On the bisimulation invariant fragment of monadic Σ_1 in the finite

Anuj Dawar[1] and David Janin[2]

[1] University of Cambridge Computer Laboratory, Cambridge CB3 0FD, UK, anuj.dawar@cl.cam.ac.uk.
[2] LaBRI, Université Bordeaux I, 33405 Talence, France, janin@labri.fr

Abstract. We investigate the expressive power of existential monadic second-order logic (monadic Σ_1) on finite transition systems. In particular, we look at its power to express properties that are invariant under forms of bisimulation and compare these to properties expressible in corresponding fixed-point modal calculi. We show that on finite unary transition systems the bisimulation invariant fragment of monadic Σ_1 is equivalent to bisimulation-invariant monadic second order logic itself or, equivalently, the mu-calculus. These results contrast with the situation on infinite structures. Although we show that these results do not extend directly to the case of arbitrary finite transition systems, we are still able to show that the situation there contrasts sharply with the case of arbitrary structures. In particular, we establish a partial expressiveness result by means of tree-like tiling systems that does not hold on infinite structures.

1 Introduction

The second author and Walukiewicz [5] showed in 1996 that any sentence of monadic second-order logic (MSO) whose models are invariant under bisimulation is equivalent to a sentence of Kozen's modal μ-calculus (L_μ).

The importance of the theorem lies, on the one hand, in the fact that monadic second-order logic is seen as a natural upper limit on the reasonable expressive power of languages for the specification of behaviours of concurrent systems. Indeed, almost all logics used in practice, such as LTL and CTL* are fragments of this logic. On the other hand, bisimulation is a natural relation describing the behavioral equivalence of processes. In speaking of behavioral specifications expressed in MSO, it seems natural to restrict oneself to those that are invariant under bisimulation. The theorem of Janin and Walukiewicz provides a syntactic characterization of the properties that are bisimulation invariant. Looked at from the other side, the theorem is also seen as an expressive completeness result for the μ-calculus.

The methodology used in the proof of this theorem is based on automata on infinite trees. Every transition system is equivalent by bisimulation to a tree and, on trees, the evaluation of MSO formulas can be expressed as the evaluation of alternating tree automata. By considering trees that are, in a precise sense, *saturated* one can transform automata to show that these formulas are equivalent to formulas of the μ-calculus. This essential methodology has proved productive in establishing variants of the original result. It is known, for instance, that sentences of MSO that are invariant under *counting bisimulation* are equivalent to C_μ—the modal fixed-point calculus with counting modalities [12, 4]. It has also been shown that the existential fragment of MSO (which we denote monadic Σ_1) is, for bisimulation invariant properties, expressively equivalent to N_1—the fragment of the μ-calculus with only greatest fixed points [4].

However, it remains an open question whether a version of this expressive completeness result is true if we restrict ourselves to finite structures. That is, is it the case that every sentence of MSO that is bisimulation-invariant *on finite structures* is equivalent, *again on finite structures* to a sentence of L_μ?

This statement has a weaker hypothesis and conclusion than the original theorem and is therefore not a consequence of it. It has been the subject of much recent investigation. The corresponding finite versions of the equivalence between monadic Σ_1 and N_1 for bisimulation invariant properties and of MSO and C_μ for counting bisimulation also remain open. One related result that is known to carry over into the finite is the theorem of van Benthem (see [11]) that any first-order definable property that is invariant under bisimulation is definable in propositional modal logic. It has been shown by Rosen [7] that this statement is still true when we restrict ourselves to finite structures.

One reason why the question of the equivalence of these logics is so different in the finite is that, once we restrict ourselves to finite structures, we no longer have a tree model property. That is, it is no longer the case that every structure is equivalent by bisimulation to a tree. In the general case, it is possible to take the collection of all (saturated) infinite trees as a canonical class of models that intersects every bisimulation equivalence class. Thus, as one is considering formulae invariant under bisimulation, one can restrict oneself to this class and on this class there are well-behaved automata models for the logics we consider. Unfortunately, there is no class of finite structures that fulfills these conditions.

Main results

In this paper, we are mainly concerned with the study of the bisimulation invariant fragment of monadic Σ_1 in the finite.

We show that restricting ourselves to finite structures that are *unary*. i.e. in which each node has a single successor, this fragment is as expressive as (the bisimulation invariant fragment of) full monadic second order logic. In other words, we obtain a complete characterization of the expressive power of the bisimulation invariant fragment of MSO on such structures. As a corollary, the correspondence between monadic Σ_1 and N_1 that holds on arbitrary (finite and infinite) unary structures just fails in the finite.

On finite structures that are not necessarily unary, however, the situation is less clearcut. We obtain a counterexample to the equivalence of monadic Σ_1 with NC_1 (the first level of the C_μ hierarchy) on finite structures, demonstrating that this situation is distinct from the case of arbitrary (finite and infinite) systems.

We also show that monadic Σ_1 is not as expressive as bisimulation-invariant MSO, so the situation also differs from the unary case.

These two negative results leads us to consider tiling systems [10], which are known to capture monadic Σ_1 on finite structures. We show that when the properties concerned are bisimulation invariant, simple tiling systems suffice. More precisely, we show that if a sentence φ of monadic Σ_1 is invariant under bisimulation then there is a class of structures, including representatives of all bisimulation classes, on which φ is characterized by a *tree-like* tiling system of *radius one* (these terms are made precise below). One might expect that this normal form could be further refined so that the tiles are what we call *forward looking*. This would establish that bisimulation invariant properties of monadic Σ_1 can be expressed in N_1. However, such a methodology would also yield the result for the counting case, which is refuted by the counterexample obtained on unary structures.

2 Background and Definitions

Models and standard logics

The logics we consider are interpreted in *transition systems*, also called Kripke structures, or simply labeled directed graphs (in the sequel, when we use the term *graph*, we mean a labeled directed graph). Fix a set A of actions and a set *Prop* of atomic propositions. A transition system for A and *Prop* is a structure

$$\mathcal{K} = \langle V, r, \{E_a\}_{a\in A}, \{p^{\mathcal{K}}\}_{p\in Prop}\rangle$$

with universe V (whose elements are called states), a distinguished element called the *root* $r \in V$, binary relations $E_a \subseteq V \times V$ for each $a \in A$ and unary relations $p^{\mathcal{K}} \subseteq V$ for each atomic proposition $p \in Prop$. For the sake of clarity, we confine ourselves in this paper to vocabularies where A consists of a single action. We then drop the subscript a on the binary relation E. All of our results apply equally well to the more general case.

Such transition systems are usually used to interpret modal logics, which we consider below. We also interpret standard predicate logics, in particular first-order logic (FO) and monadic second-order logic (MSO) in transition systems. In the sequel, we shall write $\varphi(x_1, \cdots, x_n)$ or simply $\varphi(\bar{x})$ for an FO or MSO formula with free first-order variables among $\bar{x} = (x_1, \cdots, x_n)$ *regardless* of the free monadic predicate (or set) variables occurring in φ. More precisely, given the set $\{X_1, \cdots, X_n\}$ of all set variables occurring free in φ, we shall implicitly and whenever required interpret the formula φ on transition systems with the set of atomic proposition $Prop' = Prop \cup \{X_1, \cdots, X_n\}$.

Bisimulation and counting bisimulation

A *directed (resp. undirected) path* in a transition system $\mathcal{K}$ is a (finite or infinite) sequence of vertices such that for any two consecutive vertices v_1 and v_2 in the sequence one has $(v_1, v_2) \in E$ (resp. (v_1, v_2) or $(v_2, v_1) \in E$). The directed (resp. undirected) distance $d_d(v_1, v_2)$ (resp. $d(v_1, v_2)$) between two vertices v_1 and v_2 is the length of the shortest directed (resp. undirected) path from v_1 to v_2.

A directed (resp. undirected) cycle in $\mathcal{K}$ is a periodic infinite directed (resp. undirected) path. Given an integer k, we say that a graph $\mathcal{K}$ is *k-acyclic* if any *undirected* cyclic path in $\mathcal{K}$ contains at least $k+1$ distinct vertices. Given two transition systems $\mathcal{K} = \langle V, r, E, \{p^{\mathcal{K}}\}_{p \in Prop} \rangle$ and $\mathcal{K}' = \langle V', r', E', \{p^{\mathcal{K}'}\}_{p \in Prop} \rangle$, a *bisimulation* between $\mathcal{K}$ and $\mathcal{K}'$ is a relation $B \subseteq V \times V'$ such that, if $(v, v') \in B$ then:

- for each $p \in Prop$, $v \in p^{\mathcal{K}} \iff v' \in p^{\mathcal{K}'}$;
- for each w with $(v, w) \in E$ there is a w' with $(v', w') \in E'$ and $(w, w') \in B$; and
- for each w' with $(v', w') \in E'$ there is a w with $(v, w) \in E$ and $(w, w') \in B$.

A *counting bisimulation* between $\mathcal{K}$ and $\mathcal{K}'$ is a relation $B \subseteq V \times V'$ such that, if $(v, v') \in B$ then:

- for each $p \in Prop$, $v \in p^{\mathcal{K}} \iff v' \in p^{\mathcal{K}'}$;
- B contains a bijection between the sets $\{w : (v, w) \in E\}$ and $\{w' : (v', w') \in E'\}$.

Observe that any counting bisimulation is a bisimulation.

We say that $\mathcal{K}$ and $\mathcal{K}'$ are (counting) bisimilar if there is a (counting) bisimulation B between them with $(r, r') \in B$. More generally, we say that two states $v \in \mathcal{K}$ and $v' \in \mathcal{K}'$ (where $\mathcal{K}$ and $\mathcal{K}'$ are not necessarily distinct) are (counting) bisimilar if there is a (counting) bisimulation B between the two structures with $(v, v') \in B$.

Given a class $\mathcal{C}$ of transition systems, we say that an FO or MSO sentence φ is (counting) bisimulation invariant on $\mathcal{C}$ when, for any two (counting) bisimilar models $\mathcal{K}$ and $\mathcal{K}' \in \mathcal{C}$ one has $\mathcal{K} \models \varphi$ if, and only if, $\mathcal{K}' \models \varphi$. Accordingly, we say that φ is bisimulation invariant *in the finite* when it is bisimulation invariant on the class of finite structures.

Modal logic and the mu-calculus

The *modal propositional logic* (ML) consists of formulas built up from the propositions in *Prop* and the propositional constants *true* and *false* using the Boolean connectives and the modalities $\Box$ and $\Diamond$: i.e., for a formula α, $\Box\alpha$ and $\Diamond\alpha$ are also formulas. For the semantics, we just note that $\mathcal{K}, v \models \Diamond\alpha$ if, and only if, there is a v' with $(v, v') \in E$ such that $\mathcal{K}, v' \models \alpha$ (and dually for $\Box\alpha$).

The *modal depth* of a modal formula is defined to be the maximal depth of nesting of modalities in α, i.e. the modal depth of a modality free formula is defined to be zero; if α is of modal depth k then the modal depth of $\Diamond\alpha$ or $\Box\alpha$ is $k+1$; and the modal depth of a Boolean combination of formulas is the maximum modal depth of any one of the formulas.

The *modal μ-calculus* L_μ is obtained by extending ML with a countable collection of propositional variables X so that a variable by itself is a formula and, if α is a formula and X a variable which occurs only positively (i.e., only within the scope of an even number of negation signs) in α then $\mu X.\alpha$ and $\nu X.\alpha$ are also formulas in which the variable X is bound. For the semantics, given a structure $\mathcal{K}$ and an interpretation in $\mathcal{K}$ for all the free variables of α, we say that $\mathcal{K}, v \models \mu X.\alpha$ if v is in the least set $X \subseteq V$ such that $X \iff \alpha$. Similarly $\mathcal{K}, v \models \nu X.\alpha$ if v is in the *greatest fixed point* defined by α.

A key feature of the modal logics ML and L_μ is that the properties they express are bisimulation invariant. That is, if $\mathcal{K}$ and $\mathcal{K}'$ are bisimilar then for any formula α, $\mathcal{K} \models \alpha$ if, and only if, $\mathcal{K}' \models \alpha$.

There is a standard translation of formulas of ML into the first-order logic of transition systems. That is, for each formula α of ML, there is a formula $\varphi_\alpha(x)$ of first-order logic with one free first-order variable x (in the vocabulary with a binary relation symbol E and unary relation symbols for each $p \in \mathit{Prop}$) that defines in each $\mathcal{K}$ exactly the set of states in which α is true. Similarly, there is a straightforward translation from L_μ to monadic second-order logic.

By results of van Benthem [11] and Janin and Walukiewicz [5] we know that there are converses for these translations. That is, every property of transition systems that is expressible in FO and is invariant under bisimulation is expressible in ML and any bisimulation-invariant property that is definable in MSO is also definable in L_μ.

Using the equivalences $\Box\alpha \iff \neg\Diamond\neg\alpha$, $\nu X.\alpha \iff \neg\mu X.\neg\alpha[\neg X/X]$ and De Morgan's laws, it is possible to transform any formula of L_μ into negation normal form, where negation signs only appear before propositional atoms. We write N_1 for the collection of formulas in negation normal form in which no instance of the operator μ appears. Similarly, M_1 is the collection of formulas without ν. These are the bottom two levels of an alternation hierarchy which is known to give strictly increasing expressive power (see [2]).

It is easily seen that when we translate L_μ to MSO, formulas of N_1 yield *existential* MSO formulas (i.e., in prenex normal form, all second-order quantifiers are existential) while formulas of M_1 yield *universal* MSO formulas. By a result of Janin and Lenzi [4] we get a converse of these statements for bisimulation-invariant properties. That is, any bisimulation-invariant property definable in existential MSO (also written as monadic Σ_1) is definable in N_1.

The *counting modal logic* and the *counting μ-calculus* C_μ are defined similarly to ML and L_μ except the rules for $\Box$ and $\Diamond$ are replaced by: for each $i \in \mathbb{N}$, if α is a formula then so are $\Box_i\alpha$ and $\Diamond_i\alpha$. For the semantics, we say that $\mathcal{K}, v \models \Diamond_i\alpha$ if there are at least i distinct v' such that $(v, v') \in E$ and $\mathcal{K}, v' \models \alpha$. We write NC_1 (by analogy with N_1) for the fragment of C_μ without least fixed-points.

In the sequel, we also use *backward modalities* $\Diamond^{-1}$ and $\Box^{-1}$, and *backward counting modalities* $\Diamond_i^{-1}$ and $\Box_i^{-1}$ that are defined like the ordinary modalities but with respect to the inverse edge relation E^{-1} in place of E. In the presence of backward modalities, the standard modalities are referred to as *forward modalities*.

3 Monadic Σ_1 on finite unary graphs

In this section, we study the expressive power of monadic Σ_1 on unary graphs. We first review the straightforward relationship between (bisimulation classes of) finite unary graphs and ultimately periodic infinite words. We establish that monadic Σ_1 in the finite is expressive enough to define all ω-regular languages. Then we prove that, on finite unary graphs, the bisimulation (or counting bisimulation) invariant fragment of monadic Σ_1 is the same as the bisimulation invariant fragment of full MSO. These results contrast with the case of arbitrary (finite or infinite) unary graphs where monadic Σ_1 can only express topologically closed regular languages.

A graph $\mathcal{K}$ is a *unary graph* if every vertex in $\mathcal{K}$ has a unique successor under the relation E. Of course, the bisimulation class of a unary graph is completely characterized by the infinite word (in the alphabet $\Sigma = \mathscr{P}(\mathit{Prop})$) that is described by the path emanating from the root. Thus, we can see any bisimulation-invariant property of unary finite graphs as described by a language of eventually periodic ω-words. So, given such a language $L \subseteq \Sigma^\omega$, and a class of finite unary graphs C, we say that C is equivalent to L in the finite if:

- for any graph $\mathcal{K} \in C$, there is a word $w_\mathcal{K} \in L$ such that $w_\mathcal{K}$ is the Σ-word defined by the unique infinite path starting at the root of $\mathcal{K}$.
- for any ultimately periodic word $w \in L$ there is a graph $\mathcal{K}_w \in C$ such that w is the Σ-word defined by the infinite path starting at the root of $\mathcal{K}$.

By extension, we say that an MSO sentence φ is equivalent to L when the class C_φ of finite unary graphs it defines is equivalent to L. Note that if this is the case then φ is invariant under counting bisimulation. Note further that on the class of finite unary graphs, counting bisimulation coincides with bisimulation.

Theorem 1. *For any regular ω-language $L \subseteq \Sigma^\omega$ there is a (counting bisimulation-invariant) monadic Σ_1 formula φ_L equivalent to L in the finite.*

Proof. Let L be an ω-regular language. First, one can show that that there is a nondeterministic finite Büchi automaton $\mathcal{A}_L = \langle Q, Q_0, \delta, F\rangle$ with set of states Q, set of initial states Q_0, transition function $\delta : Q \times \Sigma \to \mathscr{P}(Q)$ and accepting states F, that recognizes L and such that, for any infinite word of L of the form $u.v^\omega$, there is an initial state $q_0 \in Q_0$ and an accepting state $q \in F$ such that, there is a path in $\mathcal{A}_L$ from state q_0 to state q reading u (with $q_0 = q$ when $u = \epsilon$), and a cycle in $\mathcal{A}_L$ from q to q reading v.

The formula φ_L can now be defined as follows: there is a collection of disjoint sets X_q ($q \in Q$), such that: (i) $r \in X_{q_0}$ for some $q_0 \in Q_0$; (ii) for each $q \in Q$ and $x \in X_q$, x has a single successor y and there is a state $q' \in \delta(q, \lambda(x))$ such that $y \in X_{q'}$, where $\lambda(x) = \{p \in \mathit{Prop} : p(x) \text{ holds}\}$; and (iii) any element with two predecessors in $\bigcup_{q\in Q} X_q$ (and the root if it has one predecessor in $\bigcup_{q\in Q} X_q$) must belong to some X_q with $q \in F$.

One can check that φ_L defined in such a way (i) is counting bisimulation invariant, (ii) does enforce that there is a unique path from the root and, (iii) the word described by this path is accepted by the automaton $\mathcal{A}_L$. $\square$

Since only topologically closed regular languages are definable in the level N_1 of the mu-calculus hierarchy, this first theorem already shows that:

Corollary 2. *There is a bisimulation invariant class of unary finite models definable in monadic Σ_1 that is not definable in N_1.*

One might expect a converse to Theorem 1 to hold. Indeed, we even prove a stronger result.

Theorem 3. *For any MSO formula φ, counting-bisimulation invariant on finite graphs and true only on unary graphs, there is a regular language $L_\varphi \subseteq \Sigma^\omega$ equivalent to φ in the finite.*

The remainder of this section is dedicated to the proof of this theorem.

A unary graph $\mathcal{K}$ is called a *lasso* if the root of $\mathcal{K}$ has no predecessor and all other vertices except one (called the *knot*) have exactly one predecessor while the knot has exactly two predecessors.

Any lasso $\mathcal{K}$ is completely characterized by the two non empty finite words u and v (in the alphabet Σ) that are described respectively by the (acyclic) path from the root to the knot of $\mathcal{K}$ (excluding the knot) and the cyclic path from the knot to itself (excluding the second occurrence of the knot). In the sequel, we write $\mathcal{K}_{u,v}$ for such a lasso.

Observe that any finite unary graph is counting bisimilar to a lasso. More precisely, it is counting bisimilar to the subgraph induced by the set of vertices reachable from the root that forms (possibly after duplicating the root so that it is distinct from the knot) a lasso.

We are now ready to start the proof of Theorem 3. Let φ be an MSO formula as in Theorem 3.

Proposition 4. *There is a finite set of pairs of regular languages $(U_i, V_i)_{i\in I}$ such that, for any two words u and $v \in \Sigma^+$, $K_{u,v} \models \varphi$ if, and only if, there is some $i \in I$ such that $u \in U_i$ and $v \in V_i$.*

Proof. The mapping that maps any pair of non empty finite words $(u, v) \in \Sigma^+ \times \Sigma^+$ to the lasso $\mathcal{K}_{u,v}$ is a FO-definable transduction. It follows, by Shelah's decomposition theorem [9] that there is a finite set of pairs of MSO formulas $\{(\varphi_i, \psi_i)\}_{i\in I}$ over finite Σ-words such that for any two words u and $v \in \Sigma^+$, $K_{u,v} \models \varphi$ if and only if there is some $i \in I$ such that $u \models \varphi_i$ and $v \models \psi_i$. By Büchi's theorem, for all $i \in I$, the MSO-formulas φ_i and ψ_i define the regular languages U_i and V_i we are looking for. ☐

Remark. One might think that Proposition 4 concludes the proof of the theorem. Indeed, if $K_{u,v} \models \varphi$, then $u.v^\omega$ belongs to some $U_i.V_i^\omega$ so one might think that φ is equivalent to the language $\bigcup_{i\in I} U_i.V_i^\omega$. However, this idea fails since, a priori, nothing ensures that when an ultimately periodic word w belongs to some $U_i.V_i^\omega$ then it is of the form $u.v^\omega$ with $u \in U_i$ and $v \in V_i$ so that $K_{u,v} \models \varphi$.

So far, we have not used the fact that φ is counting bisimulation invariant on finite graphs.

Proposition 5. *For any $i \in I$ and any $(u,v) \in U_i \times V_i$, there is a triple $t = (j,r,s) \in I \times \Sigma^+ \times \Sigma^+$ such that:*

1. *$r.s^\omega = u.v^\omega$ (hence $K_{u,v}$ and $K_{r,s}$ are counting bisimilar),*
2. *for all $n > 0$, $r.s^n \in U_j$ and $s^n \in V_j$.*

Proof. Let i, u and v be as above, so $K_{u,v} \models \varphi$. By invariance of φ, for each $k > 0$, we also have $K_{u.v^k,v^k} \models \varphi$. Hence, by Proposition 4 for each $k > 0$ there is some $i_k \in I$ such that $(u.v^k, v^k) \in U_{i_k} \times V_{i_k}$. Since I is finite, there is some $j \in I$ such that $j = i_k$ for infinitely many k. Now, since both U_j and V_j are regular languages and there are infinitely many k such that $u.v^k \in U_j$ and $v^k \in V_j$ there must be some $p > 0$ such that $u.v^{pn} \in U_j$ and $v^{pn} \in V_j$ for all $n > 0$. Taking $r = u.v^p$ and $s = v^p$ gives us the desired triple t. □

A triple $t = (j,r,s)$ as in Proposition 5 is called special. Write $\mathcal{S}$ for the set of all special triples.

To continue the proof of Theorem 3, we need some standard definitions from formal language theory. Recall that the *left congruence class* $[w]^l_L$ and the *right congruence class* $[w]^r_L$ of a finite word $w \in \Sigma^+$ with respect to a language $L \subseteq \Sigma^+$ are defined as the sets of words

$$[w]^l_L = \{w' \in \Sigma^+ : \forall u \in \Sigma^*, u.w \in L \Leftrightarrow u.w' \in L\}$$

and

$$[w]^r_L = \{w' \in \Sigma^+ : \forall v \in \Sigma^*, w.v \in L \Leftrightarrow w'.v \in L\}$$

We know that if L is regular there are only finitely many distinct sets $[w]^l_L$ and $[w]^r_L$ for $w \in \Sigma^*$ and each one is a regular language.

For any special triple $t = (j,r,s)$ we define the languages

$$D_t = [r]^r_{U_j}.([s]^l_{U_j} \cap [s]^r_{V_j}) \quad \textit{and} \quad E_t = ([s]^l_{U_j} \cap [s]^r_{V_j})$$

By construction, both languages D_t and E_t are regular. Moreover:

Proposition 6. *For any special triple $t = (j,r,s)$, $D_t \subseteq U_j$, $E_t \subseteq V_j$, $D_t.E_t^+ \subseteq D_t$ and $E_t^+ \subseteq E_t$ and, for any u and $v \in \Sigma^+$, if $u \in D_t$ and $v \in E_t$ then $\mathcal{K}_{u,v} \models \varphi$.*

Proof. Immediate consequence of the constructions, Proposition 5 and Proposition 4. □

We now conclude the proof of Theorem 3 by proving the following proposition:

Proposition 7. *The ω-regular language $L = \bigcup_{t \in \mathcal{S}} D_t.(E_t)^\omega$ is equivalent to φ.*

Proof. Assume that $\mathcal{K} \models \varphi$ for some finite model $\mathcal{K}$. By assumption, $\mathcal{K}$ is unary and counting bisimilar to some lasso $\mathcal{K}_{u,v}$. We show that $u.v^\omega$ belongs to L by applying Proposition 5. Indeed, this guarantees that there is a special triple $t = (j,s,r)$ such that $u.v^\omega = r.s^\omega$ and, by construction, $r.s^\omega \in D_t.E_t^\omega$.

For the converse, let w be an ultimately periodic word in L. By definition of L, this means that there is a special triple $t = (j,r,s)$ such that $w \in D_t.(E_t)^\omega$. In other words, $w = u_1.w_1$ with $u_1 \in U_t$ and $w_1 \in V_t^\omega$.

Now, since w is ultimately periodic so is w_1 and thus, because V_t is regular, w_1 is of the form $v_1.v_2 \dots v_n.(v_{n+1} \dots v_{n+m})^\omega$ for some $v_1, \dots, v_{n+m} \in V_t$.

Defining $u = u_1.v_1 \dots v_n$ and $v = v_{n+1} \dots v_{n+m}$, we have $w = u.v^\omega$ by construction. Hence $\mathcal{K}_w$ is counting bisimilar to $\mathcal{K}_{u,v}$. We also have $u \in D_t$ and $v \in E_t$ (applying Proposition 6) hence $\mathcal{K}_{u,v} \models \varphi$ and thus $\mathcal{K}_w \models \varphi$. $\square$

Putting Theorems 1 and 3 together gives the following corollary.

Corollary 8. *Any MSO formula counting bisimulation-invariant on finite unary graphs is equivalent to a monadic Σ_1 formula.*

Moreover, restricted to the class of unary graphs, the (counting or modal) mu-calculus can define exactly the classes corresponding to ω-regular languages. This gives us the following.

Corollary 9. *The counting bisimulation-invariant fragment of monadic Σ_1 on finite unary graphs is equivalent to L_μ.*

4 Monadic Σ_1 on arbitrary finite graphs

In this section, we aim at a characterization of the bisimulation invariant fragment of monadic Σ_1 on finite graphs. We establish two negative results that demonstrate how this case differs from both the more restricted class of finite unary graphs and the wider class of arbitrary (finite or infinite) graphs. Nonetheless, by means of a translation to tiling systems [10], we obtain a partial characterization of this fragment.

Theorem 10. *There is monadic Σ_1 counting bisimulation invariant formula φ that is not equivalent to a formula of the level NC_1 of the counting mu-calculus.*

Proof. The monadic Σ_1 formula φ_L of Theorem 1 is counting bisimulation invariant on *all* finite graphs, not just unary ones. Since any formula of NC_1 defining a regular language must define a topologically closed regular language, it suffices to take for L a language that is not closed, e.g. $L = (a + b)^*.b^\omega$. $\square$

Theorem 11. *There is a bisimulation invariant MSO formula that is not equivalent (on finite graphs) to a bisimulation invariant monadic Σ_1 formula.*

Proof. We know [1] that directed reachability, though definable in monadic Π_1 in the finite, is not definable in monadic Σ_1. Consider now the μ-calculus formula $p \wedge \mu X.(q \vee \Diamond X)$ that defines the set of vertices satisfying p from which there is a (directed) path to a vertex satisfying q. If there were an equivalent monadic Σ_1 formula we would be able to define in monadic Σ_1 the class of graphs in which a distinguished target t is reachable from a source s. We would get this by replacing p and q by formulas that define s and t respectively. $\square$

We are now left with a direct attempt to characterize the expressive power of the bisimulation invariant fragment of monadic Σ_1 in the finite.

It is known (see, for instance, [8]) that monadic Σ_1 formulas can only define *local* properties. Indeed, such formulas can be characterised by *tiling systems* [10], which are a generalization of automata operating on graphs rather than strings or trees.

Given a positive integer k, we say an FO-formula φ is *k-local* around a first-order variable x if it is equivalent to the formula obtained from φ by restricting all quantifiers in φ to the k-neighborhood of x, i.e. replacing any subformula of the form $\forall y\psi$ (resp. $\exists y\psi$) in φ by one of the form $\forall y(d(x,y) \leq k) \to \psi$ (resp. $\exists y(d(x,y) \leq k) \wedge \psi$). A *local* formula is one that is k-local for some k.

Note for any modal (or counting modal) formula α of modal depth k, the FO translation $\varphi_\alpha(x)$ is k-local around x. Indeed, it is k-local and *forward-looking*, in that we can restrict the quantifiers to the *directed* k-neighborhood by replacing $\forall y\psi$ by $\forall y(d_d(x,y) \leq k) \to \psi$, etc.

Furthermore, when a sentence is (counting) bisimulation invariant, its truth in a model only depends on the submodels induced by the vertices reachable from the root. The following proposition is a consequence.

Proposition 12. *Any (counting) bisimulation invariant sentence φ of monadic Σ_1 is equivalent, on the class of finite structures, to one of the form*

$$\exists X_1 \ldots \exists X_l \forall x \varphi$$

where φ is local.

Proof. Immediate consequence of Theorem 3.4 in [8]. □

Adapting the terminology of Thomas [10], we call a monadic Σ_1 formula of this form a *tiling system*. The local formula φ in such a tiling system is called a *tiling constraint*. When the tiling constraint is k-local, we say that k is the *radius* of the tiling system. When the tiling constraint is equivalent to a modal formula (with forward and backward modalities), we say that the tiling system is *tree-like*. One can check that when no backward modalities occur in the tiling constraint, a tiling system is just a closed (modal counting) alternating tree automaton (see [4] for a precise definition).

Now, our aim is to push the construction that transforms a (counting) bisimulation invariant tiling system into a tree automaton as far as it can go on finite structures. We show that any such tiling system is equivalent to a tree-like tiling system of radius 1 on a sufficiently rich class of graphs.

We say that a graph is k-acyclic when it contains no undirected cycle of length less than $k+1$. We first show that for any structure $\mathcal{K}$ and positive integer k, we can find a k-acyclic structure that is counting bisimilar to $\mathcal{K}$ but contains no undirected cycles of length smaller than k. The construction is similar to that of acyclic covers in [6].

Definition 13 (Powergraph). *For a finite graph $\mathcal{K} = \langle V, r, E, \{p^{\mathcal{K}}\}_{p\in Prop}\rangle$ define its* powergraph $2^{\mathcal{K}}$ *to be the graph $2^{\mathcal{K}} = \langle V', r', E', \{p^{\mathcal{K}'}\}_{p\in Prop}\rangle$ defined by $V' = V \times 2^V$ (where 2^V denotes the set of maps $V \to \{0,1\}$), $r' = (r, \bar{0})$, there is an edge E' from a vertex (v,f) to a vertex (w,g) whenever $(v,w) \in E$*

and g equals the function defined from f by taking, for each $u \in V$, $g(u) = f(u)$ when $u \neq w$ and $g(w) = 1 - f(w)$, and with, for each $p \in Prop$, $p^{\mathcal{K}'} = \{(v, \bar{b}) \in V' : v \in p^{\mathcal{K}'}\}$.

Proposition 14. *Graphs $\mathcal{K}$ and $2^{\mathcal{K}}$ are counting bisimilar and, if $\mathcal{K}$ is k-acyclic for some k then $2^{\mathcal{K}}$ is $2k$-acyclic.*

Proof. (sketch) The mapping $h : V' \to V$ that maps each vertex (v, f) in $2^{\mathcal{K}}$ to the vertex $h(v, f) = v$ in $\mathcal{K}$ induces a counting bisimulation. Now, consider an undirected cycle in the graph $2^{\mathcal{K}}$. Along any edge from (v, f) to (w, g), f and g must differ in exactly one bit. Thus, for the cycle to return to its starting point, all bits that are changed must flip at least twice. This then maps via h to a cyclic path in $\mathcal{K}$ where all vertices occur at least twice. □

Corollary 15. *For each positive integer k and every graph $\mathcal{K}$, there is a k-acyclic graph $\mathcal{K}'$ counting bisimilar to $\mathcal{K}$.*

Proof. By iterating the powergraph construction. □

Let φ be a counting bisimulation invariant monadic Σ_1 formula. By applying Proposition 12, we may assume that φ is a tiling system of the form $\varphi \equiv \exists X_1 \ldots \exists X_l \forall x \psi$ with ψ k-local. The following proposition is straightforward from definitions:

Proposition 16. *Let ψ_a be the k-local FO formula asserting that the k-neighbourhood of x is acyclic. The formula φ is equivalent, over k-acyclic graphs, to the formula, $\varphi' \equiv \exists X_1 \ldots \exists X_l \forall x(\psi \wedge \psi_a)$*

Now, we obtain the following

Theorem 17. *Formula φ is equivalent on k-acyclic graphs to a formula φ'' of the form $\varphi'' \equiv \exists Y_1 \ldots \exists Y_m \forall x \psi'$ with ψ' a 1-local tree-like constraint.*

Proof. (sketch) The proof is based on the observation that the Hintikka type (see [3]) of a tree centered on a node c is completely determined by the atomic propositions that are true at c and the Hintikka types of the subtrees rooted at the neighbours of c. Thus, by introducing a fresh set of second-order quantifiers (logarithmic in the number of Hintikka types), it is not difficult to build the formula φ''. □

As the constraint ψ' is tree-like of radius 1, it can be described by a counting modal formula with forward and backward modalities.

Remark. If this formula were equivalent to one without backward modalities, then one could show that we can obtain a formula θ of NC_1 that is equivalent to φ on k-acyclic graphs. As φ is invariant under counting bisimulation on finite structures by hypothesis and θ by definition and since the class of k-acyclic graphs contains representatives of all bisimulation classes on finite structures, it follows that θ and φ are equivalent on the class of all finite structures. Thus, we would have proved that every formula of monadic Σ_1 invariant under counting bisimulation is equivalent to a formula of NC_1. This would contradict Theorem 10.

5 Conclusions

On finite unary graphs, we provide a precise characterization of bisimulation-invariant MSO. In this case, the structure of unary graphs is simple enough so that standard techniques from mathematical logic and language theory apply. Since unary graphs are closed under counting bisimulation, this also allows us to show that on finite graphs in general, monadic Σ_1 can express more counting-bisimulation invariant properties than C_μ with only greatest fixed points.

In the general case the question of whether bisimulation-invariant MSO is equivalent on finite structures to L_μ remains a challenging open problem. By investigating this question at the first level of the monadic alternating hierarchy we have been able to show that the nature of the problem is radically different to its counterpart on infinite structures, while also being different to the restriction to unary structures.

We provide a translation of bisimulation-invariant monadic Σ_1 formulas to tree-like tiling systems on a sufficiently rich class of structures. However, it seems that the use of backward modalities in such tiling systems cannot be eliminated without passing to infinite structures. The relationship between these tiling systems and the μ-calculus needs to be investigated further out.

References

1. M. Ajtai and R. Fagin. Reachability is harder for directed rather than undirected finite graphs. *Journal of Symbolic Logic*, 55:113–150, 1990.
2. J. Bradfield. The modal mu-calculus alternation hierarchy is strict. *Theoretical Computer Science*, 195:133–153, 1998.
3. H-D. Ebbinghaus and J. Flum. *Finite Model Theory.* Springer, 2 edition, 1999.
4. D. Janin and G. Lenzi. On the logical definability of topologically closed recognizable languages of infinite trees. *Computing and Informatics*, 21:185–203, 2002.
5. D. Janin and I. Walukiewicz. On the expressive completeness of the modal mu-calculus with respect to monadic second order logic. In *Conf. on Concurrency Theory (CONCUR'96)*, pages 263–277. LNCS 1119, 1996.
6. M. Otto. Modal and guarded characterisation theorems over finite transition systems. In *Proc. of the 17th IEEE Symp. on Logic in Computer Science (LICS)*, pages 371–380, 2002.
7. E. Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6:427–439, 1997.
8. T. Schwentick and K. Barthelmann. Local normal forms for first-order logic with applications to games and automata. *Discrete Mathematics and Theoretical Computer Science*, 3:109–124, 1999.
9. S. Shelah. The monadic second order theory of order. *Annals of Mathematics*, 102:379–419, 1975.
10. W. Thomas. Automata theory on trees and partial orders. In M. Dauchet M. Bidoit, editor, *TAPSOFT'97*, pages 20–38. LNCS 1214, Springer-Verlag, 1997.
11. J. van Benthem. *Modal Logic and Classical Logic.* Bibliopolis, 1983.
12. I. Walukiewicz. Monadic second order logic on tree-like structures. In *Symp. on Theoretical Aspects of Computer Science*, 1996. LNCS 1046. Full version in *Information and Computation* 164 (2001) pp. 234-263,.