

Preservation under Extensions on Well-Behaved Finite Structures ^{*}

Albert Atserias
Universitat Politècnica de Catalunya,
Barcelona, Spain
atserias@lsi.upc.edu [†]

Anuj Dawar
University of Cambridge,
Cambridge, U.K.
anuj.dawar@cl.cam.ac.uk

Martin Grohe
Humboldt Universitat zu Berlin,
Berlin, Germany
grohe@informatik.hu-berlin.de

Abstract

A class of relational structures is said to have the extension preservation property if every first-order sentence that is preserved under extensions on the class is equivalent to an existential sentence. The class of all finite structures does not have the extension preservation property. We study the property on classes of finite structures that are better behaved. We show that the property holds of classes of acyclic structures, structures of bounded degree and more generally structures that are *wide* in a sense we make precise. We also show that the preservation property holds for the class of structures of treewidth at most k , for any k . In contrast, we show that the property fails for the class of planar graphs.

^{*}A preliminary short version of this paper appeared in the proceedings of the 32nd International Conference on Automata, Languages and Programming (ICALP), volume 3580 of Lecture Notes in Computer Science, Springer-Verlag, pages 1437-1449, 2005.

[†]The first author was supported in part by CICYT TIN2004-04343 and by the European Commission through the RTN COMBSTRU HPRN-CT2002-00278.

1 Introduction

The subject of model theory is concerned with the relationship between syntactic and semantic properties of logic. Among classical results in the subject are preservation theorems which relate syntactic restrictions on first-order logic with structural properties of the classes of structures defined. A key example is the Łoś-Tarski Theorem which asserts that a first-order formula is preserved under extensions on all structures if, and only if, it is logically equivalent to an existential formula (see [13]). One direction of this result is easy, namely that any formula that is purely existential is preserved under extensions, and this holds on any class of structures. The other direction, going from the semantic restriction to the syntactic restriction makes key use of the compactness of first-order logic and hence of infinite structures.

In the early development of finite-model theory, when it was realized that finite structures are the ones that are interesting from the point of view of studying computation, it was observed that most classical preservation theorems from model theory fail when only finite structures are allowed. In particular, the Łoś-Tarski theorem fails on finite structures [16, 12]. These results suggest that the class of finite structures is not well-behaved from the point of view of model theory. However, when one considers the computational structures that arise in practice and are used as interpretations for logical languages (for instance, program models interpreting specifications or databases interpreting queries), in many cases they are not only finite but satisfy other structural restrictions as well. This motivates the study, not just of the class of finite structures, but of well-behaved subclasses of this class. Note that classical model theory, in most of its more advanced parts, also considers restricted classes of structures such as stable, simple, o-minimal structures or specific structures that are of interest in other areas of mathematics.

There are certain restrictions on finite structures that have proved especially useful in modern graph structure theory and also from an algorithmic point of view. For instance, many intractable computational problems become tractable when restricted to planar graphs or structures of bounded treewidth [4]. This is also the case in relation to evaluation of logical formulas [9]. A common generalization of classes of bounded treewidth and planar graphs are classes of structures that exclude a minor which have also been extensively studied.

A study of preservation properties for such restricted classes of finite structures was initiated in [1]. There, the focus was on the homomorphism preservation theorem, whose status on the class of finite structures was open. It was shown that this preservation property holds on any class of structures of bounded degree, bounded treewidth or that excludes some minor (and has certain other closure properties). In the present paper, we investigate the Łoś-Tarski extension preservation property on these classes of finite structures. Note that the failure of the property on the class of all finite structures does not imply its failure on subclasses. If one considers the non-trivial direction of the preservation theorem on a class \mathcal{C} , it says that any sentence φ that is preserved under extensions *on* \mathcal{C} is equivalent *on* \mathcal{C} to an existential sentence. Thus, restricting to a subclass \mathcal{C}' of \mathcal{C} weakens both the hypothesis and the consequent of the statement.

We show that the extension preservation theorem holds on any class of finite structures closed under substructures and disjoint unions that is also *wide* in the sense that any sufficiently large structure in the class contains a large number of elements that are far apart. This includes, for instance, any class of structures of bounded degree. While classes of structures of bounded treewidth are not wide, they are nearly so in that they can be made wide by removing a small number of elements. We use this property and show that it implies the extension preservation theorem for the class \mathcal{T}_k —the class of structures of treewidth k or less (note this is not as general as saying that the property holds for all classes of bounded treewidth). Finally, although all classes defined by excluded minors are known to be *almost wide* in the same sense as \mathcal{T}_k is, we show that the construction does not extend to them. We provide a counterexample to the extension preservation property

for the class of planar graphs, and indeed, even for the class of planar graphs of treewidth at most four. This contrasts with the results obtained for the homomorphism preservation property in [1] as this property was shown to hold on all classes excluding a graph minor and closed under substructures and disjoint unions.

The main methodology in establishing the preservation property for a class of structures \mathcal{C} is to show an upper bound on the size of a minimal model of a first-order sentence φ that is preserved under extensions on \mathcal{C} . The way we do this is to show that for any sufficiently large model \mathbf{A} of φ , there is a proper substructure of \mathbf{A} and an extension of \mathbf{A} that cannot be distinguished by φ . In Section 3 we establish this for the relatively simple case of acyclic structures by means of a Hanf locality argument. Section 4 contains the main combinatorial argument for wide structures which uses Gaifman locality and an iterated construction of the substructure of \mathbf{A} . In Section 5, the combinatorial argument is adapted to the classes \mathcal{T}_k . Finally, in Section 6 we discuss the existence of a counterexample in the case of planar graphs. We begin in Section 2 with some background and definitions.

2 Preliminaries

We use standard notation and terminology from finite model theory (see [5]). Some particular definitions and notation are explained in this section.

2.1 Relational structures

A *relational vocabulary* σ is a finite set of *relation symbols*, each with a specified *arity*. A σ -*structure* \mathbf{A} consists of a *universe* A , or *domain*, and an *interpretation* which associates to each relation symbol $R \in \sigma$ of some arity r , a relation $R^{\mathbf{A}} \subseteq A^r$. A *graph* is a structure $\mathbf{G} = (V, E)$, where E is a binary relation that is symmetric and anti-reflexive. Thus, our graphs are undirected, loopless, and without parallel edges.

A σ -structure \mathbf{B} is called a *substructure* of \mathbf{A} if $B \subseteq A$ and $R^{\mathbf{B}} \subseteq R^{\mathbf{A}}$ for every $R \in \sigma$. It is called an *induced substructure* if $R^{\mathbf{B}} = R^{\mathbf{A}} \cap B^r$ for every $R \in \sigma$ of arity r . Notice the analogy with the graph-theoretical concept of *subgraph* and *induced subgraph*. A substructure \mathbf{B} of \mathbf{A} is proper if $\mathbf{A} \neq \mathbf{B}$. If \mathbf{A} is an induced substructure of \mathbf{B} , we say that \mathbf{B} is an *extension* of \mathbf{A} . If \mathbf{A} is a proper induced substructure, then \mathbf{B} is a proper extension. If \mathbf{B} is the disjoint union of \mathbf{A} with another σ -structure, we say that \mathbf{B} is a *disjoint extension* of \mathbf{A} . If $S \subseteq A$ is a subset of the universe of \mathbf{A} , then $\mathbf{A} \cap S$ denotes the *induced substructure generated by S* ; in other words, the universe of $\mathbf{A} \cap S$ is S , and the interpretation in $\mathbf{A} \cap S$ of the r -ary relation symbol R is $R^{\mathbf{A}} \cap S^r$.

The *Gaifman graph* of a σ -structure \mathbf{A} , denoted by $\mathcal{G}(\mathbf{A})$, is the (undirected) graph whose set of nodes is the universe of \mathbf{A} , and whose set of edges consists of all pairs (a, a') of distinct elements of A such that a and a' appear together in some tuple of a relation in \mathbf{A} . The *degree* of a structure is the degree of its Gaifman graph, that is, the maximum number of neighbors of nodes of the Gaifman graph.

2.2 Neighborhoods and treewidth

Let $\mathbf{G} = (V, E)$ be a graph. Moreover, let $u \in V$ be a node and let $d \geq 0$ be an integer. The *d -neighborhood* of u in \mathbf{G} , denoted by $N_d^{\mathbf{G}}(u)$, is defined inductively as follows:

1. $N_0^{\mathbf{G}}(u) = \{u\}$;
2. $N_{d+1}^{\mathbf{G}}(u) = N_d^{\mathbf{G}}(u) \cup \{v \in V : (v, w) \in E \text{ for some } w \in N_d^{\mathbf{G}}(u)\}$.

If \mathbf{A} is a σ -structure, a is a point in \mathbf{A} , and \mathbf{G} is the Gaifman graph of \mathbf{A} , we let $N_d^{\mathbf{A}}(a)$ denote the d -neighborhood of a in \mathbf{G} . Where it causes no confusion, we also write $N_d^{\mathbf{A}}(a)$ for the substructure of \mathbf{A} generated by this set.

A *tree* is an acyclic connected graph. A *tree-decomposition* of $\mathbf{G} = (V, E)$ is a pair (T, L) where T is a tree and $L : T \rightarrow \wp(V)$ is a labelling of the nodes of T by sets of vertices of \mathbf{G} such that:

1. for every edge $\{u, v\} \in E$, there is a node t of T such that $\{u, v\} \subseteq L(t)$;
2. for every $u \in V$, the set $\{t \in T : u \in L(t)\}$ forms a connected subtree of T .

The *width* of a tree-decomposition (T, L) is $\max_{t \in T} |L(t)| - 1$. The *treewidth* of \mathbf{G} is the smallest k for which \mathbf{G} has a tree-decomposition of width k . The *treewidth* of a σ -structure is the treewidth of its Gaifman graph. Note that trees have treewidth one.

2.3 First-order logic, monadic second-order logic, and types

Let σ be a relational vocabulary. The *atomic formulas* of σ are those of the form $R(x_1, \dots, x_r)$, where $R \in \sigma$ is a relation symbol of arity r , and x_1, \dots, x_r are first-order variables that are not necessarily distinct. Formulas of the form $x = y$ are also atomic.

The collection of *first-order formulas* is obtained by closing the atomic formulas under negation, conjunction, disjunction, universal and existential first-order quantification. The collection of *existential* first-order formulas is obtained by closing the atomic formulas and the negated atomic formulas under conjunction, disjunction, and existential quantification. The semantics of first-order logic is standard.

The collection of *monadic second-order formulas* is obtained by closing the atomic formulas under negation, conjunction, disjunction, universal and existential first-order quantification, and universal and existential second-order quantification over sets. The semantics of monadic second-order logic is also standard.

The quantifier rank of a formula, be it first-order or monadic second-order, is the depth of nesting of quantifiers in the formula.

Let \mathbf{A} be a σ -structure, and let a_1, \dots, a_n be points in \mathbf{A} . If $\varphi(x_1, \dots, x_n)$ is a formula with free variables x_1, \dots, x_n , we use the notation $\mathbf{A} \models \varphi(a_1, \dots, a_n)$ to denote the fact that φ is true in \mathbf{A} when x_i is interpreted by a_i . If m is an integer, the first-order m -type of a_1, \dots, a_n in \mathbf{A} is the collection of all first-order formulas $\varphi(x_1, \dots, x_n)$ of quantifier rank at most m , up to logical equivalence, for which $\mathbf{A} \models \varphi(a_1, \dots, a_n)$. The monadic second-order m -type of a_1, \dots, a_n in \mathbf{A} is the collection of all monadic second-order formulas $\varphi(x_1, \dots, x_n)$ of quantifier rank at most m , up to logical equivalence, for which $\mathbf{A} \models \varphi(a_1, \dots, a_n)$. In this definition, by quantifier rank of a monadic second-order formula we mean the *total* quantifier rank, which means that we include both first-order and second-order quantifiers in the count. We note that some definitions of monadic second-order type in the literature distinguish between first-order and second-order quantifier rank [14], but we do not need this refinement.

2.4 Preservation under extensions and minimal models

Let \mathcal{C} be a class of finite σ -structures that is closed under induced substructures. Let φ be a first-order sentence. We say that φ is *preserved under extensions* on \mathcal{C} if whenever \mathbf{A} and \mathbf{B} are structures in \mathcal{C} such that \mathbf{B} is an extension of \mathbf{A} , then $\mathbf{A} \models \varphi$ implies $\mathbf{B} \models \varphi$. We say that \mathbf{A} is a *minimal model* of φ if $\mathbf{A} \models \varphi$ and every proper induced substructure \mathbf{A}' of \mathbf{A} is such that $\mathbf{A}' \not\models \varphi$. The following Lemma states that the existential sentences are precisely those that have finitely many minimal models. Its proof is part of folklore:

Lemma 2.1 *Let \mathcal{C} be a class of finite σ -structures that is closed under induced substructures. Let φ be a first-order sentence that is preserved under extensions on \mathcal{C} . Then, the following are equivalent:*

1. φ is equivalent on \mathcal{C} to an existential sentence,
2. φ has finitely many minimal models in \mathcal{C} .

In the rest of the paper, we use several times the implication from 2. to 1. Just for completeness, this is proved by taking the disjunction of the existential closure of the atomic types of each of the finitely many minimal models.

3 Acyclic Structures

We begin with the simple case of acyclic structures, by which we mean structures whose Gaifman graph is acyclic. We show that any class of such structures satisfying certain closure properties admits the extension preservation property. Note, for structures whose Gaifman graphs are acyclic, there is no loss of generality in assuming that the vocabulary σ consists of unary and binary relations only.

The proof makes heavy use of a technique known as Hanf locality, for which we provide the necessary background first.

Let \mathbf{A} and \mathbf{B} be structures. If $\mathbf{a} \in A^m, \mathbf{b} \in B^m$ are m -tuples, we write $(\mathbf{A}, \mathbf{a}) \equiv^m (\mathbf{B}, \mathbf{b})$ to denote that the first-order m -type of \mathbf{a} in \mathbf{A} is the same as the first-order m -type of \mathbf{b} in \mathbf{B} . In particular $\mathbf{A} \equiv^m \mathbf{B}$ denotes that the structures \mathbf{A} and \mathbf{B} are not distinguished by any first-order sentence of quantifier rank m or less. The equivalence relation \equiv^m is characterized by Ehrenfeucht-Fraïssé games (see, for instance, [5]). These can be used to show that the relation is a congruence with respect to disjoint union with a multiplicity threshold of m . A precise statement of this useful property is given in the following lemma. We write $\mathbf{A} \oplus \mathbf{B}$ to denote the disjoint union of the structures \mathbf{A} and \mathbf{B} and $n\mathbf{A}$ to denote the disjoint union of n copies of \mathbf{A} (see [5, Prop. 2.3.10]).

Lemma 3.1 *Let $\mathbf{A}_1, \mathbf{A}_2, \mathbf{B}_1,$ and \mathbf{B}_2 be structures, and let m, n and n' be integers.*

1. *If $\mathbf{A}_1 \equiv^m \mathbf{B}_1$ and $\mathbf{A}_2 \equiv^m \mathbf{B}_2$ then $\mathbf{A}_1 \oplus \mathbf{A}_2 \equiv^m \mathbf{B}_1 \oplus \mathbf{B}_2$.*
2. *If $n, n' \geq m$ and $\mathbf{A} \equiv^m \mathbf{B}$ then $n\mathbf{A} \equiv^m n'\mathbf{B}$.*

A useful sufficient condition for the \equiv^m equivalence of structures is provided by Hanf locality. The *Hanf type* of radius r of a structure \mathbf{A} is the multiset of isomorphism types of r -neighborhoods of elements in \mathbf{A} . We say that two structures \mathbf{A} and \mathbf{B} are *Hanf equivalent* with radius r and threshold q , written $\mathbf{A} \simeq_{r,q} \mathbf{B}$, if, for every $a \in A$, either the number of occurrences of the isomorphism type of $N_r^{\mathbf{A}}(a)$ in the Hanf type of \mathbf{A} is the same as that in the Hanf type of \mathbf{B} , or it is at least q , and conversely for every element $b \in B$. This allows us to state the following (for a proof see for instance [14, Theo. 4.24]):

Theorem 3.2 (Hanf Locality) *For every vocabulary σ and every m there are r and q such that for any pair of σ -structures \mathbf{A} and \mathbf{B} if $\mathbf{A} \simeq_{r,q} \mathbf{B}$ then $\mathbf{A} \equiv^m \mathbf{B}$.*

As a first step towards the main result of this section, we establish a useful property of connected, acyclic structures with degree at most 2. These are structures whose Gaifman graph consists of a simple path. This is a very restricted class of structures. In particular, any class of such structures is wide, in the sense of Theorem 4.3 below. Thus, on any class of such structures, the extension preservation property holds by virtue of Theorem 4.3. However, the property in Lemma 3.3 provides a useful stepping stone in our proof for all acyclic structures and also serves as a useful warm-up for the proof in Section 4.

Lemma 3.3 *For every vocabulary σ and every $m > 0$ there is a p such that if \mathbf{A} is a σ -structure whose Gaifman graph is connected, acyclic and of degree at most 2 and $|A| > p$, then there is a disjoint extension \mathbf{B} of \mathbf{A} and a proper substructure \mathbf{A}' of \mathbf{A} such that $\mathbf{A}' \equiv^m \mathbf{B}$.*

Proof: Given m , let r and q be obtained from Theorem 3.2. We first consider the $2r$ -neighborhoods of elements of \mathbf{A} , returning later to consider r -neighborhoods when we wish to establish the Hanf types of the structures we construct. Clearly, the $2r$ -neighborhood type of an element determines its r -neighborhood type. Also note that among σ -structures whose degree is bounded (by 2) there are only finitely many isomorphism types of $2r$ -neighborhoods. Let n be the number of such types, let $l = 2r(n + 1) + 1$ and let $p = nl(q + l)$.

For t the isomorphism type of a $2r$ -neighborhood in \mathbf{A} , we say that t is *frequent* if there are at least $q + l$ elements in \mathbf{A} whose type is t . Since there are at most n types, the number of occurrences of elements whose type is not frequent is less than $n(q + l)$. Thus, in a path of length p there must be a sequence of l consecutive elements of frequent type. Let a_1, \dots, a_l be such a sequence. Among the $2rn + 1$ central elements of the sequence $a_{r+1}, \dots, a_{(2n+1)r+1}$ there must be a pair a_i, a_j which have the same type and such that $j - i > 2r$. Let \mathbf{C} be the substructure of \mathbf{A} generated by the elements a_{i+1}, \dots, a_j . We define \mathbf{B} to be $\mathbf{A} \oplus \mathbf{C}$ and \mathbf{A}' to be the substructure of \mathbf{A} generated by $\mathbf{A} \setminus \mathbf{C}$.

Our aim is to prove $\mathbf{A}' \equiv^m \mathbf{B}$ by showing that $\mathbf{A}' \simeq_{r,q} \mathbf{B}$. We do this by considering how the Hanf type changes in going from \mathbf{A} to \mathbf{A}' and also how it changes in going from \mathbf{A} to \mathbf{B} . So, for t the isomorphism type of an r -neighborhood in \mathbf{A} , we say that t is *rare* if there are fewer than q elements in \mathbf{A} whose type is t . Write D for the set of elements $\{a_{i-r+1}, \dots, a_i, a_{j+1}, \dots, a_{j+r}\}$. That is, D consists of the r elements that occur immediately before \mathbf{C} and the r elements that occur immediately after \mathbf{C} in the sequence a_1, \dots, a_l . For any element $a \in \mathbf{A}$ that is not in $C \cup D$, $N_r^{\mathbf{A}}(a) = N_r^{\mathbf{A}'}(a)$. For any element a of $C \cup D$, the multiplicity of the type t of $N_r^{\mathbf{A}}(a)$ may decrease in going from \mathbf{A} to \mathbf{A}' . However, t occurs at least $q + l$ times in \mathbf{A} and this multiplicity cannot decrease by more than l as $|C \cup D| \leq l$. Thus, t is not rare in \mathbf{A}' . Clearly the elements of D may have types in \mathbf{A}' that are different to their types in \mathbf{A} and therefore the multiplicities of these types may increase.

Similarly, for any element $a \in A$, $N_r^{\mathbf{A}}(a) = N_r^{\mathbf{B}}(a)$, thus any type t that occurs in \mathbf{A} has at least the same multiplicity in \mathbf{B} . Let $C' = \{a'_{i+1}, \dots, a'_j\}$ denote the elements in the new disjoint copy of \mathbf{C} . If $a'_k \in C'$ is such that $i + r < k \leq j - r$ then the r -neighborhood of a'_k is isomorphic to $N_r^{\mathbf{A}}(a_k)$. Since the type of a_k is frequent, adding to its multiplicity is not significant. Thus, we only need to consider the types of the elements in $D' = \{a'_{i+1}, \dots, a'_{i+r+1}, a'_{j-r+1}, \dots, a'_j\}$. For these elements, the types of their r -neighborhoods in \mathbf{B} may be new and result in an increase of the multiplicities of these types over their occurrences in \mathbf{A} . Thus, to establish our result that $\mathbf{A}' \simeq_{r,q} \mathbf{B}$ it suffices to show that there is a bijection $f : D \rightarrow D'$ such that for all $a \in D$, $N_r^{\mathbf{A}'}(a) \cong N_r^{\mathbf{B}}(f(a))$. By construction, there is an isomorphism $h : N_{2r}^{\mathbf{A}}(a_i) \rightarrow N_{2r}^{\mathbf{A}}(a_j)$ and therefore in particular, for $-r \leq k \leq r$, $N_r^{\mathbf{A}}(a_{i+k}) \cong N_r^{\mathbf{A}}(a_{j+k})$. We can now define the desired bijection f as follows: for $1 \leq k \leq r$, $f(a_{i-k+1}) = a'_{j-k+1}$ and $f(a_{j+k}) = a'_{i+k}$. \square

We now use the above lemma to obtain a similar result for connected acyclic structures without a bound on the degree. This is done by reducing the case of general degree to those with degree at most 2 by means of an appropriate translation. For the vocabulary σ , there are only finitely many first-order m -types of σ -structures. Let τ_1, \dots, τ_n be an enumeration of the possible types of a in \mathbf{A} , where \mathbf{A} is a connected, acyclic structure and $a \in A$. We refer to a as *the distinguished element* of (\mathbf{A}, a) . We define a new vocabulary σ' which has the same binary relations as σ and a unary relation T_i for each τ_i .

Let \mathbf{A} be a σ' -structure that is connected, acyclic and of degree at most 2 with the property that for each $a \in A$ there is a unique i such that $T_i(a)$. We construct from \mathbf{A} a σ -structure $\tilde{\mathbf{A}}$ as follows: each element

$a \in A$ with $T_i(a)$ is replaced by a structure \mathbf{T}_a of type τ_i . Moreover, for any binary relation R , $(b, c) \in R^{\tilde{\mathbf{A}}}$ if, and only if, *either* b and c are in the same structure \mathbf{T}_a and $(b, c) \in R^{\mathbf{T}_a}$ or b is the distinguished element of \mathbf{T}_a , c is the distinguished element of $\mathbf{T}_{a'}$ and $(a, a') \in R^{\mathbf{A}}$. The structure $\tilde{\mathbf{A}}$ is not uniquely determined by \mathbf{A} as there are, in general, many structures of type τ_i . However, the following lemma is easily established along the lines of Lemma 3.1.

Lemma 3.4 *Let \mathbf{A} and \mathbf{B} be structures connected, acyclic structures of degree at most 2 with the property that for each element there is a unique i such that T_i holds, and let m be an integer. If $\mathbf{A} \equiv^m \mathbf{B}$ then $\tilde{\mathbf{A}} \equiv^m \tilde{\mathbf{B}}$.*

We will call a structure of the form $\tilde{\mathbf{A}}$ a σ -companion of \mathbf{A} .

Lemma 3.5 *For every vocabulary σ and every $m > 0$ there is a p such that if \mathbf{A} is a structure whose Gaifman graph is connected and acyclic and which contains a path with more than p elements, then there is a disjoint extension \mathbf{B} of \mathbf{A} and a proper substructure \mathbf{A}' of \mathbf{A} such that $\mathbf{A}' \equiv^m \mathbf{B}$.*

Proof: Let σ' be the vocabulary, as above, with a unary relation for each m -type of σ -structures and let p be as in Lemma 3.3 for the vocabulary σ' . Let a_1, \dots, a_p be the path of length p in \mathbf{A} . For each i , let S_i be the set of elements that are reachable (in the Gaifman graph of \mathbf{A}) from a_i without going through a_j for any $j \neq i$ and let \mathbf{S}_i be the substructure generated by S_i . We define the σ' -structure $s\mathbf{A}$ as follows. The universe of $s\mathbf{A}$ is $\{a_1, \dots, a_p\}$; $T_k(a_i)$ holds if, and only if, a_i has type τ_k in \mathbf{S}_i ; and $(a_i, a_j) \in R^{s\mathbf{A}}$ if, and only if, $(a_i, a_j) \in R^{\mathbf{A}}$. Then, it is easily seen that \mathbf{A} is a σ -companion of $s\mathbf{A}$ (which is defined, since the Gaifman graph of \mathbf{A} is acyclic).

Let $s\mathbf{A}'$ and $s\mathbf{B}$ be the structures obtained from $s\mathbf{A}$ by Lemma 3.3. We obtain \mathbf{A}' as a σ -companion of $s\mathbf{A}'$ by replacing each element a_i by the structure (\mathbf{S}_i, a_i) . This ensures that \mathbf{A}' is a substructure of \mathbf{A} . Similarly, we obtain \mathbf{B} as a σ -companion of $s\mathbf{B}$, ensuring that \mathbf{B} is a disjoint extension of \mathbf{A} . Since $s\mathbf{A}' \equiv^m s\mathbf{B}$ by Lemma 3.3, we also have $\mathbf{A}' \equiv^m \mathbf{B}$ by Lemma 3.4. \square

Note that in both Lemmas 3.3 and 3.5 \mathbf{B} is not only a disjoint extension of \mathbf{A} , it is in fact the disjoint union of \mathbf{A} with a substructure of \mathbf{A} .

In order to prove the main theorem of this section, we need one further composition property of acyclic structures along the lines of the properties in Lemma 3.1. In order to define it, we introduce some further notation. Given an acyclic structure \mathbf{A} and an element $a \in A$, for every neighbor b of a let S_b be the set of elements in A which are reachable from b (in the Gaifman graph) without going through a and let $\text{tp}_a(b)$ denote the first-order m -type of b in \mathbf{S}_b . We define the *child-type* of b with respect to a to be the pair $(\text{at}(a, b), \text{tp}_a(b))$ where $\text{at}(a, b)$ is the atomic type of the pair (a, b) . Finally, we define the *child-type* of an element a , written $\text{ct}^{\mathbf{A}}(a)$, to be the multiset of the child-types of its neighbors with respect to a . Write $(\mathbf{A}, a) \sim_m (\mathbf{B}, b)$ to denote that every type *either* occurs the same number of times in $\text{ct}^{\mathbf{A}}(a)$ as it does in $\text{ct}^{\mathbf{B}}(b)$ or occurs at least m times in both. The following lemma is now a straightforward application of games.

Lemma 3.6 *If $(\mathbf{A}, a) \sim_m (\mathbf{B}, b)$ then $(\mathbf{A}, a) \equiv_m (\mathbf{B}, b)$.*

We are now ready for the main theorem of this section.

Theorem 3.7 *Let \mathcal{C} be a class of acyclic finite structures, closed under substructures and disjoint unions. Then, on \mathcal{C} , every first-order sentence that is preserved under extensions is equivalent to an existential sentence.*

Proof: Let φ be a such a sentence of quantifier rank m . We aim to show that there is an N such that if \mathbf{A} in \mathcal{C} is a model of φ with more than N elements then \mathbf{A} is not minimal. Let p be as in Lemma 3.5, n be the number of distinct first-order m -types of connected structures in \mathcal{C} and q be the number of distinct types of the form $(\text{at}(a, b), \text{tp}_a(b))$ where a and b are neighbors in a structure in \mathcal{C} . Let $N = mn(qm)^p$.

Now, suppose \mathbf{A} is a minimal model of φ in \mathcal{C} with more than N elements. We consider three cases.

Case 1: \mathbf{A} has more than mn distinct connected components. Then there must be some collection of more than m such components that have the same first-order m -type. Consider the structure \mathbf{A}' obtained by removing one of these components. By Lemma 3.1 $\mathbf{A}' \equiv^m \mathbf{A}$, contradicting the minimality of \mathbf{A} .

If \mathbf{A} has mn or fewer connected components, one of these components must have at least $(qm)^p$ elements. Call this component \mathbf{C} the large component.

Case 2: The large component of \mathbf{A} has a node of degree greater than qm . Call this node a . The type $\text{ct}^{\mathbf{A}}(a)$ must contain a type with more than m occurrences. Let b be a neighbor of a that has this child-type with respect to a . Let \mathbf{A}' be the substructure of \mathbf{A} obtained by removing all elements in S_b . By Lemma 3.6, we have $\mathbf{A}' \equiv^m \mathbf{A}$ again contradicting the minimality of \mathbf{A} .

Case 3: If \mathbf{C} does not contain a node of degree greater than qm , it must contain a path of length p . Thus, by Lemma 3.5, there is a proper substructure \mathbf{C}' of \mathbf{C} and a disjoint extension \mathbf{D} of \mathbf{C} such that $\mathbf{C}' \equiv^m \mathbf{D}$. Let \mathbf{A}' be the structure obtained from \mathbf{A} by replacing \mathbf{C} by \mathbf{C}' and \mathbf{B} be the structure obtained from \mathbf{A} by replacing \mathbf{C} by \mathbf{D} . Then, by Lemma 3.1, $\mathbf{A}' \equiv^m \mathbf{B}$. Note also that \mathbf{A}' and \mathbf{B} are in \mathcal{C} since it is closed under substructures and disjoint unions. Since φ is preserved under extensions on \mathcal{C} , $\mathbf{B} \models \varphi$ and hence $\mathbf{A}' \models \varphi$ again contradicting the minimality of \mathbf{A} . \square

4 Wide Structures

This section will focus on classes of structures that are *wide*, meaning that large enough structures contain many points that are pairwise far apart from each other. It was shown in [1] that the homomorphism preservation theorem holds on any wide class of structures. Here we aim to establish the analogous result for the extension preservation property.

Definition 4.1 *A set of elements B in a σ -structure \mathbf{A} is d -scattered if for every pair of distinct $a, b \in B$ we have $N_d^{\mathbf{A}}(a) \cap N_d^{\mathbf{A}}(b) = \emptyset$.*

We say that a class of finite σ -structures \mathcal{C} is wide if for every d and m there exists an N such that every structure in \mathcal{C} of size at least N contains a d -scattered set of size m .

The canonical example of a wide class of structures is the the collection of all structures of degree bounded by a constant. More generally, any class of structures whose maximum degree is bounded by $n^{o(1)}$, where n is the number of elements of the structure, is wide.

Unfortunately, the techniques and arguments of Section 3 based on Hanf locality will not be enough for our current purpose. Instead, we will have to resort to Gaifman locality, for which we provide the necessary background first.

For every integer $r \geq 0$, let $\delta(x, y) \leq r$ denote the first-order formula expressing that the distance between x and y in the Gaifman graph is at most r . Let $\delta(x, y) > r$ denote the negation of this formula. Note that the quantifier rank of $\delta(x, y) \leq r$ is bounded by r . A *basic local sentence* is a sentence of the form

$$(\exists x_1) \cdots (\exists x_n) \left(\bigwedge_{i \neq j} \delta(x_i, x_j) > 2r \wedge \bigwedge_i \psi^{N_r(x_i)}(x_i) \right), \quad (4.1)$$

where ψ is a first-order formula with one free variable. Here, $\psi^{N_r(x_i)}(x_i)$ stands for the relativization of ψ to $N_r(x_i)$; that is, the subformulas of ψ of the form $(\exists x)(\theta)$ are replaced by $(\exists x)(\delta(x, x_i) \leq r \wedge \theta)$, and the subformulas of the form $(\forall x)(\theta)$ are replaced by $(\forall x)(\delta(x, x_i) \leq r \rightarrow \theta)$. The *locality radius* of a basic local sentence is r . Its *width* is n . Its *local quantifier rank* is the quantifier rank of ψ . We will use the fact that basic local sentences are preserved under disjoint extensions. Note, however, that they may not be preserved under plain extensions since in that case the neighborhoods can grow.

The main result about basic local sentences is that they form a building block for first-order logic. This is known as Gaifman's Theorem (for a proof, see, for example, [5, Theo. 2.5.1]):

Theorem 4.2 (Gaifman Locality) *Every first-order sentence is equivalent to a Boolean combination of basic local sentences.*

The following theorem contains the main technical construction of the paper.

Theorem 4.3 *Let \mathcal{C} be a class of finite σ -structures that is wide and closed under substructures and disjoint unions. Then, on \mathcal{C} , every first-order sentence that is preserved under extensions is equivalent to an existential sentence.*

Proof: Let φ be a first-order sentence that is preserved under extensions on \mathcal{C} . By Gaifman's Theorem we may assume that $\varphi = \bigvee_{i \in I} \tau_i$, with

$$\tau_i = \bigwedge_{j \in J_i} \theta_j^i \wedge \bigwedge_{k \in K_i} \neg \theta_k^i, \quad (4.2)$$

where each θ_h^i is a basic local sentence. Now we define a list of parameters that we need in the proof (the reader may skip this list now and use it to look up the values when they are needed):

- r is the maximum of the locality radii of all θ_h^i ,
- s is the sum of all widths of all θ_h^i ,
- m is the maximum of the local quantifier ranks of all θ_h^i ,
- ℓ is the number of disjuncts in φ , so $\ell = |I|$.
- $n = (\ell + 2)s$,
- $M = m + 3r + 3$,
- $d = 2(r + 1)(\ell + 1)s + 6r + 2$,
- q is the number of monadic second-order M -types with one free variable,
- N is such that every structure in \mathcal{C} of size at least N contains a $(4dq + 2r + 1)$ -scattered set of size $(n - 1)q + s + \ell s + 1$.

Our goal is to show that the minimal models of φ have size less than N . Suppose on the contrary that \mathbf{A} is a minimal model of φ of size at least N . We define the *type* of a point $a \in A$ to be its monadic second-order M -type in $\mathbf{A} \cap N_d^{\mathbf{A}}(a)$. In other words, the type of a is the collection of all monadic second-order formulas $\psi(x)$ of quantifier-rank at most M , up to logical equivalence, for which $\mathbf{A} \cap N_d^{\mathbf{A}}(a) \models \psi(a)$. We say that a realizes its type. The reason we consider monadic second-order types, instead of first-order types, will become clear later in the proof. Let t_1, \dots, t_q be all possible types. We need a couple of definitions. Let C be a subset of A and t a type. We say that t is *covered by C* if for all realizations a of t we have $N_d^{\mathbf{A}}(a) \subseteq C$. We say that t is *free over C* if there are at least n realizations a_1, \dots, a_n of t such that $N_d^{\mathbf{A}}(a_i)$ and $N_d^{\mathbf{A}}(a_j)$ are pairwise disjoint and do not intersect C .

Claim 4.4 *There exists a radius $e \leq 2dq$ and a set D of at most $(n-1)q$ points in A such that each type is either covered by $N_e^{\mathbf{A}}(D)$ or free over $N_e^{\mathbf{A}}(D)$.*

Proof: We define D and e inductively. Let $D_0 = \emptyset$ and $e_0 = 0$. Suppose now that D_i and e_i are already defined. Let $C = N_{e_i}^{\mathbf{A}}(D_i)$. If all types are either covered by C or free over C , then let $D = D_i$ and $e = e_i$. Otherwise, let j be minimal such that type t_j is neither covered by C nor free over C . We define a set E inductively as follows. Let $E_0 = \emptyset$. Suppose now that E_t is already defined. If there is no realization of t_j outside $N_{2d}^{\mathbf{A}}(C \cup E_t)$, then let $E = E_t$ and we are done with the construction of E . Otherwise, let a_{t+1} be a realization of t_j outside $N_{2d}^{\mathbf{A}}(C \cup E_t)$ and let $E_{t+1} = E_t \cup \{a_{t+1}\}$. Note that this iteration cannot continue beyond $n-1$ steps since otherwise t_j would be free over C . This means that the iteration stops, and when it does $|E| \leq n-1$ and t_j is covered by any set that contains $N_{2d}^{\mathbf{A}}(C \cup E)$, and in particular by $N_{e_i+2d}^{\mathbf{A}}(D_i \cup E)$. Let $D_{i+1} = D_i \cup E$ and $e_{i+1} = e_i + 2d$. The construction stops after at most q steps because at each step one new type is covered and remains covered for the rest of the construction. This shows that $|D| \leq (n-1)q$ and $e \leq 2dq$, which proves the claim. \square

In the following, we fix e and D according to Claim 4.4. We say that a type t is *frequent* if it is not covered by $N_e^{\mathbf{A}}(D)$. Otherwise we say that t is *rare*.

We shall build a finite sequence of sets $S_0 \subseteq S_1 \subseteq \dots \subseteq S_p \subseteq A$, with $p \leq \ell$, so that the last set S_p in the sequence will be such that the substructure of \mathbf{A} induced by S_p is a proper substructure of \mathbf{A} that satisfies φ . This will contradict the minimality of \mathbf{A} and will prove the theorem. The sequence S_i is constructed inductively together with a second sequence of sets $C_0 \subseteq C_1 \subseteq \dots \subseteq C_p \subseteq A$ called the *centers*, and a sequence of sets of indices $I_0 \subseteq I_1 \subseteq \dots \subseteq I_p \subseteq I$ (recall that φ is the disjunction of the formulas τ_i from (4.2) for $i \in I$). Moreover, the following conditions will be preserved by the inductive construction for every $i < p$.

- (a) $S_i \subseteq N_r^{\mathbf{A}}(C_i)$.
- (b) $|C_i| \leq is$.
- (c) No disjoint extension of $\mathbf{A} \cap S_i$ satisfies $\bigvee_{j \in I_i} \tau_j$.
- (d) $N_e^{\mathbf{A}}(D)$ and $N_d^{\mathbf{A}}(C_i)$ are disjoint.
- (e) $|I_i| = i$.

Observe that it is a direct consequence of property (d) that the type of each $a \in C_i$ is frequent.

Let $S_0 = C_0 = I_0 = \emptyset$, and let us assume that S_i, C_i and I_i have already been defined with the properties above. We construct S_{i+1}, C_{i+1} , and I_{i+1} . Let \mathbf{B} be the disjoint union of \mathbf{A} with a copy of $\mathbf{A} \cap S_i$.

$$\text{Since } \mathbf{B} \text{ is an extension of } \mathbf{A}, \text{ it satisfies } \varphi. \tag{4.3}$$

Therefore, there exists an $i' \in I$ such that \mathbf{B} satisfies $\tau_{i'}$. By property (c), since the extension is disjoint, we know that $i' \notin I_i$. Let $I_{i+1} = I_i \cup \{i'\}$. For the rest of the proof, the index i' will be fixed so we drop any reference to it. For example, we will write τ instead of $\tau_{i'}$ and θ_h instead of $\theta_{i'}$. Recall that

$$\tau = \bigwedge_{j \in J} \theta_j \wedge \bigwedge_{k \in K} \neg \theta_k.$$

Since \mathbf{B} satisfies τ , in particular it satisfies the positive requirements: $\mathbf{B} \models \bigwedge_{j \in J} \theta_j$. Let W_j be a minimal set of witnesses in \mathbf{B} for the outermost existential quantifiers in θ_j , and let $W = \bigcup_{j \in J} W_j$. We

have $|W| \leq s$. Some of these witnesses may be in \mathbf{A} and some may be in the new copy of $\mathbf{A} \cap S_i$ in \mathbf{B} . Let $W_A \cup W_B = W$ be such a partition, with W_A being the witnesses in A . The following claim shows that W_A can be chosen far from C_i . This will be needed later.

Claim 4.5 *There is a set W of witnesses such that $N_{r+1}^{\mathbf{A}}(C_i) \cap N_r^{\mathbf{A}}(W_A) = \emptyset$.*

Proof: Fix a set W of witnesses so that the number of points b in W_A for which $N_{r+1}^{\mathbf{A}}(C_i)$ and $N_r^{\mathbf{A}}(b)$ are not disjoint is minimal. Suppose that this number is not zero, and let $b \in W_A$ with $N_{r+1}^{\mathbf{A}}(C_i) \cap N_r^{\mathbf{A}}(b) \neq \emptyset$. Let $a \in C_i$ be such that $N_{r+1}^{\mathbf{A}}(a) \cap N_r^{\mathbf{A}}(b) \neq \emptyset$. Then $N_r^{\mathbf{A}}(b) \subseteq N_{3r+1}^{\mathbf{A}}(a) \subseteq N_d^{\mathbf{A}}(a)$. By property (d), the type t of a is frequent. So let a' be a realization of t such that $N_{r+1}^{\mathbf{A}}(W \cup C_i)$ and $N_{3r+1}^{\mathbf{A}}(a')$ are disjoint. Such an a' exists because t is frequent and thus, by Claim 4.4, is free over $N_e^{\mathbf{A}}(D)$ and thus has

$$n > (\ell + 1)s \geq |W \cup C_i|$$

realizations whose d -neighborhoods are pairwise disjoint and disjoint from $N_e^{\mathbf{A}}(D)$.

The goal now is to find a b' such that $N_r^{\mathbf{A}}(b') \subseteq N_{3r+1}^{\mathbf{A}}(a') \subseteq N_d^{\mathbf{A}}(a')$ and such that b and b' have the same first-order m -type on $\mathbf{A} \cap N_r^{\mathbf{A}}(b)$ and $\mathbf{A} \cap N_r^{\mathbf{A}}(b')$ respectively. If we achieve this, then b' can replace b as a witness in W_A , and since $N_{r+1}^{\mathbf{A}}(W \cup C_i)$ and $N_{3r+1}^{\mathbf{A}}(a')$ are disjoint, so are $N_{r+1}^{\mathbf{A}}(C_i)$ and $N_r^{\mathbf{A}}(b')$. This will contradict the minimality of W .

In order to find b' as above, let T be the first-order m -type of b on $\mathbf{A} \cap N_r^{\mathbf{A}}(b)$, and let $\xi(x)$ be the following first-order formula:

$$(\exists y) \left((\forall z) (\delta(y, z) \leq r \rightarrow \delta(x, z) \leq 3r + 1) \wedge \bigwedge_{\chi \in T} \chi^{N_r(y)}(y) \right).$$

Note that the conjunction is finite because the first-order m -type T contains finitely many formulas up to logical equivalence, and that the quantifier rank of this formula is bounded by $3r + 3 + m \leq M$. Also $N_d^{\mathbf{A}}(a) \models \xi(a)$ because b can serve as a witness for y . Therefore, since a and a' have the same monadic second-order M -type and hence the same first-order M -type in $N_d^{\mathbf{A}}(a)$ and $N_d^{\mathbf{A}}(a')$, also $N_d^{\mathbf{A}}(a') \models \xi(a')$. Note here that we are not yet using the full power of monadic second-order type, only the fact that it contains the first-order type as a subset. Let b' be the witness to y in $N_d^{\mathbf{A}}(a') \models \xi(a')$, completing the proof. \square

In the following, we fix a set W of witnesses such that $N_{r+1}^{\mathbf{A}}(C_i) \cap N_r^{\mathbf{A}}(W_A) = \emptyset$. We let \mathbf{C} be the substructure of \mathbf{A} induced by $N_e^{\mathbf{A}}(D) \cup N_r^{\mathbf{A}}(W_A) \cup S_i$. We claim that \mathbf{C} satisfies the positive requirements of τ :

Claim 4.6 *\mathbf{C} is a substructure of \mathbf{A} such that $\mathbf{C} \models \bigwedge_{j \in J} \theta_j$.*

Proof: It is obvious that \mathbf{C} is a substructure of \mathbf{A} . The point, however, is that \mathbf{C} is in fact the disjoint union of the substructure induced by $N_e^{\mathbf{A}}(D) \cup N_r^{\mathbf{A}}(W_A)$ and the substructure induced by S_i . This is because $S_i \subseteq N_r^{\mathbf{A}}(C_i)$ and $N_{r+1}^{\mathbf{A}}(C_i)$ is disjoint from $N_e^{\mathbf{A}}(D)$ by property (d) and also disjoint from $N_r^{\mathbf{A}}(W_A)$ by Claim 4.5. It follows that the witnesses from \mathbf{B} in W_B can also be found in \mathbf{C} . Obviously, also the witnesses from \mathbf{B} in W_A can be found in \mathbf{C} . This proves that \mathbf{C} satisfies the positive requirements of τ . \square

Consider φ on \mathbf{C} . If \mathbf{C} is a model of φ , let $S_p = N_e^{\mathbf{A}}(D) \cup N_r^{\mathbf{A}}(W_A) \cup S_i$ and we are done. Notice that \mathbf{C} is a proper substructure of \mathbf{A} because \mathbf{A} contains $(n-1)q + s + \ell s + 1$ points that are $(4dq + 2r + 1)$ -scattered, but $S_p \subseteq N_{2dq+r}^{\mathbf{A}}(D \cup W_A \cup C_i)$ and

$$|D \cup W_A \cup C_i| \leq (n-1)q + s + \ell s.$$

If \mathbf{C} is not a model of φ it cannot satisfy τ . However, by Claim 4.6, \mathbf{C} satisfies the positive requirements $\bigwedge_{j \in J} \theta_j$. Therefore, \mathbf{C} does not satisfy $\bigwedge_{k \in K} \neg \theta_k$. Let $k \in K$ such that $\mathbf{C} \models \theta_k$. In the next claim we find a substructure of \mathbf{A} that extends $\mathbf{A} \cap S_i$ and forces all its disjoint extensions to satisfy θ_k .

Claim 4.7 *There exist $C_{i+1} \supseteq C_i$ and $S_{i+1} \supseteq S_i$ as required by conditions (a)–(d).*

Proof: Suppose that

$$\theta_k = (\exists x_1) \dots (\exists x_{s'}) \left(\bigwedge_{i \neq j} \delta(x_i, x_j) > 2r' \wedge \bigwedge_i \psi^{N_{r'}(x_i)}(x_i) \right)$$

for some $r' \leq r, s' \leq s$, and some formula ψ of quantifier rank $m' \leq m$. Without loss of generality we may assume that $m' = m$, and in order to simplify the notation, we will assume that $r' = r$ and $s' = s$. It will suffice to replace r by r' and s by s' in the appropriate places.

We have $\mathbf{C} \models \theta_k$. Let $V = \{a_1, \dots, a_s\}$ be a set of witnesses for the outermost existential quantifiers in θ_k . Then $N_r^{\mathbf{C}}(a_i) \cap N_r^{\mathbf{C}}(a_j) = \emptyset$ for all $i \neq j$ and $\mathbf{C} \cap N_r^{\mathbf{C}}(a_i) \models \psi^{N_r(x_i)}(a_i)$ for all i . Necessarily, the type t of some $a \in V$ is frequent. Otherwise $N_r^{\mathbf{A}}(V) \subseteq N_e^{\mathbf{A}}(D) \subseteq A$, so $\mathbf{A} \models \theta_k$, and thus $\mathbf{B} \models \theta_k$, because \mathbf{B} is a disjoint extension of \mathbf{A} . However, this is impossible because $\mathbf{B} \models \tau$.

So let $a \in V$ have frequent type t . Let Z be a set of s realizations of t such that

- (i) $N_d^{\mathbf{A}}(b) \cap N_d^{\mathbf{A}}(b') = \emptyset$ for every pair of distinct $b, b' \in Z$,
- (ii) $N_e^{\mathbf{A}}(D) \cap N_d^{\mathbf{A}}(Z) = \emptyset$,
- (iii) $N_{r+1}^{\mathbf{A}}(C_i) \cap N_r^{\mathbf{A}}(Z) = \emptyset$.

Such a set Z exists because t is frequent, $n = (\ell + 2)s$, and $|C_i| \leq \ell s$ by property (b).

Now, let $F = N_r^{\mathbf{C}}(a)$. Remember that $\mathbf{C} \cap F \models \psi^{N_r(x)}(a)$. As $F \subseteq N_r^{\mathbf{A}}(a)$, it follows that $\mathbf{A} \cap F \models \psi^{N_r(x)}(a)$. Let X be a set variable, and let $\psi^{N_r(x) \cap X}(X, x)$ denote the simultaneous relativization of $\psi(x)$ to $N_r(x)$ and X , that is, the formula obtained from ψ by replacing each subformula of the form $(\exists z)\xi$ by $(\exists z)(\delta(x, z) \leq r \wedge X(z) \wedge \xi)$, and similarly for universally quantified subformulas. Observe that the quantifier rank of $\psi^{N_r(x) \cap X}(X, x)$ is at most $m + r \leq M - 1$, where we take r has an upper bound for the quantifier rank of the formula expressing $\delta(x, z) \leq r$. Moreover, $\mathbf{A} \models \psi^{N_r(x) \cap X}(F, a)$ and hence $\mathbf{A} \models \exists X \psi^{N_r(x) \cap X}(a)$.

Next comes the place where we use the full power of monadic second-order types. Since every $b \in Z$ has the same monadic second-order M -type as a , we have $\mathbf{A} \models \exists X \psi^{N_r(x_i) \cap X}(b)$. Thus there is a set $F_b \subseteq N_r^{\mathbf{A}}(b)$ such that $\mathbf{A} \models \psi^{N_r(x_i) \cap X}(F_b, b)$. It follows that

$$\mathbf{A} \cap F_b \models \psi^{N_r(x)}(b).$$

Define $C_{i+1} = C_i \cup Z$ and

$$S_{i+1} = S_i \cup \bigcup_{b \in Z} F_b.$$

Let us prove that C_{i+1} and S_{i+1} satisfy the properties (a), (b), (c), and (d). Property (a) is clear since $F_b \subseteq N_r^{\mathbf{A}}(b)$. For property (b) we have $|C_{i+1}| = |C_i| + s \leq (i + 1)s$. Property (d) is satisfied by (ii) in our choice of Z .

Finally, for property (c) we argue as follows. First note that $\mathbf{A} \cap S_{i+1}$ is a disjoint extension of $\mathbf{A} \cap S_i$ because $N_{r+1}^{\mathbf{A}}(C_i) \cap N_r^{\mathbf{A}}(Z) = \emptyset$ by (iii) and $S_i \subseteq N_r^{\mathbf{A}}(C_i)$ by a. Therefore, no disjoint extension of

$\mathbf{A} \cap S_{i+1}$ satisfies τ_j for any $j \in I_i$. It remains to show that no disjoint extension of $\mathbf{A} \cap S_{i+1}$ satisfies τ . However, this is clear from the construction because every disjoint extension of $\mathbf{A} \cap S_{i+1}$ contains witnesses for the outermost existential quantifiers in θ_k ; namely, the elements of the set Z . Suppose that $Z = \{b_1, \dots, b_s\}$. Note that b_i have pairwise distance $> 2r$ by (i), and we have $\mathbf{A} \cap S_{i+1} \models \psi^{N_r(x_i)}(b_i)$, because $N^{\mathbf{A} \cap S_{i+1}}(b_i) = F_{b_i}$ and $\mathbf{A} \cap F_{b_i} \models \psi^{N_r(x_i)}(b_i)$. \square

Note that I_{i+1} is constructed to satisfy property (e) as well. This completes the definition of the inductive construction. All it remains to see is that the construction stops in at most ℓ steps. Because suppose for contradiction that we have constructed S_ℓ, C_ℓ , and I_ℓ satisfying (a)–(e). Then $I_\ell = I$ by (e), and by (c), no disjoint extension of $\mathbf{A} \cap S_\ell$ satisfies $\varphi = \bigvee_{i \in I} \tau_i$. However,

$$\begin{aligned} & \text{the disjoint union } \mathbf{B} \text{ of } \mathbf{A} \cap S_\ell \text{ with } \mathbf{A} \text{ is an extension of } \mathbf{A} \text{ and} \\ & \text{hence does satisfy } \varphi. \end{aligned} \tag{4.4}$$

This is a contradiction. \square

As a direct application of Theorem 4.3, let us consider the class \mathcal{D}_r of all finite σ -structures of degree bounded by r . This class is both wide and closed under substructures and disjoint unions. To see the wideness, note that when the degree of every node is at most r , for any element a , $N_d(a)$ contains at most r^d elements. Thus, if a structure has size greater than $m(r^d)$, it must contain a d -scattered set of m elements.

Theorem 4.8 *Let r be an integer. Then, on \mathcal{D}_r , every first-order sentence that is preserved under extensions is equivalent to an existential sentence.*

In the following section we show how the argument of Theorem 4.3 can be extended, in some cases, to classes of structures that are *almost wide*.

5 Bounded Treewidth Structures

The class of structures of bounded degree provide a canonical example of a wide class. On the other hand, acyclic structures (which we considered in Section 3) are not wide. Indeed, in an arbitrarily large tree of height 1 all pairs of nodes are at distance at most 2 from each other and there is therefore no large d -scattered set for any $d > 2$, yet the tree may be arbitrarily large. However, in such a structure, the removal of just one element, the root, creates a large scattered set. This motivates the definition below.

Definition 5.1 *A class of finite σ -structures \mathcal{C} is almost wide if there is a k such that for every d and m there exists an N such that every structure \mathbf{A} of size at least N in \mathcal{C} contains a set B with at most k elements such that $\mathbf{A} - B$ contains a d -scattered set of size m .*

It was shown in [1] that the homomorphism preservation property holds for almost wide classes of structures which are closed under substructures and disjoint unions. It was also established that any class of graphs that excludes a minor is almost wide.

It is not the case that the extension preservation property holds for all almost wide classes. This is shown in the next section, where we show, in particular, that it fails for the class of planar graphs. It turns out that the requirement that an almost wide class be closed under substructures and disjoint unions is not sufficient to guarantee the extension preservation property. Nevertheless, closure under unions *over a set of bottlenecks* suffices, a notion we make more precise later. In this section we show that this yields the preservation under extensions property for some particularly interesting almost wide classes. To be precise, we show that the property holds for the class \mathcal{T}_k of all finite σ -structures of treewidth less than k . In other words, we aim to prove the following result:

Theorem 5.2 *Let k be an integer. Then, on \mathcal{T}_k , every first-order sentence that is preserved under extensions is equivalent to an existential sentence.*

The proof of this result requires three ingredients. The first ingredient is a generalization of the disjoint union operation on structures by allowing some non-empty intersection. Let \mathbf{A} and \mathbf{B} be σ -structures, and let $C \subseteq A \cap B$ be such that $\mathbf{A} \cap C = \mathbf{B} \cap C$. The *union of \mathbf{A} and \mathbf{B} through C* , denoted by $\mathbf{A} \oplus_C \mathbf{B}$ is a new σ -structure defined as follows. The universe of $\mathbf{D} = \mathbf{A} \oplus_C \mathbf{B}$ is $A' \cup B' \cup C$, where A' is a disjoint copy of $A - C$ and B' is a disjoint copy of $B - C$. The relations of \mathbf{D} are defined in the obvious way: If a_1, \dots, a_r are points in A and a'_1, \dots, a'_r are the corresponding points in $A' \cup C$, then $(a'_1, \dots, a'_r) \in R^{\mathbf{D}}$ if and only if $(a_1, \dots, a_r) \in R^{\mathbf{A}}$. Similarly, if b_1, \dots, b_r are points in B and b'_1, \dots, b'_r are the corresponding points in $B' \cup C$, then $(b'_1, \dots, b'_r) \in R^{\mathbf{D}}$ if and only if $(b_1, \dots, b_r) \in R^{\mathbf{B}}$. Observe that this construction is precisely the disjoint union of \mathbf{A} and \mathbf{B} when $C = \emptyset$.

The next lemma is a straightforward generalization of the obvious fact that \mathcal{T}_k is closed under disjoint unions. The Lemma states, roughly, that \mathcal{T}_k is closed under unions through *subsets of bags of tree-decompositions*.

Lemma 5.3 *Let k be an integer. Let \mathbf{A} and \mathbf{B} be two σ -structures, let $C \subseteq A \cap B$ be such that $\mathbf{A} \cap C = \mathbf{B} \cap C$, and let (T, L) and (T', L') be tree-decompositions of width k of \mathbf{A} and \mathbf{B} , respectively. Then, if there exists nodes $u \in T$ and $u' \in T'$ such that $C \subseteq L(u) \cap L'(u')$, then the union of \mathbf{A} and \mathbf{B} through C has treewidth at most k .*

Proof: The tree-decomposition of the union is $(T'', L \cup L')$, where $T'' = T \cup T'$ with a new tree edge joining u and u' . \square

The second ingredient is the fact that the class of structures of treewidth less than k is almost wide, in the sense of Definition 5.1 that there exists a small set of vertices whose removal produces a large scattered set. Such a set is henceforth called a *bottleneck*. This was proved in [1] but here we state the stronger claim that the bottleneck can be found in a single bag of a tree-decomposition. The proof is the same as in [1] and is sketched here for completeness.

Lemma 5.4 *For every k , and for every d and m , there exists an N such that if \mathbf{A} is a σ -structure of size at least N and (T, L) is a tree-decomposition of \mathbf{A} of width k , then there exist $u \in T$ and $K \subseteq L(u)$ such that $\mathbf{A} - K$ contains a d -scattered set of size m .*

Proof sketch: Let $p = (m - 1)(2d + 1) + 1$, $M = k!(p - 1)^k$, and $N = k(m - 1)^M$ and suppose that \mathbf{A} is a structure with more than N elements. Let (T, L) be a tree decomposition of \mathbf{A} such that $L(u)$ has size at most k for all $u \in T$. Note that T has size at least $N/k + 1$. Furthermore, suppose T has a node u of degree at least m . But then it is easy to see that taking $K = L(u)$ gives a graph with at least m distinct connected components and therefore a scattered set of size m . On the other hand, if every node of T has degree less than m , then T must have a path with length greater than M . By the Sunflower Lemma of Erdős and Rado [7], it follows that we can find p distinct nodes $u_1, \dots, u_p \in T$ and a set $K \subseteq A$ such that for $i \neq j$, $L(u_i) \cap L(u_j) = K$. It can then be shown that $\mathbf{A} - K$ must contain a d -scattered set of size m . \square

The third ingredient in the proof is a first-order bi-interpretation between an almost wide structure and a wide structure. From now on we focus on graphs; the construction extends easily to the general case. Let $P_1, \dots, P_k, Q_1, \dots, Q_k$ be unary relation symbols and $\sigma = \{E, P_1, \dots, P_k, Q_1, \dots, Q_k\}$. For every graph $\mathbf{G} = (V, E^{\mathbf{G}})$ and every tuple $\mathbf{a} = (a_1, \dots, a_k) \in V^k$ we define a σ -structure $\mathbf{A} = \mathbf{A}(\mathbf{G}, \mathbf{a})$ as follows:

1. $A = V$,
2. $E^{\mathbf{A}} = E^{\mathbf{G}} - \{(a, b) \in A^2 : \{a, b\} \cap \{a_1, \dots, a_k\} \neq \emptyset\}$,
3. $P_i^{\mathbf{A}} = \{a_i\}$,
4. $Q_i^{\mathbf{A}} = \{b \in A : (a_i, b) \in E^{\mathbf{G}}\}$.

Let us call a σ -structure \mathbf{A} *derived* if $E^{\mathbf{A}}$ is a symmetric and anti-reflexive binary relation, and there are elements $a_1, \dots, a_k \in A$ such that $P_i^{\mathbf{A}} = \{a_i\}$ for $1 \leq i \leq k$ and a_i is isolated in the graph underlying \mathbf{A} ; that is, for $1 \leq i \leq k$ there is no b such that $(a_i, b) \in E^{\mathbf{A}}$. Note that for every derived structure \mathbf{A} there is a unique graph $\mathbf{G}(\mathbf{A})$ and a unique k -tuple $\mathbf{a}(\mathbf{A})$ of vertices of $\mathbf{G}(\mathbf{A})$ such that

$$\mathbf{A} = \mathbf{A}(\mathbf{G}(\mathbf{A}), \mathbf{a}(\mathbf{A})).$$

The point behind the construction of $\mathbf{A} = \mathbf{A}(\mathbf{G}, \mathbf{a})$ is that if $K = \{a_1, \dots, a_k\}$ is a bottleneck of \mathbf{G} in the sense that $\mathbf{G} - K$ contains a large scattered set, then \mathbf{A} itself has a large scattered set and maintains all the information needed to reconstruct \mathbf{G} . Indeed, $\mathbf{G}(\mathbf{A})$ is first-order interpretable in \mathbf{A} , and thus we get the following lemma:

Lemma 5.5 *For every first-order sentence φ of vocabulary $\{E\}$ there is a sentence $\tilde{\varphi}$ of vocabulary σ such that for all σ -structures \mathbf{A} we have:*

1. *If $\mathbf{A} \models \tilde{\varphi}$ then \mathbf{A} is derived.*
2. *If \mathbf{A} is derived, then $\mathbf{A} \models \tilde{\varphi}$ if and only if $\mathbf{G}(\mathbf{A}) \models \varphi$.*

This follows at once from a standard result on syntactical interpretations (cf., for example, Theorem VIII.2.2 of [6]).

Equipped with these three ingredients we are ready for the main argument.

Proof: [of Theorem 5.2] Let φ be a first-order sentence that is preserved under extensions in \mathcal{T}_k . It suffices to show that φ has finitely many minimal models. Let $\mathbf{G} = (V, E^{\mathbf{G}})$ be a graph in \mathcal{T}_k that is a minimal model of φ . Suppose for contradiction that \mathbf{G} is very large. Let (T, L) be a tree-decomposition of width k of \mathbf{G} , and let $K = \{b_1, \dots, b_k\} \subseteq V$ be a bottleneck; that is, a set such that $\mathbf{G} - K$ contains a large scattered set. By Lemma 5.4 we may assume that $K \subseteq L(u)$ for some $u \in T$. Let $\mathbf{A} = \mathbf{A}(\mathbf{G}, \mathbf{b})$, where $\mathbf{b} = (b_1, \dots, b_k)$. The idea is to work with \mathbf{A} and $\tilde{\varphi}$ instead of \mathbf{G} and φ and proceed as in the proof of Theorem 4.3. The difference is that $\tilde{\varphi}$ is *not* preserved under extensions. However, preservation under extensions is used only twice in the proof of Section 4 (in (4.3) and (4.4)), both times to prove that the disjoint union \mathbf{B} of the structure \mathbf{A} with $\mathbf{A} \cap S_i$ is a model of φ . Claim 5.6 shows that in both cases, \mathbf{B} is a model of $\tilde{\varphi}$.

Claim 5.6 *Let $C \subseteq A$ such that the type of each $a \in C$ is frequent. Let $S \subseteq N_r(C)$ and let \mathbf{B} be the disjoint union of \mathbf{A} with a disjoint copy of $\mathbf{A} \cap S$. Then \mathbf{B} is derived, \mathbf{G} is an induced subgraph of $\mathbf{G}(\mathbf{B})$, and $\mathbf{G}(\mathbf{B})$ belongs to \mathcal{T}_k .*

Proof: The bottleneck points are not in C as their type is not frequent and therefore not in $N_r(C)$ as they are isolated in \mathbf{A} . Thus, Note that \mathbf{B} is derived because the bottleneck points are not in S . Let $\mathbf{H} = \mathbf{G}(\mathbf{B})$. Clearly, \mathbf{G} is an induced subgraph of \mathbf{H} . Thus all we have to prove is that \mathbf{H} belongs to \mathcal{T}_k . Let $\mathbf{A}' = \mathbf{A} \cap (S \cup K)$, where K is the bottleneck of \mathbf{G} . Again, \mathbf{A}' is derived. Let $\mathbf{G}' = \mathbf{G}(\mathbf{A}')$.

Clearly, \mathbf{G}' is an induced subgraph of \mathbf{G} . In particular, \mathbf{G}' is in \mathcal{T}_k so it has a tree-decomposition of width k . More importantly, since $K \subseteq L(u)$, we can assume as well that K is a subset of some bag of the tree-decomposition of \mathbf{G}' . These two facts together imply that the union of \mathbf{G} and \mathbf{G}' through K , which is precisely \mathbf{H} , is in \mathcal{T}_k by Lemma 5.3. \square

This shows then that the \mathbf{B} in (4.3) and (4.4) is a model of $\tilde{\varphi}$. The proof proceeds until we construct a structure \mathbf{C} that satisfies $\tilde{\varphi}$ and is a proper substructure of \mathbf{A} . We claim that \mathbf{C} is derived. This is because all bottleneck points have rare type, so they belong to D . Let $\mathbf{H} = \mathbf{G}(\mathbf{C})$. Note now that \mathbf{H} is the union of two subgraphs \mathbf{G}_1 and \mathbf{G}_2 of \mathbf{G} through the bottleneck K . Again K is a subset of a bag of the tree-decompositions of \mathbf{G}_1 and \mathbf{G}_2 , so \mathbf{H} belongs to \mathcal{T}_k by Lemma 5.3. Moreover \mathbf{H} is a proper induced subgraph of \mathbf{G} and $\mathbf{H} \models \varphi$ by Lemma 5.5. This contradicts the minimality of \mathbf{G} , which concludes the proof. \square

This completes the proof of Theorem 5.2. Note that this does not imply that the existential preservation theorem holds on all classes of bounded treewidth. Indeed, we show in the next section that it fails, in particular, for the class of planar graphs of treewidth 4.

6 Counterexample for Planar Graphs

The aim of this section is to show that the preservation under extensions property fails on the class of planar graphs. Let us focus first on the class of planar graphs whose vertices are colored either black or white. Later we show how to remove the colors. The vocabulary contains a binary relation symbol E for the edge relation, and a unary relation symbol P for the color. Let φ be the following first-order sentence:

$$\begin{aligned}\varphi &= (\exists x)(\exists y)\left(x \neq y \wedge P(x) \wedge P(y) \wedge (\varphi_1(x, y) \rightarrow \varphi_2(x, y))\right). \\ \varphi_1(x, y) &= (\forall z)\left(z \neq x \wedge z \neq y \rightarrow \neg P(z) \wedge E(x, z) \wedge E(y, z)\right). \\ \varphi_2(x, y) &= (\forall u)\left(u \neq x \wedge u \neq y \rightarrow (\exists v)(\exists w)\left(v \neq w \wedge \neg P(v) \wedge \neg P(w) \wedge E(u, v) \wedge E(u, w)\right)\right).\end{aligned}$$

We claim that φ is preserved under extensions on the class of black/white-colored planar graphs. Before we prove this we need a technical gadget. For every $n \geq 3$, let \mathbf{G}_n be the black/white-colored planar graph displayed in Figure 1, where the number of black vertices is exactly n .

It is not hard to see that \mathbf{G}_n does not have any planar proper extension in which all other vertices are adjacent to both white vertices. Let us state this:

Lemma 6.1 *Let $n \geq 3$, and let \mathbf{H} be a black/white-colored planar graph that is a proper extension of \mathbf{G}_n . Then, no vertex in $\mathbf{H} - \mathbf{G}_n$ is adjacent to both white vertices in \mathbf{G}_n .*

Proof: Let u be a vertex in $\mathbf{H} - \mathbf{G}_n$. Suppose that u is adjacent to both white vertices in \mathbf{G}_n . Then \mathbf{H} contains a \mathbf{K}_5 minor by contracting one of the edges connecting u to a white vertex in \mathbf{G}_n , and by contracting all but two of the edges in \mathbf{G}_n that do not have a white endpoint. This contradicts the planarity of \mathbf{H} . \square

Now we are ready to show that φ is preserved under extensions on the class of black/white-colored planar graphs:




Figure 1: G_9

Lemma 6.2 *Let G and H be black/white-colored planar graphs such that H is a proper extension of G . If G is a model of φ , so is H .*

Proof: Suppose that G is a model of φ , so let a and b be two different white vertices in G . If $G \not\models \varphi_1(a, b)$, then clearly $H \not\models \varphi_1(a, b)$ because G is an induced substructure of H . In this case, H is also a model of φ and we are done. Otherwise, since $G \models \varphi$ and $G \models \varphi_1(a, b)$, we have $G \models \varphi_2(a, b)$. This means that every vertex in $G - \{a, b\}$ is adjacent to at least two other black vertices. It follows that G contains some G_n as a (not necessarily induced) subgraph with a and b as white vertices. Here $n \geq 3$. It follows then by Lemma 6.1 that some vertex in $H - G_n$ fails to be connected to both a and b . But then $H \not\models \varphi_1(a, b)$ so H is a model of φ again. \square

To complete the argument we need to show that φ is not equivalent to an existential sentence on the class of black/white-colored graphs.

Lemma 6.3 *There is no existential sentence equivalent to φ on all black/white-colored planar graphs.*

Proof: By virtue of Lemma 2.1, we only need to show that φ has infinitely many minimal models among planar graphs. It is easily seen that for all n , G_n is a minimal model of φ . Indeed, if we remove at least one of the white vertices from G_n , we would not have witnesses for the two outermost existential quantifiers in φ , and if we remove at least one of the black vertices then φ_1 remains true while φ_2 fails. \square

This shows that the preservation-under-extensions property fails for the class of black/white-colored planar graphs. Removing the colors is easy. It suffices to replace each occurrence of $P(x)$ by a formula $\varphi_3(x)$ stating that x is attached to a 4×4 -grid that is otherwise disconnected from the rest of the graph. One point to note is that a node without such a grid attached in a graph G may have a grid in an extension of G . However, this would mean that φ_1 would fail in the extension and thus φ would necessarily be true. Thus, the formula is still preserved under extensions. This shows then that the preservation-under-extensions property fails for the class of planar graphs.

Note further that for any n , the treewidth of G_n is at most 4. This implies that the existential preservation theorem fails, even for the class \mathcal{P} of planar graphs of treewidth at most 4. Indeed, the sentence φ is preserved under extensions on \mathcal{P} since it is preserved under extensions on all planar graphs. However, φ still has infinitely many minimal models in this class as each G_n is in \mathcal{P} .

7 Conclusions

We have established the extension preservation theorem for a number of interesting classes of finite structures. These include all wide classes—such as any class of structures of bounded degree—and some almost wide classes, such as \mathcal{T}_k , the class of all structures of treewidth less than k . The situation for the extension preservation theorem is quite different to that established for the homomorphism preservation theorem in [1]. In particular, the former fails on the class of planar while the latter holds on all classes that exclude a graph minor. Indeed, the methods of proof used here to establish the extension preservation property are rather different to those used in [1]. It should also be noted that Rossman [15] has recently established that the homomorphism preservation theorem holds on the class of all finite structures, which contrasts with the known failure of the extension preservation theorem.

A number of recent results in finite model theory [1, 2, 3, 8, 10, 11] indicate that classes of structure such as trees or structures of bounded treewidth, planar graphs and graphs of bounded genus, graphs with excluded minors, and graphs of bounded degree are well behaved in various ways related to their first-order model theory (in a broad sense). So far, no serious attempt has been made to identify general criteria connecting the different results. The locality of first-order logic always appears to play a crucial role, and the notion of *wideness* formally introduced here seems to be a good structural counterpart. But there is more to it than this simple observation; for example, the result of this paper holds on graphs of bounded degree, but not on planar graphs, whereas for the algorithmic results of [3] it is the other way round. The order invariance result of [2] has so far eluded all efforts to extend it beyond acyclic structures.

References

- [1] A. Atserias, A. Dawar, and Ph. G. Kolaitis. On preservation under homomorphisms and unions of conjunctive queries. *Journal of the ACM*, 53:208–237, 2006.
- [2] M. Benedikt and L. Segoufin. Towards a characterization of order-invariant queries over tame structures. In C.-H.L. Ong, editor, *Proceedings of the 19th International Workshop on Computer Science Logic*, volume 3634 of *Lecture Notes in Computer Science*, pages 276–291. Springer-Verlag, 2005.
- [3] A. Dawar, M. Grohe, S. Kreutzer, and N. Schweikardt. Approximation schemes for first-order definable optimisation problems. In *Proc. 21st IEEE Symp. on Logic in Computer Science*, pages 411–420, 2006..
- [4] R.G. Downey and M.R. Fellows. *Parametrized Complexity*. Springer-Verlag, 1999.
- [5] H-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 2nd edition, 1999.
- [6] H-D. Ebbinghaus, J. Flum, and W. Thomas. *Mathematical Logic*. Springer, 2nd edition, 1994.
- [7] P. Erdős and R. Rado. Intersection theorems for systems of sets. *J. London Mathematical Society*, 35:85–90, 1960.
- [8] J. Flum and M. Grohe. Fixed-parameter tractability, definability, and model checking. *SIAM Journal on Computing*, 31(1):113–145, 2001.
- [9] J. Flum, M Frick, and M. Grohe. Query evaluation via tree-decompositions. *Journal of the ACM*, 49:716–752, 2002.

- [10] M. Frick and M. Grohe. Deciding first-order properties of locally tree-decomposable structures. *Journal of the ACM*, 48:1184–1206, 2001.
- [11] M. Grohe and Gy. Turán. Learnability and definability in trees and similar structures. *Theory of Computing Systems*, 37(1):193–220, 2004.
- [12] Yu. Gurevich. Toward logic tailored for computational complexity. In M. Richter et al., editors, *Computation and Proof Theory*, pages 175–216. Springer Lecture Notes in Mathematics, 1984.
- [13] W. Hodges. *Model Theory*. Cambridge University Press, 1993.
- [14] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
- [15] B. Rossman. Existential positive types and preservation under homomorphisms. In *20th IEEE Symposium on Logic in Computer Science*, pages 467–476, 2005.
- [16] W. W. Tait. A counterexample to a conjecture of Scott and Suppes. *Journal of Symbolic Logic*, 24:15–16, 1959.