of any model $\mathcal{T}, v$ of the formula ω-height $\wedge\, \psi$ above.

We have already seen how to transform polynomials over $\mathfrak{N}'$ and equality $x = y$ can be expressed by *everywhere*$(X \leftrightarrow Y)$. What remains is to translate quantifiers. Let $\psi(\overline{y}) := \exists x\, \varphi(x, y_1, \dots, y_k)$ be a formula in FO$[+, \cdot, 0, 1]$. By induction we get a corresponding formula $\varphi^*(X, Y_1, \dots, Y_k)$ in 1MIC. Let

$$\psi^*(\overline{Y}) := w \wedge \mathbf{ifp}\ \ X \leftarrow (\neg w \wedge \Box X) \vee (w \wedge \varphi^*(X/\neg w \wedge \Box X)).$$

Let $\mathcal{T}, v$ be any model of ω-height $\wedge\, \psi$. We claim that for all numbers $\overline{n}$ encoded by sets $\overline{N}$, $\psi(\overline{y})$ is true in $\mathfrak{N}'$ if, and only if, $\mathcal{T}, v \models \psi^*(\overline{N})$. A simple induction on the stages proves that at stage $i < \omega$, X^i contains all a and b-nodes of height less than i. Further, X^∞ contains the root v if at some stage X, the formula φ^* becomes true at the root where the variable X has been replaced by $\neg w \wedge \Box X$, i.e. at each stage i, the variable X in φ^* is interpreted by the set of nodes of height $\leq i$. Thus, in the course of the induction, φ^* is evaluated for all sets encoding natural numbers and, by induction hypothesis, becomes true at the root if, and only if, for at least one $i \in \mathbb{N}$, $\varphi(i, \overline{n})$, and thus $\psi(\overline{n})$, becomes true. This proves the claim and finishes the proof of the theorem. □

COROLLARY 7.15. *The satisfiability problem for* 1MIC *is undecidable. In fact, it is not even in the arithmetical hierarchy.*

8. CONCLUSION

While extensions of first-order logic by inflationary fixed points have been studied for a long time in finite model theory we have introduced and investigated here for the first time an analogous extension of propositional modal logic.

The pictures that emerges from our studies is actually quite surprising. While in finite model theory, least fixed point logic (LFP) and inflationary fixed point logic (IFP) behave very similarly, both concerning expressive power and algorithmic properties, the corresponding modal fixed point logics L_μ and MIC are very different.

—The modal iteration calculus MIC is much more expressive than the modal μ-calculus L_μ. In particular, MIC cannot be embedded into monadic second-order logic and does not have the finite model property.

—On the other side, one pays a prize for this greater expressive power: MIC is algorithmically much less manageable. The satisfiability problem for MIC is highly undecidable and the model checking problem is PSPACE-complete.

—There are other structural differences. Perhaps the most interesting one is that simultaneous inflationary inductions cannot be reduced to nested simple inductions.

Beyond the theoretical interest in the comparison of monotone and non-monotone inductions, our results show that MIC is *not* a useful logic for the common application areas of L_μ like automatic verification, where the expressive power of (fragments of) L_μ suffices for most applications and where efficient model checking algorithms are essential. It will be interesting to see whether inflationary fixed points in modal logic will just remain a theoretical curiosity, or whether there will be a practical use for very powerful modal logics like MIC.

REFERENCES

ABITEBOUL, S. AND VIANU, V. 1995. Computing with first-order logic. *Journal of Computer and System Sciences 50*, 309 – 335.

ARNOLD, A. AND NIWIŃSKI, D. 2001. *Rudiments of μ-calculus.* North Holland.

BLACKBURN, P., DE RIJKE, M., AND VENEMA, Y. 2001. *Modal Logic.* Cambridge University Press.

BÜCHI, J. 1960. Weak second-order arithmetic and finite automata. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik 6*, 66–92.

DAWAR, A. 1993. Feasible computation through model theory. Ph.D. thesis, University of Pennsylvania.

DAWAR, A., GRÄDEL, E., AND KREUTZER, S. 2001. Inflationary fixed points in modal logic. In *Proc. of the 10th Conf. on Computer Science Logic (CSL).* Lecture Notes in Computer Science (LNCS), vol. 2142. Springer Verlag, 277–291.

DAWAR, A. AND HELLA, L. 1995. The expressive power of finitely many generalized quantifiers. *Information and Computation 123*, 172–184.

DAWAR, A. AND KREUTZER, S. 2002. Generalising automaticity to modal properties of finite structures. In *Foundations of Software Technology and Theoretical Computer Science (FSTTCS).* Lecture Notes in Computer Science (LNCS). Springer Verlag.

DAWAR, A., LINDELL, S., AND WEINSTEIN, S. 1995. Infinitary logic and inductive definability over finite structures. *Information and Computation 119*, 160–175.

DZIEMBOWSKI, S. 1996. Bounded-variable fixpoint queries are PSPACE-complete. In *10th Annual Conference on Computer Science Logic CSL 96. Selected papers.* Lecture Notes in Computer Science, vol. 1258. Springer, 89–105.

EBBINGHAUS, H.-D. AND FLUM, J. 1999. *Finite Model Theory*, 2nd ed. Springer.

FLUM, J. 1999. On the (infinite) model theory of fixed-point logics. In *Models, algebras, and proofs*, X. Caicedo and C. Montenegro, Eds. Number 2003 in Lecture Notes in Pure and Applied Mathematics Series. Marcel Dekker, 67–75.

GÉCSEG, F. AND STEINBY, M. 1997. Tree languages. In *Handbook of Formal Languages*, G. Rozenberg and A. Salomaa, Eds. Vol. 3. Springer, New York, 1–68.

GRÄDEL, E. 2003. Finite model theory and descriptive complexity. In *Finite Model Theory and Its Applications*. Springer-Verlag. To appear.

GRÄDEL, E. AND OTTO, M. 1999. On logics with two variables. *Theoretical Computer Science 224*, 73 – 113.

GUREVICH, Y. AND SHELAH, S. 1986. Fixed-point extensions of first-order logic. *Annals of Pure and Applied Logic 32*, 265–280.

JANIN, D. AND WALUKIEWICZ, I. 1996. On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic. In *Proceedings of 7th International Conference on Concurrency Theory CONCUR '96*. Number 1119 in Lecture Notes in Computer Science. Springer-Verlag, 263–277.

JURDZINSKI, M. 1998. Deciding the winner in parity games is in UP ∩ Co-UP. *Information Processing Letters 68*, 119–124.

KREUTZER, S. 2002. Expressive equivalence of least and inflationary fixed-point logic. In *Proc. of the 17th Symp. on Logic in Computer Science (LICS)*. 403 – 413.

OTTO, M. 1999. Bisimulation-invariant Ptime and higher-dimensional mu-calculus. *Theoretical Computer Science 224*, 237–265.

SHALLIT, J. AND BREITBART, Y. 1996. Automaticity I: Properties of a measure of descriptional complexity. *Journal of Computer and System Sciences 53*, 10–25.

VAN BENTHEM, J. 1983. *Modal Logic and Classical Logic*. Bibliopolis, Napoli.

VARDI, M. 1995. On the complexity of bounded-variable queries. In *Proc. 14th ACM Symp. on Principles of Database Systems*. 266–267.

Received February 2002; accepted September 2002