An Improved Homomorphism Preservation Theorem from Lower Bounds in Circuit Complexity

Benjamin Rossman, University of Toronto

Previous work of the author [Rossman 2008a] showed that the Homomorphism Preservation Theorem of classical model theory remains valid when its statement is restricted to finite structures. In this paper, we give a new proof of this result via a reduction to lower bounds in circuit complexity, specifically on the AC^0 formula size of the colored subgraph isomorphism problem. Formally, we show the following: if a first-order sentence Φ of quantifier-rank k is preserved under homomorphisms on finite structures, then it is equivalent on finite structures to an existential-positive sentence Ψ of quantifier-rank $k^{O(1)}$. Quantitatively, this improves the result of [Rossman 2008a], where the upper bound on the quantifier-rank of Ψ is a non-elementary function of k.

This column is an abridged version of a forthcoming paper.

1. INTRODUCTION

Preservation theorems are a family of results in classical model theory that equate semantic and syntactic properties of first-order formulas. A prominent example — and the subject of this paper — is the Homomorphism Preservation Theorem, which states that a first-order sentence is preserved under homomorphisms if and only if it is equivalent to an existential-positive sentence. (Definitions for the various terms in this theorem are given in Section 3.) Two related classical preservation theorems are the Łoś-Tarski Theorem (preserved under *injective* homomorphisms $\Leftrightarrow$ equivalent to an *existential* sentence) and Lyndon's Theorem (preserved under *surjective* homomorphism $\Leftrightarrow$ equivalent to a *positive* sentence).

In all classical preservation theorems, the "syntactic property $\Rightarrow$ semantic property" direction is straightforward, while the "semantic property $\Rightarrow$ syntactic property" direction is typically proved by an application of the Compactness Theorem.[1] In order to use compactness, it is essential that the semantic property (i.e. preservation under a certain relationship between structures) holds with respect to *all* structures, that is, both finite and infinite. One may also ask about the status of classical preservation theorems relative to a class of structures $\mathscr{C}$. So long as compactness holds in $\mathscr{C}$ (for example, whenever $\mathscr{C}$ is first-order axiomatizable), so too will all of the classical preservation theorems. The situation is less clear when $\mathscr{C}$ is the class of finite structures (or a subclass thereof), as the Compactness Theorem is easily seen to be false when restricted to finite structures.[2]

The program of classifying theorems in classical model theory according to their validity over finite structures was a major line of research, initiated by Gurevich [1984],

[1]The Compactness Theorem states that a first-order theory T (i.e. set of first-order sentences) is consistent (i.e. there exists a structure $\mathcal{A}$ which satisfies every sentence in T) if every finite sub-theory of T is consistent. (See [Hodges 1993] for background and proofs of various preservation/amalgamation/interpolation theorems in classical model theory.)

[2]Consider the theory $T = \{\Phi_n : n \in \mathbb{N}\}$ where Φ_n expresses "there exist $\geq n$ distinct elements". Every finite sub-theory of T has a finite model, but T itself does not.

in the area known as *finite model theory* (see [Ebbinghaus and Flum 1996; Grädel et al. 2007; Libkin 2004]). The status of preservation theorems in particular was systematically investigated in [Alechina and Gurevich 1997; Rosen and Weinstein 1995]. Given the failure of the Compactness Theorem on finite structures, it is not surprising that nearly all of the classical preservation theorems become false when their statements are restricted to finite structures. A counterexample of Tait [1959] showed that the Łoś-Tarski Theorem is false over finite structures, while Ajtai and Gurevich [1987] showed the demise of Lyndon's Theorem via a stronger result in circuit complexity. Namely, they showed that $\mathrm{Monotone} \cap \mathrm{AC}^0 \neq \mathrm{Monotone}\text{-}\mathrm{AC}^0$, that is, there is a (semantically) monotone Boolean function that is computable by AC^0 circuits, but not by (syntactically) monotone AC^0 circuits. The failure of Lyndon's theorem on finite structures follows via the *descriptive complexity* correspondence between AC^0 and first-order logic. (See [Immerman 1999] about the nexus between logics and complexity classes.)

Given the failure of both the Łoś-Tarski and Lyndon Theorems, it might be expected that the Homomorphism Preservation Theorem also fails over finite structures (as it seems to live at the intersection of Łoś-Tarski and Lyndon). On the contrary, however, previous work of the author [Rossman 2008a] showed that the Homomorphism Preservation Theorem remains valid over finite structures. The technique of [Rossman 2008a] is model-theoretic: its starting point is a new compactness-free proof of the classical theorem, which is then adapted to finite structures. (A summary of the argument is included in Section 8.) In this column, we explain a new and completely different proof of this result — which moreover obtains a quantitative improvement — via a reduction to lower bounds in circuit complexity. The proof relies on a recent result (of independent interest) that the AC^0 formula size of the colored G-subgraph isomorphism problem is $n^{\Omega(\mathbf{tree\text{-}depth}(G)^{\varepsilon})}$ for an absolute constant $\varepsilon > 0$.

Related Work. Prior to [Rossman 2008a], the status of the Homomorphism Preservation Theorem on finite structures was investigated by Feder and Vardi [2003], Grädel and Rosen [1999], and Rosen [1995], who resolved special cases of the question for restricted classes of first-order sentences. Another special case is due to Atserias [2008] in the context of CSP dualities. (See [Rossman 2008a] for a discussion of these results.) A different — and incomparable — line of results [Atserias et al. 2006; Dawar 2010; Nesetril and De Mendez 2014] proves versions of the Homomorphism Preservation Theorems restricted to various *sparse* classes of finite structures (see Ch. 10 of [Nešetřil and de Mendez 2012], as well as [Atserias et al. 2008] related to the Łoś-Tarski Theorem). See Stolboushkin [1995] for an alternative counterexample showing that Lyndon's Theorem fails on finite structures, which is simpler than Ajtai and Gurevich [1987] (but doesn't extend to show $\mathrm{Monotone}\text{-}\mathrm{AC}^0 \neq \mathrm{Monotone} \cap \mathrm{AC}^0$).

Outline. The rest of this column is organized as follows. Because our narrative jumps between logic, graph theory and circuit complexity, for readability sake the various preliminaries — which may be familiar (at least in part) to many readers — are presented in separate sections as needed.

2. PRELIMINARIES, I

2.1. Structures and Homomorphisms

Throughout this column, let σ be a fixed finite relational signature, that is, a list of relation symbols $R^{(r)}$ (where $r \in \mathbb{N}$ denotes the arity of R). A *structure* $\mathcal{A}$ consists of a set A (called the universe of $\mathcal{A}$) together with interpretations $R^{\mathcal{A}} \subseteq A^r$ for each relation symbol $R^{(r)}$ in σ. A priori, structures may be finite or infinite.

A *homomorphism* from a structure $\mathcal{A}$ to a structure $\mathcal{B}$ is a map $f : A \to B$ such that $(a_1, \dots, a_r) \in R^{\mathcal{A}} \Longrightarrow (f(a_1), \dots, f(a_r)) \in R^{\mathcal{B}}$ for every $R^{(r)} \in \sigma$ and $(a_1, \dots, a_r) \in A^r$. Notation $\mathcal{A} \to \mathcal{B}$ asserts the existence of a homomorphism from $\mathcal{A}$ to $\mathcal{B}$.

2.2. First-Order Logic

First-order formulas (in the relational signature σ) are constructed out of atomic formulas (of the form $x_1 = x_2$ or $R(x_1, \dots, x_r)$ where $R^{(r)} \in \sigma$ and x_i's are variables) via boolean connectives ($\varphi \wedge \psi$, $\varphi \vee \psi$, and $\neg\varphi$) and universal and existential quantification ($\forall x\ \varphi(x)$ and $\exists x\ \varphi(x)$). For a structure $\mathcal{A}$ and a first-order formula $\varphi(x_1, \dots, x_k)$ and a tuple of elements $\vec{a} \in A^k$, notation $\mathcal{A} \models \varphi(\vec{a})$ is the statement that $\mathcal{A}$ satisfies φ with $\vec{a}$ instantiating the free variables $\vec{x}$. First-order formulas with no free variables are called *sentences* and represented by capital Greek letters Φ and Ψ.

A first-order sentence (or formula) is said to be:

— *positive* if it does not contain any negations (that is, it has no sub-formula of the form $\neg\varphi$),
— *existential* if it contains only existential quantifiers (that is, it has no universal quantifiers) and has no negations outside the scope of any quantifier, and
— *existential-positive* if it is both existential and positive.

Two important parameters of first-order sentences are quantifier-rank and variable-width. *Quantifier-rank* is the maximum nesting depth of quantifiers. *Variable-width* is the maximum number of free variables in a sub-formula. As we will see in Section 6, under the descriptive complexity characterization of first-order logic in terms of AC^0 circuits, variable-width corresponds to AC^0 circuit size and quantifier-rank corresponds to AC^0 formula size (or, more accurately, AC^0 formula depth when fan-in is restricted to $O(n)$).

Note that first-order sentences are not assumed to be in prenex form. For example, the formula $(\exists x\ P(x)) \vee (\exists y\ \neg Q(y))$ is existential (but not positive) and has quantifier-rank 1 and variable-width 1.

3. THE HOMOMORPHISM PRESERVATION THEOREM

Definition 1. A first-order sentence Φ is *preserved under homomorphisms* [*on finite structures*] if $(\mathcal{A} \models \Phi$ and $\mathcal{A} \to \mathcal{B}) \Longrightarrow \mathcal{B} \models \Phi$ for all [finite] structures $\mathcal{A}$ and $\mathcal{B}$. The notions of *preserved under injective homomorphisms* and *preserved under surjective homomorphisms* are defined similarly.

We now formally state the three classical preservations mentioned in the introduction.

THEOREM 2. (Łoś-Tarski / Lyndon / Homomorphism Preservation Theorems [Lyndon 1959]) *A first-order sentence is preserved under* [*injective / surjective / all*] *homomorphisms if, and only if, it is equivalent to an* [*existential / positive / existential-positive*] *sentence.*

As discussed in the introduction, Łoś-Tarski and Lyndon's Theorems become false when restricted to finite structures.

THEOREM 3. (Failure of Łoś-Tarski and Lyndon Theorems on Finite Structures [Ajtai and Gurevich 1987; Tait 1959]) *There exists a first-order sentence that is preserved under* [*injective / surjective*] *homomorphisms on finite structures, but is not equivalent on finite structures to any* [*existential / positive*] *sentence.*

In contrast, the Homomorphism Preservation Theorem remains valid over finite structures.

THEOREM 4. (Homomorphism Preservation Theorem on Finite Structures [Rossman 2008a]) *If a first-order sentence of quantifier-rank k is preserved under homomorphisms on finite structures, then it is equivalent on finite structures to an existential-positive sentence of quantifier-rank $\beta(k)$, for some computable function $\beta : \mathbb{N} \to \mathbb{N}$.*

We will refer to $\beta : \mathbb{N} \to \mathbb{N}$ in Theorem 4 as the "quantifier-rank blow-up". (Formally, there is one computable function $\beta_\sigma : \mathbb{N} \to \mathbb{N}$ for each finite relational signature σ.) We remark that the upper bound on $\beta(k)$ given by the proof of Theorem 4 is a non-elementary function of k (i.e. it is grows faster than any bounded-height tower of exponentials). In contrast, a second result in [Rossman 2008a] shows that the optimal bound $\beta(k) = k$ holds in the classical Homomorphism Preservation Theorem.

THEOREM 5. ("Equi-rank" Homomorphism Preservation Theorem [Rossman 2008a]) *If a first-order sentence of quantifier-rank k is preserved under homomorphism, then it is equivalent to an existential-positive sentence of quantifier-rank k.*

Due to reliance on the Compactness Theorem, the original proof of the classical Homomorphism Preservation Theorem gives no computable upper bound whatsoever on the quantifier-rank blow-up. Theorem 5 is proved by a constructive, compactness-free argument (see Section 8). In [Rossman 2008a], I conjectured that this stronger "equi-rank" theorem is valid over finite structures. However, new techniques were clearly needed to improve the non-elementary upper bound on $\beta(k)$.

The main result described here is a completely new proof of Theorem 4, which moreover gives a polynomial upper bound on $\beta(k)$.

THEOREM 6. ("Poly-rank" Homomorphism Preservation Theorem on Finite Structures) *If a first-order sentence of quantifier-rank k is preserved under homomorphisms on finite structures, then it is equivalent on finite structures to an existential-positive sentence of quantifier-rank $k^{O(1)}$.*

The proof of Theorem 6 involves a reduction to the AC^0 formula size of SUB_G, the colored G-subgraph isomorphism problem. This reduction transforms lower bounds on the AC^0 formula size of SUB_G into upper bounds on the quantifier-rank blow-up $\beta(k)$ in Theorem 4. In Section 7.1, we derive an exponential upper bound $\beta(k) \le 2^{O(k)}$ from an existing lower bound of [Rossman 2014] on the AC^0 formula size of SUB_{P_k} (also known as the distance-k connectivity problem). Two further steps, described in Section 7.2, are required for the polynomial upper bound $\beta(k) \le k^{O(1)}$ of Theorem 6. The first is a new result in graph minor theory from [Kawarabayashi and Rossman 2016], which gives a "polynomial excluded-minor approximation" of tree-depth, analogous to the Polynomial Grid-Minor Theorem of Chekuri and Chuzhoy [2014]. The second ingredient, in a forthcoming paper of the author [Rossman 2016], is a lower bound on AC^0 formula size of SUB_G in the special case where G is a complete binary tree.

4. PRELIMINARIES, II

4.1. Circuit Complexity

We consider *Boolean circuits* with unbounded fan-in AND and OR gates and negations on inputs. That is, inputs are labelled by variables x_i or negated variables $\overline{x}_i$ (where i comes from some finite index set, typically $\{1, \dots, n\}$). We measure *size* by the number of gates and *depth* by the maximum number of gates on an input-to-output path. Boolean circuits with fan-out 1 (i.e. tree-like Boolean circuits) are called *Boolean formulas*. (Boolean formulas are precisely the same as quantifier-free first-order formulas.)

The *depth-d circuit/formula size* of a Boolean function f is the minimum size of a depth-d circuit/formula that computes f. AC^0 refers to constant-depth, $\mathrm{poly}(n)$-size sequences of Boolean circuits/formula on $\mathrm{poly}(n)$ variables. For a sequence (f_n) of Boolean functions on $\mathrm{poly}(n)$ variables and a constant $c > 0$, we say that "(f_n) has AC^0 circuit/formula size $O(n^c)$ (resp. $\Omega(n^c)$)" if for some d (resp. for all d), the depth-d circuit/formula size of f_n is $O_d(n^c)$ (resp. $\Omega_d(n^c)$) for all n.

One slightly unusual complexity measure (which arises in the descriptive complexity correspondence between AC^0 and first-order logic in Section 6) is *fan-in n depth*, that is, the minimum depth required to compute a Boolean function by AC^0 circuits with fan-in restricted to n. Note that AC^0 formula size lower bounds imply fan-in n depth lower bounds: if f has AC^0 formula size $\omega(n^c)$, then its fan-in n formula depth is at least c (for sufficiently large n). (This follows from the observation that every depth-d formula with fan-in n is equivalent to a depth-d formula of size at most n^d.)

4.2. Monotone Projections

Definition 7 (*Monotone-Projection Reductions*). For Boolean functions $f : \{0,1\}^I \to \{0,1\}$ and $g : \{0,1\}^J \to \{0,1\}$, a *monotone-projection reduction* from f to g is a map $\rho : J \to I \cup \{0,1\}$ such that $f(x) = g(\rho^*(x))$ for all $x \in \{0,1\}^I$ where $\rho^*(x) \in \{0,1\}^J$ is defined by

$$(\rho^*(x))_j = \begin{cases} x_i & \text{if } \rho(j) = i \in I, \\ 0 & \text{if } \rho(j) = 0, \\ 1 & \text{if } \rho(j) = 1. \end{cases}$$

(Properly speaking, the "reduction" from f to g is the map $\rho^* : \{0,1\}^I \to \{0,1\}^J$ induced by ρ.) Notation $f \le_{\mathrm{mp}} g$ denotes the existence of a monotone-projection reduction from f to g.

When describing monotone-projection reductions later on, it will be natural to speak in terms of indexed sets of Boolean variables $\{X_i\}_{i \in I}$ and $\{Y_j\}_{j \in J}$, rather than sets I and J themselves. Thus, a monotone-projection reduction $\rho : J \to I \cup \{0,1\}$ associates each variable Y_j with either a constant (0 or 1) or some variable X_i.

Note that $\le_{\mathrm{mp}}$ is a partial order on Boolean functions. This is the simplest kind of reduction in complexity theory. It has the nice property that every standard complexity measure on Boolean functions is monotone under $\le_{\mathrm{mp}}$. For instance, letting $L_d(f)$ denote the depth-d formula size of f, we have $f \le_{\mathrm{mp}} g \implies L_d(f) \le L_d(g)$.

4.3. Tree-Width and Tree-Depth

Graphs in this column are finite simple graphs. (In contrast to the previous discussion of infinite structures, we assume finiteness whenever we speak of graphs.) Formally, a graph G is a pair $(V(G), E(G))$ where $V(G)$ is a finite set and $E(G) \subseteq \binom{V(G)}{2}$ is a set of unordered pairs of vertices.

Four specific graphs that arise in this column: for $k \ge 1$, let K_k denote the complete graph of order k, let P_k denote the path of order k, let B_k denote the complete binary tree of height k (where every leaf-to-root path has order k), and let $Grid_{k \times k}$ denote the $k \times k$ grid graph. (In the case $k = 1$, all four of these graphs are a single vertex.)

We recall the definitions of two structural parameters, tree-width and tree-depth, which play an important role in this column. A *tree decomposition* of a graph G consists of a tree T and a family $\mathcal{W} = \{W_t\}_{t \in V(T)}$ of sets $W_t \subseteq V(G)$ satisfying

— $\bigcup_{t \in V(T)} W_t = V(G)$ and every edge of G has both ends in some W_t, and
— if $t, t', t'' \in V(T)$ and t' lies on the path in T between t and t'', then $W_t \cap W_{t''} \subseteq W_{t'}$.

The *tree-width* of G, denoted $\mathbf{tw}(G)$, is the minimum of $\max_{t \in V(T)} |W_t| - 1$ over all tree decompositions $(T, \mathcal{W})$ of G.

The *tree-depth* of G, denoted $\mathbf{td}(G)$, is the minimum height of a rooted forest F such that $V(F) = V(G)$ and every edge of G has both ends in some branch in F (i.e. for every $\{v, w\} \in E(G)$, vertices v and w have an ancestor-descendant relationship in F). There is also an inductive characterization of tree-depth: if G has connected components $G_1, \dots, G_t$, then

$$\mathbf{td}(G) = \begin{cases} 1 & \text{if } |V(G)| = 1, \\ 1 + \min_{v \in V(G)} \mathbf{td}(G - v) & \text{if } t = 1 \text{ and } |V(G)| > 1, \\ \max_{i \in \{1,\dots,t\}} \mathbf{td}(G_i) & \text{if } t > 1. \end{cases}$$

These two structural parameters, tree-width and tree-depth, are related by inequalities:

$$\mathbf{tw}(G) \leq \mathbf{td}(G) - 1 \leq \mathbf{tw}(G) \cdot \log |V(G)|. \tag{1}$$

Tree-depth is also related to the length of the longest path in G, denoted $\mathbf{lp}(G)$:

$$\log(\mathbf{lp}(G) + 1) \leq \mathbf{td}(G) \leq \mathbf{lp}(G). \tag{2}$$

(See Ch. 6 of [Nešetřil and de Mendez 2012] for background on tree-depth and proofs of these inequalities.)

Graph parameters $\mathbf{tw}(\cdot)$ and $\mathbf{td}(\cdot)$, as well as $\mathbf{lp}(\cdot)$, are easily seen to be monotone under the graph-minor relation. A reminder what this means: recall that a graph H is a *minor* of a graph G, denoted $H \preceq G$, if H can be obtained from G by a sequence of edge contractions and vertex/edge deletions. A graph parameter $f : \{\text{graphs}\} \to \mathbb{N}$ is said to be *minor-monotone* if $H \preceq G \implies f(H) \leq f(G)$ for all graphs H and G.

5. THE COLORED G-SUBGRAPH ISOMORPHISM PROBLEM

In this section, we introduce the colored G-subgraph isomorphism problem and state the known upper and lower bounds on its complexity with respect to AC^0 circuits and formulas.

Definition 8. For a graph G and $n \in \mathbb{N}$, the *blow-up* $G^{\uparrow n}$ is the graph defined by

$$V(G^{\uparrow n}) = V(G) \times [n],$$
$$E(G^{\uparrow n}) = \big\{\{(v, a), (w, b)\} : \{v, w\} \in E(G),\ a, b \in [n]\big\}.$$

For $\alpha \in [n]^{V(G)}$, let $G^{(\alpha)}$ denote the subgraph of $G^{\uparrow n}$ defined by

$$V(G^{(\alpha)}) = \big\{(v, \alpha_v) : v \in V(G)\big\},$$
$$E(G^{(\alpha)}) = \big\{\{(v, \alpha_v), (w, \alpha_w)\} : \{v, w\} \in E(G)\big\}.$$

(Note that each $G^{(\alpha)}$ is an isomorphic copy of G.)

Definition 9. For any fixed graph G, the *colored G-subgraph isomorphism problem* asks, given a subgraph $X \subseteq G^{\uparrow n}$, to determine whether or not there exists $\alpha \in [n]^{V(G)}$ such that $G^{(\alpha)} \subseteq X$. For complexity purposes, we view this problem as a Boolean function $\mathrm{SUB}_{G,n} : \{0,1\}^{|E(G)| \cdot n^2} \to \{0,1\}$ with variables $\{X_e\}_{e \in E(G^{\uparrow n})}$. We write SUB_G for the sequence of Boolean functions $\{\mathrm{SUB}_{G,n}\}_{n \in \mathbb{N}}$.

5.1. Minor-Monotonicity

The following observation appears in [Li et al. 2014].

PROPOSITION 10. *If H is a minor of G, then* $\mathrm{SUB}_H \leq_{\mathrm{mp}} \mathrm{SUB}_G$ *(i.e.* $\mathrm{SUB}_{H,n} \leq_{\mathrm{mp}} \mathrm{SUB}_{G,n}$ *for all* $n \in \mathbb{N}$*).*

PROOF. By transitivity of $\leq_{\mathrm{mp}}$, it suffices to consider the two cases where H is obtained from G via deleting or contracting a single edge $\{v, w\} \in E(G)$. In both cases, the monotone projection maps each variable $X_{\{(v',a),(w',b)\}}$ of SUB_G with $\{v', w'\} \neq \{v, w'\}$ to the corresponding variable $Y_{\{(v',a),(w',b)\}}$ of SUB_H. In the deletion case, we set the variable $X_{\{(v,a),(w,b)\}}$ to the constant 1 for all $a, b \in [n]$. In the contraction case, we set $X_{\{(v,a),(w,b)\}}$ to 1 if $a = b$ and to 0 if $a \neq b$. (This "planted perfect matching" has the effect of gluing the v-fibre and the w-fibre for instances of SUB_H.) $\square$

Proposition 10 implies that the graph parameter $G \mapsto \mu(\mathrm{SUB}_G)$ is minor-monotone for any standard complexity measure $\mu : \{\text{Boolean functions}\} \to \mathbb{N}$ (e.g. depth-d AC^0 formula size). It also implies:

COROLLARY 11. *For all graphs G,* $\mathrm{SUB}_{P_{\mathbf{td}(G)}} \leq_{\mathrm{mp}} \mathrm{SUB}_G$.

PROOF. Recall that $\mathbf{td}(G) \leq \mathbf{lp}(G)$ by inequality (2). That is, every graph G contains a path of length $\mathbf{td}(G)$. Since subgraphs are minors, we have $P_{\mathbf{td}(G)} \preceq G$ and therefore $\mathrm{SUB}_{P_{\mathbf{td}(G)}} \leq_{\mathrm{mp}} \mathrm{SUB}_G$ by Proposition 10. $\square$

5.2. Upper Bounds

The obvious "brute-force" way of solving SUB_G has running time $O(n^{|V(G)|})$: given an input $X \subseteq G^{\uparrow n}$, check if $G^{(\alpha)} \subseteq X$ for each $\alpha \in [n]^{V(G)}$. A better upper bound comes from tree-width: based on an optimal tree-decomposition $(T, \mathcal{W})$, there is a dynamic-programming algorithm with running time $n^{\mathbf{tw}(G)+O(1)}$ [Plehn and Voigt 1990]. This algorithm can be implemented by AC^0 circuits of size $n^{\mathbf{tw}(G)+O(1)}$ and depth $O(|V(G)|)$.[3]

Unlike circuits, formulas cannot faithfully implement dynamic-programming algorithms. The fastest known formulas for SUB_G are tied to tree-depth: based on a minimum-height rooted forest F witnessing $\mathbf{td}(G)$, there are AC^0 formulas of size $n^{\mathbf{td}(G)+O(1)}$ solving SUB_G (which come from AC^0 circuits of depth $\mathbf{td}(G)+O(1)$ and fan-in $O(n)$). For future reference, these upper bounds are stated in the following proposition.[4]

PROPOSITION 12. *For all graphs G, SUB_G is solvable by AC^0 circuits of size $n^{\mathbf{tw}(G)+O(1)}$, as well as by AC^0 formulas of size $n^{\mathbf{td}(G)+O(1)}$.*

5.3. Lower Bounds: AC^0 Circuit Size

Previous work of the author [Rossman 2008b] showed that the AC^0 circuit size of SUB_{K_k} (a.k.a. the (colored) k-CLIQUE problem) is $n^{\Omega(k)}$ for every $k \in \mathbb{N}$. Generalizing the technique of [Rossman 2008b], Amano [2010] gave a lower bound on the AC^0 circuit size of SUB_G for arbitrary graphs G. In particular, he showed that the AC^0 circuit size of $\mathrm{SUB}_{Grid_{k\times k}}$ is $n^{\Omega(k)}$. This result, combined with the recent Polynomial Grid-Minor

[3]It may be possible to achieve running times of $n^{\delta\cdot\mathbf{tw}(G)+O(1)}$ for constants $\delta < 1$ using fast matrix multiplication algorithms (cp. [Williams 2014]). However, these algorithms appear to require logarithmic-depth circuits. For unrestricted Boolean circuits, no upper bound better than $n^{O(\mathbf{tw}(G))}$ is known, and in fact Marx [2010] has shown that the Strong Exponential Time Hypothesis rules out circuits smaller than $n^{O(\mathbf{tw}(G)/\log \mathbf{tw}(G))}$.

[4]For the *uncolored* G-subgraph isomorphism graph, one gets essentially the same upper bounds via a reduction to SUB_G using the "color-coding" technique of Alon et al. [1995]. Amano [2010] observed that this uncolored-to-colored reduction can be implemented by AC^0 circuits.

Theorem[5] of Chekuri and Chuzhoy [2014], implies that the AC^0 circuit size of SUB_G is $n^{\Omega(\mathbf{tw}(G)^{\varepsilon})}$ for an absolute constant $\varepsilon > 0$. An even stronger lower bound was subsequently proved by Li et al. [2014] (without appealing to the Polynomial Grid-Minor Theorem).

THEOREM 13. *For all graphs G, the AC^0 circuit size of SUB_G is $n^{\Omega(\mathbf{tw}(G)/\log \mathbf{tw}(G))}$.*

This result is nearly tight, as it matches the upper bound of Proposition 12 up to the $O(\log \mathbf{tw}(G))$ factor in the exponent.

5.4. Lower Bounds: AC^0 Formula Size

For the main result (Theorem 6), we require a lower bound on the AC^0 formula size of SUB_G (or in fact on the fan-in $O(n)$ depth of SUB_G). Since formulas are a subclass of circuits, Theorem 13 implies that the AC^0 formula size of SUB_G is at least $n^{\Omega(\mathbf{tw}(G)/\log \mathbf{tw}(G))}$. However, this does not match the $n^{\mathbf{td}(G)+O(1)}$ lower bound of Proposition 12, since $\mathbf{td}(G)$ may be larger than $\mathbf{tw}(G)$ (by up to a $\log|V(G)|$ factor). In particular, the path P_k has tree-width 1 and tree-depth $\lceil \log(k+1) \rceil$. Although Theorem 13 gives no non-trivial lower bound on the AC^0 formula size of SUB_{P_k}, a nearly optimal lower bound was proved in different work of the author [Rossman 2014]:

THEOREM 14. *The AC^0 formula size of SUB_{P_k} is $n^{\Omega(\log k)}$. More precisely, the depth-d formula size of $\mathrm{SUB}_{P_k,n}$ is $n^{\Omega(\log k)}$ for all $k, d, n \in \mathbb{N}$ with $k \leq \log\log n$ and $d \leq \frac{\log n}{(\log\log n)^3}$.*

Via the relationship between AC^0 formula size and fan-in $O(n)$ circuit depth, Theorem 14 implies:

COROLLARY 15. *Circuits with fan-in $O(n)$ computing SUB_{P_k} have depth $\Omega(\log k)$.*

Remark 16. We mention a few other lower bounds related to Corollary 15. A recent paper of Chen et al. [2015] gives a nearly optimal size-depth trade-off for AC^0 circuits computing SUB_{P_k}. Namely, they prove that the depth-d circuit size of $\mathrm{SUB}_{P_k,n}$ is $n^{\Omega(d^{-1}k^{1/(d-1)})}$ for all $k \leq n^{1/5}$. (This result is incomparable to Theorem 14.) As a corollary, this shows that circuits with fan-in $O(n)$ computing SUB_{P_k} have depth $\Omega(\log k/\log\log k)$ (a slightly weaker bound than Corollary 15). Previous size-depth trade-offs due to Beame et al. [1998] and Ajtai [1989] imply lower bounds of $\Omega(\log\log k)$ and $\Omega(\log^* k)$ respectively on the fan-in $O(n)$ depth of SUB_{P_k}.

In Section 7.1, we use Corollary 15 (together with Corollary 11) to prove a weak version of our main result, Theorem 6, with an exponential upper bound $\beta(k) \leq 2^{O(k)}$ on the quantifier-rank blow-up. We remark that the lower bound of Chen et al. implies a slightly weaker upper bound of $k^{O(k)}$, while the very first non-trivial lower bound of Ajtai implies a *non-elementary* upper bound on $\beta(k)$ (similar to the original proof of Theorem 4). For the polynomial upper bound $\beta(k) \leq k^{O(1)}$, we require a stronger $n^{\Omega(\mathbf{td}(G)^{\varepsilon})}$ lower bound on the AC^0 formula size of SUB_G for arbitrary graphs G, as we explain in Section 7.2.

6. PRELIMINARIES, III

In this section, we state a few needed lemmas on the relationship between first-order logic and AC^0 formula size. As before, let σ be a fixed finite relational signature. However, we now stipulate that ***all structures in Sections 6 and 7 are finite***. That is,

[5]This states every graph G of tree-width k contains an $\Omega(k^{\varepsilon}) \times \Omega(k^{\varepsilon})$ grid minor for an absolute constant $\varepsilon > 0$.

we drop the adjective "finite" everywhere since it is assumed. Asymptotic notation in these sections ($O(\cdot)$, etc.) implicitly depends on σ (although, essentially without loss of generality, it suffices to prove our results in the special case $\sigma = \{R^{(2)}\}$ of a single binary relation).

6.1. Descriptive Complexity: $\mathrm{FO} = \mathrm{AC}^0$

Definition 17 (*Gaifman Graphs, Encodings,* MODEL_Φ).

— For a structure $\mathcal{A}$, we denote by $\mathrm{Gaif}(\mathcal{A})$ the *Gaifman graph* of $\mathcal{A}$. This is the graph whose vertex set is the universe of $\mathcal{A}$ and whose edges are pairs $\{v, w\}$ such that $v \neq w$ and v, w appear together in a tuple of any relation of $\mathcal{A}$.
— If $\mathcal{A}$ has universe $[n]$, then we denote by $\mathrm{Enc}(\mathcal{A}) \in \{0,1\}^{\widehat{n}}$ the standard bit-string encoding of $\mathcal{A}$ where $\widehat{n} = \sum_{R^{(t)} \in \sigma} n^t$ $(= n^{O_\sigma(1)})$. That is, each bit of $\mathrm{Enc}(\mathcal{A})$ is the indicator for a tuple of some relation of $\mathcal{A}$. (Note that $\mathrm{Enc}(\cdot)$ is a bijection between structures with universe $[n]$ and strings in $\{0,1\}^{\widehat{n}}$.)
— For a first-order sentence Φ and $n \in \mathbb{N}$, let $\mathrm{MODEL}_{\Phi,n} : \{0,1\}^{\widehat{n}} \to \{0,1\}$ be the Boolean function defined, for structures $\mathcal{A}$ with universe $[n]$, by

$$\mathrm{MODEL}_{\Phi,n}(\mathrm{Enc}(\mathcal{A})) = 1 \stackrel{\text{def}}{\iff} \mathcal{A} \models \Phi.$$

We write MODEL_Φ for the sequence of Boolean functions $\{\mathrm{MODEL}_{\Phi,n}\}_{n \in \mathbb{N}}$.

The next lemma gives one-half of the descriptive complexity correspondence between first-order logic and AC^0:

LEMMA 18 ("$\mathrm{FO} \subseteq \mathrm{AC}^0$"). *For all $1 \leq w \leq k$, if Φ is a first-order sentence of quantifier-rank k and variable-width w, then MODEL_Φ is computable by AC^0 circuits of depth k and fan-in $O(n)$ and size $O(n^w)$. These circuits are equivalent with AC^0 formulas of depth k and size $O(n^k)$.*

(To be completely precise, each of these $O(\cdot)$ terms is really $O_{\sigma,k}(\cdot)$, that is, with constants that depend on k as well as the signature σ.) We remark that Lemma 18 has a converse ("$\mathrm{AC}^0 \subseteq \mathrm{FO}$") with respect to both the uniform and non-uniform versions of AC^0. We omit the statement of these results, since the description of AC^0 circuits via first-order sentences is not needed here (see [Immerman 1999] for details).

6.2. Minimal Cores of a Hom-Preserved Class

We say that a class of structures $\mathscr{C}$ (i.e. a class of finite structures) is *hom-preserved* if, whenever $\mathcal{A} \in \mathscr{C}$ and $\mathcal{A} \to \mathcal{B}$, we have $\mathcal{B} \in \mathscr{C}$. To prove our main result in the next section, we require the following lemma. (See [Hell and Nešetřil 1992; 2004; Rossman 2008a] for more details.)

LEMMA 19. *For every hom-preserved class $\mathscr{C}$, there exists a subset $\mathrm{MinCores}(\mathscr{C}) \subseteq \mathscr{C}$ with the following properties:*

(*1*) *$\mathcal{A} \in \mathscr{C}$ if and only if there exists $\mathcal{M} \in \mathrm{MinCores}(\mathscr{C})$ such that $\mathcal{M} \to \mathcal{A}$.*
(*2*) *Every homomorphism between structures in $\mathrm{MinCores}(\mathscr{C})$ is an isomorphism.*
(*3*) *$\mathscr{C}$ is definable (i.e. within the class of all finite structures) by an existential-positive sentence of quantifier-rank k if and only if $\mathbf{td}(\mathrm{Gaif}(\mathcal{M})) \leq k$ for all $\mathcal{M} \in \mathrm{MinCores}(\mathscr{C})$.*

7. PROOF OF THEOREM 6

In this section, we finally prove our main result, the "Poly-rank" Homomorphism Preservation Theorem on Finite Structures (Theorem 6, stated in Section 3). We begin in Section 7.1 by proving a weaker version of the result with an exponential upper bound $\beta(k) \leq 2^{O(k)}$. In Section 7.2, we describe the improvement to $\beta(k) \leq k^{O(1)}$, which involves new results from circuit complexity and graph minor theory.

7.1. Preliminary Bound: $\beta(k) \leq 2^{O(k)}$

For simplicity's sake, we will assume that σ consists of binary relations only.

Let Φ be a first-order sentence of quantifier-rank k, let $\mathscr{C}$ be the set of finite models of Φ, and assume that $\mathscr{C}$ is hom-preserved (that is, Φ is preserved under homomorphisms on finite structures). Our goal is to show that Φ is equivalent to an existential-positive sentence of quantifier-rank $2^{O(k)}$. By Lemma 19(3), it suffices to show that $\mathbf{td}(\mathrm{Gaif}(\mathcal{M})) \leq 2^{O(k)}$ for all $\mathcal{M} \in \mathrm{MinCores}(\mathscr{C})$.

Consider any $\mathcal{M} \in \mathrm{MinCores}(\mathscr{C})$. Let G be the Gaifman graph of $\mathcal{M}$, and let m be the size of the universe of $\mathcal{M}$. (Note that $m = |V(G)|$.) The following claim is key to showing $\mathbf{td}(G) \leq 2^{O(k)}$.

CLAIM 20. *For all $n \in \mathbb{N}$, there exists a monotone-projection reduction* $\mathrm{SUB}_{G,n} \leq_{\mathrm{mp}} \mathrm{MODEL}_{\Phi,mn}$.

In order to define this monotone-projection reduction, let us identify $[mn]$ with the set $V(G^{\uparrow n})$ ($= V(G) \times [n]$). Variables X_e of $\mathrm{SUB}_{G,n}$ are indexed by potential edges $e \in E(G^{\uparrow n})$ in a subgraph $X \subseteq G^{\uparrow n}$. Variables Y_i of $\mathrm{MODEL}_{\Phi,mn}$ are indexed by the set

$$I := \big\{(R,(v,a),(w,b)) : R^{(2)} \in \sigma,\ (v,a),(w,b) \in V(G^{\uparrow n})\big\}.$$

(That is, I is the set of potential 2-tuples of relations of structures with universe $V(G^{\uparrow n})$.) Define the monotone projection $\rho : \{Y_i\}_{i \in I} \to \{X_e\}_{e \in E(G^{\uparrow n})} \cup \{0,1\}$ by

$$\rho : Y_{(R,(v,a),(w,b))} \mapsto \begin{cases} X_{\{(v,a),(w,b)\}} & \text{if } (v,w) \in R^{\mathcal{M}} \text{ and } v \neq w,\\ 1 & \text{if } (v,w) \in R^{\mathcal{M}} \text{ and } v = w,\\ 0 & \text{otherwise.} \end{cases}$$

We must show that the corresponding map

$$\rho^* : \{\text{subgraphs of } G^{\uparrow n}\} \to \{\text{structures with universe } V(G^{\uparrow n})\}$$

is in fact a reduction from $\mathrm{SUB}_{G,n}$ to $\mathrm{MODEL}_{\Phi,mn}$. That is, we must show that for any $X \subseteq G^{\uparrow n}$,

$$\mathrm{SUB}_{G,n}(X) = 1 \iff \mathrm{MODEL}_{\Phi,mn}(\rho^*(X)) = 1. \tag{3}$$

For the $\Longrightarrow$ direction of (3): Assume $\mathrm{SUB}_{G,n}(X) = 1$. Then $G^{(\alpha)} \subseteq X$ for some $\alpha \in [n]^{V(G)}$. The definition of ρ ensures that the map $v \mapsto (v, \alpha_v)$ is a homomorphism from $\mathcal{M}$ to the structure $\rho^*(X)$. Since $\mathscr{C}$ is hom-preserved, it follows that $\rho^*(X) \in \mathscr{C}$ and therefore $\mathrm{MODEL}_{\Phi,mn}(\rho^*(X)) = 1$.

For the $\Longleftarrow$ direction of (3): Assume $\mathrm{MODEL}_{\Phi,mn}(\rho^*(X)) = 1$, that is, $\rho^*(X) \in \mathscr{C}$. By Lemma 19(1) there exist $\mathcal{N} \in \mathrm{MinCores}(\mathscr{C})$ and a homomorphism $\gamma : \mathcal{N} \to \rho^*(X)$. The definition of ρ ensures that the map $\pi : (v,i) \mapsto v$ is a homomorphism from $\rho^*(X)$ to $\mathcal{M}$. The composition $\pi \circ \gamma$ is a homomorphism from $\mathcal{N}$ to $\mathcal{M}$. By Lemma 19(2), it is an isomorphism. Therefore, without loss of generality, we may assume that $\mathcal{M} = \mathcal{N}$ and $\pi \circ \gamma$ is the identity map on the universe $V(G)$ of $\mathcal{M}$. This means that $\pi(v) \in \{(v,a) : a \in [n]\}$ for all $v \in V(G)$. We may now define $\alpha \in [n]^{V(G)}$ as the unique element such that $\gamma : v \mapsto (v, \alpha_v)$ for all $v \in V(G)$. From the definition of ρ and the fact that $G = \mathrm{Gaif}(\mathcal{M})$, we infer that $G^{(\alpha)} \subseteq X$. We conclude that $\mathrm{SUB}_{G,n}(X) = 1$, finishing the proof of Claim 20.

We proceed to show that $\mathbf{td}(G) \le 2^{O(k)}$. By Corollary 11, we have $\mathrm{SUB}_{P_{\mathbf{td}(G)},n} \le_{\mathrm{mp}} \mathrm{SUB}_{G,n}$. By Claim 20 and transitivity of $\le_{\mathrm{mp}}$, it follows that $\mathrm{SUB}_{P_{\mathbf{td}(G)},n} \le_{\mathrm{mp}} \mathrm{MODEL}_{\Phi,kn}$. Therefore, $\mu(\mathrm{SUB}_{P_{\mathbf{td}(G)},n}) \le \mu(\mathrm{MODEL}_{\Phi,kn})$ for every standard complexity measure $\mu : \{\text{Boolean functions}\} \to \mathbb{N}$ (in particular, depth-k formula size). By Lemma 18 (the simulation of first-order logic by AC^0), there exist depth-k formulas of size $O((mn)^k)$ which compute $\mathrm{MODEL}_{\Phi,mn}$. Therefore, there exist depth-k formulas of size $O((mn)^k)$ which compute $\mathrm{SUB}_{P_{\mathbf{td}(G)},n}$. On the other hand, by Theorem 14, the depth-k formula size of $\mathrm{SUB}_{P_{\mathbf{td}(G)},n}$ is $n^{\Omega(\log \mathbf{td}(G))}$ for all sufficiently large n such that $k < \log\log n$. Therefore, we have $n^{\Omega(\log \mathbf{td}(G))} \le O((mn)^k)$ for all sufficiently large n. Since m $(= |V(G)|)$ is constant, it follows that $k \ge \Omega(\log \mathbf{td}(G))$, that is, $\mathbf{td}(G) \le 2^{O(k)}$. This completes the proof that $\beta(k) \le 2^{O(k)}$ for binary signatures σ.

Remark 21. In this argument, as an alternative to *depth-k formula size*, we may instead consider *fan-in $O(n)$ depth* (i.e. *fan-in cn depth* for a sufficiently large constant c) and appeal to Corollary 15 instead of Theorem 14.

7.2. Improvement to $\beta(k) \le k^{O(1)}$

The upper bound $\beta(k) \le 2^{O(k)}$ in the previous section relies on the *exponential approximation* of tree-depth in terms of the longest path, that is, $\log(\mathbf{lp}(G)+1) \le \mathbf{td}(G) \le \mathbf{lp}(G)$ (inequality (2)). To achieve a polynomial upper bound on $\beta(k)$, we require a *polynomial approximation* of tree-depth in terms of a few manageable classes of "excluded minors". This realization led to a conjecture of the author, which was thereafter proved in forthcoming work [Kawarabayashi and Rossman 2016].

THEOREM 22. *Every graph G of tree-depth k satisfies one (or more) of the following conditions for $\ell = \widetilde{\Omega}(k^{1/5})$:*

(*1*) $\mathbf{tw}(G) \ge \ell$,
(*2*) *G contains a path of length 2^ℓ, or*
(*3*) *G contains a B_ℓ-minor.*

This result is analogous to the Polynomial Grid-Minor Theorem [Chekuri and Chuzhoy 2014], which can be used to replace condition (i) with the condition that G contains an $\Omega(k^\varepsilon) \times \Omega(k^\varepsilon)$ grid minor for an absolute constant $\varepsilon > 0$. In cases (i) and (ii), Theorems 13 and 14 respectively imply that SUB_G has AC^0 formula size $n^{\widetilde{\Omega}(\mathbf{td}(G)^{1/5})}$. This leaves only case (iii), where forthcoming work [Rossman 2016] shows the following (via a generalization of the "pathset complexity" framework of [Rossman 2014]).

THEOREM 23. *The AC^0 formula size of SUB_{B_k} is $n^{\Omega(k^\varepsilon)}$ for an absolute constant $\varepsilon > 0$.*

Together Theorems 22 and 23 imply:

THEOREM 24. *For all graphs G, the AC^0 formula complexity of SUB_G is $n^{\Omega(\mathbf{td}(G)^\varepsilon)}$ for an absolute constant $\varepsilon > 0$.*

Plugging Theorem 24 into the argument in the previous subsection directly yields the polynomial upper bound $\beta(k) \le k^{O(1)}$ of Theorem 4. (In fact, we get $\beta(k) \le k^{1/\varepsilon}$ for the constant $\varepsilon > 0$ of Theorem 24.)

8. COMPARISON WITH THE METHOD IN (R. 2008)

In this section, for the sake of comparison, we summarize the model-theoretic approach of the original proof of Theorem 4 in [Rossman 2008a]. The starting point in [Rossman

2008a] was a new compactness-free proof of the classical Homomorphism Preservation Theorem, which moreover yields the stronger "equi-rank" version (Theorem 5). The proof is based on an operation mapping each structure $\mathcal{A}$ to an infinite co-retract $\Gamma(\mathcal{A})$. (We drop the assumption of the last two sections that structures are finite by default.) In order to state the key property of this operation, we introduce notation $\mathcal{A} \equiv_{\mathrm{FO}(k)} \mathcal{B}$ (resp. $\mathcal{A} \equiv_{\exists^+\mathrm{FO}(k)} \mathcal{B}$) denoting that $\mathcal{A}$ and $\mathcal{B}$ satisfy the same first-order sentences (resp. existential-positive sentences) of quantifier-rank k.

THEOREM 25. *There is an operation $\Gamma : \{$structures$\} \to \{$structures$\}$ associating every structure $\mathcal{A}$ with a co-retract $\Gamma(\mathcal{A}) \overset{\supseteq}{\to} \mathcal{A}$ such that, for all structures $\mathcal{A}$ and $\mathcal{B}$ and $k \in \mathbb{N}$,*

$$\mathcal{A} \equiv_{\exists^+\mathrm{FO}(k)} \mathcal{B} \implies \Gamma(\mathcal{A}) \equiv_{\mathrm{FO}(k)} \Gamma(\mathcal{B}).$$

There is a straightforward proof that Theorem 25 implies Theorem 5 (see [Rossman 2008a]). The structure $\Gamma(\mathcal{A})$ is the Fraïssé limit of the class of co-finite co-retracts of $\mathcal{A}$ (that is, structures $\mathcal{A}'$ such that $\mathcal{A}' \overset{\supseteq}{\to} \mathcal{A}$ and $A' \setminus A$ is finite). We remark that $\Gamma(\mathcal{A})$ is infinite, even when $\mathcal{A}$ is finite. For this reason, Theorem 25 says nothing in the setting of finite structures.

The Homomorphism Preservation Theorem on Finite Structures (Theorem 4) is proved in [Rossman 2008a] by considering a sequence of finitary "approximations" of $\Gamma(\mathcal{A})$. (This is somewhat analogous to the sense in which large random graph $G(n, 1/2)$ "approximate" the infinite Rado graph.)

THEOREM 26. *There is a computable function $\beta : \mathbb{N} \to \mathbb{N}$ and a sequence $\{\Gamma_k\}_{k\in\mathbb{N}}$ of operations $\Gamma_k : \{$finite structures$\} \to \{$finite structures$\}$ associating every finite structure $\mathcal{A}$ with a sequence $\{\Gamma_k(\mathcal{A})\}_{k\in\mathbb{N}}$ of finite co-retracts $\Gamma_k(\mathcal{A}) \overset{\supseteq}{\to} \mathcal{A}$ such that, for all finite structures $\mathcal{A}$ and $\mathcal{B}$ and $k \in \mathbb{N}$,*

$$\mathcal{A} \equiv_{\exists^+\mathrm{FO}(\beta(k))} \mathcal{B} \implies \Gamma_k(\mathcal{A}) \equiv_{\mathrm{FO}(k)} \Gamma_k(\mathcal{B}).$$

Theorem 4 follows directly from Theorem 26, inheriting the same quantifier-rank blow-up $\beta(k)$. The proof of Theorem 26 in [Rossman 2008a] implies a non-elementary upper bound on $\beta(k)$. While the work we describe improves the upper bound $\beta(k) \leq k^{O(1)}$ in Theorem 4, we remark that this it does not imply any improvement to $\beta(k)$ in Theorem 26.

9. SYNTAX VS. SEMANTICS IN CIRCUIT COMPLEXITY

We conclude by stating some consequences of our results in circuit complexity. Let $\mathrm{HomPreserved}$ denote the class of all homomorphism-preserved graph properties (for example, $\{G : \mathrm{girth}(G) \leq 20$ or $\mathrm{clique\text{-}number}(G) \geq 10\}$). This is a semantic class, akin to the class $\mathrm{Monotone}$ of all monotone languages. The new proof of the Homomorphism Preservation Theorem on Finite Structures using AC^0 lower bounds also implies the following "Homomorphism Preservation Theorem for (non-uniform) AC^0":

$$\mathrm{HomPreserved} \cap \mathrm{AC}^0 = \exists^+\mathrm{FO} \quad (\subseteq \{\text{poly-size monotone DNFs}\}).$$

In other words, every homomorphism-preserved graph property in AC^0 is definable (over finite graphs) by an existential-positive first-order sentence and, therefore, also by a polynomial-size monotone DNF (moreover, with constant bottom fan-in). As a consequence, for every integer $d \geq 2$, we get a collapse of the AC^0 depth hierarchy with respect to homomorphism-preserved properties:

$$\mathrm{HomPreserved} \cap \mathrm{AC}^0[\text{depth } d] = \mathrm{HomPreserved} \cap \mathrm{AC}^0[\text{depth } d+1].$$

In contrast, it is known that $\mathrm{AC}^0[\text{depth } d] \neq \mathrm{AC}^0[\text{depth } d+1]$ by the Depth Hierarchy Theorem [Håstad 1986].

These results have an opposite nature to the "syntactic monotonicity $\neq$ semantic monotonicity" counterexamples of Ajtai and Gurevich [1987] and Razborov [1985] (as well as Tardos [1988]), which respectively show that

$$\mathrm{Monotone} \cap \mathrm{AC}^0 \neq \text{Monotone-AC}^0 \quad \textbf{and} \quad \mathrm{Monotone} \cap \mathrm{P} \neq \text{Monotone-P}.$$

In light of the results just presented, I feel that questions of syntax vs. semantics in circuit complexity are worth re-examining. For instance, so far as I know, there is no known separation between the *uniform average-case* monotone vs. non-monotone complexity of any monotone function in any well-studied class of Boolean circuits (AC^0, NC^1, etc.) It is plausible that syntactic monotonicity = semantic monotonicity in the average-case. Evidence for this viewpoint comes from considering the *slice distribution* (that is, the uniform distribution on inputs of Hamming weight exactly $\lfloor n/2 \rfloor$). With respect to the slice distribution, it is known that monotone and non-monotone complexity are equivalent within a $\mathrm{poly}(n)$ factor by a classic result of Berkowitz [1982].

As for an even stronger "Homomorphism Preservation Theorem" in circuit complexity, we can state the following: if for every k, SUB_{P_k} requires unbounded-depth formula size $n^{\Omega(\log k)}$ (which is widely conjectured to be true) or even $n^{\omega_{k\to\infty}(1)}$, then $\mathrm{HomPreserved} \cap \mathrm{NC}^1 = \exists^+\mathrm{FO}$. Therefore, I strongly believe in a "Homomorphism Preservation Theorem for NC^1". On the other hand, the homomorphism-preserved property of being 2-colorable a.k.a. non-bipartite ($= \{G : C_k \to G$ for any odd $k\}$) is in $\mathrm{Logspace}$ (this follows from Reingold's theorem [Reingold 2008]), yet it is not $\exists^+\mathrm{FO}$-definable. Therefore, we may assert that $\mathrm{HomPreserved} \cap \mathrm{Logspace} \neq \exists^+\mathrm{FO}$.

Acknowledgements. The author is supported by NSERC and the JST ERATO Kawarabayashi Large Graph Project. This column was partially written at the National Institute of Informatics in Tokyo and during a visit to IMPA, the National Institute for Pure and Applied Mathematics in Rio de Janeiro.

REFERENCES

Miklos Ajtai. 1989. First-order definability on finite structures. *Annals of Pure and Applied Logic* 45, 3 (1989), 211–225.

M. Ajtai and Y. Gurevich. 1987. Monotone versus positive. *J. ACM* 34 (1987), 1004–1015.

N. Alechina and Y. Gurevich. 1997. Syntax vs. semantics on finite structures. In *Structures in Logic and Computer Science*, J. Mychielski, G. Rozenberg, and A. Salomaa (Eds.). Springer-Verlag, 14–33.

Noga Alon, Raphael Yuster, and Uri Zwick. 1995. Color-coding. *J. ACM* 42, 4 (1995), 844–856.

Kazuyuki Amano. 2010. k-Subgraph isomorphism on AC^0 circuits. *Computational Complexity* 19, 2 (2010), 183–210.

Albert Atserias. 2008. On digraph coloring problems and treewidth duality. *European Journal of Combinatorics* 29, 4 (2008), 796–820.

Albert Atserias, Anuj Dawar, and Martin Grohe. 2008. Preservation under extensions on well-behaved finite structures. *SIAM J. Comput.* 38, 4 (2008), 1364–1381.

A. Atserias, A. Dawar, and Ph.G. Kolaitis. 2006. On preservation under homomorphisms and unions of conjunctive queries. *J. ACM* 53, 2 (2006), 208–237.

Paul Beame, Russell Impagliazzo, and Toniann Pitassi. 1998. Improved depth lower bounds for small distance connectivity. *Computational Complexity* 7, 4 (1998), 325–345.

Stephen Bellantoni, Toniann Pitassi, and Alasdair Urquhart. 1992. Approximation and small-depth Frege proofs. *SIAM J. Comput.* 21, 6 (1992), 1161–1179.

S Berkowitz. 1982. *On some relationships between monotone and nonmonotone circuit complexity*. Technical Report. Technical report, Department of Computer Science, University of Toronto, Canada, Toronto, Canada.

Chandra Chekuri and Julia Chuzhoy. 2014. Polynomial bounds for the grid-minor theorem. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*. ACM, 60–69.

Xi Chen, Igor C Oliveira, Rocco A Servedio, and Li-Yang Tan. 2015. Near-optimal small-depth lower bounds for small distance connectivity. *arXiv preprint arXiv:1509.07476* (2015).

Anuj Dawar. 2010. Homomorphism preservation on quasi-wide classes. *J. Comput. System Sci.* 76, 5 (2010), 324–332.

H.-D. Ebbinghaus and J. Flum. 1996. *Finite Model Theory*. Springer-Verlag.

T. Feder and M.Y. Vardi. 2003. Homomorphism closed vs. existential positive. In *Proceedings of the 18th IEEE Symposium on Logic in Computer Science*. 310–320.

E. Grädel, P.G. Kolaitis, L. Libkin, M. Marx, J. Spencer, M.Y. Vardi, Y. Venema, and S. Weinstein. 2007. *Finite Model Theory and its Applications*. Springer.

E. Grädel and E. Rosen. 1999. On preservation theorems for two-variable logic. *Math. Logic Quart.* 45 (1999), 315–325.

Martin Grohe. 2007. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM* 54, 1 (2007), 1–24.

Y. Gurevich. 1984. Toward logic tailored for computational complexity. In *Computation and Proof Theory*, M. M. Richter et al. (Ed.). Springer Lecture Notes in Mathematics, 175–216.

Johan Håstad. 1986. Almost optimal lower bounds for small depth circuits. In *STOC '86: Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*. 6–20.

P. Hell and J. Nešetřil. 1992. The core of a graph. *Discrete Math.* 109 (1992), 117–126.

P. Hell and J. Nešetřil. 2004. *Graphs and Homomorphisms*. Oxford University Press.

W. Hodges. 1993. *Model Theory*. Cambridge University Press.

N. Immerman. 1999. *Descriptive Complexity Theory*. Springer, New York.

Kenichi Kawarabayashi and Benjamin Rossman. 2016. An Excluded-Minor Approximation of Tree-Depth. (2016). manuscript.

Yuan Li, Alexander Razborov, and Benjamin Rossman. 2014. On the AC0 Complexity of Subgraph Isomorphism. In *Found. of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*. IEEE, 344–353.

L. Libkin. 2004. *Elements of Finite Model Theory*. Springer-Verlag.

R.C. Lyndon. 1959. Properties preserved under homomorphism. *Pacific J. Math.* 9 (1959), 129–142.

Dániel Marx. 2010. Can You Beat Treewidth? *Theory of Computing* 6 (2010), 85–112.

Jaroslav Nesetril and Patrice Ossona De Mendez. 2014. On first-order definable colorings. (2014). https://arxiv.org/abs/1403.1995.

J. Nešetřil and P. Ossona de Mendez. 2006. Tree depth, subgraph coloring and homomorphism bounds. *European J. Combin.* 27, 6 (2006), 1022–1041.

J. Nešetřil and P. Ossona de Mendez. 2012. Sparsity (graphs, structures, and algorithms), Algorithms and Combinatorics, vol. 28. (2012).

Jürgen Plehn and Bernd Voigt. 1990. Finding minimally weighted subgraphs. In *International Workshop on Graph-Theoretic Concepts in Computer Science*. Springer, 18–29.

Alexander A Razborov. 1985. Lower bounds on the monotone complexity of some Boolean functions. In *Dokl. Akad. Nauk SSSR*, Vol. 281. 798–801.

Omer Reingold. 2008. Undirected connectivity in log-space. *Journal of the ACM (JACM)* 55, 4 (2008), 17.

E. Rosen. 1995. *Finite model theory and finite variable logic*. Ph.D. Dissertation. University of Pennsylvania.

E. Rosen and S. Weinstein. 1995. Preservation theorems in finite model theory. In *Logic and Computational Complexity*, D. Leivant (Ed.). Spring-Verlag, 480–502.

Benjamin Rossman. 2008a. Homomorphism preservation theorems. *Journal of the ACM (JACM)* 55, 3 (2008), 15.

Benjamin Rossman. 2008b. On the constant-depth complexity of k-clique. In *40th Annual ACM Symposium on Theory of Computing (STOC)*. 721–730.

Benjamin Rossman. 2014. Formulas vs. circuits for small distance connectivity. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*. ACM, 203–212.

Benjamin Rossman. 2016. Lower Bounds for Subgraph Isomorphism. (2016). manuscript.

A. Stolboushkin. 1995. Finite monotone properties. In *Proceedings of the 10th IEEE Symposium on Logic in Computer Science*. 324–330.

W. Tait. 1959. A counterexample to a conjecture of Scott and Suppes. *J. Symbolic Logic* 24 (1959), 15–16.

Éva Tardos. 1988. The gap between monotone and non-monotone circuit complexity is exponential. *Combinatorica* 8, 1 (1988), 141–142.

Ryan Williams. 2014. Faster decision of first-order graph properties. In *Proc. 29th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*.