Deciding first-order properties of nowhere dense graphs

Martin Grohe
RWTH Aachen University
grohe@informatik.rwth-aachen.de

Stephan Kreutzer
Technical University Berlin
stephan.kreutzer@tu-berlin.de

Sebastian Siebertz
Technical University Berlin
sebastian.siebertz@tu-berlin.de

Abstract

Nowhere dense graph classes, introduced by Nešetřil and Ossona de Mendez [29], form a large variety of classes of "sparse graphs" including the class of planar graphs, actually all classes with excluded minors, and also bounded degree graphs and graph classes of bounded expansion.

We show that deciding properties of graphs definable in first-order logic is fixed-parameter tractable on nowhere dense graph classes. At least for graph classes closed under taking subgraphs, this result is optimal: it was known before that for all classes $\mathcal{C}$ of graphs closed under taking subgraphs, if deciding first-order properties of graphs in $\mathcal{C}$ is fixed-parameter tractable, then $\mathcal{C}$ must be nowhere dense (under a reasonable complexity theoretic assumption).

As a by-product, we give an algorithmic construction of sparse neighbourhood covers for nowhere dense graphs. This extends and improves previous constructions of neighbourhood covers for graph classes with excluded minors. At the same time, our construction is considerably simpler than those.

Our proofs are based on a new game-theoretic characterisation of nowhere dense graphs that allows for a recursive version of locality-based algorithms on these classes. On the logical side, we prove a "rank-preserving" version of Gaifman's locality theorem.

1 Introduction

Algorithmic meta theorems attempt to explain and unify algorithmic results by proving tractability not only for individual problems, but for whole classes of problems. These classes are typically defined in terms of logic. The meaning of "tractability" varies; for example, it may be linear or polynomial time solvability, fixed-parameter tractability, or polynomial time approximability to some ratio. The prototypical example of an algorithmic meta theorem is Courcelle's Theorem [4], stating that all properties of graphs of bounded tree-width that are definable in monadic second-order logic are decidable in linear time. Another well-known example is Papadimitriou and Yannakakis's [31] result that all optimisation problems in the class MAXSNP, which is defined in terms of a fragment of existential second-order logic, admit constant-ratio polynomial time approximation algorithms. By now, there is a rich literature on algorithmic meta theorems (see, for example, [2, 5, 6, 7, 8, 14, 18, 25, 26, 32] and the surveys [20, 22, 24]). While the main motivation for proving such meta theorems may be to understand the "essence" and the scope of certain algorithmic techniques by abstracting from problem-specific details, sometimes meta theorems are also crucial for obtaining new algorithmic results. A recent example is the quadratic time algorithm for a structural decomposition of graphs with excluded minors from [21], which builds on Courcelle's Theorem in an essential way. Furthermore, meta theorems often give a quick and easy way to see that certain problems can be solved efficiently (in principle), for example in linear time on graphs of bounded tree-width. Once this has been established, a
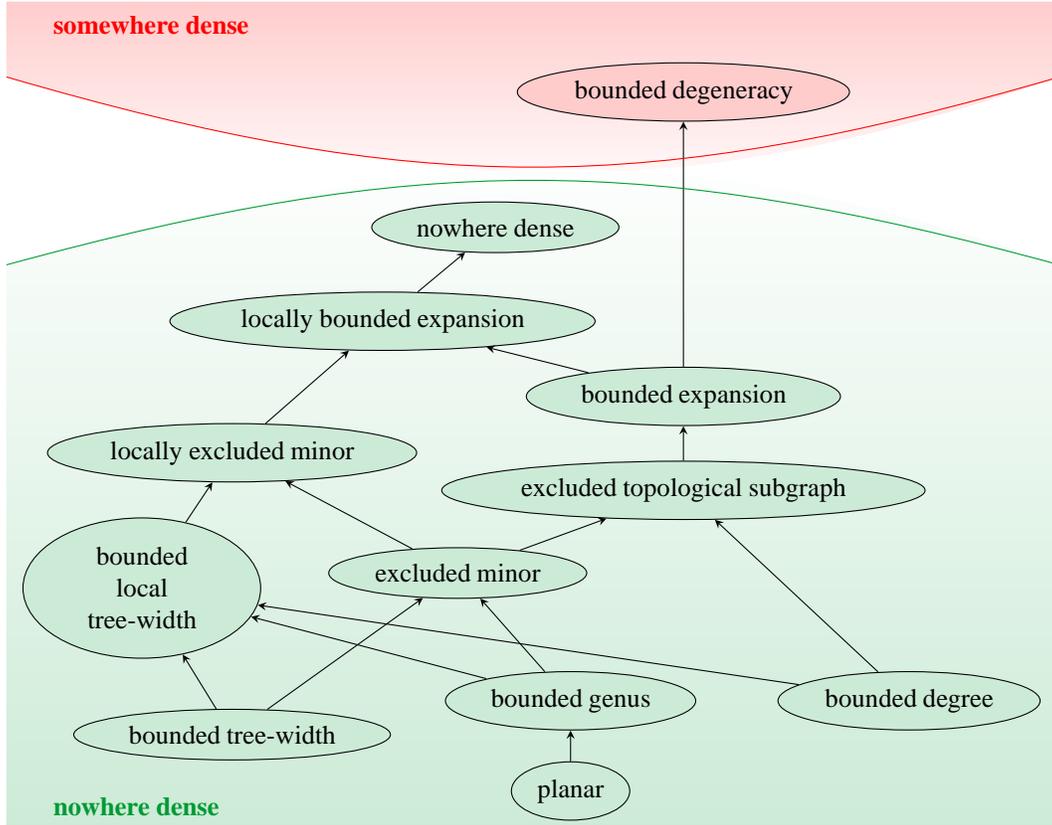
Figure 1: Sparse graph classes

problem specific analysis may yield better algorithms – even though implementations of, for instance, Courcelle's theorem have shown that the direct application of meta theorems can yield competitive algorithms for common problems such as the dominating set problem (see [27]).

In this paper, we prove a new meta theorem for first-order logic on nowhere dense classes of graphs. These classes were introduced by Nešetřil and Ossona de Mendez [28, 29] as a formalisation of classes of "sparse" graphs. All familiar examples of sparse graph classes, like the class of planar graphs, classes of bounded tree-width, classes of bounded degree, and indeed all classes with excluded topological subgraphs are nowhere dense. Figure 1 shows the containment relations between these and other sparse graph classes.[1] "Nowhere density" turns out to be a very robust concept with several seemingly unrelated natural characterisations (see [28, 29]). Furthermore, Nešetřil and Ossona de Mendez [29] established a clear-cut dichotomy between nowhere dense and somewhere dense graph classes. The exact definition of nowhere dense graph classes is technical and we defer it to Section 3.

Theorem 1.1 *For every nowhere dense class $\mathcal{C}$ and every $\varepsilon > 0$, every property of graphs definable in first-order logic can be decided in time $\mathcal{O}(n^{1+\varepsilon})$ on $\mathcal{C}$.*

---

[1]Notably, classes of bounded average degree or bounded degeneracy are not necessarily nowhere dense. To be precise: for every $k \geq 2$ the class of all graphs of degeneracy at most k is somewhere dense. This is reasonable, because every graph can be turned into a graph of degeneracy 2 by simply subdividing every edge once. Recall that a graph has *degeneracy* at most d if every subgraph has a vertex of degree at most d. Degeneracy at most d implies that the graph and all its subgraphs have average degree at most $2d$ and hence have a linear number of edges. Contrarily, graphs in nowhere dense classes can have an edge density of $n^{1+\varepsilon}$ and are therefore not necessarily degenerate.

In particular, deciding first-order properties is fixed-parameter tractable on nowhere dense graph classes.[2] Deciding first-order properties of arbitrary graphs is known to be complete for the parameterized complexity class AW[∗] and thus unlikely to be fixed-parameter tractable [12].

Nešetřil and Ossona de Mendez [28] already proved that deciding properties definable in existential first-order logic is fixed-parameter tractable on nowhere dense graphs. Dawar and Kreutzer [9] showed that dominating set (parameterized by the size of the solution) is fixed-parameter tractable on nowhere dense graphs. Our theorem implies new fixed-parameter tractability results on nowhere dense graphs for many other standard parameterized problems, for example, connected dominating set and digraph kernel (both parameterized by the size of the solution), Steiner tree (parameterized by the size of the tree) and circuit satisfiability (parameterized by the depth of the circuit and the Hamming weight of the solution). The last result requires the generalisation of our theorem from graphs to arbitrary relational structures, which is straightforward.

Our theorem can be seen as the culmination of a long line of meta theorems for first order logic. The starting point is Seese's [32] result that first-order properties of bounded degree graphs can be decided in linear time. Frick and Grohe [18] gave linear time algorithms for planar graphs and all apex-minor-free graph classes and $\mathcal{O}(n^{1+\varepsilon})$ algorithms for graphs of bounded local tree-width. Flum and Grohe [16] proved that deciding first-order properties is fixed-parameter tractable on graph classes with excluded minors, and Dawar, Grohe, and Kreutzer [7] extended this to classes of graphs locally excluding a minor. Finally, Dvořák, Král, and Thomas [14] proved that first-order properties can be decided in linear time on graph classes of bounded expansion and in time $\mathcal{O}(n^{1+\varepsilon})$ on classes of locally bounded expansion. All these classes are nowhere dense, and there are nowhere dense classes that do not belong to any of these classes. For example, the class of all graphs whose girth is larger than the maximum degree is nowhere dense, but has unbounded expansion. If to every graph in this class we add one vertex and connect it with all other vertices, we obtain a class of graphs that is still nowhere dense, but does not even have locally bounded expansion. However, what makes our theorem interesting is not primarily that it is yet another extension of the previous results, but that it is optimal for classes $\mathcal{C}$ closed under taking subgraphs: under the standard complexity theoretic assumption FPT $\neq$ W[1], Kreutzer [24] and Dvořák et al. [14] proved that if a class $\mathcal{C}$ closed under taking subgraphs is somewhere dense (that is, not nowhere dense), then deciding first-order properties of graphs in $\mathcal{C}$ is not fixed-parameter tractable. Note that all classes considered in the previous results are closed under taking subgraphs. Hence our result supports the intuition that nowhere dense classes are the natural limit for many algorithmic techniques for sparse graph classes.

Technically, we neither use the structural graph theory underlying [7, 16] nor the quantifier elimination techniques employed by [14]. Our starting point is the locality based technique introduced in [18]. In a nutshell, this technique works as follows. Using Gaifman's theorem, the problem to decide whether a general first-order formula φ is true in a graph can be reduced to testing whether a formula is true in r-neighbourhoods in the graph, where the radius r only depends on φ, and solving a variant of the (distance d) independent set problem. Hence, if $\mathcal{C}$ is a class of graphs where r-neighbourhoods have a simple structure, such as the class of planar graphs or classes of bounded local tree-width, this method gives an easy way for deciding properties definable in first-order logic.

Applying this technique to nowhere dense classes of graphs immediately runs into problems, as r-neighbourhoods in nowhere dense graphs do not necessarily have a simple structure that can be exploited algorithmically. We therefore iterate the locality based approach. Using locality we reduce the first-order model-checking problem to the problems of evaluating formulas in r-neighbourhoods and solving a variant of the independent set problem. We then show that r-neighbourhoods N in nowhere dense graphs can be split by deleting a set W of only a few vertices into smaller neighbourhoods. We apply the locality argument again and transform our formula into formulas to be evaluated in r-neighbourhoods in $N - W$ and solving

---

[2]There is a minor issue regarding non-uniform vs uniform fixed-parameter tractability, see Remark 3.2.

the independent set problem on $N - W$. We show that on nowhere dense classes of graphs this process terminates after a constant number of steps.

The three main steps of our proof, each of which may be of independent interest, are the following.

- An algorithmic construction of *sparse neighbourhood covers* for nowhere dense graphs (Section 6). The parameters are surprisingly good: we can cover all r-neighbourhoods with sets (called *clusters*) of radius $2r$ such that each vertex is contained in $n^{o(1)}$ clusters. For classes of bounded expansion (see Figure 1), we even get such covers where each vertex is only contained in a constant number of clusters. In particular, the small radius of the clusters substantially improves known results for planar graphs and graphs with excluded minors [1, 3], which all have bounded expansion.

- A new characterisation of nowhere dense graph classes in terms of a game, the *Splitter Game* (Section 4). We use this game to formalise the process of localising and splitting described above and showing that it terminates on nowhere dense graphs. It turns out that it only terminates on nowhere dense graphs, thus providing a necessary and sufficient condition for nowhere density.

- A *Rank-Preserving Locality Theorem* (Section 7), strengthening Gaifman's well-known locality theorem for first-order logic by translating first-order formulas into local formulas of the same rank. The key innovation here is a new, discounted rank measure for first-order formulas.

We describe the main algorithm proving Theorem 1.1 in Section 8.

2 Preliminaries

We assume familiarity with basic concepts of graph theory and refer to [10] for background. We denote the set of positive integers by $\mathbb{N}$. For $k \in \mathbb{N}$ we write $[k]$ for the set $\{1, \ldots, k\}$. We will often write $\bar{a}$ for a k-tuple $(a_1, \ldots, a_k)$ and $a \in \bar{a}$ for $a \in \{a_1, \ldots, a_k\}$.

In this section, we will review the necessary background from graph theory and parameterized complexity theory. We will provide some background on logic in Section 7.

Background from graph theory. All graphs in this paper are finite and simple, i.e., they do not have loops or multiple edges between the same pair of vertices. Whenever we speak of a graph we mean an undirected graph and we will explicitly mention when we deal with directed graphs.

If G is a graph then $V(G)$ denotes its set of vertices and $E(G)$ its set of edges. We write $n := |V(G)|$ for the *order* of G.

An *orientation* of G is a directed graph $\vec{G}$ on the same vertex set, which is denoted $V(\vec{G})$, such that for each edge $\{u, v\} \in E(G)$ the set of arcs $E(\vec{G})$ contains exactly one of the arcs (u, v) or (v, u). For $v \in V(\vec{G})$, the set $N^-(v) := \{u : (u, v) \in E(\vec{G})\}$ denotes the *in-neighbours* of v and $N^+(v) := \{w : (v, w) \in E(\vec{G})\}$ denotes the *out-neighbours* of v. The *indegree* $d^-(v)$ of a vertex v is the number in-neighbours of v. We denote the *maximum indegree* of $\vec{G}$ by $\Delta^-(\vec{G})$. For any directed graph $\vec{G}$ we denote the underlying undirected graph by G.

We assume that all graphs are represented by adjacency lists so that the total size of the representation of a graph is linear in the number of edges and vertices. In fact we will often store an orientation $\vec{G}$ of a graph G and use one adjacency list for the in-neighbours and one adjacency list for the out-neighbours of each vertex. This representation allows to check adjacency of vertices in time $\mathcal{O}(\Delta^-(\vec{G}))$.

For a set $X \subseteq V(G)$ we write $G[X]$ for the subgraph of G induced by X and we let $G \setminus X := G[V(G) \setminus X]$. For $k \in \mathbb{N}$, G is k-*degenerate* if for each $X \subseteq V(G)$ the graph $G[X]$ contains a vertex of degree at most k. If a graph G is k-degenerate then G contains at most $k \cdot n$ edges and an orientation $\vec{G}$ of G with $\Delta^-(G) \leq k$ can be computed in time $\mathcal{O}(k \cdot n)$ by a simple greedy algorithm.

The *distance* $\mathrm{dist}^G(u, v)$ between two vertices $u, v \in V(G)$ is the length of a shortest path from u to v if such a path exists and ∞ otherwise. The *radius* $\mathrm{rad}(G)$ of G is $\min_{u \in V(G)} \max_{v \in V(G)} \mathrm{dist}^G(u, v)$. A vertex $u \in V(G)$ such that $\max_{v \in V(G)} \mathrm{dist}^G(u, v) = \mathrm{rad}(G)$ is called a *centre vertex* of G.

By $N_r^G(v)$ we denote the *r-neighbourhood* of v in G, i.e., the set of vertices of distance at most r from v in G. A set $W \subseteq V(G)$ is *r-independent in G* if $\mathrm{dist}^G(u, v) > r$ for all distinct $u, v \in W$. A 1-independent set is simply called independent. A set $W \subseteq V(G)$ is *r-scattered in G* if $N_r^G(u) \cap N_r^G(w) = \emptyset$ for all distinct $u, w \in W$, i.e., if it is $2r$-independent.

A graph H is a *minor* of a graph G, written $H \preceq G$, if H can be obtained from a subgraph of G by contracting edges. Equivalently, H is a minor of G if there is a map that associates with every vertex $v \in V(H)$ a tree $T_v \subseteq G$ such that T_u and T_v are disjoint for $u \neq v$ and whenever there is an edge $\{u, v\} \in E(H)$ there is an edge in G between some node in T_u and some node in T_v. The subgraphs T_v are called *branch sets*.

Let $r \in \mathbb{N}$. H is a *depth-r minor* of G, denoted $H \preceq_r G$, if H is a minor of G and this is witnessed by a collection of branch sets $\{T_v : v \in V(H)\}$, each of which is a tree of radius at most r.

For $s \geq 1$ we denote the complete graph on s vertices by K_s.

Parameterized complexity. The complexity theoretical framework we use in this paper is parameterized complexity theory, see [11, 17]. A *parameterized problem* is a pair (P, χ), where P is a decision problem and χ is a polynomial time computable function that associates with every instance w of P a positive integer, called the *parameter*. The *model-checking problem* for first-order logic on a class $\mathcal{C}$ of graphs is the following decision problem. Given an FO-sentence and a graph $G \in \mathcal{C}$, decide whether G satisfies φ, written $G \models \varphi$. The parameter is $|\varphi|$. We say that the model-checking problem on a class $\mathcal{C}$ is *fixed-parameter tractable*, or in the complexity class FPT, if there is an algorithm that decides on input (G, φ) whether $G \models \varphi$, in time $f(|\varphi|) \cdot |V(G)|^{\mathcal{O}(1)}$ for some computable function $f : \mathbb{N} \to \mathbb{N}$. The model-checking problem for first-order logic on the class of all graphs is known to be complete for the parameterized complexity class AW[*], which is widely believed to strictly contain the class FPT. Thus, it is widely believed that model-checking for first-order logic is not fixed-parameter tractable.

3   Nowhere Dense Classes of Graphs

Nowhere dense classes of graphs were introduced by Nešetřil and Ossona de Mendez [28, 29] as a formalisation of classes of "sparse" graphs.

Definition 3.1 (Nowhere dense classes) *A class $\mathcal{C}$ of graphs is* nowhere dense *if for every r there is a graph H_r such that $H_r \not\preceq_r G$ for all $G \in \mathcal{C}$.*

It is immediate from the definition that if $\mathcal{C}$ excludes a minor then it is nowhere dense. But note that excluding some graph as a depth-r minor is a "local" condition that is much weaker than excluding it "globally" as a minor.

Remark 3.2 We call a class $\mathcal{C}$ *effectively nowhere dense* if there is a computable function f such that $K_{f(r)} \not\preceq_r G$ for all $G \in \mathcal{C}$. All natural nowhere dense classes are effectively nowhere dense, but it is possible to construct artificial classes that are nowhere dense, but not effectively so.

The way Theorem 1.1 is stated in the introduction only asserts that deciding first-order properties of nowhere dense graphs is *non-uniformly* fixed-parameter tractable. That is, for every $\varepsilon > 0$ and every sentence φ of first-order logic there is an algorithm deciding the property defined by φ in time $\mathcal{O}(n^{1+\varepsilon})$. This allows for the algorithms for different sentences to be unrelated. For effectively nowhere dense classes $\mathcal{C}$, we obtain uniform fixed-parameter tractability, that is, a single algorithm that, given an n-vertex graph $G \in \mathcal{C}$, $\varepsilon > 0$ and a sentence φ of first-order logic, decides whether φ holds in G in time $f(|\varphi|, \varepsilon) \cdot n^{1+\varepsilon}$, for some computable function f. $\dashv$

"Nowhere density" turns out to be a very robust concept with several seemingly unrelated natural characterisations (see [28, 29]). We will use several different characterisations, each supporting different algorithmic techniques. In the rest of this section we will recall the required equivalences.

The following characterization relates nowhere density to sparsity, albeit sparsity in the liberal sense that the number of edges of an n-vertex graph is $n^{1+o(1)}$.

Lemma 3.3 (Nešetřil-Ossona de Mendez [29]) *A class $\mathcal{C}$ of graphs is nowhere dense if, and only if, for every $r \in \mathbb{N}$*

$$\limsup_{n \to \infty} \left\{ \frac{\log |E(H)|}{\log |V(H)|} \ \middle| \ H \preceq_r G \text{ with } |V(H)| \geq n, G \in \mathcal{C} \right\} \leq 1. \tag{3.1}$$

Here we take $\frac{\log |E(H)|}{\log |V(H)|}$ to be $-\infty$ if $E(H) = \emptyset$, and we take the supremum to be 0 if the set is empty, that is, if $\mathcal{C}$ contains no graphs of order at least n.

Note that the supremum in (3.1) always exists, because $\frac{\log |E(H)|}{\log |V(H)|} \leq 2$ for all H. The lemma states that, as n gets large, the number of edges in all r-shallow minors of n-vertex graphs in $\mathcal{C}$, is $n^{1+o(1)}$. Thus the graphs in $\mathcal{C}$ are very uniformly sparse: not only the graphs and all their subgraphs are sparse, but even all graphs that can be obtained from subgraphs by "local" contractions are. As a further justification of why nowhere dense classes are inherently interesting as a "limit of sparse graph classes", Nešetřil-Ossona de Mendez proved a trichotomy stating that for all graph classes $\mathcal{C}$, the limit in (3.1) approaches 0 or 1 or 2 as r goes to infinity. This means that if a class $\mathcal{C}$ is not nowhere dense, then in the limit it is really dense.

For our algorithmic purpose, we state the result in a different form which follows immediately from the proof of Lemma 3.3.

Lemma 3.4 *A class $\mathcal{C}$ of graphs is nowhere dense if, and only if, there is a function f such that for every $r \in \mathbb{N}$ and every $\varepsilon > 0$, every depth-r minor H of a graph $G \in \mathcal{C}$ with $n \geq f(r, \varepsilon)$ vertices satisfies $|E(H)| \leq n^{1+\varepsilon}$. Furthermore, $\mathcal{C}$ is effectively nowhere dense if, and only if, the function f is computable.*

We close the section with stating another characterisation of nowhere dense classes that will be used below.

Definition 3.5 (Uniformly quasi-wide classes) *A class $\mathcal{C}$ of graphs is* uniformly quasi-wide *with margin $s : \mathbb{N} \to \mathbb{N}$ and $N : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ if for all $r, k \in \mathbb{N}$, if $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| > N(r, k)$, then there is a set $S \subseteq V(G)$ with $|S| < s(r)$, such that W contains an r-scattered set of size at least k in $G \setminus S$.*

We call $\mathcal{C}$ effectively uniformly quasi-wide if the margins s and N are computable functions.

Lemma 3.6 (Nešetřil-Ossona de Mendez [29]) *A class $\mathcal{C}$ of graphs is (effectively) nowhere dense if, and only if, it is (effectively) uniformly quasi-wide.*

4 Game theoretic characterisation of nowhere dense classes

We now provide a new characterisation of nowhere dense classes in terms of a game.

Definition 4.1 (Splitter game) *Let G be a graph and let $\ell, m, r > 0$. The (ℓ, m, r)-splitter game* on G *is played by two players, "Connector" and "Splitter", as follows. We let $G_0 := G$. In round $i + 1$ of the game, Connector chooses a vertex $v_{i+1} \in V(G_i)$. Then Splitter picks a subset $W_{i+1} \subseteq N_r^{G_i}(v_{i+1})$ of size at most m. We let $G_{i+1} := G_i[N_r^{G_i}(v_{i+1}) \setminus W_{i+1}]$. Splitter wins if $G_{i+1} = \emptyset$. Otherwise the game continues at G_{i+1}. If Splitter has not won after ℓ rounds, then Connector wins.*

A strategy *for Splitter is a function f that associates to every partial play* $(v_1, W_1, \ldots, v_s, W_s)$ *with associated sequence* $G_0, \ldots, G_s$ *of graphs and move* $v_{s+1} \in V(G_s)$ *by Connector a set* $W_{s+1} \subseteq N_r^{G_s}(v_{s+1})$ *of size at most m. A strategy f is a* winning *strategy for Splitter in the (ℓ, m, r)-splitter game on G if Splitter wins every play in which he follows the strategy f. If Splitter has a winning strategy, we say that he* wins *the (ℓ, m, r)-splitter game on G.*

Theorem 4.2 *Let $\mathcal{C}$ be a nowhere dense class of graphs. Then for every $r > 0$ there are $\ell, m > 0$, such that for every $G \in \mathcal{C}$, Splitter wins the (ℓ, m, r)-splitter game on G.*

If $\mathcal{C}$ is effectively nowhere dense, then ℓ and m can be computed from r.

Proof. As $\mathcal{C}$ is nowhere dense, it is also uniformly quasi-wide. Let $s_\mathcal{C}$ and $N_\mathcal{C}$ be the margin of $\mathcal{C}$. Let $r > 0$ and let $\ell := N_\mathcal{C}(r, 2s_\mathcal{C}(r))$ and $m := \ell \cdot (r + 1)$. Note that both ℓ and m only depend on $\mathcal{C}$ and r. We claim that for any $G \in \mathcal{C}$, Splitter wins the (ℓ, m, r)-splitter game on G.

Let $G \in \mathcal{C}$ be a graph. In the (ℓ, m, r)-splitter game on G, Splitter uses the following strategy. In the first round, if Connector chooses $v_1 \in V(G_0)$, where $G_0 := G$, then Splitter chooses $W_1 := \{v_1\}$. Now let $i > 1$ and suppose that $v_1, \ldots, v_i, G_1, \ldots, G_i, W_1, \ldots, W_i$ have already been defined. Suppose Connector chooses $v_{i+1} \in V(G_i)$. We define W_{i+1} as follows. For each $1 \le j \le i$, choose a path $P_{j,i+1}$ in $G_{j-1}[N_r^{G_{j-1}}(v_j)]$ of length at most r connecting v_j and v_{i+1}. Such a path must exist as $v_{i+1} \in V(G_i) \subseteq V(G_j) \subseteq N_r^{G_{j-1}}(v_j)$. We let $W_{i+1} := \bigcup_{1 \le j \le i} V(P_{j,i+1}) \cap N_r^{G_i}(v_{i+1})$. Note that $|W_{i+1}| \le i \cdot (r + 1)$ (the paths have length at most r and hence consist of $r + 1$ vertices). It remains to be shown is that the length of any such play is bounded by ℓ.

Assume towards a contradiction that Connector can play on G for $\ell' = \ell + 1$ rounds. Let $(v_1, \ldots, v_{\ell'}, G_1, \ldots, G_{\ell'}, W_1, \ldots, W_{\ell'})$ be the play. As $\ell' > N_\mathcal{C}(r, 2s_\mathcal{C}(r))$, for $W := \{v_1, \ldots, v_{\ell'}\}$ there is a set $S \subseteq V(G)$ with $|S| < s_\mathcal{C}(r)$, such that W contains an r-scattered set I of size $t := 2s_\mathcal{C}(r)$ in $G \setminus S$. Suppose that $I = \{u_1, \ldots, u_t\}$, where $u_j = v_{i_j}$ for indices $1 \le i_1 < i_2 < \ldots < i_t \le \ell'$.

We now consider the pairs (u_{2j-1}, u_{2j}) for $1 \le j \le s(r)$. By construction, $P_j := P_{i_{2j-1}, i_{2j}}$ is a path of length at most r from u_{2j-1} to u_{2j} in $G_{i_{2j-1}-1}$. Any path P_j must necessarily contain a vertex $s_j \in S$, as otherwise the path would exist in $G \setminus S$, contradicting the fact that I is r-scattered in $G \setminus S$. We claim that for $i \ne j$, $s_i \ne s_j$, but this is not possible, as there are strictly less than $s_\mathcal{C}(r)$ vertices in S. The claim follows easily from the following observation. Assume $i > j$. Then $V(P_j) \cap V(G_{2j-1} \subseteq W_{2j}$, thus $V(P_j) \cap V(G_{2j+1}) = \emptyset$, and $V(P_i) \subseteq V(G_{2j+1}) \subseteq V(G_{2i})$. Thus $V(P_i) \cap V(P_j) = \emptyset$ for $i \ne j$. $\square$

Remark 4.3 In the proof of our main theorem, we will also have to compute Splitter's winning strategy efficiently in the following sense.

Suppose that we are in a play $v_1, W_1, \ldots, v_i, W_i$, *and let* $G_0, G_1, \ldots, G_i$ *be the graphs associated with the play (that is, $G_0 = G$ and $G_{j+1} = G_j[N_r^{G_j}(v_{j+1}) \setminus W_{j+1}]$). For $1 \le j \le i$, let T_j be a breadth-first search tree of depth r in G_{j-1} with root v_j.*

Then, given $v_1, W_1, \ldots, v_i, W_i, v_{i+1}$ *and* $T_1, \ldots, T_j$ *and Connector's move v_{i+1} in round $(i + 1)$, we can compute Splitter's answer W_{i+1} according to her winning strategy in time $\mathcal{O}(ri|V(G_i)| + |E(G_i)|)$.*

To see this, recall that $W_{i+1} := \bigcup_{1 \le j \le i} V(P_{j,i+1}) \cap N_r^{G_i}(v_{i+1})$, where $P_{j,i+1}$ can be any shortest path from v_j to v_{i+1} in G_{j-1}. We choose the path from v_{i+1} to v_j in the tree T_j. We can compute this path in time $\mathcal{O}(r)$ and thus all paths in time $\mathcal{O}(ri)$. We can compute $N_r^{G_i}(v_{i+1})$ in time $\mathcal{O}(|V(G_i)| + |E(G_i)|)$ and the intersection in time $\mathcal{O}(ri|V(G_i)|)$.

Remark 4.4 If Splitter wins the (ℓ, m, r)-splitter game on a graph G, then he also wins if we remove in each step of the game a superset of his chosen set W.

We implicitly use this remark when sometimes in a graph G_i reached after i rounds of the game and after choices v_{i+1}, W_{i+1} in the next round we do not continue the game the graph $G_{i+1} = G_i[N_r^{G_i}(v_{i+1}) \setminus W_{i+1}]$, but in a subgraph of G_{i+1}.

We close the section by observing the converse of Theorem 4.2 and hence show that the splitter game provides another characterisation of nowhere dense classes of graphs.

Theorem 4.5 *Let $\mathcal{C}$ be a class of graphs. If for every $r > 0$ there are $\ell, m > 0$ such that for every graph $G \in \mathcal{C}$, Splitter wins the (ℓ, m, r)-splitter game, then $\mathcal{C}$ is nowhere dense.*

Proof. We show that if $\mathcal{C}$ is not nowhere dense, i.e., $\mathcal{C}$ contains all graphs as depth-r minors at some depth r, then for all $\ell, m > 0$ there is a graph $G \in \mathcal{C}$ such that Connector wins the $(\ell, m, 4r + 1)$-splitter game.

Let $\ell, m > 0$. We choose $G \in \mathcal{C}$ such that G contains the complete graph $K := K_{\ell m+1}$ as a depth-r minor. Connector uses the following strategy to win the $(\ell, m, 4r)$-splitter game. Connector chooses any vertex from the branch set of a vertex of K. The $4r + 1$-neighbourhood of this vertex contains the branch sets of all vertices of K. Splitter removes any m vertices. We actually allow him to remove the complete branch sets of all m vertices he chose. In round 2 we may thus assume to find the complete graph $K_{(\ell-1)m+1}$ as a depth-r minor and continue to play in this way until in round ℓ at least the branch set of a single vertex remains. □

5 Independent Sets in Nowhere Dense Classes of Graphs

In this section we use the splitter game to show that the DISTANCE INDEPENDENT SET problem, which is NP-complete in general, is fixed-parameter tractable on nowhere dense classes of graphs. This will be used later in the proof of our main theorem but is also of independent interest. Recall from Section 2 that, for $r \geq 0$, a set of vertices in a graph is *r-independent* if their mutual distance is greater than r.

Theorem 5.1 *Let $\mathcal{C}$ be a nowhere dense class of graphs. There is a function f such that for every $\varepsilon > 0$ the following problem can be solved in time $f(\varepsilon, r, k) \cdot |V(G)|^{1+\varepsilon}$.*

> DISTANCE INDEPENDENT SET
> *Input:* Graph $G \in \mathcal{C}$, $W \subseteq V(G)$, $k, r \in \mathbb{N}$.
> *Problem:* Determine whether G contains an r-independent set of size k.

Furthermore, if $\mathcal{C}$ is effectively nowhere dense, then f is computable.

We will show that we can solve a coloured version of the problem, called the RAINBOW DISTANCE INDEPENDENT SET problem, and reduce the original distance independent set problem to the rainbow distance independent set problem. We first give a formal definition of rainbow sets.

Definition 5.2 *A coloured graph $(G, C_1, \ldots, C_t)$ is a graph G together with relations $C_1, \ldots C_t \subseteq V(G)$, called* colours*, such that $C_i \cap C_j = \emptyset$ for all $i \neq j$. A vertex $v \notin \bigcup_{1 \leq i \leq t} C_i$ is called* uncoloured*. A set $X \subseteq V(G)$ is a* rainbow set *if all of its elements have distinct colours (and no vertex is uncoloured).*

The RAINBOW DISTANCE INDEPENDENT SET *problem on a class $\mathcal{C}$ of graphs is the following problem.*

> RAINBOW DISTANCE INDEPENDENT SET (RAINBOW DIS)
> *Input:* Graph $G \in \mathcal{C}$, $C_1, \ldots, C_t \subseteq V(G)$, $k, r \in \mathbb{N}$.
> *Problem:* Determine whether G contains a rainbow r-independent set of size k.

Before we describe the algorithm for solving the RAINBOW DISTANCE INDEPENDENT SET problem, let us show how the plain DISTANCE INDEPENDENT SET problem can be reduced to the rainbow version.

The *lexicographic product $G \bullet H$* of two graphs G and H is defined by $V(G \bullet H) = V(G) \times V(H)$ and $E(G \bullet H) = \big\{ \{(x, y), (x', y')\} : \{x, x'\} \in E(G)$ or $\big(x = x'$ and $\{y, y'\} \in E(H)\big) \big\}$. The graph $G \bullet H$ has

a natural coloured version $G \circ H$: we associate a colour with every vertex of H and colour every vertex of $G \bullet H$ by its projection on H. That is, the colour of (x, y) is y (or the colour associated with y). It is easy to see that a graph G has an r-independent set of size k if and only if $G \circ K_k$ has a rainbow r-independent of size k. This gives us the reduction from distance independent sets to their rainbow variant. Furthermore, observe that if Splitter wins the (l, m, r)-splitter game on a graph G, for some $r, l, m \geq 0$, then he also wins the $(l, k \cdot m, r)$-splitter game on $G \bullet K_k$, for all k. As a consequence, together with Theorem 4.2 and Theorem 4.5 this implies a different and very simple proof of the following result by Nešetřil and Ossona de Mendez (Theorem 13.1 of [28]) that nowhere dense classes of graphs are preserved by taking lexicographic products in the following sense.

Corollary 5.3 *If $\mathcal{C}$ is a nowhere dense class of graphs then for every $k \geq 0$, $\{G \bullet K_k : G \in \mathcal{C}\}$ is also nowhere dense.*

Note, however, that the reduction above reduces DISTANCE INDEPENDENT SET on a class $\mathcal{C}$ of graphs to RAINBOW DISTANCE INDEPENDENT SET on the class $\bigcup_{k \geq 1} \mathcal{C} \bullet K_k$, where $\mathcal{C} \bullet K_k := \{G \bullet H : G \in \mathcal{C}\}$. For the non-uniform version of our results, this is no problem, because by the previous result, if $\mathcal{C}$ is nowhere dense then $\mathcal{C} \bullet K_k$ is nowhere dense as well, and in the nonuniform setting we only have to deal with fixed k. We need to be slightly more careful for the uniform version. The key insight is that we can easily translate a winning strategy for Splitter in the (ℓ, m, r)-splitter game on a graph G to a winning strategy in the (ℓ, km, r)-splitter game on $G \bullet K_k$.

We are now ready to use this reduction to complete the proof of Theorem 5.1. Let $\varepsilon > 0$ and let ℓ, m be chosen according to Theorem 4.2 such that Splitter has a winning strategy for the $(\ell, m, 4k^2 r)$-splitter game on every graph in $\mathcal{C}$. Choose $n_0 = n_0(\varepsilon)$ according to Theorem 3.4 such that every graph $G \in \mathcal{C}$ of order $n \geq n_0$ has at most $n^{1+\varepsilon}$ many edges.

Suppose we are given an instance G, k, r, W of DISTANCE INDEPENDENT SET, where $G \in \mathcal{C}$. We first compute the coloured graph $G' := G \circ K_k$. Let $C_1, \ldots, C_t$, where $t := k$, be the colours of G'. As explained above, Splitter wins the $(l, mk, 4k^2 r)$-splitter game on G' and his winning strategy can easily be computed from any winning strategy for the $(\ell, m, 4k^2 r)$-splitter game on G.

We need to decide if $(G', C_1, \ldots, C_t)$ has a rainbow r-independent set of size k. If $n = |V(G)| \leq n_0$, we test whether this set exists by brute force. In this case the running time is bounded by a function of r, k and ε. So let us assume $n \geq n_0$.

Let $G_1 := G'$. We compute an inclusion-wise maximal rainbow r-independent set $I_1 = \{x_1^1, \ldots, x_1^{k_1}\}$ of size $k_1 \leq k$ by a greedy algorithm. If $k_1 = k$, we are done and return the independent set. Otherwise, we may assume without loss of generality that x_i^j has colour j. Let $X_1 := N_r(I_1)$. Then all elements with colours $k_1 + 1, \ldots, t$ are contained in X_1. Let $Y_1 := N_r(X_1)$. Then all paths of length at most r between elements of colour $k_1 + 1, \ldots, k$ lie inside Y_1. Let $G_2 := G_1 \setminus Y_1$.

We continue by computing an inclusion-wise maximal rainbow r-independent set in G_2. Denote this set by $I_2 = \{x_2^1, \ldots, x_2^{k_2}\}$. Note that all occurring colours are among $1, \ldots, k_1$ and in particular we have $k_2 \leq k_1$ because no other colours occur in $G_1 \setminus Y_1$. Again we may assume without loss of generality that x_2^i has colour i. Let $X_2 := N_r(I_2)$. Then we find all elements with colours $k_2 + 1, \ldots, t$ in $X_1 \cup X_2$. We let $Y_2 := N_r(X_2)$. Let $G_3 := G_2 \setminus Y_2$.

We repeat this construction until $k_s = k_{s+1}$ or until $G_{s+1} = \emptyset$. Note that $s \leq k$, because $k_1 < k$. In the first case we have constructed $s + 1$ sets $I_i = \{x_i^1, \ldots, x_i^{k_i}\}$, X_i and Y_i such that x_i^j has colour j for $1 \leq i \leq s + 1$, $1 \leq j \leq k_s$. Furthermore, the colours $k_s + 1, \ldots, t$ occur only in $X_1 \cup \ldots \cup X_s$ and all paths of length at most r between vertices of these colours lie in $Y_1 \cup \ldots \cup Y_s$. By construction, no vertex of colour $k_s + 1, \ldots, t$ has distance at most r to any vertex of I_{s+1}. Hence we may assume that any rainbow r-independent set includes the vertices $x_{s+1}^1, \ldots, x_{s+1}^{k_s}$ of colour $1, \ldots, k_s$. It remains to solve the rainbow r-independent set problem with parameter $k' := k - k_s$ and colours $k_s + 1, \ldots, t$ on $G' := G[Y_1 \cup \ldots \cup Y_s]$.

In the other case ($G_{s+1} = \emptyset$) we also let $G'' := G[Y_1 \cup \ldots Y_s]$. The only difference is that we have to solve the original problem with parameter $k' = k$.

If G'' is not connected, let $U_1, \ldots, U_c \subseteq G''$ be the components of G''. For all possible partitions of the set $C_1, \ldots, C_t$ of colours into parts $\mathcal{V}_1, \ldots, \mathcal{V}_c$ we proceed as follows. For all $1 \leq i \leq c$ we delete all colours from U_i not in $\mathcal{V}_i$, i.e. work in the coloured graph $(U_i, \mathcal{V}_i)$. We then solve the problem separately for all components $(U_i, \mathcal{V}_i)$ and for each component determine the maximal value $k'' \leq k'$ so that $(U_i, \mathcal{V}_i)$ contains a rainbow r-independent set. We then simply check whether for some partition $(V_1, \ldots, V_c)$ of the colours the maximal values for the individual components sum up to at least k'.

Hence, we can assume that G'' is connected. Then G''' has diameter at most $4k^2 \cdot r$ (there are at most $\sum_{i=1}^{k} i \leq k^2$ many vertices in the independent sets surrounded by their $2r$-neighbourhoods of diameter at most $4r$). Hence the radius of G'' is at also at most $4k^2 \cdot r$.

Let v be a centre vertex of G''. We let v be Connector's choice in the $(\ell, km, 4k^2 r)$-splitter game and let M be Splitter's answer. Without loss of generality we assume that $M = \{m_1, \ldots, m_m\} \neq \emptyset$. We let $G''' := G'' \setminus M$ and continue with a different colouring of G''' as follows. Let $X \subseteq M$ be a rainbow r-independent set in G'', possibly $X = \emptyset$ (we test for all possible sets $X \subseteq M$ whether they are rainbow r-independent sets and recurse with every possible such set). We remove the colours occurring in X completely from the graph and furthermore we remove the colour of vertices from $N_r^{G''}(X)$.

We now change the colours of G''' as follows. For every colour C_i, with $1 \leq i \leq t$, and every *distance vector* $\bar{d} := (d_1, \ldots, d_m)$, where $d_i \in \{1, \ldots, r, \infty\}$, we add a new colour $C_{i,\bar{d}}$ and set $C_{i,\bar{d}}$ to be the set of all vertices $w \in C_i$ such that $\text{dist}_{G''}(w, m_i) = d_i$, for all $1 \leq i \leq r$, where we define $\text{dist}_{G''}(w, m_i) = \infty$ if the distance is bigger then r. Note that the number of colours added in this way is only $t \cdot d'$, where $d' := (r+1)^m$ is the number of distance vectors, and hence only depends on the number of original colours and r and m. We call a subset $C_{i_1, \bar{d}_1}, \ldots, C_{i_{t''}, \bar{d}_{t''}}$ of the colours a *valid sub-colouring* if the colours satisfy the following constraints:

1. If $C_{i_j, \bar{d}_j} \neq \emptyset$ for a colour which states that the distance to some element $m \in M$ is $r' < r$, then $D_{i_{j'}, \bar{d}_{j'}} = \emptyset$ for all colours which state that the distance to m is at most $r - r'$.

2. If $C_{i_j, \bar{d}_j}$ and $C_{i_{j'}, \bar{d}_{j'}}$ are colours such that $i_j = i_{j'}$ and $\bar{d}_j \neq \bar{d}_{j'}$ then $C_{i_j, \bar{d}_j} = \emptyset$ or $C_{i_{j'}, \bar{d}_{j'}} = \emptyset$.

We now check for all possible sub-colourings $C_{i_1, \bar{d}_{i_1}}, \ldots, C_{i_{t''}, \bar{d}_{t''}}$ of G''' whether they are valid and for each valid sub-colouring we recursively call the algorithm on G''' with colouring $C_{i_1, \bar{d}_{i_1}}, \ldots, C_{i_{t''}, \bar{d}_{t''}}$ and parameter $k'' := k' - |X|$. The number of valid sub-colourings only depends on the original number of colours and on m and r.

We claim that this procedure correctly decides whether G'' contains a rainbow r-independent set of size k'. If there exists such a set Z, let $X := M \cap Z$. Then X will be considered as one of the potential sets to be extended by the algorithm. No vertex from $Z \setminus X$ may have a colour of X, hence we may remove these colours completely from the graph. Furthermore, $Z \cap N_r(X) = X$, hence we may remove the colours from $N_r(X)$. Also, if $u \in Z$ with $\text{dist}_{G''}(u, m) = r' < r$ for some $m \in M$, then $v \notin Z$ for all v with $\text{dist}_{G''}(v, m) \leq r - r'$. Hence we will find Z in the graph where all colours which state that the distance to m is at most $r - r'$ are removed. Conversely assume that the algorithm has chosen a rainbow r-independent set I in G''' of size $k' - |X|$ for some $X \subseteq M$ and some valid sub-colouring of a colouring which is consistent with X. By Condition (1) of valid sub-colourings, I is also an r-independent set in G''. By Condition (2) of valid sub-colourings, I is also rainbow in G''.

We now analyse the running time of the algorithm. First observe that in a recursive call the parameters r and m are left unchanged and k can only decrease. Moreover it follows from the definition of G''' that Splitter has a winning strategy for the $(\ell - 1, km, 4k^2 r)$-splitter game on G'''. Thus in each recursive call we can reduce the parameter ℓ by 1. Once we have reached $\ell = 0$, the graph G''' will be empty and the algorithm terminates.

10

There is one more issue we need to attend to, and that is how we compute Splitter's winning strategy, that is, the sets M. We use Remark 4.3. This means that to compute M in some recursive call, we need the whole history of the game (in a sense, the whole call stack). In addition, we need a breadth-first search tree in all graphs that appeared in the game before. It is no problem to compute a breadth-first search tree once when we first need it and then store it with the graph; this only increases the running time by a constant factor.

Let us first describe the running time of the algorithm on level j of the recursion. The time for computing k maximal r-independent sets of size at most k and their $2r$-neighbourhoods can be bounded by time $c_0 \cdot n^{1+\varepsilon}$. The factor $n^{1+\varepsilon}$ stems from the breadth-first searches we have to perform in order to find the sets $Y(i)$ and Splitter's strategy and c_0 is a constant depending only on r, k, ε and $\mathcal{C}$.

As the initial number of colours was k and the number of colours in every recursive step increases by a factor depending only on r and m (which depends only on r, k and $\mathcal{C}$), the total number of colours depends only on r, k and $\mathcal{C}$. Hence the number of rainbow r-independent subsets X of an occurring set M is bounded by a constant c_1 depending only on r, k and $\mathcal{C}$. The number of valid sub-colourings in any recursive step is bounded by a constant c_2 depending only on r, k and $\mathcal{C}$.

Furthermore, for $n \leq n_0$ the running time can be bounded by a constant c_3 that only depends on k, r, ε and $\mathcal{C}$. For $j = 0$, the running time can be bounded by a constant c_4 depending only on k, r, ε and $\mathcal{C}$. We obtain the following recurrence for T.

$$
\begin{aligned}
T(0) &\leq c_3 + c_4, \\
T(j) &\leq c_3 + c_0 \cdot n^{1+\varepsilon} + c_1 \cdot c_2 \cdot T(j-1) \qquad\qquad \text{for all } j \geq 1.
\end{aligned}
$$

We conclude that there is a constant c depending only on k, r, ε and $\mathcal{C}$ such that $T(\ell) \leq c \cdot n^{1+\varepsilon}$.

This completes the proof of Theorem 5.1. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ □

6 Sparse Neighbourhood Covers

Neighborhood covers of small radius and small size play a key role in the design of many data structures for distributed systems. Such covers will also form the basis of the data structure constructed in our first-order model-checking algorithm on nowhere dense classes of graphs. In this section we will show that nowhere dense classes of graphs admit sparse neighbourhood covers of small radius and small size and present an fpt-algorithm for computing such covers.

Definition 6.1 *For $r \in \mathbb{N}$, an r-neighbourhood cover $\mathcal{X}$ of a graph G is a set of connected subgraphs of G called* clusters, *such that for every vertex $v \in V(G)$ there is some $X \in \mathcal{X}$ with $N_r(v) \subseteq X$.*

The radius $\mathrm{rad}(\mathcal{X})$ *of a cover $\mathcal{X}$ is the maximum radius of any of its clusters. The* degree $d^{\mathcal{X}}(v)$ *of v in $\mathcal{X}$ is the number of clusters that contain v. The* maximum degree $\Delta(\mathcal{X})$ *of $\mathcal{X}$ is $\Delta(\mathcal{X}) = \max_{v \in V(G)} d^{\mathcal{X}}(v)$. The size of $\mathcal{X}$ is $\|\mathcal{X}\| = \sum_{X \in \mathcal{X}} |X| = \sum_{v \in V(G)} d^{\mathcal{X}}(v)$.*

The main result of this section is the following theorem.

Theorem 6.2 *Let $\mathcal{C}$ be a nowhere dense class of graphs. There is a function f such that for all $r \in \mathbb{N}$ and $\varepsilon > 0$ and all graphs $G \in \mathcal{C}$ with $n \geq f(r, \varepsilon)$ vertices, there exists an r-neighbourhood cover of radius at most $2r$ and maximum degree at most n^{ε} and this cover can be computed in time $f(r, \varepsilon) \cdot n^{1+\varepsilon}$. Furthermore, if $\mathcal{C}$ is effectively nowhere dense, then f is computable.*

To prove the theorem we use the concept of generalised colouring numbers introduced by Kierstead and Yang in [23]. For a graph G, let $\Pi(G)$ be the set of all linear orderings of $V(G)$. For $u, v \in V(G)$

and $k \in \mathbb{N}$, we say that u is *weakly k-accessible* from v with respect to $< \in \Pi(G)$ if $u < v$ and there is a $u-v$-path P of length at most k such that for all $w \in V(P)$ we have $u \le w$. We write $\le$ for the reflexive ordering induced by $<$. Let $\mathrm{WReach}_k(G, <, v)$ be the set of vertices that are weakly k-accessible from v and let $\mathrm{WReach}_k[G, <, v] := \mathrm{WReach}_k(G, <, v) \cup \{v\}$. The *weak k-colouring number* $\mathrm{wcol}_k(G)$ of G is defined as

$$\mathrm{wcol}_k(G) = \min_{< \in \Pi(G)} \max_{v \in V(G)} |\mathrm{WReach}_k[G, <, v]|.$$

Zhu [33] (and in fact also Kierstead and Yang but they were not aware of the depth-r minor terminology) showed that general colouring numbers and densities of depth-r minors are strongly related. From this, Nešetřil and Ossona de Mendez conclude that the weak colouring number on nowhere dense classes is small.

Lemma 6.3 ([33, 29]) *Let $\mathcal{C}$ be a nowhere dense class of graphs. Then there is a function f such that for every $r \in \mathbb{N}$, every $\varepsilon > 0$, every graph $G \in \mathcal{C}$ with $n \ge f(r, \varepsilon)$ vertices satisfies $\mathrm{wcol}_r(G) \le n^\varepsilon$. Furthermore, if $\mathcal{C}$ is effectively nowhere dense, then f is computable.*

For our purpose, we need an efficient algorithm for ordering the vertices of G in an order witnessing $\mathrm{wcol}_r(G) \le n^\varepsilon$. Dvořák [13] conjectures that in general computing $\mathrm{wcol}_r(G)$ is NP-complete. We are able to prove his conjecture for all $r \ge 3$. He provides an approximation algorithm to solve the problem, but its running time is $\mathcal{O}(r \cdot n^3)$ which is too expensive for our purpose. We propose a more efficient approximation algorithm, based on Nešetřil and Ossona de Mendez's transitive fraternal augmentation technique and an argument from Zhu's proof.

In the following we will work with ordered representations of graphs where each vertex stores an adjacency list for its in-neighbours and an adjacency list for its out-neighbours.

Definition 6.4 *Let $\vec{G}$ be a directed graph. A* tight 1-transitive fraternal augmentation *of $\vec{G}$ is a directed graph $\vec{H}$ on the same vertex set such that for all distinct vertices u, v, w*

- *if $(u, v) \in E(\vec{G})$, then $(u, v) \in E(\vec{H})$.*

- *if $(u, w), (w, v) \in E(\vec{G})$, then $(u, v) \in E(\vec{H})$,*

- *if $(u, w), (v, w) \in E(\vec{G})$, then (u, v) or (v, u) are arcs of $\vec{H}$ and*

- *for all $(u, v) \in E(\vec{H})$, either $(u, v) \in E(\vec{G})$ or there is some w such that $(u, w), (w, v) \in E(\vec{G})$ or $(u, w), (v, w) \in E(\vec{G})$.*

We write $\mathrm{aug}(\vec{G}, 1)$ for any tight 1-transitive fraternal augmentation of $\vec{G}$ and for $r > 1$ we write $\mathrm{aug}(\vec{G}, r)$ for $\mathrm{aug}(\mathrm{aug}(\vec{G}, r-1), 1)$. We call $\mathrm{aug}(\vec{G}, r)$ a tight r-transitive fraternal augmentation of $\vec{G}$. We will often write $\mathrm{aug}(G, r)$ and speak of an r-transitive fraternal augmentation of G instead of $\mathrm{aug}(\vec{G}, r)$ and an r-transitive fraternal augmentation of an orientation $\vec{G}$ of G.

In [30], Nešetřil-Ossona de Mendez show how to efficiently compute tight transitive fraternal augmentations. They state the result in terms of average densities of depth-r minors, for our purpose it suffices to state their result for nowhere dense classes. All functions $f(r, \varepsilon)$ in the following lemmas are computable if $\mathcal{C}$ is effectively nowhere dense.

Lemma 6.5 (Nešetřil-Ossona de Mendez [30], Corollary 4.2, Theorem 4.3) *Let $\mathcal{C}$ be a nowhere dense class of graphs. There is a function f such that for all $r \in \mathbb{N}$ and $\varepsilon > 0$ and all graphs $G \in \mathcal{C}$ with $n \ge f(r, \varepsilon)$ vertices, there exists an r-transitive fraternal augmentation $\vec{H} = \mathrm{aug}(G, r)$ of G such that $\Delta^-(\vec{H}) \le n^\varepsilon$. Furthermore, $\vec{H}$ can be computed from G in time $f(r, \varepsilon) \cdot n^{1+\varepsilon}$.*

We will write $\text{aug}(G, r, \varepsilon)$ for an augmentation $\vec{H} = \text{aug}(G, r)$ such that $\Delta^-(\vec{H}) \leq n^\varepsilon$.

The following property of transitive fraternal augmentations is noted in the proof of Theorem 5.1 in [30].

Lemma 6.6 ([30]) *Let G be a graph and let $r \in \mathbb{N}$. Let $\vec{H} = \text{aug}(G, r)$ be an r-transitive fraternal augmentation of G. Let $v \in V(G)$ and $w \in N_r^G(v)$. Let $v = v_1, v_2, \ldots, v_l = w$ be a path of length at most r from v to w in G. Then either $(v, w) \in E(\vec{H})$ or $(w, v) \in E(\vec{H})$ or there is some v_i such that $(v_i, v), (v_i, w) \in E(\vec{H})$.*

In fact, for the results in the previous lemma it would suffice to use an $\left\lceil \log_{3/2} r \right\rceil + 1$-augmentation. While this would make the construction more efficient, we refrain from doing so for ease of presentation.

We now show how to approximate $\text{wcol}_r(G)$ with the help of r-transitive fraternal augmentations.

Lemma 6.7 *Let G be a graph and let $r > 0$. Let $\vec{H} = \text{aug}(G, r)$ be an r-transitive fraternal augmentation of G such that $\Delta^-(\vec{H}) \leq d$. Then $\text{wcol}_r(G) \leq 2(d + 1)^2$.*

Proof. As $\Delta^-(\vec{H}) \leq d$, the underlying undirected graph H is $2d$-degenerate and we can order the vertices of H such that each vertex has at most $2d$ smaller neighbours. Denote this order by $<$. For each vertex $v \in V(G)$ we count the number of end-vertices of paths of length at most r from v such that the end-vertex is the smallest vertex of the path. This number bounds $|\text{WReach}_r[G, <, v)]|$.

By Lemma 6.6, for each such path with end-vertex w, we either have an edge (v, w) or an edge (w, v) or there is u on the path and we have edges $(u, v), (u, w)$ in H. By construction of the order there are at most $2d$ edges (v, w) or (w, v) such that $w < v$. Furthermore, we have at most d edges (u, v), as v has indegree at most d and for each such u there are at most $2d$ edges (u, w) such that $w < u$ by construction of the order. These are exactly the pairs of edges we have to consider, as no vertex on the path from v to w may be smaller than w. Hence in total we have $|\text{WReach}_r[G, <, v]| \leq 2d + 2d^2 + 1 \leq 2(d + 1)^2$. $\square$

Corollary 6.8 *Let $\mathcal{C}$ be a nowhere dense class of graphs. There is a function f such that for all $r \in \mathbb{N}$ and $\varepsilon > 0$ and every $G \in \mathcal{C}$ with $n \geq f(r, \varepsilon)$ vertices, we can order the vertices of G in order $<$ such that $|\text{WReach}_r[G, <, v]| \leq n^\varepsilon$ for all $v \in V(G)$ in time $f(r, \varepsilon) \cdot n^{1+\varepsilon}$. Furthermore, if $\mathcal{C}$ is effectively nowhere dense, then f is computable.*

Proof. Let $\delta := \varepsilon/4$. We compute an r-transitive fraternal augmentation $\vec{H} = \text{aug}(G, r, \delta)$ of G in time $g(r, \delta) \cdot n^{1+\delta}$ by Lemma 6.5, where g is the function from the lemma. We can order the vertices as in the proof of Theorem 6.7 by a simple greedy algorithm in time $\mathcal{O}(n^{1+\delta})$ and obtain an order witnessing $\text{wcol}_r(G) \leq 2(n^\delta + 1)^2 \leq n^\varepsilon$. $\square$

In the next lemma we use the weak colouring number to prove the existence of sparse neighbourhood covers in nowhere dense classes of graphs.

Definition 6.9 *Let G be a graph, let $<$ be an ordering of $V(G)$ and let $r > 0$. For a vertex $v \in V(G)$ we define*
$$X_r[G, <, v] := \{w \in V(G) : v \in \text{WReach}_r[G, <, w]\}.$$

Lemma 6.10 *Let G be a graph such that $\text{wcol}_{2r}(G) \leq s$ and let $<$ be an order witnessing this. Then $\mathcal{X} = \{X_{2r}[G, <, v] : v \in V(G)\}$ is an r-neighbourhood cover of G with radius at most $2r$ and maximum degree at most s.*

Proof. Clearly the radius of each cluster is at most $2r$, because if v is weakly $2r$-accessible from w then $w \in N_{2r}(v)$. Furthermore, every r-neighbourhood lies in some cluster. To see this, let $v \in V(G)$. Let u be the minimum of $N_r(v)$ with respect to $<$. Then u is weakly $2r$-accessible from every $w \in N_r(v) \setminus \{u\}$ as there

13

is a path from w to u which uses only vertices of $N_r(v)$ and has length at most $2r$ and u is the minimum element of $N_r(v)$. Thus $N_r(v) \subseteq X_{2r}[G, <, u]$. Finally observe that for every $v \in V(G)$,

$$d^{\mathcal{X}}(v) = |\{u \in V(G) : v \in X_{2r}[G, <, u]\}|$$
$$= |\{u \in V(G) : u \in \mathrm{WReach}_{2r}[G_<, v]\}| = |\mathrm{WReach}_{2r}[G_<, v]| \leq s.$$

$\square$

Proof of Theorem 6.2. Let $\delta := \varepsilon/2$. We order the vertices of G in order $<$ as in Corollary 6.8, where δ plays the role of ε in the corollary, such that $\mathrm{WReach}_{2r}[G, <, v] \leq n^\delta$ for all $v \in V(G)$ in time $g(r, \delta) \cdot n^{1+\delta}$, where g is the function from the corollary.

Let us first note the following observation.

Claim 1. For $v \in V(G)$ let $S(v) := \{u : u < v\}$. Then $X_{2r}[G, <, v] = N_{2r}^{G \setminus S(v)}(v)$.

Our algorithm computes the sets $X_{2r}[G, <, v]$ in ascending order. To do so, it chooses the smallest vertex v, performs $2r$ levels of a breadth-first search and then deletes v from the graph. Correctness of the algorithm follows immediately from Claim 1. Let us analyse the running time.

We construct the following representation of G which is easily seen to be computable in time $\mathcal{O}(n^{1+\delta})$. We split the edges of G into edges going to larger elements and into edges going to smaller elements with respect to the ordering. For each $v \in V(G)$ we write $N_>(v)$ (resp. $N_<(v)$) for the neighbours of v that are larger (resp. smaller) than v. We write $d_>(v)$ for $|N_>(v)|$ and $d_<(v)$ for $|N_<(v)|$. Note that we have $d_<(v) \leq n^\delta$ for each $v \in V(G)$, as $d_<(v) \leq |\mathrm{WReach}_{2r}[G, <, v]|$.

Let G' be a subgraph of G with n' vertices. We can count the edges of G' by counting the sum of $d_<(v)$ over all $v \in V(G')$, hence G' has at most $n' \cdot n^\delta$ many edges. We can thus perform each breadth-first search to compute $X_{2r}[G, <, v]$ in time $\mathcal{O}(|X_{2r}[G, <, v]| \cdot n^\delta)$ for each vertex $v \in V(G)$. Furthermore, we have the following overhead in the breadth-first search for deleting edges that point to v, which must be deleted. As we store the edges of each vertex in separate lists, for each vertex $w \in N_>(v)$ (this is the first level of the breadth-first search), we have to access only the edges to vertices of $N_<(w)$. No other vertex is connected to v in $G \setminus S(v)$. Hence, the deletion of v from the adjacency list of w can be done in time $d_<(w) \leq n^\delta$. The number of such vertices w is $d_>(v)$, which at the time of deletion of v is bounded by $|X_{2r}[G, <, v]|$.

For ease of presentation let $X_v := X_{2r}[G, <, v]$ and let us drop any constant factors in the following estimation. We get a total running time of

$$\sum_{v \in V(G)} \left(|X_v| \cdot n^\delta + \sum_{w \in N_>(v)} d_<(w) \right)$$

$$= \sum_{v \in V(G)} |X_v| \cdot n^\delta + \sum_{v \in V(G)} \sum_{w \in N_>(v)} d_<(w)$$

$$\leq \sum_{v \in V(G)} |X_v| \cdot n^\delta + \sum_{v \in V(G)} |X_v| \cdot n^\delta$$

$$= 2n^\delta \sum_{v \in V(G)} |X_v|$$

$$\leq 2n^{1+2\delta} =: f(r, \varepsilon) \cdot n^{1+\varepsilon}$$

$\square$

Remark 6.11 *By definition, an r-neighbourhood cover $\mathcal{X}$ of a graph G contains for each $v \in V(G)$ a cluster $X \in \mathcal{X}$ such that $N_r^G(v) \subseteq X$. For the algorithmic applications below it will be useful to store*

along with the neighbourhood cover a function $f_{\mathcal{X}} : V(G) \to \mathcal{X}$ which associates with every vertex v such a cluster X containing its r-neighbourhood.

The proof of the previous theorem can easily be modified to compute such a function along with the neighbourhood cover as follows: we associate with $v \in V(G)$ the set $X_{2r}[G, <, u]$ for the $<$-minimal $u \in V(G)$ such that $v \in N_r^{G \setminus S(u)}(u)$, where $S(u)$ is defined as in Claim 1 in the proof of Theorem 6.2. As the sets $X_{2r}[G, <, u]$ are computed in increasing order, this can be done at no extra cost.

We remark that our construction also yields very good covers for other restricted classes of graphs, in particular for classes with excluded minors and classes of graphs of bounded expansion, where we can replace the maximum degree n^{ε} of the neighbourhood cover by a constant. See the conclusions (Section 9) for further comments.

7 Locality of First-Order Logic

In this chapter, we prove the "rank-preserving" version of Gaifman's locality theorem stated in the introduction.

7.1 Background on First-Order Logic

We start with a brief review of first-order logic. For background, we refer the reader to [15]. A *(relational) vocabulary* is a finite set of relation symbols, each with a prescribed arity. Throughout this paper, we let σ be a vocabulary. A *σ-structure A* consist of a (not necessarily finite) set $V(A)$, called the *universe* or *vertex set* of A, and for each k-ary relation symbol $R \in \sigma$ a k-ary relation $R(A) \subseteq V(A)^k$. A structure A is *finite* if its universe is.

For example, graphs may be viewed as $\{E\}$-structures, where E is a binary relation symbol.

Let A be a σ-structure. For a subset $X \subseteq V(A)$, the *induced substructure* of A with universe X is the σ-structure $A[X]$ with $V(A[X]) = X$ and $R(A[X]) = R(A) \cap X^k$ for every k-ary $R \in \sigma$. For a vocabulary $\sigma' \subseteq \sigma$, the *σ'-restriction* of A is the σ'-structure A' with $V(A') = V(A)$ and $R(A') = R(A)$ for all $R \in \sigma'$. Conversely, A is a *σ-expansion* of a σ'-structure A' if A' is the σ'-restriction of A.

First-order formulas of vocabulary σ are formed from atomic formulas $x = y$ and $R(x_1, \ldots, x_k)$, where $R \in \sigma$ is a k-ary relation symbol and $x, y, x_1, \ldots, x_k$ are variables (we assume that we have an infinite supply of variables) by the usual Boolean connectives $\neg$ (negation), $\wedge$ (conjunction), and $\vee$ (disjunction) and existential and universal quantification $\exists x, \forall x$, respectively. The set of all first-order formulas of vocabulary σ is denoted by FO$[\sigma]$, and the set of all first-order formulas by FO. The free variables of a formula are those not in the scope of a quantifier, and we write $\varphi(x_1, \ldots, x_k)$ to indicate that the free variables of the formula φ are among $x_1, \ldots, x_k$. A *sentence* is a formula without free variables. The *quantifier rank* qr(φ) of a formula φ is the nesting depth of quantifiers in φ, defined recursively in the obvious way. A formula without any quantifiers is called *quantifier-free*.

To define the semantics, we inductively define a satisfaction relation $\models$, where for a σ-structure A, a formula $\varphi(x_1, \ldots, x_k)$, and elements $a_1, \ldots, a_k \in V(A)$,

$$A \models \varphi(a_1, \ldots, a_k)$$

means that A satisfies φ if the free variables $x_1, \ldots, x_k$ are interpreted by $a_1, \ldots, a_k$, respectively. If $\varphi(x_1, \ldots, x_k) = R(x_1, \ldots, x_k)$ is atomic, then $A \models \varphi(a_1, \ldots, a_k)$ if $(a_1, \ldots, a_k) \in R(A)$. The meaning of the equality symbol, the Boolean connectives, and the quantifiers is the usual one.

For example, consider the formula $\varphi(x_1, x_2) = \forall y (x_1 = y \vee x_2 = y \vee E(x_1, y) \vee E(x_2, y))$ in the vocabulary $\{E\}$ of graphs. For every graph G and vertices $v_1, v_2 \in V(G)$ we have $G \models \varphi(v_1, v_2)$ if any

only if $\{v_1, v_2\}$ is a dominating set of G. Thus G satisfies the sentence $\exists x_1 \exists x_2 \varphi(x_1, x_2)$ if, and only if, it has a (nonempty) dominating set of size at most 2.

Whenever a σ-structure occurs as the input of an algorithm, we implicitly assume that it is finite and encoded in a suitable way. Similarly, we assume that formulas φ appearing as input are encoded suitably. By $|\varphi|$, we denote the length of the encoding of φ.

A formula $\varphi(x_1, \ldots, x_k) \in \mathrm{FO}[\sigma]$ is *valid* if for all σ-structures A and all elements $a_1, \ldots, a_k \in V(A)$ it holds that $A \models \varphi(a_1, \ldots, a_k)$. The Completeness Theorem for First-Order Logic implies that the set of valid formulas is recursively enumerable. Two formulas $\varphi(x_1, \ldots, x_k), \psi(x_1, \ldots, x_k) \in \mathrm{FO}[\sigma]$ are *equivalent* if for all σ-structures A and all elements $a_1, \ldots, a_k \in V(A)$ we have $A \models \varphi(a_1, \ldots, a_k) \iff A \models \psi(a_1, \ldots, a_k)$.

Up to logical equivalence, for all k, q there are only finitely many FO-formulas $\varphi(x_1, \ldots, x_k)$ of quantifier-rank at most q. Indeed, by systematically renaming the bound variables, bringing Boolean combinations into conjunctive normal form, and deleting duplicate entries from the disjunctions and conjunctions, we can normalise FO-formulas in such a way that every formula can be effectively translated into an equivalent normalised formula of the same quantifier rank, and for all k, q the set $\Phi(\sigma, k, q)$ of all normalised FO-formulas $\varphi(x_1, \ldots, x_k)$ of quantifier rank at most q is finite and computable.

The *Gaifman graph* G_A of a σ-structure A is the graph with vertex set $V(A)$ and an edge between $a_1, a_2 \in V(A)$ if a_1, a_2 appear together in some tuple of some relation in A. The *distance* $\mathrm{dist}^A(a, b)$, or just $\mathrm{dist}(a, b)$, between two elements $a, b \in V(A)$ in A is the length of the shortest path from a to b in G_A, and the *r-neighbourhood* of a in A is the set $N_r^A(a)$, or just $N_r(a)$, of all $b \in V(A)$ such that $\mathrm{dist}(a, b) \leq r$. For a tuple $\bar{a} = (a_1, \ldots, a_k)$, we let $N_r(\bar{a}) = \bigcup_{i=1}^{k} N_r(a_i)$.

A first-order formula $\psi(\bar{x})$ is called *r-local* if its truth value at a tuple $\bar{a}$ of vertices in a structure A only depends on the r-neighbourhood of $\bar{a}$ in A, that is, $A \models \varphi(\bar{a}) \iff A[N_r(\bar{a})] \models \varphi(\bar{a})$. For all $d \geq 0$ there is an FO-formula $\delta_{\leq d}(x, y)$ stating that the distance between x and y is at most d. We write $\delta_{>d}(x, y)$ instead of $\neg \delta_{\leq d}(x, y)$. A *basic local sentence* is a first-order sentence of the form

$$\exists x_1 \ldots \exists x_k \big(\bigwedge_{1 \leq i < j \leq k} \delta_{>2r}(x_i, x_j) \wedge \bigwedge_{i=1}^{k} \varphi(x_i) \big), \tag{7.1}$$

where φ is r-local.

Theorem 7.1 (Gaifman's Locality Theorem [19]) *Every first-order sentence is equivalent to a Boolean combination of basic local sentences.*

The algorithm of Frick and Grohe [18] for deciding first-order properties on graph classes of bounded local tree width relies on Gaifman's theorem. Unfortunately, we cannot use Gaifman's theorem here, at least not directly, because it does not give us sufficient control over the quantifier rank of the basic local sentences we translate a sentence to. As we intend to apply the theorem repeatedly, such control will be crucial. To get around these difficulties, we need a discounted rank measure, which does not charge the full quantifier rank to distance formulas, and a refined version of Gaifman's theorem.

7.2 The Logic FO$^+$

We define an extension FO$^+$ of first-order logic by adding new atomic formulas $\mathrm{dist}(x, y) \leq d$, for all variables x, y and all $d \in \mathbb{N}$. We call these formulas *distance atoms*. The meaning of the distance atoms is obvious. Note that every FO$^+$-formula φ is equivalent to an FO-formula φ^- obtained from φ by replacing each distance atom $\mathrm{dist}(x, y) \leq d$ by the FO-formula $\delta_{\leq d}(x, y)$. Thus FO$^+$ is only a syntactic extension of FO. However, the quantifier rank of $\delta_{\leq d}(x, y) \in \mathrm{FO}$ is at least $\lceil \log d \rceil$, whereas by definition the quantifier

rank of the atomic FO$^+$-formula $\mathrm{dist}(x,y) \leq d$ is 0. With this definition as one of the base steps, we can define the quantifier rank $\mathrm{qr}(\varphi)$ for FO$^+$-formulas φ recursively as for FO-formulas.

We now define the discounted rank measure. Let $q \in \mathbb{N}$.

We say that φ has *q-rank* at most ℓ if φ has quantifier-rank at most ℓ and if each distance atom $\mathrm{dist}(x,y) \leq d$ in the scope of $i \leq \ell$ quantifiers satisfies $d \leq (4q)^{q+\ell-i}$.

For example, the sentence

$$\exists x \exists y \Big(\mathrm{dist}(x,y) \leq 12^5 \wedge \exists z \big(\mathrm{dist}(x,z) \leq 12^6 \wedge \forall z'(\neg \mathrm{dist}(z,z') \leq 12^4 \vee \mathrm{dist}(z',y) \leq 12^4)) \big) \Big)$$

has 3-rank 6, because for the distance atom $\mathrm{dist}(x,z) \leq 12^6$ in the scope of 3 quantifiers we have $12^6 = (4 \cdot 3)^{3+6-3}$. Note that the quantifier-rank of this formula is 4 and hence $\leq \ell = 6$.

For convenience, we let

$$f_q(\ell) := (4q)^{q+\ell}. \tag{7.2}$$

This is is the largest value of d which may occur in a distance atom $\mathrm{dist}(x,y) \leq d$ of a formula of q-rank ℓ.

The definition of the q-rank arises from the necessities of the proof of Theorem 7.5. Note that this rank measure makes it cheaper to define distances as in FO-formulas: with an FO$^+$-formula of q-rank q we can define distances up to $(4q)^{2q}$, which is much more than the distance 2^q we can define with an FO-formula of quantifier rank q. Also note that defining distances becomes more expensive in the scope of quantifiers.

Up to logical equivalence, for all k, q, ℓ there are only finitely many FO$^+[\sigma]$-formulas $\varphi(x_1, \ldots, x_k)$ of q-rank at most ℓ. As FO-formulas, we can normalise FO$^+$ formulas such that every formula can be effectively translated into an equivalent normalised formula of the same rank, and for all k, q, ℓ the set $\Phi^+(\sigma, k, q, \ell)$ of all normalised FO$^+$-formulas $\varphi(x_1, \ldots, x_k)$ of q-rank at most ℓ is finite and computable.

7.3 An Ehrenfeucht-Fraïssé Game for FO$^+$

For σ-structures A, B and tuples $\bar{a} = (a_1, \ldots, a_k) \in V(A)^k, \bar{b} = (b_1, \ldots, b_k) \in V(B)^k$ we write $(A, \bar{a}) \equiv^+_{q,\ell} (B, \bar{b})$ (and say that $(A, \bar{a})$ and $(B, \bar{b})$ are $(q,\ell)^+$-equivalent) if for all $\varphi(\bar{x}) \in$ FO$^+$ of q-rank at most ℓ we have $A \models \varphi(\bar{a}) \iff B \models \varphi(\bar{b})$. Observe that $(A, \bar{a}) \equiv^+_{q,\ell} (B, \bar{b})$ implies for all $i, j \in [k]$ that either $\mathrm{dist}(a_i, a_j) = \mathrm{dist}(b_i, b_j)$ or $\mathrm{dist}(a_i, a_j) > f_q(\ell)$ and $\mathrm{dist}(b_i, b_j) > f_q(\ell)$.

We generalise the well-known characterisation of first-order equivalence by means of the Ehrenfeucht-Fraïssé (EF) game (see, for example, [15]) to the logic FO$^+$ parameterized by q-ranks. A *partial d-isomorphism* between two structures A, B is a mapping p with domain $\mathrm{dom}(p) \subseteq V(A)$ and range $\mathrm{rg}(p) \subseteq V(B)$ that is an isomorphism between the induced substructure $A[\mathrm{dom}(p)]$ and the induced substructure $B[\mathrm{rg}(p)]$ and in addition, preserves distances up to d, that is, for all $a, a' \in \mathrm{dom}(p)$ either $\mathrm{dist}(a, a') = \mathrm{dist}(p(a), p(a'))$ or $\mathrm{dist}(a, a') > d$ and $\mathrm{dist}(p(a), p(a')) > d$.

Definition 7.2 (EF^+_q-game) Let A, B be σ-structures, $\bar{a} = (a_1, \ldots, a_k) \in V(A)^k, \bar{b} = (b_1, \ldots, b_k) \in V(B)^k$ and $q \in \mathbb{N}$. Let $0 \leq \ell \leq q$. The ℓ-*round EF^+_q-game* on $(A, \bar{a}, B, \bar{b})$ is played by two players, called *Spoiler* and *Duplicator*. The game is played for ℓ rounds. In round i, Spoiler picks an element $a_{k+i} \in V(A)$ or an element $b_{k+i} \in V(B)$. If Spoiler picks $a_{k+i} \in V(A)$, then Duplicator must choose an element $b_{k+i} \in V(B)$ and if Spoiler picks $b_{k+i} \in V(B)$, then Duplicator must choose an element $a_{k+i} \in V(A)$. Duplicator wins the game if for $0 \leq i \leq \ell$, the mapping $a_j \mapsto b_j$ for $1 \leq j \leq k+i$ is a partial $f_q(\ell - i)$-isomorphism.

Theorem 7.3 *For all* $q, 0 \le \ell \le q$, A, B *and* $\bar{a} \in V(A)^k, \bar{b} \in V(B)^k$, *the following are equivalent.*

1. *Duplicator has a winning strategy for the* ℓ*-round* EF_q^+ *game on* $(A, \bar{a}, B, \bar{b})$.

2. $(A, \bar{a}) \equiv_{q,\ell}^+ (B, \bar{b})$.

The proof of Theorem 7.3 requires some familiarity with logic. It is similar to the proof that equivalence in first-order logic is characterised by the standard Ehrenfeucht-Fraïssé game (see, for example, [15]).

For $\bar{a} = (a_1, \ldots, a_k) \in V(A)^k$ and $a \in V(A)$, write $\mathrm{dist}(\bar{a}, a) =_{q,\ell} \bar{d} \in (\{0, \ldots, f_q(\ell)\} \cup \{\infty\})^k$ if for all $i \in [k]$ we have $\mathrm{dist}(a_i, a) = d_i \le f_q(\ell)$ or $\mathrm{dist}(a_i, a) > f_q(\ell)$ and $d_i = \infty$. Note that we can easily write a quantifier-free FO^+-formula of q-rank ℓ expressing $\mathrm{dist}(\bar{x}, x) =_{q,\ell} \bar{d}$.

We can rephrase the existence of a winning strategy for Duplicator in the ℓ-round EF_q^+ game on $(A, \bar{a}, B, \bar{b})$ as follows.

- Duplicator has a winning strategy for the 0-round FO_q^+-game on $(A, \bar{a}, B, \bar{b})$ if, and only if, $\bar{a} \mapsto \bar{b}$ is a partial $f_q(0)$-isomorphism.

- For $0 < \ell \le q$, Duplicator has a winning strategy for the ℓ-round FO_q^+-game on $(A, \bar{a}, B, \bar{b})$ if, and only if,

  (1) $\bar{a}$ and $\bar{b}$ satisfy the same distance formulas up to $f_q(\ell)$ and

  (2) for every $a \in V(A)$ there is a $b \in V(B)$ such that Duplicator has a winning strategy for the $\ell - 1$-round FO_q^+-game on $(A, \bar{a}a)$ and $(B, \bar{b}b)$ and

  (3) for every $b \in V(B)$ there is an $a \in V(A)$ such that Duplicator has a winning strategy for the $\ell - 1$-round FO_q^+-game on $(A, \bar{a}a)$ and $(B, \bar{b}b)$.

This description of winning strategies can be defined in FO^+ as follows. Let A and $q \in \mathbb{N}$ be given. For $\bar{a} = (a_1, \ldots, a_k) \in V(A)^k$, $\bar{x} := (x_1, \ldots, x_k)$ and $0 \le \ell \le q$, let

$$\vartheta_{\bar{a}}^{q,\ell}(\bar{x}) := \bigwedge_{\substack{a_i, a_j \in \bar{a} \\ \mathrm{dist}(a_i, a_j) = d \le f_q(\ell)}} \mathrm{dist}(x_i, x_j) = d \quad \wedge \bigwedge_{\substack{a_i, a_j \in \bar{a} \\ \mathrm{dist}(a_i, a_j) > f_q(\ell)}} \mathrm{dist}(x_i, x_j) > f_q(\ell).$$

For $\ell = 0$, let

$$\varphi_{\bar{a}}^{q,0}(\bar{x}) := \vartheta_{\bar{a}}^{q,0}(\bar{x}) \quad \wedge \bigwedge_{\substack{\varphi(\bar{x}) \in \Phi(\sigma, k, 0) \\ A \models \varphi(\bar{a})}} \varphi(\bar{x}).$$

Recall that $\Phi(\sigma, k, 0)$ denotes the (finite) set of all quantifier free normalised $FO[\sigma]$-formulas $\varphi(\bar{x})$. For $1 \le \ell \le q$, let

$$\varphi_{\bar{a}}^{q,\ell}(\bar{x}) := \vartheta_{\bar{a}}^{q,\ell}(\bar{x}) \wedge \bigwedge_{a \in V(A)} \exists x_{k+1} \varphi_{\bar{a}a}^{q,\ell-1}(\bar{x}, x_{k+1}) \wedge \forall x_{k+1} \bigvee_{a \in V(A)} \varphi_{\bar{a}a}^{q,\ell-1}(\bar{x}, x_{k+1}).$$

If we remove repeated entries from the big conjunction and the big disjunction in the definition of $\varphi_{\bar{a}}^{q,\ell}(\bar{x})$, we obtain a well-defined finite formula even for infinite structures A. Moreover, it is easy to see that the q-rank of this formula is ℓ. The following lemma implies Theorem 7.3.

Lemma 7.4 *Given* $q, 0 \le \ell \le q$, A, B *and* $\bar{a} \in V(A)^k, \bar{b} \in V(B)^k$, *the following are equivalent.*

1. *Duplicator has a winning strategy for the* ℓ*-round* EF_q^+ *game on* $(A, \bar{a}, B, \bar{b})$.

2. $B \models \varphi_{\bar{a}}^{q,\ell}(\bar{b})$.

3. $(A, \bar{a}) \equiv_{q,\ell}^{+} (B, \bar{b})$.

Proof. Assertion (3) implies assertion (2), as the q-rank of $\varphi_{\bar{a}}^{q,\ell}$ is ℓ and $A \models \varphi_{\bar{a}}^{q,\ell}(\bar{a})$.

Let $q \in \mathbb{N}$. We prove the equivalence of (1) and (2) by induction on ℓ.

For $\ell = 0$, $(A, \bar{a}) \equiv_{q,\ell}^{+} (B, \bar{b})$ if, and only if, $\bar{a} \mapsto \bar{b}$ is a partial $f_q(0)$-isomorphism. This is exactly the meaning of $\varphi_{\bar{a}}^{q,0}$.

For $\ell > 0$,

  Duplicator has a winning strategy for the ℓ-round EF_q^{+} game on $(A, \bar{a}, B, \bar{b})$

$\Longleftrightarrow$ $\bar{a}$ and $\bar{b}$ satisfy the same distance formulas up to $f_q(\ell)$ and

- for every $a \in V(A)$ there is a $b \in V(B)$ such that Duplicator has a winning strategy for the $\ell - 1$-round FO_q^{+}-game on $(A, \bar{a}a)$ and $(B, \bar{b}b)$ and

- for every $b \in V(B)$ there is an $a \in V(A)$ such that Duplicator has a winning strategy for the $\ell - 1$-round FO_q^{+}-game on $(A, \bar{a}a)$ and $(B, \bar{b}b)$

$\Longleftrightarrow$ $\bar{a}$ and $\bar{b}$ satisfy the same distance formulas up to $f_q(\ell)$ and

- for every $a \in V(A)$ there is a $b \in V(B)$ such that $B \models \varphi_{\bar{a}a}^{q,\ell-1}(\bar{b}b)$ and

- for every $b \in V(B)$ there is an $a \in V(A)$ such that $B \models \varphi_{\bar{a}a}^{q,\ell-1}(\bar{b}b)$ (by induction hypothesis)

$\Longleftrightarrow$ $B \models \varphi_{\bar{a}}^{q,\ell}(\bar{b})$ (by construction of $\varphi_{\bar{a}}^{q,\ell}$).

It remains to show that (1) implies (3). The proof is by induction on ℓ. Case $\ell = 0$ is handled as above. Let $\ell > 0$ and suppose that Duplicator has a winning strategy for the ℓ-round EF_q^{+} game starting in position $(A, \bar{a}, B, \bar{b})$. Then the truth of atomic formulas and distances up to $f_q(\ell)$ in $\bar{a}$ and $\bar{b}$ are preserved. Clearly, the set of formulas whose truth values are preserved is closed under negation and disjunction. Suppose that $\varphi(\bar{x}) = \exists y \psi(\bar{x}, y)$ and φ is of rank at most (q, ℓ). Assume, for instance, $A \models \varphi(\bar{a})$. Then there is $a \in V(A)$ such that $A \models \varphi(\bar{a}, a)$. By assumption Duplicator has a winning strategy for the ℓ-round EF_q^{+} game starting in position $(A, \bar{a}, B, \bar{b})$ and thus there is $b \in V(B)$ such that Duplicator has a winning strategy for the $\ell - 1$-round EF_q^{+} game starting in position $(A, \bar{a}a, B, \bar{b}b)$. Since the q-rank of ψ is at most $\ell - 1$, the induction hypothesis yields $B \models \psi(\bar{b}, b)$ and hence $B \models \varphi(\bar{b})$. $\square$

7.4 The Rank-Preserving Locality Theorem

We expand σ-structures A by adding definable information about neighbourhoods to every vertex. Let $\mathcal{X}$ be an r-neighbourhood cover of A. For every $a \in V(G)$, we fix some cluster $\mathcal{X}(a) \in \mathcal{X}$ such that $N_r(a) \subseteq \mathcal{X}(a)$. Actually, we view this assignments of clusters to the vertices as being given with the neighbourhood cover. Formally, we thus view an r-neighbourhood cover $\mathcal{X}$ as a mapping that associates with every vertex $a \in V(G)$ a set $\mathcal{X}(a) \subseteq V(G)$ such that $N_r(a) \subseteq \mathcal{X}(a)$. For all $q \in \mathbb{N}$, let $\sigma \star q$ be the vocabulary obtained from σ by adding a fresh unary relation symbol P_φ for each $\varphi = \varphi(x) \in \Phi^+(\sigma, 1, q, q)$. For a σ-structure A, let $A \star_{\mathcal{X}} q$ be the $\sigma \star q$-expansion of A in which P_φ is interpreted by the set of all $a \in V(A)$ such that $A\big[\mathcal{X}(a)\big] \models \varphi(a)$. We let $\sigma \star^0 q := \sigma$ and $A \star_{\mathcal{X}}^0 q := A$. For $i \geq 0$, we let $\sigma \star^{i+1} q := (\sigma \star^i q) \star q$ and $A \star_{\mathcal{X}}^{i+1} q := \big(A \star_{\mathcal{X}}^i q\big) \star_{\mathcal{X}} q$.

A (q, r)-*independence sentence* is a sentence of the form

$$\exists x_1 \dots \exists x_q \Big(\bigwedge_{1 \leq i < j \leq q} \mathrm{dist}(x_i, x_j) > 2r \wedge \bigwedge_{1 \leq i \leq q} \varphi(x_i) \Big)$$

for a quantifier-free first-order formula $\varphi(x_i)$. Note that the independence sentences have the same form as the basic local sentences in Gaifman's Theorem, except that the formula $\varphi(x)$ is required to be quantifier-free, which implies that it is s-local for every $s \geq 0$. We denote the set of all (q, r)-independence sentences of vocabulary σ by $\Psi(\sigma, q, r)$.

Theorem 7.5 (Rank-Preserving Locality Theorem) *Let $q \in \mathbb{N}$ and $r = f_q(q)$. For every* $\text{FO}[\sigma]$*-formula $\varphi(x)$ of quantifier rank q there is an* $\text{FO}^+[\sigma \star^{q+1} q]$*-formula $\widehat{\varphi}(x)$, which is a Boolean combination of $(q + 1, r)$-independence sentences and atomic formulas, such that for every σ-structure A, every r-neighbourhood cover $\mathcal{X}$ of A, and every $a \in V(A)$,*

$$A \models \varphi(a) \iff A \star_{\mathcal{X}}^{q+1} q \models \widehat{\varphi}(a).$$

Furthermore, $\widehat{\varphi}$ is computable from φ.

Even though we need the theorem in this general form, it may be worthwhile to state, as a corollary, a version that does not refer to any neighbourhood cover. It is obtained by applying the theorem to the generic r-neighbourhood cover $\mathcal{X} = \{N_r(v) \mid v \in V(G)\}$. We omit the index $\mathcal{X}$ in the $\star$-notation when we refer to this neighbourhood cover. As a further simplification, we only state the corollary for sentences.

Corollary 7.6 *Let $q \in \mathbb{N}$ and $r = f_q(q)$. For every* $\text{FO}[\sigma]$*-sentence φ of quantifier rank q there is an* $\text{FO}^+[\sigma \star^{q+1} q]$*-sentence $\widehat{\varphi}$, which is a Boolean combination of $(q + 1, r)$-independence sentences, such that for every σ-structure A and every $a \in V(A)$,*

$$A \models \varphi \iff A \star^{q+1} q \models \widehat{\varphi}.$$

Furthermore, $\widehat{\varphi}$ is computable from φ.

To prove the theorem, it will be convenient to introduce the language of types. The (q, ℓ)-*type* of a tuple $\bar{a} \in V(A)^k$ in a σ-structure A is the set $\text{tp}_{q,\ell}^+(A, \bar{a})$ of all formulas $\varphi(\bar{x}) \in \Phi^+(\sigma, k, q, \ell)$ (normalised $\text{FO}^+[\sigma]$-formulas of q-rank at most ℓ) such that $A \models \varphi(\bar{a})$. Note that

$$(A, \bar{a}) \equiv_{q,\ell}^+ (B, \bar{b}) \iff \text{tp}_{q,\ell}^+(A, \bar{a}) = \text{tp}_{q,\ell}^+(B, \bar{b}).$$

We call $\text{atp}_q^+(A, \bar{a}) := \text{tp}_{q,0}^+(A, \bar{a})$ the *atomic q-type* of $\bar{a}$ in A. We denote the set of all (q, ℓ)-types of k-tuples in σ-structures by $T(\sigma, k, q, \ell)$.

The (q, r)-*independence type* of a structure A is the set $\text{itp}_{q,r}^+(A)$ of all (q', r')-independence sentences for $q' \leq q$ and $r' \leq r$ that are satisfied by A. The set of all (q, r)-independence types of σ-structures is denoted by $I(\sigma, q, r)$.

Lemma 7.7 *Let $q \in \mathbb{N}$ and $r := f_q(q)$. Let A, B be σ-structures and $\mathcal{X}, \mathcal{Y}$ r-neighbourhood covers of A, B, respectively. Let $a_0 \in V(A), b_0 \in V(B)$ such that*

$$\text{itp}_{q+1,r}^+ \left(A \star_{\mathcal{X}}^q q \right) = \text{itp}_{q+1,r}^+ \left(B \star_{\mathcal{Y}}^q q \right)$$
$$\text{and } \text{atp}_q^+ (A \star_{\mathcal{X}}^{q+1} q, a_0) = \text{atp}_q^+ (B \star_{\mathcal{Y}}^{q+1} q, b_0).$$

Then $(A, a_0) \equiv_{q,q}^+ (B, b_0)$.

Proof. We start by fixing some notation. For $0 \leq k \leq q$, we let $\sigma_k := \sigma \star^{q-k} q$ and $A_k := A \star_{\mathcal{X}}^{q-k} q$ and $B_k := B \star_{\mathcal{Y}}^{q-k} q$ and $r_k := f_q(q - k)$. Throughout the proof, $\bar{x}$ always denotes a tuple $(x_0, \ldots, x_k)$ (for varying k), and similarly $\bar{a}, \bar{b}$ denote tuples $(a_0, \ldots, a_k)$ and $(b_0, \ldots, b_k)$. We write $J \sqsubseteq H$ to denote that J

20

is a connected component of a graph H. Furthermore, if $V(H) = \{0, \ldots, k\}$ and $J \sqsubseteq H$, then $\bar{x}_J$ denotes the sub-tuple of $\bar{x}$ with entries x_j for $j \in V(J)$, and $\bar{a}_J, \bar{b}_J$ denote the corresponding sub-tuples of $\bar{a}, \bar{b}$.

We shall prove that Duplicator has a winning strategy for the q-round EF_q^+ game on (A, a_0, B, b_0). We describe a winning strategy for Duplicator satisfying the following conditions for every position $p = (A, \bar{a}, B, \bar{b})$, where $\bar{a} = (a_0, a_1, \ldots, a_k)$ and $\bar{b} = (b_0, b_1, \ldots, b_k)$, of the game that can be reached if Duplicator plays according to this strategy. Let H_p be the graph with vertex set $V(H_p) = \{0, \ldots, k\}$ and edge set

$$E(H_p) := \{ij : \mathrm{dist}(a_i, a_j) \le r_k \text{ or } \mathrm{dist}(b_i, b_j) \le r_k\}.$$

Then for every component $J \sqsubseteq H_p$ there are induced substructures $A_J \subseteq A_k$, $B_J \subseteq B_k$ such that the following conditions are satisfied:

(i) $N_{r_k}(a_j) \subseteq V(A_J)$ and $N_{r_k}(b_j) \subseteq V(B_J)$ for all $j \in V(J)$;

(ii) $(A_J, \bar{a}_J) \equiv_{q, q-k}^+ (B_J, \bar{b}_J)$.

Note that this implies that $\bar{a} \mapsto \bar{b}$ is a partial $f_q(q - k)$-isomorphism.

The proof is by induction on k. For the base step $k = 0$, note that the graph $H := H_p$ is the one-vertex graph, which is connected. We let $A_H := A_0[\mathcal{X}(a_0)]$ and $B_H := B_0[\mathcal{Y}(b_0)]$. Then (i) holds, because $\mathcal{X}, \mathcal{Y}$ are r-neighbourhood covers and $r = r_0$. By the assumption of the lemma, we have $\mathrm{atp}_q^+(A_0 \star_{\mathcal{X}}^{q+1} q, a_0) = \mathrm{atp}_q^+(B_0 \star_{\mathcal{Y}}^{q+1} q, b_0)$. In particular, for every formula $\varphi(x) \in \Phi^+(\sigma_0, 1, q, q)$ we have $A_0 \star_{\mathcal{X}}^{q+1} q \models P_\varphi(a_0) \iff B_0 \star_{\mathcal{Y}}^{q+1} q \models P_\varphi(b_0)$, which implies $A_H \models \varphi(a_0) \iff B_H \models \varphi(b_0)$ by the definition of the $\star$-operator. As every $\mathrm{FO}^+[\sigma_0]$-formula $\varphi(x_0)$ of q-rank at most q is equivalent to a formula in $\Phi^+(\sigma_0, 1, q, q)$, this implies $(A_H, a_0) \equiv_{q,q}^+ (B_H, b_0)$, that is, assertion (ii).

For the inductive step, suppose that we are in a position $p = (A, \bar{a}, B, \bar{b})$, where $\bar{a} = (a_0, a_1, \ldots, a_k)$ and $\bar{b} = (b_0, b_1, \ldots, b_k)$ for some $k < q$. Again, let $H := H_p$. Suppose that in the $(k+1)$st round of the game, Spoiler picks $a_{k+1} \in V(A)$.

Case 1: $\mathrm{dist}(a_{k+1}, a_i) \le r_k$ for some $i \in \{0, \ldots, k\}$.

Let $I \sqsubseteq H$ be the connected component of i, and let $A_I \subseteq A_k$, $B_I \subseteq B_k$ be substructures satisfying (i) and (ii). By (i), $a_{k+1} \in V(A_I)$. By (ii), $(A_I, \bar{a}_I) \equiv_{q, q-k}^+ (B_I, \bar{b}_I)$, and thus Duplicator has a winning strategy for the $q - k$-round EF_q^+-game on $(A_I, \bar{a}_I, B_I, \bar{b}_I)$. Let b_{k+1} be Duplicator's answer if Spoiler picks a_{k+1} in this game. Then

$$(A_I, \bar{a}_I a_{k+1}) \equiv_{q, q-k-1}^+ (B_I, \bar{b}_I b_{k+1}). \tag{7.3}$$

This implies $\mathrm{atp}_q^+(A_k, a_{k+1}) = \mathrm{atp}_q^+(B_k, b_{k+1})$ and thus

$$(A_{k+1}[\mathcal{X}(a_{k+1})], a_{k+1}) \equiv_{q, q-k-1}^+ (B_{k+1}[\mathcal{Y}(b_{k+1})], b_{k+1}). \tag{7.4}$$

We choose b_{k+1} as Duplicator's answer in the game on A, B. Thus the new position is

$$p' := (A, \bar{a} a_{k+1}, B, \bar{b} b_{k+1}).$$

Let $H' := H_{p'}$.

Case 1a: $\mathrm{dist}(a_{k+1}, a_i) \le r_{k+1}$ for some $i \in \{0, \ldots, k\}$.

Then

$$N_{r_{k+1}}(a_{k+1}) \subseteq N_{r_k}(\bar{a}) \subseteq V(A_I), \tag{7.5}$$

because $r_k \ge 2 r_{k+1}$, and

$$N_{r_{k+1}}(b_{k+1}) \subseteq N_{r_k}(\bar{b}) \subseteq V(B_I), \tag{7.6}$$

21

because $(q, q - k - 1)^+$-equivalence preserves distances up to r_{k+1}.

Let $J' \sqsubseteq H'$. Then there is a $J \sqsubseteq H$ such that $V(J') \cap \{0, \ldots, k\} \subseteq V(J)$. To see this, just note that if $j(k+1) \in E(H')$ and $(k+1)j' \in E(H')$ then $jj' \in E(H)$, because $2r_{k+1} \leq r_k$. Thus, whenever there is a path between two vertices $j, j' \in \{0, \ldots, k\}$ in H' there also is a path in H. We let $A_{J'} \subseteq A_{k+1}$ be the restriction of $A_J \subseteq A_k$ to σ_{k+1} and $B_{J'} \subseteq B_{k+1}$ the restriction of $B_J \subseteq A_k$ to σ_{k+1}. Then if $J = I$ and hence $k+1 \in V(J')$, (i) for p' and $A_{J'}, B_{J'}$ follows from (7.5) and (7.6), and (ii) follows from (7.3). If $J \neq I$, then (i) and (ii) for p' and $A_{J'}, B_{J'}$ are inherited from (i) and (ii) for p and A_J, B_J.

Case 1b: $\mathrm{dist}(a_{k+1}, a_i) > r_{k+1}$ for all $i \in \{0, \ldots, k\}$.

Let $J' \sqsubseteq H'$. Then either $V(J') = \{k+1\}$, or there is a $J \sqsubseteq H$ such that $V(J') \subseteq V(J)$. If $V(J') = \{k+1\}$, we let $A_{J'} := A_{k+1}\big[\mathcal{X}(a_{k+1})\big]$ and $B_{J'} := B_{k+1}\big[\mathcal{Y}(b_{k+1})\big]$. Then (i) holds because $\mathcal{X}$ and $\mathcal{Y}$ are r-neighbourhood covers, and (ii) follows from (7.4). If there is a connected component J of H such that $V(J') \subseteq V(J)$, we let $A_{J'} \subseteq A_{k+1}$ be the restriction of $A_J \subseteq A_k$ to σ_{k+1} and $B_{J'} \subseteq B_{k+1}$ the restriction of $B_J \subseteq B_k$ to σ_{k+1}. Then (i) and (ii) for p' and $A_{J'}, B_{J'}$ are inherited from (i) and (ii) for p and A_J, B_J.

Case 2: $\mathrm{dist}(a_{k+1}, a_i) > r_k$ for all $i \in \{0, \ldots, k\}$.

Let $t := \mathrm{atp}_q^+(A_k, a_{k+1})$. We will prove the existence of a $b_{k+1} \in V(B)$ with $\mathrm{atp}_q^+(B_k, b_{k+1}) = t$ and $\mathrm{dist}(b_{k+1}, b_i) > r_{k+1}$ for all $i \in \{0, \ldots, k\}$. We can then argue as in Case 1b. Assume towards a contradiction that

(A) there is no $b \in V(B)$ with $\mathrm{atp}_q^+(B_k, b) = t$ and $\mathrm{dist}(b, b_i) > r_{k+1}$ for all $i \in \{0, \ldots, k\}$.

The first step is to construct d, D, ℓ such that $2r_{k+1} \leq d \leq D - 4r_{k+1}$ and $D \leq r_k$ and $\ell \leq k$ and there are elements $a^0, \ldots, a^\ell \in V(A)$ with $\mathrm{atp}_q^+(A_k, a^i) = t$ and $\mathrm{dist}(a^i, a^j) > D$ for $i \neq j \in \{0, \ldots, \ell\}$, but no elements $a_*^0, \ldots, a_*^{\ell+1} \in V(A)$ with $\mathrm{atp}_q^+(A_k, a_*^i) = t$ and $\mathrm{dist}(a_*^i, a_*^j) > d$ for $i \neq j \in \{0, \ldots, \ell+1\}$.

We let $d_0 := 2r_{k+1}$, and we let ℓ_0 be maximal such that there are $a_0^0, \ldots, a_0^{\ell_0}$ with $\mathrm{atp}_q^+(A_k, a_0^i) = t$ for all $i \in \{0, \ldots, \ell_0\}$ and $\mathrm{dist}(a_0^i, a_0^j) > d_0$ for all $i \neq j \in \{0, \ldots, \ell_0\}$. Suppose first that $\ell_0 > k$. As A and B satisfy the same $(k+1, d_0/2)$-independence sentences (note that d_0 is even), there are elements $b_0^0, \ldots, b_0^{k+1} \in V(B)$ with $\mathrm{atp}_q^+(B_k, b_0^i) = t$ for all $i \in \{0, \ldots, k+1\}$ and $\mathrm{dist}(b_0^i, b_0^j) > d_0$. By (A), for every $i \in \{0, \ldots, k+1\}$ there is a $j(i) \in \{0, \ldots, k\}$ such that $\mathrm{dist}(b_0^i, b_{j(i)}) \leq r_{k+1} = d_0/2$. As $\mathrm{dist}(b_0^i, b_0^j) > d_0$, we have $j(i) \neq j(i')$ for $i \neq i' \in \{0, \ldots, k+1\}$. This is a contradiction, which proves that $\ell_0 \leq k$.

Now suppose that d_h, ℓ_h are defined for some $h \geq 0$. Let $d_{h+1} := d_h + 4r_{k+1}$, and let ℓ_{h+1} be maximal such that there are $a_{h+1}^0, \ldots, a_{h+1}^{\ell_{h+1}}$ with $\mathrm{atp}_q^+(A_k, a_{h+1}^i) = t$ for all $i \in \{0, \ldots, \ell_{h+1}\}$ and $\mathrm{dist}(a_{h+1}^i, a_{h+1}^j) > d_{h+1}$ for all $i \neq j \in \{0, \ldots, \ell_{h+1}\}$. Then $\ell_{h+1} \leq \ell_h$. If $\ell_{h+1} = \ell_h$ for the first time, we stop the construction. Then $h \leq k$ and thus $d_{h+1} = (4(h+1) - 2)r_{k+1} \leq r_k$. We let $d := d_h$ and $D := d_{h+1}$ and $\ell := \ell_h = \ell_{h+1}$.

As A and B satisfy the same $(k+1, D/2)$-independence sentences, there are elements $b^0, \ldots, b^\ell \in V(B)$ with $\mathrm{atp}_q^+(B_k, b^i) = t$ and $\mathrm{dist}(b^i, b^j) > D$. Then for every $i \in \{0, \ldots, \ell\}$ there is a $j(i) \in \{0, \ldots, k\}$ such that $\mathrm{dist}(b^i, b_{j(i)}) \leq r_{k+1}$. The $j(i)$ are mutually distinct, because $\mathrm{dist}(b^i, b^j) > 2r_{k+1}$ for $i \neq j$. To simplify the notation, let us assume that $j(i) = i$ for all $i \in \{0, \ldots, \ell\}$. As $\mathrm{dist}(b^i, b^j) > D$, we have $\mathrm{dist}(b_i, b_j) > D - 2r_{k+1}$. Then it follows from (ii) that $\mathrm{dist}(a_i, a_j) > D - 2r_{k+1}$, because $D - 2r_{k+1} \leq r_k$. It also follows from (ii) that for all $i \in \{0, \ldots, \ell\}$ there is an a_*^i such that $\mathrm{dist}(a_*^i, a_i) \leq r_{k+1}$ and $\mathrm{atp}_q^+(A_k, a_*^i) = t$. Then for $i \neq j$ we have $\mathrm{dist}(a_*^i, a_*^j) > D - 4r_{k+1} \geq d$. Furthermore, we have $\mathrm{dist}(a_{k+1}, a_*^i) > r_k - r_{k+1} \geq d$. Letting $a_*^{\ell+1} := a_{k+1}$, we

have found $a_*^1, \ldots, a_*^{\ell+1} \in V(A)$ with $\mathrm{atp}_q^+(A_k, a_*^i) = t$ and $\mathrm{dist}(a_*^i, a_*^j) > d$. This is a contradiction.

$\square$

We will show next how the Rank Preserving Locality Theorem follows from this lemma by standard techniques from logic.

Proof of the Rank Preserving Locality Theorem. Let $\varphi(x) \in \mathrm{FO}[\sigma]$ be a first-order formula of quantifier rank q. Let $r := f_q(q)$ and $\sigma_I := \sigma \star^q q$ and $\sigma_T := \sigma \star^{q+1} q$. Furthermore, let $I := I(\sigma_I, q+1, r)$ and $T := T(\sigma_T, 1, q, 0)$. A pair $(\eta, \theta) \in I \times T$ is *satisfiable* if there are a σ-structure A and an r-neighbourhood cover $\mathcal{X}$ of A and an $a \in V(A)$ such that $\mathrm{itp}_{q+1,r}^+(A \star_{\mathcal{X}}^q q) = \eta$ and $\mathrm{atp}_q^+(A \star_{\mathcal{X}}^{q+1} q, a) = \theta$.

It follows from Lemma 7.7 that for all satisfiable pairs $(\eta, \theta) \in I \times T$ the following two statements are equivalent.

(A) There are a σ-structure A and an r-neighbourhood cover $\mathcal{X}$ of A and an $a \in V(A)$ such that $\mathrm{itp}_{q+1,r}^+(A\star_{\mathcal{X}}^q q) = \eta$ and $\mathrm{atp}_q^+(A \star_{\mathcal{X}}^{q+1} q, a) = \theta$ and $A \models \varphi(a)$.

(B) For all σ-structures A and r-neighbourhood covers $\mathcal{X}$ of A and $a \in V(A)$, if $\mathrm{itp}_{q+1,r}^+(A \star_{\mathcal{X}}^q q) = \eta$ and $\mathrm{atp}_q^+(A \star_{\mathcal{X}}^{q+1} q, a) = \theta$, then $A \models \varphi(a)$.

Thus there is a subset $S_\varphi \subseteq I \times T$ such that for all σ-structures A, all r-neighbourhood covers $\mathcal{X}$ of A, and all $a \in V(A)$,

$$A \models \varphi(a) \iff \exists (\eta, \theta) \in S_\varphi : \mathrm{itp}_{q+1,r}^+(A \star_{\mathcal{X}}^q q) = \eta \text{ and } \mathrm{atp}_q^+(A \star_{\mathcal{X}}^q q, a) = \theta. \qquad (7.7)$$

Recall that every $(q+1, r)$-independence type $\eta \in I$ is a subset of the finite set $\Psi(\sigma_I, q+1, r)$, and for every σ_I-structure A we have

$$\mathrm{itp}_{q+1,r}^+(A) = \eta \iff A \models \bigwedge_{\psi \in \eta} \psi \wedge \bigwedge_{\psi \in \Psi(\sigma_I, q+1, r) \setminus \eta} \neg \psi.$$

We denote the sentence $\bigwedge_{\psi \in \eta} \psi \wedge \bigwedge_{\psi \in \Psi(\sigma_I, q+1, r) \setminus \eta} \neg \psi$ by $\widetilde{\eta}$ and say that it *defines* the type η. But we can actually define $\widetilde{\eta}$ for every subset $\eta \subseteq \Psi(\sigma_I, q+1, r)$. Then either $\widetilde{\eta}$ is unsatisfiable or there is some σ_I-structure A such that $\mathrm{itp}_{q+1,r}^+(A) = \eta$.

Similarly, every atomic type $\theta \in T(\sigma_T, 1, q, 0)$ is a subset of the finite set $\Phi^+(\sigma_T, 1, q, 0)$, and for every σ_T-structure A and every $a \in V(A)$ we have

$$\mathrm{atp}_q^+(A, a) = \theta \iff A \models \bigwedge_{\zeta(x) \in \theta} \zeta(a) \wedge \bigwedge_{\zeta(x) \in \Phi(\sigma_T, 1, q, 0) \setminus \theta} \neg \zeta(a).$$

We denote the formula $\bigwedge_{\zeta(x) \in \theta} \zeta(x) \wedge \bigwedge_{\zeta(x) \in \Phi(\sigma_T, 1, q, 0) \setminus \theta} \neg \zeta(x)$ by $\widetilde{\theta}(x)$. Again, we can define $\widetilde{\theta}(x)$ for every subset $\theta \subseteq \Phi^+(\sigma_T, 1, q, 0)$. Then either $\widetilde{\theta}(x)$ is unsatisfiable, or there is some σ_T-structure A and $a \in V(A)$ such that $\mathrm{atp}_q^+(A, a) = \theta$.

It follows from (7.7) that for all σ-structures A, all r-neighbourhood covers $\mathcal{X}$ of A, and all $a \in V(A)$,

$$A \models \varphi(a) \iff A \star_{\mathcal{X}}^{q+1} q \models \bigvee_{(\eta, \theta) \in S_\varphi} (\widetilde{\eta} \wedge \widetilde{\theta}(a)). \qquad (7.8)$$

Here we use that the σ_T-structure $A \star_{\mathcal{X}}^{q+1} q$ is an expansion of the σ_I-structure $A \star_{\mathcal{X}}^q q$.

We could let $\widehat{\varphi}(x) = \bigvee_{(\eta, \theta) \in S_\varphi} (\widetilde{\eta} \wedge \widetilde{\theta}(x))$. Clearly, this formula has the desired syntactic form, and by (7.8) satisfies the assertion of the theorem. However, we want $\widehat{\varphi}(x)$ to be computable from $\varphi(x)$, and with

23

this definition, it is not, because the choice of S_φ is not unique and, so far, arbitrary. However, we will prove that we can compute some set S_φ satisfying (7.8).

We need to incorporate the r-neighbourhood covers into the logical framework. Let R be a fresh binary relation symbol and $\sigma_R := \sigma \cup \{R\}$. For every σ-structure A and every mapping $\mathcal{X} : V(A) \to 2^{V(A)}$, we let $A^{\mathcal{X}}$ be the $\sigma \cup \{R\}$-expansion of A with

$$R(A^{\mathcal{X}}) = \{ab \mid b \in \mathcal{X}(a)\}.$$

Recall that we view r-neighbourhood covers of A as mappings $\mathcal{X} : V(A) \to 2^{V(A)}$ where $N_r(a) \subseteq \mathcal{X}(a)$ for each $a \in V(A)$. We let $\gamma := \forall x \forall y (\mathrm{dist}(x,y) \leq r \longrightarrow R(x,y))$. Then $\mathcal{X}$ is an r-neighbourhood cover of A if, and only if, $A^{\mathcal{X}} \models \gamma$. It is not hard to see that the structure $A \star_{\mathcal{X}} q$ is definable within $A^{\mathcal{X}}$, which means that for every (unary) relation symbol $P \in (\sigma \star q) \setminus \sigma$ there is a $\sigma \cup \{R\}$-formula $\chi_P(x)$ such that $P(A \star_{\mathcal{X}} q) = \{a \in V(A) \mid A^{\mathcal{X}} \models \chi_P(a)\}$. By the so-called Lemma on Syntactical Interpretations (see [15]), this implies that for every $\sigma \star q$-formula $\psi(x)$ there is a $\sigma \cup \{R\}$-formula $\psi_R(x)$ such that $A \star_{\mathcal{X}} q \models \psi(a) \iff A^{\mathcal{X}} \models \psi_R(a)$. Using this, we can inductively prove that $A \star_{\mathcal{X}}^\ell q$ is definable within $A^{\mathcal{X}}$ and that for every $\sigma \star^\ell q$-formula $\psi(x)$ there is a $\sigma \cup \{R\}$-formula $\psi_R(x)$ such that $A \star_{\mathcal{X}}^\ell q \models \psi(a) \iff A^{\mathcal{X}} \models \psi_R(a)$. In particular, for every $\eta \subseteq \Psi(\sigma_I, q+1, r)$ there is a σ_R-sentence $\widetilde{\eta}_R$ such that $A \star_{\mathcal{X}}^{q+1} q \models \widetilde{\eta} \iff A^{\mathcal{X}} \models \widetilde{\eta}_R$ and for every $\theta(x) \subseteq \Phi(\sigma_T, 1, q, 0)$ there is a σ_R-sentence $\widetilde{\theta}_R(x)$ such that $A \star_{\mathcal{X}}^{q+1} q \models \widetilde{\theta}(a) \iff A^{\mathcal{X}} \models \widetilde{\theta}_R(a)$.

It follows from (7.8) that for all σ-structures A, all r-neighbourhood covers $\mathcal{X}$ of A, and all $a \in V(A)$,

$$A \models \varphi(a) \iff A^{\mathcal{X}} \models \bigvee_{(\eta,\theta) \in S_\varphi} \left(\widetilde{\eta}_R \wedge \widetilde{\theta}_R(a) \right). \tag{7.9}$$

As $A^{\mathcal{X}}$ is an expansion of A, on the left-hand side of (7.9) we can replace A by $A^{\mathcal{X}}$ and thus rewrite (7.9) as

$$A^{\mathcal{X}} \models \varphi(a) \longleftrightarrow \bigvee_{(\eta,\theta) \in S_\varphi} \left(\widetilde{\eta}_R \wedge \widetilde{\theta}_R(a) \right). \tag{7.10}$$

Recalling that a σ_R-structure A_R equals $A^{\mathcal{X}}$ for some r-neighbourhood cover $\mathcal{X}$ of a σ-structure A if any only if $A_R \models \gamma$, for all σ_R-structures A_R and all $a \in V(A_R)$ we thus have

$$A_R \models \gamma \longrightarrow \left(\varphi(a) \longleftrightarrow \bigvee_{(\eta,\theta) \in S_\varphi} \left(\widetilde{\eta}_R \wedge \widetilde{\theta}_R(a) \right) \right). \tag{7.11}$$

For every subset $S \subseteq I \times T$, let

$$\alpha_S(x) = \gamma \longrightarrow \left(\varphi(x) \longleftrightarrow \bigvee_{(\eta,\theta) \in S} \left(\widetilde{\eta}_R \wedge \widetilde{\theta}_R(x) \right) \right).$$

By (7.11), the formula $\alpha_{S_\varphi}(x)$ is valid. Note that so far we thought of $\alpha_S(x)$ as an FO$^+$-formula, but we can directly translate every FO$^+$-formula into an equivalent FO-formula by substituting appropriate distance formulas for the distance atoms. This changes the rank, but at this point we no longer care about the rank. Thus we view $\alpha_S(x)$ as an FO$[\sigma_R]$-formula.

The set of all valid FO$[\sigma_R]$-formulas is recursively enumerable. We start an enumeration algorithm and wait for the first formula $\alpha_S(x)$ it produces. This will happen eventually, because we know that $\alpha_{S_\varphi}(x)$ is valid. The set $S \subseteq I \times T$ of the first formula $\alpha_S(x)$ returned by enumeration algorithm is not necessarily

the same as the set S_φ we started with. However, by retracing our construction backwards, it is easy to see that S satisfies (7.8), that is, for all σ-structures A, all r-neighbourhood covers $\mathcal{X}$ of A, and all $a \in V(A)$,

$$A \models \varphi(a) \iff A \star_\mathcal{X}^{q+1} q \models \bigvee_{(\eta,\theta)\in S} \left(\widetilde{\eta} \wedge \widetilde{\theta}(a)\right).$$

We define $\widehat{\varphi}(x) := \bigvee_{(\eta,\theta)\in S} \left(\widetilde{\eta} \wedge \widetilde{\theta}(x)\right)$. As argued above, this formula satisfies the conditions of the theorem, and by construction it is computable from $\varphi(x)$.

Note that if, given a formula φ, we first compute an equivalent normalised formula φ' and then apply the procedure above to φ', then we can compute an upper bound for the running time. $\qquad\square$

8 The Main Algorithm

We are now ready to prove our main result, Theorem 1.1. We actually prove a slightly more general theorem. A *coloured-graph vocabulary* consists of the binary relation symbol E and possibly finitely many unary relation symbols. In particular, if σ is a coloured-graph vocabulary then $\sigma \star q$ (as defined in Section 7.4) is a coloured graph vocabulary. A *σ-coloured graph* is a σ-structure whose $\{E\}$-restriction is a simple undirected graph.[3] We call the $\{E\}$-restriction of a σ-colored graph the *underlying graph* of G.

Theorem 8.1 *For every nowhere dense class $\mathcal{C}$, every $\varepsilon > 0$, every coloured graph vocabulary σ, and every first-order formula $\varphi(x) \in \mathrm{FO}[\sigma]$, there is an algorithm that, given a σ-coloured graph G whose underlying graph is in $\mathcal{C}$, computes the set of all $v \in V(G)$ such that $G \models \varphi(v)$ in time $\mathcal{O}(n^{1+\varepsilon})$.*

Furthermore, if $\mathcal{C}$ is effectively nowhere dense, then there is a computable function f and an algorithm that, given $\varepsilon > 0$, a formula $\varphi(x) \in \mathrm{FO}[\sigma]$ for some coloured-graph vocabulary σ, and a σ-coloured graph G, computes the set of all $v \in V(G)$ such that $G \models \varphi(v)$ in time $f(|\varphi|, \varepsilon) \cdot n^{1+\varepsilon}$.

Clearly, this implies Theorem 1.1.

We need one more lemma for the proof. It describes a standard reduction that allows us to remove a bounded number of elements from a structure in which we want to evaluate a formula.

Lemma 8.2 *Let σ be a coloured-graph vocabulary and $k, \ell, m, q \in \mathbb{N}$ with $0 \le \ell \le q$. Then there are*

1. *a coloured-graph vocabulary $\sigma' \supseteq \sigma$,*

2. *for every $\mathrm{FO}^+[\sigma]$-formula $\varphi(x_1, \ldots, x_k, y_1, \ldots, y_m)$ of q-rank ℓ and every atomic q-type $\theta \in T(\sigma, m, q, 0)$ an $\mathrm{FO}^+[\sigma']$ formula $\varphi^\theta(x_1, \ldots, x_k)$ of q-rank at most ℓ,*

3. *for every σ-coloured graph G and all $w_1, \ldots, w_m \in V(G)$ a σ'-expansion G' of $G \setminus \{w_1, \ldots, w_m\}$,*

such that if $\mathrm{atp}_q^+(G, w_1, \ldots, w_m) = \theta$ then for all $v_1, \ldots, v_k \in V(G) \setminus \{w_1, \ldots, w_m\}$

$$G \models \varphi(v_1, \ldots, v_k, w_1, \ldots, w_m) \iff G' \models \varphi^\theta(v_1, \ldots, v_k).$$

Furthermore, φ^θ is computable from φ and θ, and G' is computable from G and $w_1, \ldots, w_m$ in time $f(\ell, m, q) \cdot (|V(G)| + |E(G)|)$.

Proof. We use a game theoretic argument similar to (but simpler than) the proof of the rank preserving locality theorem.

---

[3] To see that this is consistent with the definition of coloured graphs in Section 5, we may define the *colour* of a vertex v in a σ-coloured graph G to be the set of all unary relation symbols $P \in \sigma$ such that $v \in P(G)$.

For $1 \leq i \leq f_q(\ell)$ and $1 \leq j \leq m$, we let Q_{ij} be a fresh unary relation symbol, and we let σ' be the union of σ with all these Q_{ij}. For every σ-coloured graph G and all $w_1, \ldots, w_m \in V(G)$ we let G' be the σ'-expansion of $G \setminus \{w_1, \ldots, w_m\}$ with

$$Q_{ij}(G') = \{v \in V(G) \setminus \{w_1, \ldots, w_m\} \mid \mathrm{dist}^G(v, w_j) = i\}.$$

Clearly, G' can be computed from G in time $f(\ell, m, q) \cdot (|V(G)| + |E(G)|)$, for some function f.

Claim 2. Let G_1, G_2 be σ-coloured graphs and $v_{11}, \ldots, v_{1k}, w_{11}, \ldots, w_{1m} \in V(G_1), v_{21}, \ldots, v_{2k}, w_{21}, \ldots, w_{2m} \in V(G_2)$ such that

$$\mathrm{atp}_q^+(G_1, w_{11}, \ldots, w_{1m}) = \mathrm{atp}_q^+(G_2, w_{21}, \ldots, w_{2m})$$

and

$$G_1', (v_{11}, \ldots, v_{1k}) \equiv_{(q,\ell)}^+ G_2', (v_{21}, \ldots, v_{2k}).$$

Then

$$G_1, (v_{11}, \ldots, v_{1k}, w_{11}, \ldots, w_{1m}) \equiv_{q,\ell}^+ G_2, (v_{21}, \ldots, v_{2k}, w_{21}, \ldots, w_{2m}).$$

Proof. It is easy to see that Duplicator has a winning strategy for the ℓ-round EF_q^+-game on $(G, (v_{11}, \ldots, v_{1k}, w_{11}, \ldots, w_{1m}), G_2, (v_{21}, \ldots, v_{2k}, w_{21}, \ldots, w_{2m}))$: she simply plays according to a winning strategy for the ℓ-round EF_q^+-game on $(G_1', (v_{11}, \ldots, v_{1k}), G_2', (v_{21}, \ldots, v_{2k}))$, and whenever Spoiler selects a w_{ij} she answers by selecting $w_{(3-i)j}$. ⊣

The claim implies that there is a set $S_{\varphi,\theta} \subseteq T(\sigma', k, q, \ell)$ such that

$$G \models \varphi(v_1, \ldots, v_k, w_1, \ldots, w_m) \iff G' \models \bigvee_{\eta \in S_{\varphi,\theta}} \bigwedge_{\psi(x_1, \ldots, x_k) \in \eta} \psi(v_1, \ldots, v_k).$$

It remains to prove that we can compute such a set $S_{\varphi,\theta}$ from φ and θ. We use an argument based on the recursive enumerability of the valid first-order sentences similar to the one in the proof of the Rank Preserving Locality Theorem. □

Proof of Theorem 8.1. Let $\mathcal{C}$ be a nowhere dense class of graphs and $\varepsilon > 0$. Without loss of generality we may assume that $\varepsilon \leq 1/2$, which implies $\varepsilon^2 \leq \varepsilon/2$, and that $\mathcal{C}$ is closed under taking subgraphs.

The input to our algorithm is an $\varepsilon \leq 1/2$, a σ-coloured graph G whose $\{E\}$-restriction is in $\mathcal{C}$ and an $FO^+[\sigma]$-formula $\varphi(x)$, for some coloured-graph vocabulary σ. Our algorithm will compute the set of all $v \in V(G)$ such that $G \models \varphi(v)$ in time $\mathcal{O}(n^{1+\varepsilon})$.

We start by fixing a few parameters. We choose q such that the q-rank of φ is at most q and let $r = f_q(q)$. By the Rank-Preserving Locality Theorem, we can find an $FO^+[\sigma \star^{q+1} q]$-formula $\widehat{\varphi}(x)$, which is a Boolean combination of $(q+1, r)$-independence sentences and atomic formulas, such that for all σ-coloured graphs G, all r-neighbourhood covers $\mathcal{X}$ of G, and all $v \in V(G)$ we have $G \models \varphi(v) \iff G \star_{\mathcal{X}}^{q+1} q \models \widehat{\varphi}(v)$. We choose ℓ, m according to Theorem 4.2 such that Splitter has a winning strategy for the $(\ell, m, 2r)$-splitter game on every graph in $\mathcal{C}$. Note that q, r, ℓ, m and $\widehat{\varphi}$ only depend on φ and the class $\mathcal{C}$, but not on ε or the input graph G. Now ε comes into play. Let $\delta = \varepsilon/(2\ell)$. Choose $n_0 = n_0(\delta, r)$ according to Theorem 6.2 such that every graph $G \in \mathcal{C}$ of order $n \geq n_0$ has an r-neighbourhood cover of radius at most $2r$ and maximum degree at most n^δ. Choose $n_1 \geq n_0$ such that $n_1^{\delta/2} \geq 2$ and that every graph $G \in \mathcal{C}$ of order $n \geq n_1$ has at most $n^{1+\delta}$ edges. The existence of such an n_1 follows from Lemma 3.3. All the parameters and the formula $\widehat{\varphi}(x)$ can be computed from φ, ε and the nowhere-density parameters of $\mathcal{C}$ if $\mathcal{C}$ is effectively nowhere dense.

Now consider the σ-coloured input graph G. If $n = |V(G)| < n_1$, we compute the set of all $v \in V(G)$ such that $G \models \varphi(v)$ by brute force; in this case the running time can be bounded in terms of φ, ε, and $\mathcal{C}$. So let us assume that $n \geq n_1$. We compute an r-neighbourhood cover $\mathcal{X}$ of G of radius $2r$ and maximum degree n^δ. The main task of our algorithm will be to compute $G \star_{\mathcal{X}}^{q+1} q$. Before we describe how to do this, let us assume that we have computed $G \star_{\mathcal{X}}^{q+1} q$ and describe how the algorithm proceeds from there. The next step is to evaluate all (q, r)-independence sentences in the Boolean combination $\widehat{\varphi}(x)$ in $G \star_{\mathcal{X}}^{q+1} q$. Consider such a sentence

$$\psi = \exists x_1 \ldots \exists x_q \Big(\bigwedge_{1 \leq i < j \leq q} \mathrm{dist}(x_i, x_j) > 2r \wedge \bigwedge_{1 \leq i \leq q} \chi(x_i) \Big).$$

Remember that $\chi(x_i)$ is an atomic formula. Thus we can easily compute the set U of all $v \in V(G)$ such that $G \star_{\mathcal{X}}^{q+1} q \models \chi(v)$. Then we can use the algorithm of Theorem 5.1 to decide if U has k elements of pairwise distance greater than $2r$. This is the case if and only if $G \star_{\mathcal{X}}^{q+1} q \models \psi$. This way, we decide which (q, r)-independence sentences in $\widehat{\varphi}(x)$ are satisfied in $G \star_{\mathcal{X}}^{q+1} q$. It remains to evaluate the atomic formulas in $\widehat{\varphi}(x)$ and combine the results to evaluate the Boolean combination. Both tasks are easy.

Let us now turn to computing $G \star_{\mathcal{X}}^{q+1} q$. We inductively compute $G \star_{\mathcal{X}}^{i} q$ for $0 \leq i \leq q + 1$. The base step $i = 0$ is trivial, because $G \star_{\mathcal{X}}^{0} q = G$. As each $G \star_{\mathcal{X}}^{i} q$ is a σ' coloured graph for some σ' (to be precise, $\sigma' = \sigma \star^i q$), it suffices to show how to compute $G \star_{\mathcal{X}} q$ from G. To do this, for each formula $\xi(x) \in \Phi^+(\sigma, 1, q, q)$ we need to compute the set $P_\xi(G \star_{\mathcal{X}} q)$ of all $v \in V(G)$ such that $G[\mathcal{X}(v)] \models \xi(v)$. Let us fix a formula $\xi(x) \in \Phi^+(\sigma, 1, q, q)$.

For every $X \in \mathcal{X}$, let $v_X \in X$ be a "centre" of $G[X]$, that is, a vertex with $X \subseteq N_{2r}(v_X)$. Such a v_X exists because the radius of $G[X]$ is at most $2r$. Let $W_X \subseteq N_{2r}^G$ be Splitter's response if Connector chooses v_X in the first round of the $(\ell, m, 2r)$-splitter game on G. Without loss of generality we assume that $W_X \neq \emptyset$. Let $w_1, \ldots, w_m$ be an enumeration of W_X. We apply Lemma 8.2 with $k = 1$, $\ell = q$, and m, q to the formulas $\xi_0(x_1, y_1 \ldots, y_m) = \xi(x_1)$ and $\xi_j(x_1, y_1 \ldots, y_m) = \xi(y_j)$ for $j = 1, \ldots, m$. Let σ' be the vocabulary obtained by Lemma 8.2 (1), and let G_X be the graph obtained from G and $w_1, \ldots, w_m$ by Lemma 8.2 (3). (Neither σ' nor G_X depend on the formula.) For $0 \leq j \leq m$, let $\xi'_j(x_1)$ be the formula obtained from ξ_j by Lemma 8.2 (2). We recursively evaluate the formulas $\xi'_0, \ldots, \xi'_1$ in G_X. This gives us the set Ξ_X of all $v \in V(G)$ such that $G[X] \models \xi(v)$. Doing this for all $X \in \mathcal{X}$, we can compute the set

$$P_\xi(G \star_{\mathcal{X}} q) = \{ v \in V(G) \mid G[\mathcal{X}(v)] \models \xi(v) \} = \bigcup_{X \in \mathcal{X}} \big(\Xi_x \cap \{ v \in V(G) \mid \mathcal{X}(v) = X \} \big).$$

The crucial observation to ensure that the algorithm terminates is that in a recursive call with input G_X, ξ'_j the parameters q and hence $r = f_q(q)$ can be left unchanged. Moreover, it follows from the definition of G_X that Splitter has a winning strategy for the $(\ell - 1, m, 2r)$-splitter game on G_X. Thus we can reduce the parameter ℓ by 1. Once we have reached $\ell = 0$, the graph G_X will be empty, and the algorithm terminates.

There is one more issue we need to attend to, and that is how we compute Splitter's winning strategy, that is, the sets W_X. We use Remark 4.3. This means that to compute W_X in some recursive call, we need the whole history of the game (in a sense, the whole call stack). In addition, we need a breadth-first search tree in all graphs that appeared in the game before. It is no problem to compute a breadth-first search tree once when we first need it and then store it with the graph; this only increases the running time by a constant factor.

This completes the description of the algorithm.

Let us analyse the running time. The crucial parameters are the order n of the input graph and the level j of the recursion. As argued above, we have $j \leq \ell$. We write the running time as a function T of j and n. We first observe that the time used by the algorithm without the recursive calls can be bounded by $c_1 n^{1+\delta}$ for a suitable constant c_1 depending on the input sentence φ, the parameter ε, and the class $\mathcal{C}$, but not on n

or j. Furthermore, for $n < n_1$ the running time can be bounded by a constant c_2 that again only depends on φ, ε, and $\mathcal{C}$, and for $j = 0$ the running time can be bounded by c_3. Furthermore, there is a c_4 such that for each $X \in \mathcal{X}$ at most c_4 recursive calls are made to the graph G_X. Let $n_X = |V(G_X)| \le |X|$ and $c = \max\{c_1, c_2, c_3, c_4\}$. We obtain the following recurrence for T:

$$T(0, n) \le c,$$
$$T(j, n) \le c \qquad\qquad\qquad\qquad\qquad\qquad \text{for all } n < n_1,$$
$$T(j, n) \le \sum_{X \in \mathcal{X}} cT(j-1, n_X) + cn^{1+\delta} \qquad\qquad \text{for all } j \ge 1, n \ge n_1$$

We claim that for all $n \ge 1$ and $0 \le j \le \ell$ we have

$$T(j, n) \le c^j n^{1+2j\delta} = c^\ell n^{1+\varepsilon}. \tag{8.1}$$

As c and ℓ are bounded in terms of $\varphi, \varepsilon, \mathcal{C}$, this proves the theorem.

(8.1) can be proved by a straightforward induction. The crucial observation is

$$\sum_{X \in \mathcal{X}} n_X = \sum_{v \in V(G)} |\{X \in \mathcal{X} \mid v \in X\}| \le nn^\delta = n^{1+\delta}. \tag{8.2}$$

The base steps $j = 0$ and $n < n_1$ are trivial. In the inductive step, we have

$$\begin{aligned}
T(j, n) &\le \sum_{X \in \mathcal{X}} cT(j-1, n_X) + cn^{1+\delta} \\
&\le \sum_{X \in \mathcal{X}} cc^{j-1} n_X^{1+2(j-1)\delta} + cn^{1+\delta} && \text{(Induction Hypothesis)} \\
&\le c^j \left(\sum_{X \in \mathcal{X}} n_X \right)^{1+2(j-1)\delta} + cn^{1+\delta} \\
&\le c^j n^{(1+\delta)(1+2(j-1)\delta)} + cn^{1+\delta} && \text{(by (8.2))} \\
&\le c^j \left(n^{1+(2j-1)\delta + 2(j-1)\delta^2} + n^{1+\delta} \right) \\
&\le c^j \left(\frac{n^{1+2j\delta} + n^{1+(3/2)\delta}}{n^{\delta/2}} \right) && \text{(because } 2(j-1)\delta^2 \le \frac{\varepsilon^2}{2\ell} \le \delta/2\text{)} \\
&\le c^j n^{2j\delta} && \text{(because } n^{\delta/2} \ge 2\text{).}
\end{aligned}$$

$\square$

9 Conclusion

We prove that deciding first-order properties is fixed-parameter tractable on nowhere dense graph classes. This generalises a long list of previous algorithmic meta theorems for first-order logic. Furthermore, it is optimal on classes of graphs closed under taking subgraphs. It remains open to find an optimal meta theorem for first-order properties on classes that are not closed under taking subgraphs, but only satisfy some weaker closure condition like being closed under taking induced subgraphs.

Our theorem underlines that nowhere dense graph classes have very favourable algorithmic properties. As opposed to Robertson and Seymour's structure theory underlying most algorithms on graph classes with excluded minors, the graph theory behind our algorithms does not cause enormous hidden constants in the running time.

A particularly interesting property of nowhere dense classes and classes of bounded expansion that we uncover here for the first time is that they have simple sparse neighbourhood covers with very good parameters. We have focussed on the radius of the covering sets and have not tried to optimise the degree of the cover, that is, the number of covering sets a vertex may be contained in. As the graph theory underlying our result is not very complicated, we believe that it is possible to obtain good degree bounds as well, probably much better than those obtained through graph minor theory [1, 3] (even though the classes we consider are much larger). However, this remains future work.

References

[1] I. Abraham, C. Gavoille, D. Malkhi, and U. Wieder. Strong-diameter decompositions of minor free graphs. In *Proceedings of the nineteenth annual ACM Symposium on Parallel Algorithms and Architectures*, pages 16–24, 2007.

[2] H.L. Bodlaender, F.V. Fomin, D. Lokshtanov, E. Penninkx, S. Saurabh, and D.M. Thilikos. (Meta) Kernelization. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 629–638, 2009.

[3] C. Busch, R. LaFortune, and S. Tirthapura. Improved sparse covers for graphs excluding a fixed minor. In *Proceedings of the twenty-sixth annual ACM Symposium on Principles of Distributed Computing*, pages 61–70, 2007.

[4] B. Courcelle. Graph rewriting: An algebraic and logic approach. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 194–242. Elsevier Science Publishers, 1990.

[5] B. Courcelle, J.A. Makowsky, and U. Rotics. Linear time solvable optimization problems on graphs of bounded clique width. *Theory of Computing Systems*, 33(2):125–150, 2000.

[6] B. Courcelle, J.A. Makowsky, and U. Rotics. On the fixed-parameter complexity of graph enumeration problems definable in monadic second-order logic. *Discrete Applied Mathematics*, 108(1–2):23–52, 2001.

[7] A. Dawar, M. Grohe, and S. Kreutzer. Locally excluding a minor. In *Proceedings of the 22nd IEEE Symposium on Logic in Computer Science*, pages 270–279, 2007.

[8] A. Dawar, M. Grohe, S. Kreutzer, and N. Schweikardt. Approximation schemes for first-order definable optimisation problems. In *Proceedings of the 21st IEEE Symposium on Logic in Computer Science*, pages 411–420, 2006.

[9] A. Dawar and S. Kreutzer. Domination problems in nowhere-dense classes of graphs. In *Foundations of software technology and theoretical computer science (FSTTCS)*, pages 157–168, 2009.

[10] R. Diestel. *Graph Theory*. Springer-Verlag, 3rd edition, 2005.

[11] R. Downey and M. Fellows. *Parameterized Complexity*. Springer, 1999.

[12] R.G. Downey, M.R. Fellows, and U. Taylor. The parameterized complexity of relational database queries and an improved characterization of W[1]. In D.S. Bridges, C. Calude, P. Gibbons, S. Reeves, and I.H. Witten, editors, *Combinatorics, Complexity, and Logic*, volume 39 of *Proceedings of DMTCS*, pages 194–213. Springer-Verlag, 1996.

[13] Z. Dvořák. Constant-factor approximation of the domination number in sparse graphs. *Eur. J. Comb.*, 34(5):833–840, 2013.

[14] Z. Dvořák, D. Král, and R. Thomas. Deciding first-order properties for sparse graphs. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 133–142, 2010.

[15] H.-D. Ebbinghaus, J. Flum, and W. Thomas. *Mathematical Logic*. Springer, 2nd edition, 1994.

[16] J. Flum and M. Grohe. Fixed-parameter tractability, definability, and model checking. *SIAM Journal on Computing*, 31(1):113–145, 2001.

[17] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer Verlag, 2006.

[18] M. Frick and M. Grohe. Deciding first-order properties of locally tree-decomposable structures. *Journal of the ACM*, 48:1184–1206, 2001.

[19] H. Gaifman. On local and non-local properties. In J. Stern, editor, *Herbrand Symposium, Logic Colloquium '81*, pages 105 – 135. North Holland, 1982.

[20] M. Grohe. Logic, graphs, and algorithms. In J. Flum, E. Grädel, and T. Wilke, editors, *Logic and Automata – History and Perspectives*, volume 2 of *Texts in Logic and Games*, pages 357–422. Amsterdam University Press, 2007.

[21] M. Grohe, K. Kawarabayashi, and B. Reed. A simple algorithm for the graph minor decomposition – logic meets structural graph theory–. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2013.

[22] M. Grohe and S. Kreutzer. Methods for algorithmic meta theorems. In M. Grohe and J.A. Makowsky, editors, *Model Theoretic Methods in Finite Combinatorics*, volume 558 of *Contemporary Mathematics*, pages 181–206. American Mathematical Society, 2011.

[23] H.A. Kierstead and D. Yang. Orderings on graphs and game coloring number. *Order*, 20(3):255–264, 2003.

[24] S. Kreutzer. Algorithmic meta-theorems. In J. Esparza, C. Michaux, and C. Steinhorn, editors, *Finite and Algorithmic Model Theory*, London Mathematical Society Lecture Note Series, chapter 5, pages 177–270. Cambridge University Press, 2011.

[25] S. Kreutzer and S. Tazari. Lower bounds for the complexity of monadic second-order logic. In *Proceedings of the 25th IEEE Symposium on Logic in Computer Science*, 2010.

[26] S. Kreutzer and S. Tazari. On brambles, grid-like minors, and parameterized intractability of monadic second-order logic. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, 2010.

[27] A. Langer, F. Reidl, P. Rossmanith, and S. Sikdar. Evaluation of an mso-solver. In *ALENEX*, pages 55–63, 2012.

[28] J. Nešetřil and P. Ossona de Mendez. *Sparsity*. Springer, 2012.

[29] J. Nešetřil and P. Ossona de Mendez. On nowhere dense graphs. *European Journal of Combinatorics*, 32(4):600–617, 2011.

[30] J. Nešetřil and P. Ossona de Mendez. Grad and classes with bounded expansion ii. algorithmic aspects. *arXiv*, math/0508324v2, 2005.

[31] C.H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *Journal of Computer and System Sciences*, 43:425–440, 1991.

[32] D. Seese. Linear time computable problems and first-order descriptions. *Mathematical Structures in Computer Science*, 6:505–526, 1996.

[33] X. Zhu. Colouring graphs with bounded generalized colouring number. *Discrete Mathematics*, 309:5562–5568, 2009.