The Monadic Quantifier Alternation Hierarchy over Graphs is Infinite

Oliver Matz and Wolfgang Thomas
Institut für Informatik und Praktische Mathematik
Christian-Albrechts-Universität zu Kiel, D-24098 Kiel
email : {oma|wt}@informatik.uni-kiel.de

Abstract

We show that in monadic second-order logic over finite directed graphs, a strict hierarchy of expressiveness is obtained by increasing the (second-order) quantifier alternation depth of formulas. Thus, the "monadic analogue" of the polynomial hierarchy is found to be strict, which solves a problem of Fagin. The proof is based on automata theoretic concepts (rather than Ehrenfeucht-Fraïssé games) and starts from a restricted class of graph-like structures, namely finite two-dimensional grids. We investigate monadic second-order definable sets of grids where the width of grids is a function of the height. In this context, the infiniteness of the quantifier alternation hierarchy is witnessed by n-fold exponential functions for increasing n. It is notable that these witness sets of the monadic hierarchy all belong to the complexity class NP, the first level of the polynomial hierarchy.

1 Introduction

The subject of this paper is monadic second-order logic over graphs. In this logic, one can quantify over vertices and sets of vertices; typical properties which are formalizable in monadic logic are k-colourability and connectivity.

Monadic second-order graph properties can be classified by the alternation depth of set quantifiers which occur in defining formulas. One speaks of a Σ_n-formula (Π_n-formula) if its prenex normal form has a prefix of n alternating blocks of set quantifiers starting with an existential (universal) block, followed by a first-order kernel.

Fagin [7] raised the question whether for increasing n the monadic Σ_n-formulas allow to define more and more properties of finite graphs. This question is the "monadic analogue" of the problem whether the polynomial hierarchy (of complexity theory) is infinite; in the context of finite graphs the n-th level of the polynomial hierarchy is given by the graph properties which are definable by general Σ_n-formulas where second-order quantifiers may range also over relations of higher arity than 1. (For more background we refer the reader to [4].)

As partial results on Fagin's problem we mention his result [6] (see also [8]) that connectivity of graphs is a monadic property which is not Σ_1 (i.e., not definable by a Σ_1-formula). In [17], Schwentick extended this result to graphs with built-in order. Recent work of Ajtai, Fagin, and Stockmeyer ([1]) announces that monadic logic allows to define complete problems for each level of the polynomial hierarchy. Within the range of monadic Σ_1-formulas, Otto [15] showed that the length of the (single) block of leading existential set quantifiers induces a strict hierarchy of properties. From automata theory (cf. Büchi [2], Elgot [5], and subsequent work), it is also known that over words and trees (considered as special labelled graphs), all monadic second-order properties are in fact Σ_1-properties; in [18] it was observed that even a single existential set quantifier suffices.

In the present paper we show that over finite graphs in general, the Σ_n-formulas induce a strict hierarchy of graph properties, thus answering Fagin's question affirmatively.

Our proof method does not use Ehrenfeucht-Fraïssé games (as in the above-mentioned papers) but has a more automata theoretic and arithmetical flavour.

A first idea is to consider rectangular grids instead of graphs. (At a later point we transfer the hierarchy theorem from the domain of grids back to the domain of graphs.) Grids are, roughly speaking, finite graphs whose vertices are arranged as elements of a matrix and which have two edge relations corresponding to vertical and horizontal successors. A grid can be identified with its size (m, n), indicating that there are m rows and n columns; in this way a grid property determines a binary relation over $\mathbb{N}_{\geq 1}$. We shall consider mainly those grid properties where this relation is a function, i.e. where for every m there is exactly one grid of height m.

A second idea is to study the growth rates of functions which can be realized with Σ_n-formulas. A key lemma states that with a Σ_{2n+3}-formula one can define an n-fold exponential function, but the rate of growth of a function realizable by a Σ_n-formula can be at most n-fold exponential in m.

For the latter we use an automata theoretic view of formulas as developed in [11]; this in turn allows to apply pumping arguments. The approach is a generalization of a method due to Giammarresi [9] which argues against Σ_1-formulas as follows: Consider grids of a fixed height m; then such a formula, say with a single leading existential quantifier, will be equivalent to a nondeterministic finite automaton (NFA) which works through the given grid, column by column from left to right. A "state" of the automaton is a column vector of labels from $\{0,1\}$, and a "run" is thus an assignment of boolean values to the grid's vertices. Since for grids of fixed height m the number of states of the automaton is 2^m, the shortest grid cannot be longer than 2^m. Thus the growth rate of functions realized by such an automaton (and hence by corresponding formulas) is at most 2^m. Using the powerset and projection construction for NFAs, one proves that the set of grids satisfying a Σ_n-formula is accepted by an NFA whose number of states (on grids of height m) is n-fold exponential in m. The width of shortest accepted grids is then bounded accordingly.

The grid properties showing the infiniteness of the monadic hierarchy can be defined with very moderate use of set quantifiers when the use of transitive closure operators is allowed. Let us denote by $\Sigma_1^{TC(1)}$ the set of properties definable by a formula with a prefix of existential monadic quantifiers and a kernel formula with first-order quantifications and transitive-closure operators ranging over binary relations. $\Sigma_1^{TC(1)}$ is a subset of NP. The proofs given below show that the properties used for the hierarchy proof are all in $\Sigma_1^{TC(1)}$. Invoking the announced results of [1] we thus obtain that monadic second-order logic allows to define grid properties arbitrarily high in the polynomial hierarchy, but that the infiniteness of the monadic hierarchy can be realized already within NP. In Section 5, we show that by existential monadic formulas with transitive-closure operators applied to 4-ary relations in the kernel ($\Sigma_1^{TC(2)}$-formulas in our notation), one can realize even non-elementary functions.

As a final step, we transfer the hierarchy result from the domain of grids to the domain of graphs. From the infinity of the quantifier alternation hierarchy over grids we can conclude that over graphs the corresponding hierarchy is *strict*. As a result, the levels Σ_n, Π_n, and Δ_n (for $n = 1, 2, \ldots$) of the monadic hierarchy over graphs are related precisely as in the classical case of the arithmetical hierarchy of recursion theory. (Δ_n contains those properties which are both Σ_n and Π_n.)

Let us mention that also over the domain of grids this strictness of the hierarchy has meanwhile been established. After communication of a preliminary version of the present paper, Nicole Schweikardt [16] showed (by a refinement of the proof of Section 4 below) that a certain n-fold exponential function is definable by a Σ_n- and a Π_n-formula, rather than just by a Σ_{2n+3}-formula as presented here. Together with the results of Section 3 of this paper, this implies an even finer separation than that between Σ_n- and Σ_{n+1}-properties: The class of boolean combinations of Σ_n-properties of grids is strictly included in the class of Δ_{n+1}-properties. Moreover, this can be transferred to the domain of graphs, where this class of boolean combinations is then located strictly between the levels Δ_{n+1} and Σ_n, respectively Π_n.

We thank Yuri Gurevich for helpful and inspiring discussions on the subject of this paper in its early stage, and Nicole Schweikardt for her contributions and remarks in the late stage.

2 Notation and Main Result

2.1 Rectangles, Grids, and Pictures

By $[1..n]$ we denote, for any $n \in \mathbb{N}$, the set $\{1, \ldots, n\}$. A *rectangle* is a set of the form $[1..m] \times [1..n]$, where $m, n \geq 1$. With the words *row, column, top, bottom* etc. we refer to the usual formation of indices inside a matrix, e.g. $(1, n)$ is the rightmost position of the top row of the rectangle $[1..m] \times [1..n]$.

For any rectangle R, the *grid* $\underline{R}$ is the structure (R, S_1, S_2) with universe R, where $S_1 \subseteq R \times R$ is the vertical successor relation containing all pairs $((i,j),(i+1,j))$ of elements of R, and $S_2 \subseteq R \times R$ is the horizontal successor relation containing all pairs $((i,j),(i,j+1))$ of elements of R.

A *picture of size* (m,n) (where $m, n \geq 1$) over an alphabet Γ is a $m \times n$-matrix over Γ, i.e. a mapping $[1..m] \times [1..n] \longrightarrow \Gamma$. The set of pictures of size (m,n) over Γ is denoted by $\Gamma^{m,n}$. The set of pictures of arbitrary size over Γ is denoted by $\Gamma^{+,+}$. A subset of $\Gamma^{+,+}$ is called a *picture language over* Γ.

2.2 Monadic Second-Order Formulas

We will consider formulas over three different classes of structures, namely the class $\mathcal{R}$ of grids, the class $\mathcal{P}$ of pictures, and the class $\mathcal{G}$ of graphs. Here, we will sketch our notations for formulas over grids.

We use $\boldsymbol{u}, \boldsymbol{v}, \boldsymbol{x}, \boldsymbol{y}, \ldots$ as first-order variables and $\boldsymbol{X}, \boldsymbol{Y}, \ldots$ as monadic second-order variables, the latter officially numbered as $\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots$. MSO-*formulas* are built up as usual from atomic formulas $\boldsymbol{x} = \boldsymbol{y}$, $\boldsymbol{x} S_1 \boldsymbol{y}$, $\boldsymbol{x} S_2 \boldsymbol{y}$, and $\boldsymbol{X}(\boldsymbol{y})$ by propositional connectives and first- and second-order quantifications. FO-*formulas* are MSO-formulas in which no second-order quantifier occurs.

By $\mathrm{FO}^{\leq}$ we denote the set of formulas in which we allow also $\boldsymbol{x} \leq_1 \boldsymbol{y}$ and $\boldsymbol{x} \leq_2 \boldsymbol{y}$, which state that $\boldsymbol{x}$ is above, resp. left from $\boldsymbol{y}$ in a grid, as well as $\boldsymbol{x} \equiv_d \boldsymbol{y}$, which asserts that the sum of the two coordinates is the same for $\boldsymbol{x}$ and $\boldsymbol{y}$, i.e. $\boldsymbol{x}$ and $\boldsymbol{y}$ are on the same counter-diagonal.

By $\mathrm{FO}^{TC(k)}$ we denote the set of FO-formulas in which also the transitive closure operator ranging over $2k$-ary relations is allowed. (For a definition see e.g. [4]). Every formula in $\mathrm{FO}^{\leq}$ is equivalent to an $\mathrm{FO}^{TC(1)}$-formula. For example, $\boldsymbol{x} \leq_1 \boldsymbol{y}$ is equivalent to $TC(\boldsymbol{u}, \boldsymbol{u}', \boldsymbol{u}S_1\boldsymbol{u}')(\boldsymbol{x}, \boldsymbol{y})$.

Satisfaction of these formulas is defined as usual. For an MSO-formula $\varphi(\boldsymbol{X}_1, \dots, \boldsymbol{X}_m, \boldsymbol{x}_1, \dots, \boldsymbol{x}_n)$ with free variables $\boldsymbol{X}_1, \dots, \boldsymbol{X}_m$ and $\boldsymbol{x}_1, \dots, \boldsymbol{x}_n$, we write $\underline{R} \models \varphi[X_1, \dots, X_m, x_1, \dots, x_n]$ iff the assignment f that maps $\boldsymbol{X}_i$ to the set $X_i \subseteq R$ and $\boldsymbol{x}_j$ to the position $x_j \in R$ (for all $i \leq m$ and $j \leq n$) makes φ true.

For an MSO-sentence φ we write $Mod_{\mathcal{R}}(\varphi)$ for the set of grids that satisfy φ. Two sentences φ, ψ are called equivalent over $\mathcal{R}$ if $Mod_{\mathcal{R}}(\varphi) = Mod_{\mathcal{R}}(\psi)$. We identify a grid $\underline{[1..m]\times[1..n]}$ with the pair (m, n), thus we have for every MSO-sentence φ:

$$Mod_{\mathcal{R}}(\varphi) = \{(m, n) \mid \underline{[1..m]\times[1..n]} \models \varphi\}.$$

A relation $M \subseteq (\mathbb{N}_{\geq 1})^2$ is MSO-*definable* if there is an MSO-sentence φ such that $Mod_{\mathcal{R}}(\varphi) = M$. If an MSO-definable relation is a function, we call it MSO-definable function. Do not confuse this notion of definability with the one used in arithmetic.

Definition 1 (Quantifier alternation hierarchy)

Let Σ_0 be the set of FO-formulas. For every $n \in \mathbb{N}$, we denote by Σ_{n+1} the set of all formulas of the form $\exists \boldsymbol{X}_1 \dots \exists \boldsymbol{X}_m \neg\varphi$, where φ is in Σ_n.

By $\underline{\Sigma_n}[\mathcal{R}]$ we denote the class of sets of grids which are Σ_n-definable (i.e. defined by some formula in Σ_n). $\underline{\Pi_n}[\mathcal{R}]$ denotes the class of sets of grids whose complement is in $\underline{\Sigma_n}[\mathcal{R}]$. Furthermore, $\underline{\Delta_n}[\mathcal{R}] = \underline{\Sigma_n}[\mathcal{R}] \cap \underline{\Pi_n}[\mathcal{R}]$, and $\underline{\mathrm{B}(\Sigma_n)}[\mathcal{R}]$ denotes the class of sets of grids which are some boolean combination of sets of grids in $\underline{\Sigma_n}[\mathcal{R}]$. $\square$

By $\Sigma_1^{TC(k)}$ we denote the set of formulas of the form $\exists \boldsymbol{X}_1 \dots \exists \boldsymbol{X}_m \varphi$, where φ is in $\mathrm{FO}^{TC(k)}$.

Formulas over the class $\mathcal{G}$ of finite directed graphs are considered in the usual signature with the binary symbol E. By $\underline{\Sigma_n}[\mathcal{G}], \underline{\Pi_n}[\mathcal{G}]$ etc. we denote the classes of sets of graphs definable by corresponding formulas.

2.3 Main Result

In the Sections 3, 4, and 5 we will consider the class of grids. We will investigate the rate of growth of MSO-definable functions on integers. For this purpose we introduce, for every $n \geq 0$, the n-fold exponential functions $f_n : \mathbb{N}_{\geq 1} \to \mathbb{N}_{\geq 1}$ and $s_n : \mathbb{N}_{\geq 1} \to \mathbb{N}_{\geq 1}$ defined by

$$f_0(i) = i, \quad f_{n+1}(i) = f_n(i)2^{f_n(i)}$$
$$s_0(i) = i, \quad s_{n+1}(i) = 2^{s_n(i)}.$$

Then we have

Theorem 2 1. For every $n \in \mathbb{N}$, the function $m \mapsto f_n(m) + 1$ is $\Sigma_1^{TC(1)}$-definable and Σ_{2n+3}-definable.

2. The function $m \mapsto f_{m-1}(m) + 1$ is $\Sigma_1^{TC(2)}$-definable.

3. For every $n \in \mathbb{N}$, every Σ_n-definable function $m \mapsto f(m)$ is $s_n(\mathcal{O}(m))$. $\square$

The first claim will be proven in Section 4, the second one in Section 5, and the third one in the next section. The following is a consequence of Theorem 2 because $f_{n+1}(m)$ is not $s_n(\mathcal{O}(m))$.

Corollary 3 1. $\underline{\Sigma_n}[\mathcal{R}] \subsetneq \underline{\Sigma_{2n+5}}[\mathcal{R}]$ for every $n \in \mathbb{N}_{\geq 1}$, i.e. the monadic second-order quantifier alternation hierarchy $\underline{\Sigma_1}[\mathcal{R}], \underline{\Sigma_2}[\mathcal{R}], \dots$ over grids is infinite.

2. $\underline{\Sigma_1^{TC(1)}}[\mathcal{R}] \not\subseteq \underline{\Sigma_n}[\mathcal{R}]$ for every $n \in \mathbb{N}$.

$\square$

In Section 6 we conclude the strictness of the monadic alternation hierarchy over graphs:

Theorem 4 $\underline{\Sigma_n}[\mathcal{G}] \subsetneq \underline{\Sigma_{n+1}}[\mathcal{G}]$ for $n \geq 0$, i.e. the monadic quantifier alternation hierarchy over graphs is strict. Moreover, $\underline{\Sigma_n}[\mathcal{G}]$ and $\underline{\Pi_n}[\mathcal{G}]$ are incomparable (by inclusion) for every $n \geq 1$. $\square$

As a consequence, $\underline{\Delta_{n+1}}[\mathcal{G}]$ properly contains $\underline{\Sigma_n}[\mathcal{G}]$, $\underline{\Pi_n}[\mathcal{G}]$ and is properly contained in $\underline{\Sigma_{n+1}}[\mathcal{G}], \underline{\Pi_{n+1}}[\mathcal{G}]$.

3 Bounds on MSO-definable Functions

In this section we will prove the third claim of Theorem 2. For this aim we will develop upper bounds for the rate of growth of MSO-definable functions.

In order to argue against formulas by combinatorial arguments, we take an automata theoretic view of formulas, here by using the notion of *local domino systems (DS)*. DS are a formalism for the definition of *picture languages* (sets of pictures).

Definition 5 Let P be a picture. We denote by $\hat{P}$ the picture that results from P by surrounding it with the (fresh) boundary symbol #.

A *local domino system (DS)* over an alphabet Γ is a triple $(\Gamma, \Delta_1, \Delta_2)$, where $\Delta_1 \subseteq (\Gamma \cup \{\#\})^{2,1}$ and $\Delta_2 \subseteq (\Gamma \cup \{\#\})^{1,2}$. (So Δ_1 contains "vertical dominoes" of size 2×1 and Δ_2 contains "horizontal dominoes" of size 1×2.)

The language recognized by a local domino system $\mathcal{T} = (\Gamma, \Delta_1, \Delta_2)$ is given by the set of pictures P such that all 2×1-subblocks and all 1×2-subblocks of $\hat{P}$ are in $\Delta_1 \cup \Delta_2$.

A picture language $L \subseteq \Gamma^{+,+}$ is *domino-local* iff there is a local domino system that recognizes L. □

The above definition generalizes the definition of "local word languages". Recall that a word language is local iff it is fully characterized by the set of letters that may occur at the beginning of words, the set of letters that may occur at the end of words, and the set of infixes of length 2 that may occur inside words.

It is well known that every local word language over an alphabet A is recognized by some deterministic finite automaton (DFA) whose number of states is $1 + |A|$.

Now we introduce some notations for the special case where the alphabet is of the form $\{0,1\}^k$.

Definition 6 Let $k \geq 0$. For an $m\times n$-picture P over the alphabet $\{0,1\}^k$ and $l \in \{1,\ldots,k\}$, let X_l^P be the subset of $[1..m]\times[1..n]$ of those positions (i,j) for which the l-th component of $P(i,j)$ is 1.

Let $\varphi(\boldsymbol{X}_1,\ldots,\boldsymbol{X}_k)$ be an MSO-formula. For $P \in (\{0,1\}^k)^{+,+}$, we shall allow the notation $P \models_{\mathcal{P},k} \varphi$ for $\underline{\mathrm{dom}(P)} \models \varphi[X_1^P,\ldots,X_k^P]$. The *picture language defined by* $\varphi(\boldsymbol{X}_1,\ldots,\boldsymbol{X}_k)$ is the set $Mod_{\mathcal{P},k}(\varphi) = \{P \in (\{0,1\}^k)^{+,+} \mid P \models_{\mathcal{P},k} \varphi\}$. □

The meaning $Mod_{\mathcal{R}}(\varphi)$ of a formula φ without free variables has been defined in Section 2. It coincides with the above definition for case $k=0$ if one identifies an $m\times n$-picture over a singleton alphabet (namely $\{0,1\}^0$) with the pair (m,n) and thus with a grid.

The following fact is shown in [10], [13], and [14], using [11]. It generalizes the well-known result from word language theory that every recognizable word language is the projection of some local word language. A projection is here a mapping from one alphabet to another, and is lifted to words, respectively pictures, and to languages.

Theorem 7 Every Σ_1-definable picture language is a projection of some domino-local picture language. □

For the proof of this theorem one uses the fact ([11]) that every Σ_1-definable picture language is a projection of a *local* picture language. (A picture language is *local* iff membership of a picture P in L is completely determined by the set of 2×2-subblocks of the picture $\hat{P}$.) It is then simple to pass from local picture languages to projections of domino-local picture languages ([10, 13, 14]).

In order to state the key theorem 9, we need the following.

Definition 8 For a picture language $L \subseteq \Gamma^{+,+}$ and an integer $m \geq 1$, we denote by $L(m)$ the following word language over $\Gamma^{m,1}$:

$$L(m) = \left\{ \begin{pmatrix} a_{11} \\ \vdots \\ a_{m1} \end{pmatrix} \cdots \begin{pmatrix} a_{1n} \\ \vdots \\ a_{mn} \end{pmatrix} \;\middle|\; \begin{matrix} a_{11} \cdots a_{1n} \\ \vdots \quad\quad \vdots \\ a_{m1} \cdots a_{mn} \end{matrix} \in L \right\}$$

□

Recall also the sequence of functions (s_n), where s_0 is the identity on $\mathbb{N}_{\geq 1}$ and $s_{n+1}(m) = 2^{s_n(m)}$.

Theorem 9 For $n \geq 1$, $k \geq 0$, and every formula $\varphi(\boldsymbol{X}_1,\ldots,\boldsymbol{X}_k) \in \Sigma_n$ there exists $c \in \mathbb{N}$ such that for all $m \geq 1$ there is an NFA with $s_{n-1}(c^m)$ states that recognizes the word language $Mod_{\mathcal{P},k}(\varphi)(m)$ over $(\{0,1\}^k)^{m,1}$. □

Proof The proof is by induction over n. For the case $n=1$ one constructs an NFA with the characterization of Theorem 7: Let $k \geq 0$ and $\varphi(\boldsymbol{X}_1,\ldots,\boldsymbol{X}_k) \in \Sigma_1$. There exists a domino system $\mathcal{T} = (\Gamma, \Delta_1, \Delta_2)$ and a projection $\alpha : \Gamma \to \{0,1\}^k$ such that $\alpha(\mathcal{L}(\mathcal{T})) = Mod_{\mathcal{P},k}(\varphi)$.

Let $c := |\Gamma| + 1$. For every $m \geq 1$, consider the word language $\mathcal{L}(\mathcal{T})(m)$ over the alphabet $\Gamma^{m,1}$ (of columns of height m). It is easy to see that, for every $m \geq 1$, the word language $\mathcal{L}(\mathcal{T})(m)$ is local and therefore recognized by some DFA with $|\Gamma^{m,1}|+1$ states. Since $\alpha(\mathcal{L}(\mathcal{T})(m)) = Mod_{\mathcal{P},k}(\varphi)(m)$ for every m, the well-known construction for projections of word languages gives, for every m, an NFA with $|\Gamma^{m,1}| + 1 \leq c^m$ states that recognizes $Mod_{\mathcal{P},k}(\varphi)(m)$. This completes the induction basis.

For the induction step, assume we have proven the claimed implication for some n. Let $\varphi(\boldsymbol{X}_1,\ldots,\boldsymbol{X}_k)$ be a Σ_{n+1}-formula. There is some $l \in \mathbb{N}$ and a Σ_n-formula $\psi(\boldsymbol{X}_1,\ldots,\boldsymbol{X}_{k+l})$ such that φ equals $\exists \boldsymbol{X}_{k+1}\ldots\exists \boldsymbol{X}_{k+l}\neg\psi$, i.e. $Mod_{\mathcal{P},k}(\varphi)$ is a projection of the complement of the picture language $Mod_{\mathcal{P},k+l}(\psi)$. So for every $m \geq 1$, the word language $Mod_{\mathcal{P},k}(\varphi)(m)$ is a projection of the complement of the word language $Mod_{\mathcal{P},k+l}(\psi)(m)$.

By induction hypothesis there is a $c \in \mathbb{N}$ such that for every $m \geq 1$ there is an NFA with $s_{n-1}(c^m)$ states that recognizes the word language $Mod_{\mathcal{P},k+l}(\psi)(m)$. Hence (by the well-known constructions for complementation and projection) the word language $Mod_{\mathcal{P},k}(\varphi)(m)$ is recognized by an NFA with $2^{s_{n-1}(c^m)} = s_n(c^m)$ states.

This completes the induction step. □

Now we can apply this theorem to sentences in order to show the third claim of Theorem 2, namely that every Σ_n-definable function $m \mapsto f(m)$ is $s_n(\mathcal{O}(m))$.

Proof (third claim of Theorem 2): Let $n \geq 1$ and φ be a Σ_n-sentence that defines a function $f : \mathbb{N}_{\geq 1} \to \mathbb{N}_{\geq 1}$. By Theorem 9 there is a c such that for all $m \geq 1$ there is an NFA with $s_{n-1}(c^m)$ states and singleton alphabet that accepts only the word of length $f(m)$. This implies that $f(m) \leq s_{n-1}(c^m) = s_n(m \log_2 c)$ for all m. $\square$

The preceding fact has been shown in [9, 10] for the special case $n = 1$ using essentially the same argument.

In the remainder of this section we strengthen this result to cover the functions which are definable by boolean combinations of Σ_n-sentences, rather than just Σ_n-sentences. This becomes relevant in the presence of the result of N. Schweikardt (see the Introduction) that some n-fold exponential function is definable both by a Σ_n- and a Π_n-formula: With Corollary 14 below it follows that $\mathrm{B}(\Sigma_n)[\mathcal{R}] \subsetneq \Delta_{n+1}[\mathcal{R}]$.

We need some preparations on word languages over a singleton alphabet. For a one-letter alphabet a word is identified with its length and hence a word language is identified with a subset of $\mathbb{N}$. Consequently, if $R \subseteq \mathbb{N} \times \mathbb{N}$ and $m \in \mathbb{N}$, then we let $R(m) = \{n \in \mathbb{N} \mid (m, n) \in R\}$.

Definition 10 Let $k \in \mathbb{N}$. A set N of integers is *periodic at* k (short: k*-periodic*) iff there is an integer $p \in \mathbb{N}$ such that $\forall l \geq k(l \in N \leftrightarrow l + p \in N)$. $\square$

We remark that every boolean combination of k-periodic sets of integers is also k-periodic. (As period p take the least common multiple of the periods of the given sets.) It is well-known that the regular languages over a one-letter alphabet are exactly those sets of integers that are k-periodic for some k.

The following Theorem of Chrobak states that any NFA with n states over a singleton alphabet can be transformed into an equivalent DFA whose transition structure is a loop together with a path of length $\leq (n + 2)^2$ from the initial state to this loop.

Theorem 11 (see [3, Theorem 4.4]) Let $N \subseteq \mathbb{N}$ be recognizable by some n-state NFA. Then there are some $k \leq (n + 2)^2$ and an integer p such that N is recognized by a DFA $\mathcal{A}$ with states $0, \ldots, (k - 1) + p$ such that $\mathcal{A}$ reaches the state $k + ((l - k) \mod p)$ after reading an input of length $l \geq k$. $\square$

Corollary 12 Let $N \subseteq \mathbb{N}$ be recognizable by some n-state NFA. Then N is $(n + 2)^2$-periodic.

Proof Let $\mathcal{A}$, k, p be as in the above theorem. For every $l \geq k$, the automaton reaches the same state on inputs of length either l or $l + p$. Thus $l \in N \iff l + p \in N$ for every $l \geq k$. $\square$

Now we apply Theorem 9 to sentences and hence to singleton alphabets.

Theorem 13 Let $n \geq 1$ and φ be a boolean combination of Σ_n-sentences. There is a $c \in \mathbb{N}$ such that for every $m \geq 1$ the set $\mathit{Mod}_{\mathcal{R}}(\varphi)(m) = \{n \mid \underline{[1..m] \times [1..n]} \models \varphi\}$ is $s_{n-1}(c^m)$-periodic. $\square$

Proof Assume that φ is a boolean combination of the Σ_n-sentences $\varphi_1, \ldots, \varphi_r$. By the Theorem 9, there are $c_1, \ldots, c_r$ such that for every $m \geq 1$ and every $i \leq r$, the set $\mathit{Mod}_{\mathcal{R}}(\varphi_i)(m) = \{n \mid \underline{[1..m] \times [1..n]} \models \varphi_i\}$ is recognized by some NFA with $s_{n-1}(c_i^m)$ states. Let $c = \max\{(c_1 + 2)^2, \ldots, (c_r + 2)^2\}$.

Since $(s_{n-1}(c_i^m) + 2)^2 \leq s_{n-1}((c_i + 2)^{2m}) \leq s_{n-1}(c^m)$ for every $m \geq 1$ and every $i \leq r$, Corollary 12 implies that $\mathit{Mod}_{\mathcal{R}}(\varphi_i)(m)$ is $s_{n-1}(c^m)$-periodic for every i. For every $m \geq 1$, the set $\mathit{Mod}_{\mathcal{R}}(\varphi)(m)$ is a boolean combination of the sets $\mathit{Mod}_{\mathcal{R}}(\varphi_1)(m), \ldots, \mathit{Mod}_{\mathcal{R}}(\varphi_r)(m)$ and hence also $s_{n-1}(c^m)$-periodic. $\square$

Now we get the third claim of Theorem 2, namely:

Corollary 14 Every function $m \mapsto f(m)$ that is definable by a boolean combination of Σ_n-sentences is $s_n(\mathcal{O}(m))$.

Proof Let f be definable by a boolean combination of Σ_n-sentences. By the above theorem, there is a c such that for all m, the singleton $\{f(m)\}$ is $s_{n-1}(c^m)$-periodic. This implies that $f(m) \leq s_{n-1}(c^m) = s_n(m \log_2 c)$ for all m. $\square$

4 Fast Growing MSO-Definable Functions

In this section we show how to define, for each $k \geq 0$, the k-fold exponential function $m \mapsto f_k(m) + 1$.

The definitions are obtained inductively. The idea for the definition of $m \mapsto f_{k+1}(m) + 1$ is to describe, over a grid of height m, a counting process from 0 up to $2^{f_k(m)} - 1$. We shall generate the binary notation of these numbers, each of length $f_k(m)$, on the top row of the grid, and at the end add one further bit. Such a 0-1-sequence defines a subset of (the top row) of the grid, denoted Y_k in the sequel. Its length is $f_k(m) \cdot 2^{f_k(m)} + 1 = f_{k+1}(m) + 1$, as desired. An auxiliary subset X_k of the top row is given by the 0-1-sequence which fixes the beginnings of the binary numbers, i.e. which is built up by concatenating copies of the word $10^{f_k(m)-1}$. For describing the counting process we also introduce formulas which allow to compare two number representations starting at two points x and y: The formulas $\sigma_k^{=}$, σ_k^{+1}, $\sigma_k^{1,0}$ will say, respectively, that the two numbers are equal, are in succession, or ihave the representations $1 \ldots 1$ and $0 \ldots 0$ of length $f_k(m)$.

Let us describe this more precisely.

Definition 15 For every $j, m \in \mathbb{N}$ such that $j < 2^m$, we denote by $\mathrm{BIN}(j, m)$ the string in $\{0, 1\}^m$ that is the binary representation of the number j (the lowest bit coming last).

For a picture P over $\{0,1\}$ and $f : \mathbb{N}_{\geq 1} \to \mathbb{N}_{\geq 1}$ we say that *row i of P is f-well-numbered* iff the i-th row of P is in $(\mathrm{BIN}(0, f(m)) \ldots \mathrm{BIN}(2^{f(m)}-1, f(m)))^*\{0,1\}$. $\square$

Note that if some row of an $m\times n$-picture is f-well-numbered, then n is congruent to 1 modulo $f(m)2^{f(m)}$.

If X is a subset of some rectangle $[1..m]\times[1..n]$, then the *characteristic $m\times n$-picture of X* is the $m\times n$-picture P over $\{0,1\}$ that has 1's exactly at the positions in X. If $x = (i,j) \in [1..m]\times[1..n]$ and $0 \leq l \leq n-j$, then we write $X[x, x+l)$ for the string of length l that is on the l positions right from x (inclusively) in the characteristic picture of X.

A formula which defines $m \mapsto f_k(m)+1$ has to assert for a grid R that the universe is a rectangle of size $(m, f_k(m)+1)$. This is expressed using the condition that there is a set Y_{k-1} such that the top row of the characteristic picture of Y_{k-1} is f_{k-1}-well-numbered.

In the following lemma, formulas are given which describe this property (inductively over k).

Lemma 16 For $k \in \mathbb{N}$ there are $\mathrm{FO}^{\leq}$-formulas $\varphi_k(\boldsymbol{X}_0, \ldots, \boldsymbol{X}_k, \boldsymbol{Y}_0, \ldots, \boldsymbol{Y}_{k-1})$, $\sigma_k^{=}(\boldsymbol{X}_0, \ldots, \boldsymbol{X}_k, \boldsymbol{Y}_0, \ldots, \boldsymbol{Y}_k, \boldsymbol{x}, \boldsymbol{y})$, $\sigma_k^{+1}(\boldsymbol{X}_0, \ldots, \boldsymbol{X}_k, \boldsymbol{Y}_0, \ldots, \boldsymbol{Y}_k, \boldsymbol{x}, \boldsymbol{y})$, and $\sigma_k^{1,0}(\boldsymbol{X}_0, \ldots, \boldsymbol{X}_k, \boldsymbol{Y}_0, \ldots, \boldsymbol{Y}_k, \boldsymbol{x}, \boldsymbol{y})$ such that for every m and every grid R of height m and all subsets $X_0, \ldots, X_k$, $Y_0, \ldots, Y_k$ of its top row the following holds:

$$R \models \bigwedge_{i=0}^{k} \varphi_i[X_0, \ldots, X_i, Y_0, \ldots, Y_{i-1}]$$

$\iff$ (*) The top row of the characteristic picture of X_i is in $(10^{f_i(m)-1})^*1$ for every $i \leq k$; and the top row of the characteristic picture of Y_i is f_i-well-numbered for every $i < k$;

and if R is a rectangle of height m such that (*) holds and a (resp. b) are the numbers represented binarily on the $f_k(m)$ positions right from x (resp. y) in the characteristic picture of Y_k, then the formulas $\sigma_k^{=}$, σ_k^{+1}, and $\sigma_k^{1,0}$ assert that $a = b$, resp. $a+1 = b$, resp. $(a,b) = (2^{f_k(m)}-1, 0)$.

Besides, φ_k, $\sigma_k^{=}$, σ_k^{+1}, and $\sigma_k^{1,0}$ are equivalent to Σ_{2k+3}-formulas for every k. $\square$

Proof We only give the proof idea. For $n = 0$ one uses the following formula, written in shorthand, to fix corresponding positions $\boldsymbol{u}$, $\boldsymbol{w}$ in two m-digit number representations which are starting at $\boldsymbol{x}$ and $\boldsymbol{y}$ on the top row of a grid of height m.

$$\exists \boldsymbol{x}_1 \exists \boldsymbol{y}_1 (\boldsymbol{x}_1 \geq_1 \boldsymbol{x} \leq_2 \boldsymbol{u} \equiv_d \boldsymbol{x}_1 \leq_2 \boldsymbol{y}_1 \equiv_d \boldsymbol{w} \geq_2 \boldsymbol{y} \leq_1 \boldsymbol{y}_1).$$

For example, the formula $\sigma_0^{+1}(\boldsymbol{X}_0, \boldsymbol{Y}_0, \boldsymbol{x}, \boldsymbol{y})$ compares $Y_0[x, x+m)$ and $Y_0[y, y+m)$ by asserting that there is one pivot bit at corresponding positions that is zero in the first sequence and one in the latter, and that corresponding bits before this pivot are equal, and that the bits after the pivot are ones in the first sequence and zeros in the latter.

In the induction step from k to $k+1$, the numbers derived from the Y_k-colouring serve as addresses of sections of length $f_k(m)$ in the characteristic picture of Y_{k+1}.

For example, the formula φ_{k+1} asserts that the top-left and the top-right position are in X_{k+1} and that the following holds: For every two consecutive X_{k+1}-positions x, x' in the top row, x, x' are also in X_k and the Y_k-colourings of $f_k(m)$-sequences of every two consecutive X_k-positions between x, x' are binary number representations of consecutive integers (here the formula σ_k^{+1} from the induction hypothesis is needed), and the first (resp. last) of these sequences is $0^{f_k(m)}$ (resp. $1^{f_k(m)}$).

The formula σ_{k+1}^{+1} that is to compare $Y_{k+1}[x, x+f_{k+1}(m))$ with $Y_{k+1}[y, y+f_{k+1}(m))$ works similar to the formula σ_0^{+1} from the induction basis, but σ_{k+1}^{+1} compares *corresponding sequences of length $f_k(m)$* instead of *corresponding bits*. Now corresponding $f_k(m)$-sequences can be determined by their equal Y_k-colouring, using the formula $\sigma_k^{=}$ from the induction hypothesis.

This completes our sketch of the proof of Lemma 16. $\square$

With this lemma, the $\Sigma_1^{TC(1)}$-definability and the Σ_{2k+3}-definability of every function f_k (i.e. the first claim of Theorem 2) is immediate: Namely, for every $k \in \mathbb{N}$ the formula

$$\exists \boldsymbol{X}_0 \ldots \exists \boldsymbol{X}_k \exists \boldsymbol{Y}_0 \ldots \exists \boldsymbol{Y}_{k-1} \Big(\Big(\textstyle\bigwedge_{i=0}^{k} \varphi_i(\boldsymbol{X}_0, \ldots, \boldsymbol{X}_i, \boldsymbol{Y}_0, \ldots, \boldsymbol{Y}_{i-1}) \Big) \wedge \psi(\boldsymbol{X}_k) \Big)$$

is equivalent to a $\Sigma_1^{TC(1)}$-formula (and to a Σ_{2k+3}-formula) that defines f_k, provided $\psi(\boldsymbol{X}_k)$ is an FO-formula that asserts that in the top row exactly the leftmost and rightmost position are in X_k.

5 Beyond Monadic Second-Order Logic

For the proof of the second claim of Theorem 2, i.e. that the function $m \mapsto f_{m-1}(m)+1$ is $\Sigma_1^{TC(2)}$-definable, we need the following lemma.

Lemma 17 There is an $\mathrm{FO}^{TC(2)}$-formula $\varphi(\boldsymbol{X}, \boldsymbol{Y})$ such that for all grids R of height m and all sets $X, Y \subseteq R$ we have that $R \models \varphi[X, Y]$ holds iff the i-th row of the characteristic picture of Y is f_{i-1}-well-numbered for every $i \in \{1, \ldots, m-1\}$ and row i of the characteristic picture of X is in $(10^{f_{i-1}(m)-1})^*1$ for every $i \in \{1, \ldots, m\}$. $\square$

Proof With the binary *TC*-operator one can easily obtain an $\mathrm{FO}^{TC(2)}$-formula $\sigma(\boldsymbol{Y}, \boldsymbol{x}, \boldsymbol{y})$ that asserts that the number represented binarily in the characteristic picture of Y on the positions between x (inclusively) and y (exclusively) is one less than the number represented binarily by the equally long sequence of positions right from y (inclusively).

Using this formula, one can find an $\mathrm{FO}^{TC(2)}$-formula $\pi(\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{w}, \boldsymbol{w}')$ that asserts, for given positions w and w' in some row i such that in the characteristic picture of X the row section from w to w' (inclusively) is in $(10^{f_{i-1}(m)-1})^*1$, that this section of the characteristic picture of Y is $\mathrm{BIN}(0, f_{i-1}(m)) \ldots \mathrm{BIN}(2^{f_{i-1}(m)} - 1, f_{i-1}(m))\{0,1\}$ (and hence of length $f_{i-1}(m)2^{f_{i-1}(m)} + 1 = f_i(m) + 1$). For this aim, π asserts that $\sigma(\boldsymbol{Y}, \boldsymbol{x}, \boldsymbol{y})$ holds for all consecutive X-positions x, y between w and w', and that in the characteristic picture of Y the sequence between w and the next X-position is in 0^* and the sequence between w' and the preceding X-position is in 1^*.

The next step is to define a formula $\varphi'(\boldsymbol{X}, \boldsymbol{Y})$ that asserts that for every $2 \leq i \leq m$ (where m is the height of the grid) the following holds: The ith row of the characteristic picture of Y is f_{i-1}-well-numbered and the $(i+1)$-th row of the characteristic picture of X is in $(10^{f_i(m)-1})^*1$, provided that row i of the characteristic picture of X is in $(10^{f_{i-1}(m)-1})^*1$. For this aim, φ' asserts that for any two (horizontally) consecutive X-positions (except for those in the top row) there are vertical predecessors w, w' such that $\pi(X, Y, w, w')$ holds.

Recall the formula $\varphi_0(\boldsymbol{X}, \boldsymbol{Y})$, which was defined in the last section and asserts that row 1 of the characteristic picture of X is in $1(0^{f_0(m)-1})^*1$. Let $\varphi(\boldsymbol{X}, \boldsymbol{Y}) := \varphi_0(\boldsymbol{X}, \boldsymbol{Y}) \wedge \varphi'(\boldsymbol{X}, \boldsymbol{Y})$. A simple inductive argument shows that $\varphi(\boldsymbol{X}, \boldsymbol{Y})$ has indeed the desired property, which completes the proof of Lemma 17. $\square$

With Lemma 17, the second claim of Theorem 2 is clear because the formula $\exists \boldsymbol{Y} \exists \boldsymbol{X} (\varphi(\boldsymbol{Y}, \boldsymbol{X}) \wedge \psi(\boldsymbol{X}))$ defines $m \mapsto f_{m-1}(m) + 1$, where ψ is an FO-formula that asserts that the bottom row of the characteristic picture of X is in 10^*1.

6 From Grids to Pictures and Graphs

In this section we pass to the domains of pictures and graphs, over which the monadic second-order hierarchy is shown to be strict. The key ideas are the following: Over pictures, equality between the levels n and $n+1$ of the monadic hierarchy would extend to all higher levels, contradicting the infiniteness of the hierarchy over grids. Furthermore, pictures can be encoded by graphs in such a way that every Σ_n-formula that talks about pictures can be transformed into a Σ_n-sentence that says the same about those graphs that are encodings of pictures.

In addition to the classes $\mathcal{R}$ and $\mathcal{G}$ for grids and graphs, respectively, we consider for every $k \geq 0$ the class $\mathcal{P}^k = (\{0,1\}^k)^{+,+}$ of pictures over $\{0,1\}^k$, and the class $\mathcal{P} = \bigcup_{k \geq 0} \mathcal{P}^k$. As mentioned in Section 3 we identify $\mathcal{P}^0$ with $\mathcal{R}$.

For defining sets of grids and graphs we use sentences in the signatures $\tau_{\mathcal{R}} := \{S_1, S_2\}$ and $\tau_{\mathcal{G}} := \{E\}$, respectively. For the description of sets of pictures over $\{0,1\}^k$ we adopt the conventions of Definition 6 and use formulas $\varphi(\boldsymbol{X}_1, \ldots, \boldsymbol{X}_k)$ of signature $\tau_{\mathcal{R}}$. Recall that for $P \in \mathcal{P}^k$ and $\varphi = \varphi(\boldsymbol{X}_1, \ldots, \boldsymbol{X}_k)$ we write the satisfaction relation as $P \models_{\mathcal{P},k} \varphi$ (or $P \in Mod_{\mathcal{P},k}(\varphi)$).

We use $\underline{\Sigma_n}[\mathcal{P}^k]$ and similar adaptations of the notions of Definition 1 with their obvious meaning.

We can conclude the following from the results of the previous section:

Theorem 18 For every $n \geq 0$ there is a k such that $\underline{\Sigma_n}[\mathcal{P}^k] \subsetneq \underline{\Sigma_{n+1}}[\mathcal{P}^k]$.

For every $n \geq 1$ there is a k such that $\underline{\Sigma_n}[\mathcal{P}^k]$ and $\underline{\Pi_n}[\mathcal{P}^k]$ are incomparable. $\square$

Proof Suppose the converse of the first claim, i.e. there is an n such that $\underline{\Sigma_n}[\mathcal{P}^k] = \underline{\Sigma_{n+1}}[\mathcal{P}^k]$ for every k. This implies that every Σ_{n+1}-formula over $\tau_{\mathcal{R}}$ is equivalent to a Σ_n-formula over $\tau_{\mathcal{R}}$.

Every Σ_{n+2}-sentence φ over $\tau_{\mathcal{R}}$ can be written in the form $\exists \boldsymbol{X}_1 \ldots \exists \boldsymbol{X}_k \neg\psi$ for a Σ_{n+1}-formula ψ. The latter is equivalent (over $\mathcal{P}^k$) to a Σ_n-formula ψ', so that φ is equivalent (over $\mathcal{R}$) to the Σ_{n+1}-sentence $\exists \boldsymbol{X}_1 \ldots \exists \boldsymbol{X}_k \neg\psi'$.

This shows that $\underline{\Sigma_{n+2}}[\mathcal{R}] = \underline{\Sigma_{n+1}}[\mathcal{R}]$. By repeating this argument one shows that $\underline{\Sigma_{n+i}}[\mathcal{R}] = \underline{\Sigma_n}[\mathcal{R}]$ for every i, which means that the monadic hierarchy collapses at level n. But this contradicts Theorem 2.

The second claim is shown similarly. $\square$

The abovementioned recent result of Nicole Schweikardt ([16]) shows that the first claim of Theorem 18 remains true when one fixes the value of k to zero, or, in other words, when one considers sentences over the domain of grids rather than formulas with an unbounded number of free variables.

The incomparability of $\underline{\Sigma_1}[\mathcal{P}^1]$ and $\underline{\Pi_1}[\mathcal{P}^1]$ has been shown in [11].

We wish to transfer Theorem 18 to graphs. Here is the key lemma.

Lemma 19 Let $k \geq 0$, and $\Gamma = \{0,1\}^k$. There is a mapping $\alpha : \Gamma^{+,+} \longrightarrow \mathcal{G}$ such that for every n and every Σ_n-formula $\varphi_1(\boldsymbol{X}_1, \ldots \boldsymbol{X}_k)$ over the signature $\tau_{\mathcal{R}}$ there is a Σ_n-sentence φ_2 over the signature $\tau_{\mathcal{G}}$ and vice versa, such that

$$P \models_{\mathcal{P},k} \varphi_1 \iff \alpha(P) \models \varphi_2 \quad \text{for all } P \in \Gamma^{+,+}, \qquad (1)$$

and thus $Mod_{\mathcal{P},k}(\varphi_1) = \alpha^{-1}(Mod_{\mathcal{G}}(\varphi_2))$. $\square$

Proof Let α be a mapping that maps a $m \times n$-picture P over Γ to the graph graph (V, E), where the components are defined as follows: $V = (\mathrm{dom}P \times \{-1, 0, \ldots, k\})$.

E contains

1. for every $v \in \mathrm{dom}P$ the edges $((v,-1),(v,-1))$, $((v,0),(v,-1))$;
2. for every $v \in \mathrm{dom}P$ and every $l \in [0..k]$ the edges $((v,l-1),(v,l))$;
3. for every $v \in \mathrm{dom}P$, $l \in [1..k]$ with $v \in X_l^P$ the edge $((v,l),(v,l))$.
4. for every $(u,v) \in S_1$ the edge $((u,-1),(v,-1))$;
5. for every $(u,v) \in S_2$ the edge $((u,0),(v,-1))$.

Each vertex set $\{v\} \times \{-1, \ldots, k\}$ forms together with the first three types of vertices a "gadget" whose structure allows to extract the "colour" $P(v)$ of v. There are FO-formulas that allow to fix and distinguish the nodes of one such gadget, i.e. there is for every $l \in \{-1, \ldots, k\}$ an FO-formula $\psi_l(\boldsymbol{x})$ over $\tau_{\mathcal{G}}$ that asserts for every picture P that the node x of $\alpha(P)$ is in $\mathrm{dom}P \times \{l\}$.

Now it is quite simple to translate a Σ_n-formula φ_1 into a Σ_n-sentence φ_2 such that the assertion (1) holds.

Conversely, when translating a Σ_n-sentence φ_2 over $\tau_{\mathcal{G}}$ into a Σ_n-formula $\varphi_1(\boldsymbol{X}_1, \ldots, \boldsymbol{X}_k)$ over $\tau_{\mathcal{R}}$, one encounters the problem that in the α-image of a picture P there are $(k+2)$-times as many nodes than in the universe of P. Therefore one has to introduce second-order variables $\boldsymbol{Y}_{-1}, \ldots \boldsymbol{Y}_k$ for every free second-order variable $\boldsymbol{Y}$ in φ_1. Here Y_l (for $-1 \leq l \leq k$) is intended to contain those positions v in P for which (v,l) is in Y, where Y is some subset of the vertex set of $\alpha(P)$.

For the translation of the first-order kernel formula we need the following auxiliary definition: For every picture P and every mapping κ from the set of free variables of a formula ψ over $\tau_{\mathcal{G}}$ into the set $\{-1, \ldots, k\}$, we say that an assignment f into $\alpha(P)$ is *of type* κ iff $f(\boldsymbol{x}) \in \mathrm{dom}P \times \{\kappa(\boldsymbol{x})\}$ for every free variable $\boldsymbol{x}$ in ψ. The translation of $\tau_{\mathcal{G}}$-formulas with free first-order variables depends on this κ and works only for assignments of type κ.

Thus when we translate an existential first-order quantification inside the FO-kernel of φ_2 with respect to a given type κ, we have to proceed as follows: A subformula of the form $\exists \boldsymbol{y} \sigma_2$ becomes

$$\bigvee_{\kappa'} (\sigma_1^{\kappa'} \wedge \psi_{\kappa'(\boldsymbol{y})}(\boldsymbol{y})),$$

where the disjunction ranges over all κ' that extend the mapping κ (defined on the set of free variables of $\exists \boldsymbol{y} \sigma_2$) to the additional free variable $\boldsymbol{y}$, and $\sigma_1^{\kappa'}$ is the translation of σ_2 with respect to κ' for each of these κ'.

Atomic subformulas of φ_2 such as $\boldsymbol{x} E \boldsymbol{y}$ and $\boldsymbol{Y}(\boldsymbol{x})$ can be translated with respect to κ, depending on the values $\kappa(\boldsymbol{x})$ and $\kappa(\boldsymbol{y})$.

The details are left to the reader. $\square$

Now we turn to the proof of Theorem 4, i.e. the claim $\underline{\Sigma_n}[\mathcal{G}] \subsetneq \underline{\Sigma_{n+1}}[\mathcal{G}]$ (for $n \geq 0$) and the incomparability of $\underline{\Sigma_n}[\mathcal{G}]$ and $\underline{\Pi_n}[\mathcal{G}]$ for $n \geq 1$.

Proof of Theorem 4: We start with the first claim. Let $n \geq 0$. By Theorem 18 there is some k and some Σ_{n+1}-formula $\varphi(\boldsymbol{X}_1, \ldots, \boldsymbol{X}_n)$ over $\tau_{\mathcal{R}}$ that is not equivalent (over $\mathcal{P}^k$) to any Σ_n-formula. Now choose a Σ_{n+1}-sentence φ' over $\tau_{\mathcal{G}}$ according to Lemma 19, i.e. $\mathit{Mod}_{\mathcal{P},k}(\varphi) = \alpha^{-1}(\mathit{Mod}_{\mathcal{G}}(\varphi'))$.

If φ' was equivalent to some Σ_n-sentence ψ' over $\tau_{\mathcal{G}}$, then (again by Lemma 19) there would be a Σ_n-formula $\psi(\boldsymbol{X}_1, \ldots, \boldsymbol{X}_k)$ such that $\mathit{Mod}_{\mathcal{P},k}(\psi) = \alpha^{-1}(\mathit{Mod}_{\mathcal{G}}(\psi')) = \alpha^{-1}(\mathit{Mod}_{\mathcal{G}}(\varphi')) = \mathit{Mod}_{\mathcal{P},k}(\varphi)$, i.e. ψ would be equivalent to φ, which contradicts the choice of φ. Thus we have shown that $\mathit{Mod}_{\mathcal{G}}(\varphi')$ is in $\underline{\Sigma_{n+1}}[\mathcal{G}]$ but not in $\underline{\Sigma_n}[\mathcal{G}]$.

The incomparability of $\underline{\Sigma_n}[\mathcal{G}]$ and $\underline{\Pi_n}[\mathcal{G}]$ for $n \geq 1$ is shown similarly. $\square$

7 Conclusion

We have shown that the monadic quantifier alternation hierarchy over grids is infinite and have concluded the strictness of this hierarchy over the domain of pictures (coloured grids) and over the domain of graphs. Over grids, the hierarchy was established by considering, for each $n \geq 1$, the set of grids of size $(m, f_n(m)+1)$, where f_n is a certain n-fold exponential function.

Over the domain $\mathcal{G}$ of graphs, the structure of the monadic quantifier alternation hierarchy is summarized in Figure 1.

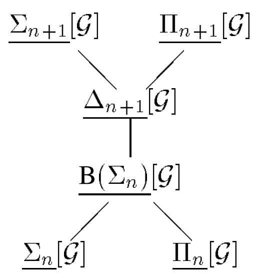

Figure 1. The Monadic Hierarchy over Graphs

Here lines indicate proper inclusions. As explained before (in the Introduction and Section 3), the central proper inclusion between $\underline{\Delta_{n+1}}[\mathcal{G}]$ and $\underline{\mathrm{B}(\Sigma_n)}[\mathcal{G}]$ relies on the results of [16] (together with Section 3 above), while the others are shown in this paper. The inclusions displayed in the figure hold as well over the domain of pictures (using alphabets of unbounded finite size).

In the hierarchy proof over the class of grids, the witness sets can be defined even with existential MSO-formulas when the kernel is either first-order in the signature $\{\leq_1, \leq_2, \equiv_d\}$ or (equivalently) in the extension of first-order logic over $\{S_1, S_2\}$ by the unary transitive closure operator (see Lemma 16). The corresponding sets of grids are projections of $\mathrm{FO}^{TC(1)}$-definable sets. Invoking [12], this coincides with projections of grid sets in NLOG. Hence the examples showing the infiniteness of the monadic hiereracby are all in NP and thus in the lowest level of the polynomial hierarchy. We have even given an example of a set of grids which is $\Sigma_1^{TC(2)}$-definable and hence in NP, but not MSO-definable. This indicates that the monadic hierarchy is essentially different from the polynomial hierarchy, and showing the infiniteness of the monadic hierarchy does not mean any progress in answering the question whether the polynomial hierarchy is also infinite.

Let us close with some (more accessible) open questions. Whereas over graphs and pictures all proper inclusions have been verified following the pattern of the classical hierarchies (of recursion theory), some separation claims are still unsettled over grids. We know the following inclusion chain: $\underline{\Sigma_n}[\mathcal{R}] \subseteq \underline{\mathrm{B}(\Sigma_n)}[\mathcal{R}] \subsetneq \underline{\Delta_{n+1}}[\mathcal{R}] \subseteq \underline{\Sigma_{n+1}}[\mathcal{R}]$, where the strictness of the first and last inclusion (i.e. the closure of $\underline{\Sigma_n}[\mathcal{R}]$ under complement) is open. By the results of Section 3 (Corollary 14) together with Theorem 2 (Claim 3), these classes do not differ with respect to the growth rate of functions which can be defined by respective formulas.

Finally, one can try to sharpen the hierarchy proof by allowing stronger kernels in the Σ_n-formulas. It is open whether the monadic alternation hierarchy over grids is infinite when these kernels are first-order formulas in the signature $\{\leq_1, \leq_2\}$ (or even $\{\leq_1, \leq_2, \equiv_d\}$), or when the transitive closure operator is adjoined to first-order logic.

References

[1] M. Ajtai, R. Fagin, and L. Stockmeyer. private communication, 1996.

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

[3] M. Chrobak. Finite automata and unary languages. *Theoretical Computer Science*, 47:149–158, 1986.

[4] H. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer-Verlag, New York, 1995.

[5] C. Elgot. Decision problems of finite automata design and related arithmetics. *Trans. Amer. Math. Soc.*, 98:21–52, 1961.

[6] R. Fagin. Monadic generalized spectra. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 21:89–96, 1975.

[7] R. Fagin. Problem proposed in Oberwolfach Meeting 1995; published in: Collection of open problems in finite model theory, see `http://www.informatik.rwth-aachen.de/WWW-math/FMT.html`, 1995.

[8] R. Fagin, L. Stockmeyer, and M. Vardi. On monadic NP vs. monadic co-NP. *Information and Computation*, 120:78–92, 1995.

[9] D. Giammarresi. Two-dimensional languages and recognizable functions. In G. Rozenberg and A. Salomaa, editors, *Developments in Language Theory, Proceedings of the Conference, Turku (Finnland) '93*, pages 290–301. world scientific, Singapore, 1994.

[10] D. Giammarresi and A. Restivo. Two-dimensional languages. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Language Theory*, volume III. Springer-Verlag, New York, 1996.

[11] D. Giammarresi, A. Restivo, S. Seibert, and W. Thomas. Monadic second-order logic and recognizability by tiling systems. *Information and Computation*, 125:32–45, 1996.

[12] N. Immerman. Languages that capture complexity classes. *SIAM J. Comput.*, 16(4), 1988.

[13] M. Latteux and D. Simplot. Recognizable picture languages and domino tiling. Internal Report IT-94-264, Laboratoire d'Informatique Fondamentale de Lille, Université de Lille, France, 1994.

[14] O. Matz. *Klassifizierung von Bildsprachen mit rationalen Ausdrücken, Grammatiken und Logik-Formeln*. Diploma thesis, Christian-Albrechts-Universität Kiel, 1995. (German).

[15] M. Otto. Note on the number of monadic quantifiers in monadic Σ_1^1. *Inf. Process. Lett.*, 53:337–339, 1995.

[16] N. Schweikardt. Diploma thesis, Universität Mainz. (in preparation).

[17] T. Schwentick. Graph connectivity and monadic NP. In *Proceedings of the 35th IEEE Symposium on Foundations of Computer Science*, pages 614–622, 1994.

[18] W. Thomas. Classifying regular events in symbolic logic. *J. Comput. Syst. Sci.*, 25:360–376, 1982.