Characterising Bounded Expansion by Neighbourhood Complexity

Felix Reidl Fernando Sánchez Villaamil
Konstantinos Stavropoulos

RWTH Aachen University
{reidl,fernando.sanchez,stavropoulos}@cs.rwth-aachen.de.

Abstract

We show that a graph class $\mathcal{G}$ has *bounded expansion* if and only if it has bounded *r-neighbourhood complexity*, *i.e.* for any vertex set X of any subgraph H of $G \in \mathcal{G}$, the number of subsets of X which are exact r-neighbourhoods of vertices of H on X is linear in the size of X. This is established by bounding the r-neighbourhood complexity of a graph in terms of both its *r-centred colouring number* and its *weak r-colouring number*, which provide known characterisations to the property of bounded expansion.

1 Introduction

Graph classes of *bounded expansion* (and their further generalisation, nowhere dense classes) have been introduced by Nešetřil and Ossona de Mendez [23, 24, 25] as a general model of *structurally sparse* graph classes. They include and generalise many other natural sparse graph classes, among them all classes of bounded degree, classes of bounded genus, and classes defined by excluded (topological) minors. Nowhere dense classes even include classes that locally exclude a minor, which in turn generalises graphs with locally bounded treewidth.

The appeal of this notion and its applications stems from the fact that bounded expansion has turned out to be a very robust property of graph classes with various seemingly unrelated characterisations (see [19, 25]). These include characterisations through the density of shallow minors [23], *quasi-wideness* [3] *low treedepth colourings* [23], and *generalised colouring numbers* [31]. The latter two are particularly relevant towards algorithmic applications, as we will discuss in the sequel. Furthermore, there is good evidence that real-world graphs (often dubbed 'complex networks') might exhibit this notion of structural sparseness [6, 27], whereas stricter notions (planar, bounded degree, excluded (topological) minors, *etc.*) do not apply.

It seems unlikely that bounded-expansion and nowhere dense classes admit global Robertson-Seymour style decompositions as they are available for classes excluding a fixed minor [28], a topological minor [21], an immersion [30], or an odd minor [5]. However, Nešetřil and Ossona de Mendez showed [24] that bounded-expansion and nowhere dense classes admit a 'local' decomposition, a so-called *low r-treedepth colouring*, in the following sense: for every integer r, every graph G from a bounded expansion (nowhere dense) class can be coloured with $f(r)$ (respectively $O(n^{o(1)})$) colours such that every union of $p < r$ colour classes induces a graph of treedepth at most p. We denote by $\chi_r(G)$ the minimal number of colours needed for a low r-treedepth colouring of G. These types of colourings generalise the star-colouring number [25] introduced by Fertin, Raspaud, and Reed [12]. In that context, low r-treedepth colourings are usually called *r-centred colourings*[1] (the precise definition of which we defer to Section 2).

This 'decomposition by colouring' has direct algorithmic implications. For example, counting how often an h-vertex graph appears in a host graph G as a subgraph, induced subgraph or homomorphism is possible in linear time [24] through the application of low r-centred (r-treedepth) colourings. A more precise bound for the running time of $O(|c(G)|^{2h}6^h h^2 \cdot |G|)$ was shown by Demaine *et al.* [6] if an appropriate low treedepth colouring c is provided as input. Low r-centred (r-treedepth) colourings can be further used to check whether an existential first-order sentence is true [25] or to approximate the problems $\mathcal{F}$-Deletion and Induced-$\mathcal{F}$-Deletion to within a factor that only depends on the precise bounded expansion graph class G belongs to and the set $\mathcal{F}$ [27].

Another characterisation of bounded expansion is obtained via the *weak r-colouring numbers*, denoted by $\mathrm{wcol}_r(G)$. The name 'colouring number' reflects the fact that the weak 1-colouring number is sometimes also called the *colouring number* of the graph, which only differs to the *degeneracy* of a graph by one. Roughly, the weak colouring number describes how well the vertices of a graph can be linearly ordered such that for any vertex v, the number of vertices that can reach v via short paths that use higher-order vertices is bounded. We postpone the precise definition of weak r-colouring

[1] Depending on the way r-treedepth colourings are defined, r-centred colourings might appear in the literature as $r-1$-treedepth colourings, as for example in [25]. For convenience, here we define them in a way so that the gap in the depth r is alleviated.

numbers to Section 2, but let us emphasise their utility: Grohe, Kreutzer, and Siebertz [20] used weak r-colouring numbers to prove the milestone result that first-order formulas can be decided in almost linear time for nowhere-dense classes (improving upon a result by Dvořák, Král, and Thomas for bounded expansion classes [11] and the preceding work for smaller sparse classes [4, 13, 17, 29]).

Our work here centres on a new characterisation, motivated by recent progress in the area of kernelisation. This field, a subset of parametrised complexity theory, formalises polynomial-time preprocessing of computationally hard problems. For an introduction to kernelisation we refer the reader to the seminal work by Downey and Fellows [8]. Gajarský *et al.* [18] extended the meta-kernelisation framework initiated by Bodlaender *et al.* [2] for bounded-genus graphs to nowhere-dense classes (notable intermediate results where previously obtained for excluded-minor classes [14] and classes excluding a topological minor [22]). In a largely independent line of research, Drange *et al.* recently provided a kernel for Dominating Set on nowhere-dense classes [9]. Previous results showed kernels for planar graphs [1], bounded-genus graphs [2], apex- minor-free graphs [14], graphs excluding a minor [15] and graphs excluding a topological minor [16].

A feature exploited heavily in the above kernelisation results for bounded expansion classes is that for any graph G from such a class, every subset $X \subseteq G$ has the property that the number of ways vertices from $V(G) \setminus X$ connect to X is linear in the size of X. Formally, we have that

$$|\{N(v) \cap X\}_{v \in V(G)}| \leqslant c \cdot |X|$$

where c only depends on the graph class from which G was drawn. One wonders whether this property of bounded expansion classes can be turned into a characterisation. It is, however, missing one important ingredient present in all known notions related to bounded expansion: a notion of *depth* via an appropriate distance-parameter. This brings us to the central notion of our work: If we denote by $N^r[\circ]$ the closed r-neighbourhood around a vertex, we define the *r-neighbourhood complexity* as

$$\nu_r(G) := \max_{H \subseteq G, \varnothing \neq X \subseteq V(H)} \frac{|\{N^r[v] \cap X\}_{v \in H}|}{|X|}.$$

That is, the value ν_r tells in how many different ways vertices can be joined to a vertex set X via paths of length at most r. Note that we define the value over all possible subgraphs: otherwise uniform dense graphs (*e.g.* complete graphs) would yield very low values[2].

[2]While this might be an interesting measure in and of itself, in this work we want to develop a measure for sparse graph classes and therefore choose the above definition.

The main result of this paper is the following characterisation of bounded expansion through neighbourhood complexity. We say that a graph class $\mathcal{G}$ has *bounded neighbourhood complexity* if there exists a function f such that for every r it holds that $\nu_r(\mathcal{G}) \leqslant f(r)$.

Theorem 1. *A graph class $\mathcal{G}$ has bounded expansion if and only if it has bounded neighbourhood complexity.*

Specifically, we show that the following relations between the r-neighbourhood complexity ν_r, the r-centred colouring number χ_r, and the weak r-colouring number wcol_r of a graph.

Theorem 2. *For all graphs G and all non-negative integers r it holds that*

$$\nu_r(G) \leqslant (r+1)2^{\chi_{2r+2}(G)^{r+2}}.$$

Theorem 3. *For every graph G and all non-negative integers r it holds that*

$$\nu_r(G) \leqslant \frac{1}{2}(2r+2)^{\mathrm{wcol}_{2r}(G)}\mathrm{wcol}_{2r}(G) + 1.$$

The characterisation of bounded expansion through generalised colouring numbers in [31] was provided by relating r-centred colourings to generalised colouring numbers. We believe that this interaction of the two notions is also highlighted in this paper, in the sense that when one can use one of the two notions as a direct proof tool, it might often be the case that the other might also serve as a direct proof tool, the most appropriate to be chosen depending on the occasion. As we believe it is also the case with neighbourhood complexity, it is still, as a consequence, useful to have access to a result through both parameters, since the general known bounds relating r-centred colourings and generalised colouring numbers seem to be very loose and most probably not optimal. For example, it is still unclear to our knowledge if one is always smaller than the other. Moreover, bounds for both parameters are not in general known for all kinds of specific graph glasses. It can then be the case that for different questions and different graph classes, r-centred colourings are more appropriate than generalised colouring numbers or vice versa.

2 Preliminaries

The main challenge is to prove that graphs from a graph class of bounded expansion have low neighbourhood complexity. To this end, some definitions will be necessary to prove Theorems 2 and 3.

2.1 Graphs and Signatures

For an integer n we write $[n] = \{1, \dots, n\}$. All logaritms in this paper are of base 2 and we only write $\log x$ instead of $\log_2 x$. We only consider non-empty, finite and simple graphs. For a graph G we write $V(G)$ and $E(G)$ to denote vertices and edges of G, respectively. We use the notations $|G| = |V(G)|$ and $\|G\| = |E(G)|$. Following the notation of Diestel [7], we denote an edge between two nodes $u, v \in V(G)$ by uv. In the following we will sometimes use the symbol $\circ$ to denote an arbitrary vertex and it should be understood that each occurence of $\circ$ can denote a different vertex. The statement 'there exist two edges $v\circ$,$w\circ$' therefore means 'there exist two edge va_1,va_2' and we will prefer the former if a_1, a_2 are not referenced later.

For a vertex $v \in V(G)$, we denote by $N^r_G(v) := \{u \in V(G) \mid \mathrm{dist}_G(u, v) = r\}$ the *r-th neighbourhood around v* for $r \geqslant 0$. Analogously, the r-th *closed* neighbourhood around v is defined as $N^r_G[v] := \bigcup_{i=0}^{r} N^i(v)$. In particular, $N^0_G(v) = N^0_G[v] = \{v\}$. We usually omit the subscript G if the context is clear.

A *signature* σ over a universe U is a sequence of elements $(u_i)_{1 \leqslant i \leqslant \ell}, u_i \in U$ where ℓ is the *length* of the signature, also denoted by $|\sigma|$. Accordingly, an *ℓ-signature* is simply a signature of length ℓ. We use the notation $\sigma[i] := u_i$ to signify the i-th element of σ. A signature is *proper* if all its elements are distinct. We assume that the elements of U are ordered. We assume an order on all signatures (say, lexicographic). Thus for a set $\mathcal{S}$ of signatures and a function $f\colon \mathcal{S} \to A$ for an arbitrary set A, we employ the notation $(f(\sigma))_{\sigma \in \mathcal{S}}$ to obtain sequences over elements of A derived from that ordering. For example, $(|\sigma|)_{\sigma \in \{\sigma_a, \sigma_b, \sigma_c\}}$ is shortand for the sequence $(|\sigma_a|, |\sigma_b|, |\sigma_c|)$ if $\sigma_a \leqslant \sigma_b \leqslant \sigma_c$ according to our (arbitrary) total order.

For a path $P = x_1 \dots x_\ell$ we write $P[x_i, x_j] = x_i \dots x_j$ to denote the subpath of P starting at x_i and ending at x_j. As such, we treat paths as ordered. Similarly, for an integer $1 \leqslant i \leqslant |P|$ we denote by $P[i]$ the i-th vertex on the path and we call i the *index* of that vertex on P. Hence, for non-empty paths, $P[1]$ is the start and $P[|P|]$ the end of the path. If G is a graph coloured by $c\colon V(G) \to [\xi]$ for some $\xi \in \mathbb{N}$ and P is a path in G, then we write σ_P to denote the $|P|$-signature over $[\xi]$ with $\sigma_P[i] = c(P[i])$. For a fixed signature σ, we say that P is a *σ-path* if $\sigma_P = \sigma$.

For a fixed signature σ over $[\xi]$, we define the *σ-neighbourhood* of a vertex v in G as

$$N^\sigma(v) := \{w \in V(G) \mid \exists vPw \text{ such that } \sigma_{vPw} = \sigma\}$$

Note that $N^\sigma(v) \subseteq N^{|\sigma|}(v)$ and that $N^\sigma(v) = \varnothing$ whenever $\sigma[1] \neq c(v)$. We use the following extension to vertex sets $X \in V(G)$ and sets of signatures $\mathcal{S}$

over $[\xi]$:

$$N^{\mathcal{S}}(v) := \bigcup_{\sigma\in\mathcal{S}} N^\sigma(v) \qquad N^\sigma(X) := \bigcup_{v\in X} N^\sigma(v) \qquad N^{\mathcal{S}}(X) := \bigcup_{v\in X}\bigcup_{\sigma\in\mathcal{S}} N^\sigma(v)$$

Similarly, the *σ-in-neighbourhood* of a vertex v is defined

$$N^{-\sigma}(v) := \{w \in V(G) \mid \exists wPv \text{ such that } \sigma_{wPv} = \sigma\}$$

and we extend this notation to vertex and signature sets in the same manner as above:

$$N^{-\mathcal{S}}(v) := \bigcup_{\sigma\in\mathcal{S}} N^{-\sigma}(v) \qquad N^{-\sigma}(X) := \bigcup_{v\in X} N^{-\sigma}(v) \qquad N^{-\mathcal{S}}(X) := \bigcup_{v\in X}\bigcup_{\sigma\in\mathcal{S}} N^{-\sigma}(v)$$

The following basic fact about σ-neighbourhoods for proper signatures σ is easy to verify.

Observation 1. *Let $u, v \in V(G)$ be distinct vertices and $uP\circ$, $vP\circ$ be two σ-paths for some proper signature σ. Then for any $x \in u\sigma\circ \cap v\sigma\circ$ it holds that x has the same index on both $u\sigma\circ$ and $v\sigma\circ$ and that xs colour appears exactly once in $uP\circ \cup vP\circ$.*

Finally the *lexicographic product* $G_1 \bullet G_2$ is the graph with vertices $V(G_1) \times V(G_2)$, where two nodes (u, x) and (v, y) are connected by an edge iff either a) $uv \in E(G_1)$ or b) $u = v$ and $xy \in E(G_2)$.

2.2 Grad and Expansion

The property of *bounded expansion* was introduced by Nešetřil and Ossona de Mendez using the notion of shallow minors [23, 24]: the basic idea is to exclude different minors depending on how 'local' the contracted portions of the graph is. Building on Dvořák's work [10], Nešetřil, Ossona de Mendez, and Wood later introduced an equivalent definition via shallow *topological* minors [26]. This seem surprising at first, since graphs defined via (unrestricted) forbidden minors are vastly different objects than graphs defined via forbidden *topological* minors. We will only introduce the topological variant here.

Definition 1 (Topological minor embedding)**.** A *topological minor embedding* of a graph H into a graph G is a pair of functions $\phi_V\colon V(H) \to V(G)$, $\phi_E\colon E(H) \to 2^{V(H)}$ where ϕ_V is injective and for every $uv \in H$ we have that

1. $\phi_E(uv)$ is a path in G with endpoints $\phi_V(u), \phi_V(v)$ and

2. for every $u'v' \in H$ with $u'v' \neq uv$ the two paths $\phi_E(uv)$, $\phi_E(u'v')$ are internally vertex-disjoint.

We define the *depth* of the topological minor embedding ϕ_V, ϕ_E as the half-integer $(\max_{uv \in H} |\phi_E(uv)| - 1)/2$, *i.e.* an embedding of depth r will map the edges of H onto paths in G of length at most $2r + 1$.

Accordingly, if H has a topological minor embedding of depth r into G we say that H is a an *r-shallow topological minor* of G and write $H \preccurlyeq_t^r G$. Note that this relationship is monotone in the sense that an r-shallow topological minor of G is also an $r+1$-shallow topological minor of G.

Definition 2 (Grad and bounded expansion)**.** For a graph G and an integer $r \geqslant 0$, we define the *topologically greatest reduced average density (top-grad) at depth r* as

$$\widetilde{\nabla}_r(G) = \max_{H \preccurlyeq_t^r G} \frac{\|H\|}{|H|}.$$

We extend this notation to graph classes as $\widetilde{\nabla}_r(\mathcal{G}) = \sup_{G \in \mathcal{G}} \widetilde{\nabla}_r(G)$. A graph class $\mathcal{G}$ then has *bounded expansion* if there exists a function $f \colon \mathbb{N} \to \mathbb{R}$ such that for all r we have that $\widetilde{\nabla}_r(\mathcal{G}) \leqslant f(r)$.

2.3 r-Centred Colourings and Weak r-Colouring Number

Equivalent definitions for classes of bounded expansion are related to the r-centred colouring number and the weak r-colouring number of graphs.

Definition 3 (r-centred colourings)**.** An r-centred colouring of a graph G is a vertex colouring such that, for any (induced) connected subgraph H, either some colour $c(H)$ colours exactly one node (a *centre*) in H or H gets at least r colours.

The minimum number of colours of an r-centred colouring of G is denoted by $\chi_r(G)$. Let us see the characterisation of bounded expansion via χ_r.

Proposition 1 (Nešetřil, Ossona de Mendez [23])**.** *Let $\mathcal{G}$ be a graph class of bounded expansion. Then there exists a function f_c such that for every $r \in \mathbb{N}$ and every $G \in \mathcal{G}$ it holds that $\chi_r(G) \leqslant f_c(r)$.*

Let $\Pi(G)$ be the set of linear orders on $V(G)$ and let $\preceq \in \Pi(G)$. We represent $\preceq$ as an injective function $L \colon V(G) \to \mathbb{N}$ with the property that $v \preceq w$ if and only if $L(v) \leqslant L(w)$.

A vertex u is *weakly r-reachable* from v with respect to the order L, if there is a path P of length at most r from v to u such that $L(u) \leqslant L(w)$ for all $w \in V(P)$. Let $\mathrm{WReach}_r[G, L, v]$ be the set of vertices that are weakly r-reachable from v with respect to L. The *weak r-colouring number* $\mathrm{wcol}_r(G)$ is now defined as

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

For a set of vertices $X \subseteq V(G)$, we let

$$\mathrm{WReach}_r[G, L, X] = \bigcup_{v \in X} \mathrm{WReach}_r[G, L, v].$$

Zhu [31] showed that a graph class has bounded expansion if and only if the weak r-colouring number wcol_r of every member is bounded by a function that only depends on r.

2.4 Neighbourhood Complexity

Definition 4 (Neighbourhood complexity)**.** For a graph G the *r-neighbourhood complexity* is a function ν_r defined via

$$\nu_r(G) := \max_{H \subseteq G, \varnothing \neq X \subseteq V(H)} \frac{|\{N^r[v] \cap X\}_{v \in V(H)}|}{|X|}.$$

We extend this definition to graph classes $\mathcal{G}$ via $\nu_r(\mathcal{G}) := \sup_{G \in \mathcal{G}} \nu_r(G)$.

Alternatively, we can define the neighbourhood complexity via the index of an equivalence relation. This turns out to be a useful perspective in the subsequent proofs. For $r \in \mathbb{N}$ and $X \subseteq V(G)$, we define the *(X, r)-twin equivalence* over $V(G)$ as

$$u \simeq_r^{G,X} v \iff N^r[u] \cap X = N^r[v] \cap X$$

which gives rise to the alternative definition

$$\nu_r(G) = \max_{H \subseteq G, \varnothing \neq X \subseteq V(G)} \frac{|V(H)/{\simeq_r^{H,X}}|}{|X|}.$$

We will usually fix a graph in the following and hence omit the superscript G of this relation. Recall that we say that a graph class $\mathcal{G}$ has *bounded neighbourhood complexity* if there exists a function f such that for every r it holds that $\nu_r(\mathcal{G}) \leqslant f(r)$.

3 Neighbourhood Complexity and r-Centred Colourings

This section is dedicated to proving the following relation between the r-neighbourhood complexity and the $(2r+2)$-centred colouring number of a graph.

Theorem 2. *For all graphs G and all non-negative integers r it holds that*

$$\nu_r(G) \leqslant (r+1)2^{\chi_{2r+2}(G)^{r+2}}.$$

For the remainder of this section, we fix a graph G, a subset of vertices $\varnothing \neq X \subseteq V(G)$, an integer r and a $(2r+2)$-centred colouring $c\colon V(G) \to [\xi]$ where $\xi = \chi_{2r+2}(G)$. We will assume that G and X are chosen such that $|V(G)/{\simeq_r^{G,X}}| = \nu_r(G) \cdot |X|$. For readability we will drop the superscript G from $\simeq_r^{G,X}$ in the following.

In the following we introduce a sequence of equivalence relations over $V(G)$ and prove that they successively refine $\simeq_r^X$. To that end, define $\mathcal{S}_{\leqslant r}$ to be the set of all signatures over $[\xi]$ of length at most r. The subsequent lemmas will elucidate the connection between centred colourings and proper signatures.

Lemma 1. *For any* proper *signature $\sigma \in \mathcal{S}_{\leqslant r}$ and any vertices $u, v \in V(G)$, either $N^\sigma(u) \cap N^\sigma(v) = \varnothing$ or $N^\sigma(u) = N^\sigma(v)$.*

Proof. Assume there exists $x \in N^\sigma(u) \cap N^\sigma(v)$ but $N^\sigma(u) \neq N^\sigma(v)$. Without loss of generality, let $y \in N^\sigma(v) \setminus N^\sigma(u)$.

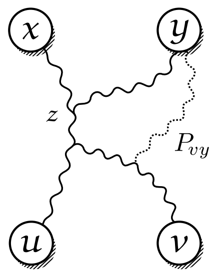

Fix a σ-path P_{ux} and a σ-path P_{vx}. Let $s \in P_{ux} \cap P_{vx}$ be the first vertex in which both paths intersect (since both paths end in x, such a vertex must exist). Further, fix a σ-path P_{vy}. Now if $P_{vy} \cap P_{ux}$ is non-empty, then y is σ-reachable from u: by Observation 1, there would be a vertex $z \in P_{vy} \cap P_{ux}$ that has the same index on both paths. Since σ is proper, the subpath of $P_{vy}[z, y]$ cannot share a vertex with $P_{ux}[u, z]$, thus we can construct a σ-path by first taking the subpath $P_{ux}[u, z]$ and then the subpath $P_{vy}[z, y]$. This path would mean that $y \in N^{\sigma}(u)$, contradicting our choice of y.

Hence, assume P_{vy} and P_{ux} do not intersect. But then the graph $P_{ux} \cup P_{vx} \cup P_{vy}$ is connected and contains every colour of σ at least twice. Since $|\sigma| \leqslant 2r + 1$ this contradicts our assumption that the colouring c is $(2r + 2)$-centred. $\square$

We see that a single proper signature σ imposes a very restricted structure on the respective σ-neighbourhoods in the graph. Even more interesting is the interaction of proper signatures with each other. To that end, let us introduce the notion of (X, σ)-*equivalence*: vertices u and v are equivalent if their respective σ-neighbourhoods in X are the same, *i.e.*

$$u \simeq_{\sigma}^{X} v \iff N^{\sigma}(u) \cap X = N^{\sigma}(v) \cap X.$$

Lemma 2. *Let σ_1, σ_2 be a pair of proper signatures. Let further $Y_{\sigma_1,\sigma_2} = N^{-\sigma_1}(X) \cap N^{-\sigma_2}(X)$ be all vertices that can reach at least one vertex in X via a σ_1-path and at least one vertex via a σ_2-path.*

Fix two arbitrary equivalence classes $C_{\sigma_1} \in Y_{\sigma_1,\sigma_2}/\simeq_{\sigma_1}^{X}$ and $C_{\sigma_2} \in Y_{\sigma_1,\sigma_2}/\simeq_{\sigma_2}^{X}$. Then either $C_{\sigma_1} \cap C_{\sigma_2} = \varnothing$, $C_{\sigma_1} \subseteq C_{\sigma_2}$, or $C_{\sigma_1} \supseteq C_{\sigma_2}$.

Proof. The statement is trivial if $\sigma_1 = \sigma_2$ or $C_{\sigma_1} = C_{\sigma_2}$. Otherwise, assume that there exist $C_{\sigma_1} \neq C_{\sigma_2}$ such that indeed C_{σ_1} and C_{σ_2} are not related in the three above ways. Since this is impossible when $|C_{\sigma_1}| = 1$ or $|C_{\sigma_2}| = 1$, we know that there exists vertices $u, v, w \in Y_{\sigma_1,\sigma_2}$ with $u \in C_{\sigma_1} \setminus C_{\sigma_2}$, $v \in C_{\sigma_2} \setminus C_{\sigma_1}$ and $w \in C_{\sigma_1} \cap C_{\sigma_2}$.

The respective membership in these classes tell us the following about the vertices u, v, w:

$$\begin{aligned} N^{\sigma_1}(u) \cap X = N^{\sigma_1}(w) \cap X \neq N^{\sigma_1}(v) \cap X \quad &\text{and} \\ N^{\sigma_2}(u) \cap X \neq N^{\sigma_2}(w) \cap X = N^{\sigma_2}(v) \cap X. & \end{aligned}$$

Using Lemma 1 we can strengthen this statement: $N^{\sigma_1}(u) \cap N^{\sigma_1}(v) = \varnothing$ and $N^{\sigma_2}(u) \cap N^{\sigma_2}(v) = \varnothing$ and since u, v, w are contained in Y_{σ_1,σ_2}, we know that all the involved neighbourhoods intersect X.

Therefore, we can pick distinct vertices $x_1, y_1, x_2, y_2 \in X$ such that $x_1 \in N^{\sigma_1}(u)$, $y_1 \in N^{\sigma_1}(v)$ and $x_2 \in N^{\sigma_2}(u)$, $y_2 \in N^{\sigma_2}(v)$.

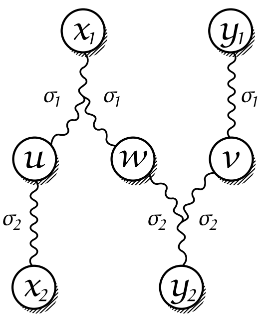

Since $N^{\sigma_1}(w) = N^{\sigma_1}(u)$, we can connect the vertices u, w with two (not necessarily disjoint) σ_1-paths $P_u^{\sigma_1}, P_w^{\sigma_1}$ that start both in x_1. Further, there exists a σ_1-path $P_v^{\sigma_1}$ from y_1 to v. If $P_v^{\sigma_1}$ would intersect either $P_u^{\sigma_1}$ or $P_w^{\sigma_1}$, we could not have that $N^{\sigma_1}(v) \cap N^{\sigma_1}(u) = \varnothing$ according to Lemma 1. We conclude that indeed $P_v^{\sigma_1}$ is disjoint from both $P_u^{\sigma_1}$ and $P_w^{\sigma_1}$.

We repeat the same construction for x_2, y_2 and the signature σ_2 to obtain paths $P_u^{\sigma_2}, P_v^{\sigma_2}, P_w^{\sigma_2}$. This time, $P_u^{\sigma_2}$ is necessarily disjoint from both $P_v^{\sigma_2}$ and $P_w^{\sigma_2}$ (*cf.* figure above). We reach a contradiction: observe that the graph induced by the paths $P_u^{\sigma_1}, P_v^{\sigma_1}, P_w^{\sigma_1}, P_u^{\sigma_2}, P_v^{\sigma_2}, P_w^{\sigma_2}$ is connected, contains every colour of $\sigma_1 \cup \sigma_2$ at least twice and in total at most $2r+1$ colours. This is impossible if c was indeed $(2r+2)$-centred. □

For the next lemma we extend the notion of (X, σ)-equivalence to sets of proper signatures $\mathcal{S}$. We define the $(X, \mathcal{S})$-equivalence relation on the vertices of G as follows:

$$u \simeq_{\mathcal{S}}^{X} v \iff \text{for all } \sigma \in \mathcal{S}, N^{\sigma}(u) \cap X = N^{\sigma}(v) \cap X$$

Lemma 3. *Let $\hat{\mathcal{S}} \subseteq \mathcal{S}_{\leqslant r}$ be a set of proper signatures and let $W_{\hat{\mathcal{S}}} = \bigcap_{\sigma \in \hat{\mathcal{S}}} N^{-\sigma}(X)$ be those vertices in G which have a non-empty σ-neighbourhood in X for every $\sigma \in \hat{\mathcal{S}}$. Then $|W_{\hat{\mathcal{S}}}/{\simeq_{\hat{\mathcal{S}}}^{X}}| \leqslant |\hat{\mathcal{S}}| \cdot |X|$.*

Proof. Define the set family $\mathcal{F} := \bigcup_{\sigma \in \hat{\mathcal{S}}} (W_{\hat{\mathcal{S}}}/{\simeq_{\sigma}^{X}})$ of the classes of all equivalence relations defined via a signature contained in $\hat{\mathcal{S}}$. By Lemma 2 and our choice of $W_{\hat{\mathcal{S}}}$, every pair $B_1, B_2 \in \mathcal{F}$ satisfies $B_1 \cap B_2 \in \{\varnothing, B_1, B_2\}$ (*i.e.* $\mathcal{F}$ is a *laminar* family).

Consider a class $B \in W_{\hat{\mathcal{S}}}/{\simeq_{\hat{\mathcal{S}}}^{X}}$. Then B is the result of a intersection of at most $|\hat{\mathcal{S}}|$ classes in $\mathcal{F}$. Since $B \neq \varnothing$ and $\mathcal{F}$ is laminar, it follows that $B \in \mathcal{F}$. We conclude that

$$|W_{\hat{\mathcal{S}}}/{\simeq_{\hat{\mathcal{S}}}^{X}}| \leqslant |\mathcal{F}| \leqslant |\hat{\mathcal{S}}| \cdot |X|,$$

where the second inequality follows from Lemma 1. □

In order to apply the above lemma it is left to bound the number of possible r-neighbourhoods in X by σ-neighbourhoods of *proper* signatures. We establish this bound by successively refining the (X, r)-twin equivalence. The following figure gives an overview over the proof (using relations yet to be introduced).

$$\begin{array}{lcc}
u \simeq^X_{r-1} v & \iff & N^{r-1}[u] \cap X = N^{r-1}[v] \cap X \\
\Uparrow \text{Lemma 4} & & \\
u \simeq^X_{\mathcal{S}_{\leqslant r}} v & \iff & \Big(N^\sigma(u) \cap X\Big)_{\sigma \in \mathcal{S}_{\leqslant r}} = \Big(N^\sigma(v) \cap X\Big)_{\sigma \in \mathcal{S}_{\leqslant r}} \\
\Uparrow \text{Lemma 5} & & \\
u \simeq^X_{\hat{\mathcal{S}}_{\leqslant r}} v & \iff & \Big(N^{\hat\sigma}_{\hat G}(u^1) \cap X^{|\hat\sigma|}\Big)_{\hat\sigma \in \hat{\mathcal{S}}_{\leqslant r}} = \Big(N^{\hat\sigma}_{\hat G}(v^1) \cap X^{|\hat\sigma|}\Big)_{\hat\sigma \in \hat{\mathcal{S}}_{\leqslant r}}
\end{array}$$

Where the last relation is defined with the help of an auxiliary graph $\hat{G}$ and signature set $\hat{\mathcal{S}}_{\leqslant r}$ whose construction is described later. The bound on the index of this last relation will prove Theorem 2.

Lemma 4. *The equivalence relation $\simeq^X_{\mathcal{S}_{\leqslant r}}$ over $V(G)$ defined via*

$$u \simeq^X_{\mathcal{S}_{\leqslant r}} v \iff \Big(N^\sigma(u) \cap X\Big)_{\sigma \in \mathcal{S}_{\leqslant r}} = \Big(N^\sigma(v) \cap X\Big)_{\sigma \in \mathcal{S}_{\leqslant r}}$$

is a refinement of $\simeq^X_{r-1}$.

Proof. Assume $u \simeq^X_{\mathcal{S}_{\leqslant r}} v$. We need to prove that $N^{r-1}[u] \cap X = N^{r-1}[v] \cap X$. The equivalence of u and v implies that

$$\begin{aligned}
w \in N^{r-1}[v] \cap X &\iff \exists \sigma \in \mathcal{S}_{\leqslant r} \colon w \in N^\sigma(v) \cap X \\
&\iff \exists \sigma \in \mathcal{S}_{\leqslant r} \colon w \in N^\sigma(u) \cap X \\
&\iff w \in N^{r-1}[u] \cap X.
\end{aligned}$$

□

We now construct an auxiliary graph and colouring as follows: Let $\hat{G} = G \bullet \overline{K}_r$. Assuming that $V(K_r) = [r]$ and hence $V(\hat{G}) = V(G) \times [r]$, we will use the shorthand $v^i = (v, i)$ for $v \in V(G), i \in [r]$ and call v^i the *ith copy* of v. Using this notation, we define a colouring $\hat{c} \colon V(\hat{G}) \to [\xi] \times [r]$ of $\hat{G}$ via $\hat{c}(v^i) = (c(v), i)$. Note that $\hat{c}$ is a $(2r+2)$-centred colouring of $\hat{G}$: any connected subgraph $\hat{H} \subseteq \hat{G}$ with less than $2r+2$ colours and no centre would directly imply that the subgraph $H \subseteq G$ with vertex set $V(H) = \bigcup_{1 \leqslant i \leqslant r} \{v \in V(G) \mid v^i \in \hat{H}\}$ contains at most $2r+2$ colours and no centre, contradicting our choice of c.

For a signature $\sigma \in \mathcal{S}_{\leqslant r}$ we define the proper signature $\hat{\sigma} = ((\sigma[i], i))_{1 \leqslant i \leqslant |\sigma|}$. Accordingly, we define the set of proper signatures $\hat{\mathcal{S}}_{\leqslant r}$ over colours $[\xi] \times [r]$ as $\hat{\mathcal{S}}_{\leqslant r} = \{\hat{\sigma} \mid \sigma \in \mathcal{S}_{\leqslant r}\}$. The following lemma connects the sigma-equivalence $\simeq^X_{\mathcal{S}_{\leqslant r}}$ over $V(G)$ with a suitable equivalence defined over the above auxiliary structure.

Lemma 5. *The equivalence relation $\simeq^X_{\hat{\mathcal{S}}_{\leqslant r}}$ over $V(G)$ defined via*

$$u \simeq^X_{\hat{\mathcal{S}}_{\leqslant r}} v \iff \left(N^{\hat{\sigma}}_{\hat{G}}(u^1) \cap X^{|\hat{\sigma}|}\right)_{\hat{\sigma} \in \hat{\mathcal{S}}_{\leqslant r}} = \left(N^{\hat{\sigma}}_{\hat{G}}(v^1) \cap X^{|\hat{\sigma}|}\right)_{\hat{\sigma} \in \hat{\mathcal{S}}_{\leqslant r}}$$

is a refinement of $\simeq^X_{\mathcal{S}_{\leqslant r}}$ where $X^i := \{v^i \mid v \in X\}$.

Proof. Assume $u \simeq^X_{\hat{\mathcal{S}}_{\leqslant r}} v$. Then for every signature $\hat{\sigma} \in \mathcal{S}_{\leqslant r}$ we have that $N^{\hat{\sigma}}_{\hat{G}}(u^1) \cap X^{|\hat{\sigma}|} = N^{\hat{\sigma}}_{\hat{G}}(v^1) \cap X^{|\hat{\sigma}|}$. Now note that if $w^{|\hat{\sigma}|}$ is $\hat{\sigma}$-reachable from u^1 in $\hat{G}$, then w is σ-reachable from u in G: if $u^1 x_2^2 \dots x_{|\hat{\sigma}|-1}^{|\hat{\sigma}|-1} w^{|\hat{\sigma}|}$ is a $\hat{\sigma}$-path in $\hat{G}$, then $u x_1 \dots x_{|\hat{\sigma}|-1} w$ is, by construction of $\hat{\sigma}$, a σ-path in G.

Accordingly $w^{|\hat{\sigma}|} \in N^{\hat{\sigma}}(u^1)$ implies that $w \in N^{\sigma}(u)$. We conclude that therefore $N^{\sigma}(u) \cap X^{|\sigma|} = N^{\sigma}(v) \cap X^{|\sigma|}$ and thus $u \simeq^X_{\mathcal{S}_{\leqslant r}} v$. ☐

Lemma 6. $|V(G)/{\simeq^X_{\hat{\mathcal{S}}_{\leqslant r}}}| \leqslant r 2^{\xi^{r+1}} \cdot |X|$.

Proof. To obtain the bound, we apply Lemma 3 to every subset of signatures $\hat{\mathcal{S}} \subseteq \hat{\mathcal{S}}_{\leqslant r}$. Let $\hat{X} \subseteq \hat{G}$ be the set containing all copies of vertices in X.

$$|V(G)/{\simeq^X_{\hat{\mathcal{S}}_{\leqslant r}}}| \leqslant |V(\hat{G})/ \simeq^{\hat{X}^r}_{\hat{\mathcal{S}}_{\leqslant r}}| \leqslant \sum_{\hat{\mathcal{S}} \subseteq \hat{\mathcal{S}}_{\leqslant r}} |\hat{\mathcal{S}}| \cdot |\hat{X}| = r 2^{\xi^{r+1}} \cdot |X| \qquad ☐$$

The proof of this section's theorem is now only a technicality.

Proof of Theorem 2. By Lemma 4 and 5 we have that

$$|V(G)/{\simeq^X_r}| \leqslant |V(G)/{\simeq^X_{\mathcal{S}_{\leqslant r+1}}}| \leqslant |V(G)/{\simeq^X_{\hat{\mathcal{S}}_{\leqslant r+1}}}|$$

Which, by Lemma 6, is at most $(r+1) 2^{\chi_{2r+2}(G)^{r+2}} \cdot |X|$ and the claim follows. ☐

4 Neighbourhood Complexity and Weak Colouring Number

Having obtained a bound for the neighbourhood complexity in terms of the r-centred colouring number, we now derive a bound in terms of the weak r-colouring number. For the next proof, we say that two vertices $u, v \in V(G)$ have *the same distances* to $Z \subseteq V(G)$ if for every $z \in Z$ we have $d_G(u, z) = d_G(v, z)$.

Theorem 3. *For every graph G and all non-negative integers r it holds that*

$$\nu_r(G) \leqslant \frac{1}{2}(2r+2)^{\mathrm{wcol}_{2r}(G)}\mathrm{wcol}_{2r}(G) + 1.$$

Proof. Fix a graph G and choose any subset $\varnothing \neq X \subseteq V(G)$. We will show in the following that

$$|V(G)/{\simeq_r^X}| \leqslant \left(\frac{1}{2}(2r+2)^{\mathrm{wcol}_{2r}(G)}\mathrm{wcol}_{2r}(G) + 1\right)|X|,$$

from which the claim immediately follows.

Let $\alpha_0 \in V(G)/{\simeq_r^X}$ be the equivalence class of $\simeq_r^X$ corresponding to the vertices of G with an empty r-neighbourhood in X and let $\mathcal{W} = \left(V(G)/{\simeq_r^X}\right) \setminus \{\alpha_0\}$. Moreover, let $L \in \Pi(G)$ be such that $\mathrm{wcol}_{2r}(G) = \max_{v\in V(G)} |\mathrm{WReach}_{2r}[G, L, v]|$. We will estimate the neighbourhood complexity of X via the neighbourhood complexity of a certain *good* subset of $\mathrm{WReach}_r[G, L, X]$.

For a vertex $v \in N^r(X)$ and a vertex $x \in N^r[v] \cap X$, let $\mathcal{P}_v^x$ be the set of all shortest (v, x)-paths (of length at most r). We define as $G^r[v]$ the graph induced by the union of the paths of all $\mathcal{P}_v^x$, namely

$$G^r[v] = G\Big[\bigcup_{x\in N^r[v]\cap X} \bigcup_{P\in\mathcal{P}_v^x} V(P)\Big].$$

By its construction, $G^r[v]$ contains, for every $x \in N^r[v] \cap X$, all shortest paths of length at most r that connect v to x.

Now, for every equivalence class $\kappa \in \mathcal{W}$, choose a representative vertex $v_\kappa \in \kappa$. Let $C = \{v_\kappa\}_{\kappa\in\mathcal{W}}$ be the set of representative vertices for all classes in $\mathcal{W}$. Using the representatives from C, we define for every class $\kappa \in \mathcal{W}$ the set (see Fig. 1)

$$Y_\kappa = \mathrm{WReach}_r[G^r[v_\kappa], L, v_\kappa] \cap \mathrm{WReach}_r[G, L, N^r[v_\kappa] \cap X]$$

and join all such sets into $Y = \bigcup_{\kappa\in\mathcal{W}} Y_\kappa$. Then,

$$Y \subseteq \bigcup_{\kappa\in\mathcal{W}} \mathrm{WReach}_r[G, L, N^r[v_\kappa] \cap X] \subseteq \mathrm{WReach}_r[G, L, X].$$

Moreover, by definition and the fact that L is an ordering achieving $\mathrm{wcol}_{2r}(G)$ (and not necessarily one achieving $\mathrm{wcol}_r(G)$), we have

$$|Y_\kappa| \leqslant |\mathrm{WReach}_r[G, L, v_\kappa]| \leqslant |\mathrm{WReach}_{2r}[G, L, v_\kappa]| \leqslant \mathrm{wcol}_{2r}(G).$$

Notice that for every $x \in N^r[v] \cap X$, the minimum vertex (according to L) of a path in $\mathcal{P}_v^x$ will always belong to Y_κ, therefore the set Y_κ intersects every path of

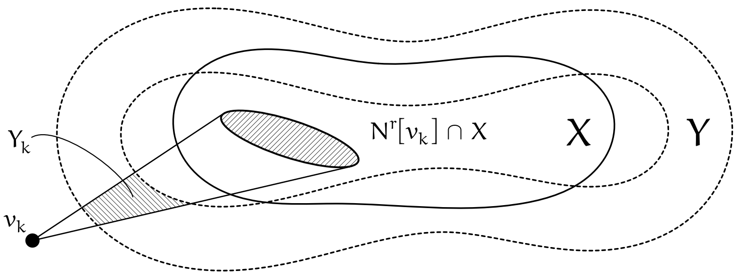

Figure 1: A set Y_κ and the set Y.

the sets $\mathcal{P}^x_{v_\kappa}$ forming $G^r[v_\kappa]$. We want to see how many different equivalence classes of $\mathcal{W}$ produce the same Y_κ set. This will allow us to bound the neighbourhood complexity of X by relating it to the number of different Y_κ's.

Suppose that $\kappa \neq \lambda$ with $Y_\kappa = Y_\lambda = Z$. Recall that Y_κ intersects all the shortest paths from v_κ to the vertices of $N^r[v_\kappa] \cap X$ and that $G^r[v_\kappa]$ is formed by all such shortest paths. Hence, if v_κ and v_λ have the same distances to Z, then we clearly get $N^r[v_\kappa] \cap X = N^r[v_\lambda] \cap X$, a contradiction. This means that if $Y_\kappa = Y_\lambda = Z$, the vertices v_κ and v_λ cannot have the same distances to Z. But there are at most $(r+1)^{|Z|}$ possible configurations of distances of the vertices of a set Z to a vertex v that has distance at most r to every vertex of Z. It follows that the number of equivalence classes of $\mathcal{W}$ that produce the same set Y_κ through their representative v_κ from C is at most $(r+1)^{|Y_\kappa|} \leqslant (r+1)^{\mathrm{wcol}_{2r}(G)}$.

Let $\mathcal{Y} := \{Y_\kappa \mid \kappa \in \mathcal{W}\}$ be the set of all (*different*) Y_κ's, and define $\gamma\colon \mathcal{Y} \to Y$ by $\gamma(Y_\kappa) = \arg\max_{y \in Y_\kappa} L(y)$. That is, $\gamma(Y_\kappa)$ is that vertex in Y_κ that comes last according to L. Observe that—by definition—every vertex in Y_κ is weakly r-reachable from v_κ. It follows that every vertex in Y_κ is weakly $2r$-reachable from $\gamma(Y_\kappa)$ via v_κ. In other words, $Y_\kappa \subseteq \mathrm{WReach}_{2r}[G, L, \gamma(Y_\kappa)]$. Consequently, for every vertex $y \in \gamma(\mathcal{Y})$, it holds that[3]

$$\bigcup \gamma^{-1}(y) \subseteq \mathrm{WReach}_{2r}[G, L, y],$$

i.e. the union $\bigcup \gamma^{-1}(y)$ of all Y_κ's that choose the same vertex y via γ has size at most $\mathrm{wcol}_{2r}(G)$. But every set in the family $\gamma^{-1}(y)$ is a subset of $\bigcup \gamma^{-1}(y)$ that contains y. Since there are at most $2^{|\bigcup \gamma^{-1}(y)|-1}$ different such subsets of $\bigcup \gamma^{-1}(y)$, the number of *different* Y_κ's for which the same vertex is chosen via γ is bounded by $2^{\mathrm{wcol}_{2r}(G)-1}$, i.e.

$$|\gamma^{-1}(y)| \leqslant 2^{\mathrm{wcol}_{2r}(G)-1}.$$

Recalling that one Y_κ corresponds to at most $(r+1)^{\mathrm{wcol}_{2r}(G)}$ equivalence classes of $\mathcal{W}$ and that $Y \subseteq \mathrm{WReach}_r[G, L, X]$, we can now bound the size of

[3] We remind the reader that this union expresses the union of a set in the set theoretical sense, i.e. the union of a set is the union of all of its elements (as sets).

$\mathcal{W}$ as follows:

$$\begin{aligned}|\mathcal{W}| &\leqslant (r+1)^{\mathrm{wcol}_{2r}(G)} \cdot |\mathcal{Y}| = (r+1)^{\mathrm{wcol}_{2r}(G)} \cdot \sum_{y \in \gamma(\mathcal{Y})} |\gamma^{-1}(y)| \\ &\leqslant (r+1)^{\mathrm{wcol}_{2r}(G)} \cdot \sum_{y \in \gamma(\mathcal{Y})} 2^{\mathrm{wcol}_{2r}(G)-1} \\ &= \frac{1}{2}(2r+2)^{\mathrm{wcol}_{2r}(G)} \cdot |\gamma(\mathcal{Y})|\end{aligned}$$

from which we obtain that

$$\begin{aligned}|V(G)/{\simeq_r^X}| &\leqslant |\mathcal{W}| + 1 \leqslant \frac{1}{2}(2r+2)^{\mathrm{wcol}_{2r}(G)} \cdot |\gamma(\mathcal{Y})| + 1 \\ &\leqslant \frac{1}{2}(2r+2)^{\mathrm{wcol}_{2r}(G)} \cdot |Y| + 1 \\ &\leqslant \frac{1}{2}(2r+2)^{\mathrm{wcol}_{2r}(G)} \cdot |\mathrm{WReach}_r[G, L, X]| + 1 \\ &\leqslant \frac{1}{2}(2r+2)^{\mathrm{wcol}_{2r}(G)} \mathrm{wcol}_{2r}(G) \cdot |X| + 1 \\ &\leqslant \left(\frac{1}{2}(2r+2)^{\mathrm{wcol}_{2r}(G)} \mathrm{wcol}_{2r}(G) + 1\right) |X|,\end{aligned}$$

as claimed. □

5 Completing the Characterisation

We have seen in the previous two sections that bounded expansion implies bounded neighbourhood complexity. Let us now prove the other direction to arrive at the full characterisation. We begin by proving that every bipartite graph with low neighbourhood complexity must have low minimum degree. To that end, we will need the following Lemma.

Lemma 7 (Nešetřil & Ossona de Mendez [25])**.** *Let $G = (A, B, E)$ be a bipartite graph and let $1 \leqslant r \leqslant s \leqslant |A|$. Assume each vertex in B has degree at least r.*

Then there exists a subset $A' \subseteq A$ and a subset $B' \subseteq B$ such that $|A'| = s$ and $|B'| \geqslant |B|/2$ and every vertex in B' has at least $r\frac{|A'|}{|A|}$ neighbours in A'.

The minimum degree and depth-one neighbourhood complexity ν_1 of a bipartite graph can now be related to each other as follows:

Lemma 8. *Let $G = (A, B, E)$ be a non-empty bipartite graph. Then*

$$\delta(G) < 4\nu_1(G)\Big(2\lceil \log \nu_1(G) \rceil + 1\Big)\Big(64\nu_1(G)^3 \lceil \log \nu_1(G) \rceil + 16\nu_1(G)^2 + 1\Big).$$

Proof. Let

$$\alpha = 4\nu_1(G)\Big(2\lceil\log\nu_1(G)\rceil + 1\Big)\Big(64\nu_1(G)^3\lceil\log\nu_1(G)\rceil + 16\nu_1(G)^2 + 1\Big)$$

and suppose that $\delta(G) \geqslant \alpha$. Assume without loss of generality that $|B| \geqslant |A|$ and let $\nu = 2^{\lceil\log\nu_1(G)\rceil}$. Observe that both $\nu, \log\nu$ are integers and that $\nu_1(G) \leqslant \nu < 2\nu_1(G)$. Therefore,

$$|B| \geqslant |A| \geqslant \delta(G) > 2\nu(2\log\nu + 1)\Big(8\nu^3\log\nu + 4\nu^2 + 1\Big).$$

Let us apply Lemma 7 on G with $r = 8\nu^3\log\nu + 4\nu^2 + 1$ and $s = \lfloor\frac{|A|}{2\nu(2\log\nu+1)}\rfloor$. Notice that this is indeed possible, because $|A| > 2\nu(2\log\nu + 1)\cdot r$ and therefore $s \geqslant r$. We obtain a subgraph $G' = (A', B', E')$ with

1. $\frac{|A|}{2\nu(2\log\nu+1)} - 1 < |A'| = s \leqslant \frac{|A|}{2\nu(2\log\nu+1)}$,
2. $|B'| \geqslant \frac{|B|}{2}$, and thus $|B'| \geqslant \frac{|A|}{2} \geqslant \nu(2\log\nu + 1)|A'|$,
3. and such that for every $v \in B'$ we have that $\deg_{G'}(v) \geqslant r \cdot \frac{|A'|}{|A|}$.

Combining the first and third property with $|A| > 2\nu(2\log\nu + 1)\cdot r$, we obtain

$$\begin{aligned}\deg_{G'}(v) &\geqslant r\cdot\frac{|A'|}{|A|} > r\left(\frac{1}{2\nu(2\log\nu+1)} - \frac{1}{|A|}\right)\\ &> r\left(\frac{1}{2\nu(2\log\nu+1)} - \frac{1}{2\nu(2\log\nu+1)\cdot r}\right)\\ &= \frac{r-1}{2\nu(2\log\nu+1)} = 2\nu^2.\end{aligned}$$

Now, note that any graph H with at least two vertices trivially has $\nu_1(H) \geqslant 2$ by taking X to be a single vertex of H. Hence, if $K_{2\nu^2, 2\log\nu+1}$ is a subgraph of G', we have that

$$\nu_1(G) \geqslant \nu_1(G') \geqslant \nu_1(K_{2\nu^2,2\log\nu+1}) \geqslant \frac{2\nu^2}{2\log\nu + 1} > \nu,$$

where the last inequality follows by the fact that $\nu \geqslant 2$, a contradiction.

So, let us call two vertices $u, v \in V(G')$ *twins* if $N^1_{G'}(u) = N^1_{G'}(v)$ and let us partition B' into twin-classes $B'_1, \dots B'_\ell$. Since each twin-class has at least $2\nu^2$ neighbours, the size of each twin-class must be bounded by $|B'_i| < 2\log\nu + 1$. Hence, the number of twin-classes is at least $\ell > \frac{|B'|}{2\log\nu+1}$. Since each twin-class has, by definition, a unique neighbourhood in A', we conclude that

$$\nu_1(G') \geqslant \frac{\ell}{|A'|} > \frac{|B'|}{2\log\nu + 1}\frac{\nu(2\log\nu+1)}{|B'|} = \nu \geqslant \nu_1(G),$$

a contradiction. □

It easily follows that every graph with low neighbourhood complexity must have low average degree.

Corollary 1. *Let G be a graph. Then $\widetilde{\nabla}_0(G) < 5445 \cdot \nu_1(G)^4 \log^2 \nu_1(G)$.*

Proof. We assume that $\widetilde{\nabla}_0(G) = \|G\|/|G|$, otherwise we restrict ourselves to a suitable subgraph of G with that property. The case where $|G| = 1$ is trivial, therefore we may assume that $|G| \geqslant 2$. It is folklore that G contains a bipartite graph H such that $\|H\| \geqslant \|G\|/2$. We can further ensure that $\delta(H) \geqslant \|H\|/|H|$ by excluding vertices of lower degree (this operation cannot decrease the density of H). Applying Lemma 8 to H, we obtain that

$$\widetilde{\nabla}_0(G) = \frac{\|G\|}{|G|} \leqslant 2\frac{\|H\|}{|H|} \leqslant 2\delta(H).$$

We apply the bound provided by Lemma 8 and relax it to the more concise polynomial $(5445/2) \cdot \nu_1(G)^4 \log^2 \nu_1(G)$, using the fact that $\nu_1(G) \geqslant 2$. □

The next theorem now leads to the full characterisation as stated in Theorem 1.

Theorem 4. *For every graph G and every half-integer r it holds that*

$$\widetilde{\nabla}_r(G) \leqslant (2r+1) \max\left\{5445\nu_1(G)^4 \log^2 \nu_1(G),\, \nu_2(G), \ldots,\, \nu_{\lceil r+1/2 \rceil}(G)\right\}.$$

Proof. Fix r and let $H \preccurlyeq^r_t G$ be an r-shallow topological minor of maximal density, *i.e.* $\widetilde{\nabla}_0(H) = \widetilde{\nabla}_r(G)$. Let further ϕ_V, ϕ_E be a topological minor embedding of H into G of depth r.

Let us label the edges of H by the respective path-length in the embedding ϕ_V, ϕ_E: an edge $uv \in H$ receives the label $\|\phi_E(uv)\|$. Let r' be the label of highest frequency and let $H' \subseteq H$ be the graph obtained from H by only keeping edges labelled with r'. Since there were up to $2r+1$ labels in H, we have that $(2r+1)\|H'\| \geqslant |H|$ and therefore

$$\widetilde{\nabla}_r(G) = \widetilde{\nabla}_0(H) \leqslant (2r+1)\frac{\|H'\|}{|H'|} \leqslant (2r+1)\widetilde{\nabla}_0(H'). \tag{1}$$

First, consider the case that $r' = 1$, *i.e.* H' is a subgraph of G. Combining (1) with Corrollary 1, we obtain

$$\begin{aligned}\widetilde{\nabla}_r(G) &\leqslant (2r+1)\widetilde{\nabla}_0(H') \leqslant (2r+1)\widetilde{\nabla}_0(G) \\ &\leqslant (2r+1) \cdot 5445\, \nu_1(G)^4 \log^2 \nu_1(G).\end{aligned}$$

Otherwise, assume that $r' \geqslant 2$, *i.e.* every edge of H' is embedded into a path of length at least 2 in G by ϕ_V, ϕ_E. Construct the subgraph $G' \subseteq G$ that contains all edges and vertices involved in the embedding of H' into G, that is, G' has vertices $\bigcup_{v \in H'} V(\phi_V(v)) \cup \bigcup_{e \in H'} V(\phi_E(e))$ and edges $\bigcup_{e \in H'} E(\phi_E(e))$.

Let $X = \bigcup_{v \in H'} V(\phi_V(v))$ and let $S \subseteq V(G')$ be a set constructed as follows: for every edge $e \in H'$ we add the middle vertex of the path $\phi_E(e)$ to S—in case r' is odd, we pick one of the two vertices that lie in the middle of $\phi_E(e)$ arbitrarily. Because X is an independent set in G' and $r' > 1$, every vertex in S has exactly two neighbours in X at distance $\lceil r'/2 \rceil$ in the graph G'. By construction, there is a one-to-one correspondence between these $\lceil r'/2 \rceil$-neighbourhoods and the edges of H'. Accordingly,

$$\|H'\| = |\{N_{G'}^{\lceil r'/2 \rceil}(v) \cap X\}_{v \in S}|$$

and therefore, using also the fact that G' is a subgraph of G,

$$\frac{\|H'\|}{|H'|} = \frac{|\{N_{G'}^{\lceil r'/2 \rceil}(v) \cap X\}_{v \in S}|}{|X|} \leqslant \nu_{\lceil r'/2 \rceil}(G') \leqslant \nu_{\lceil r'/2 \rceil}(G).$$

which, taken together with (1) and the fact that G' is a subgraph of G, yields

$$\widetilde{\nabla}_r(G) \leqslant (2r+1)\widetilde{\nabla}_0(H') \leqslant (2r+1)\nu_{\lceil r'/2 \rceil}(G).$$

Putting everything together, we finally arrive at

$$\widetilde{\nabla}_r(G) \leqslant (2r+1) \max\Big\{5445\, \nu_1(G)^4 \log^2 \nu_1(G),\ \nu_2(G), \ldots,\ \nu_{\lceil r+1/2 \rceil}(G)\Big\}.$$

proving the theorem. □

We conclude that graph classes with bounded neighbourhood complexity have bounded expansion. Theorem 1 follows by Theorems 2, 3 and 4.

6 Concluding Remarks

One should note that in Theorems 2 and 3 the derived bounds are *exponential* in the measures χ_{2r+2} and wcol_{2r}. Consequently, we cannot use neighbourhood complexity to characterise nowhere dense classes: in these classes, the quantities χ_r and wcol_r can only be bounded by $O(|G|^{o(1)})$ which only results in superpolynomial bounds for ν_r.

This constitutes an unusual phenomenon in the following sense: so far, every known characterisation of bounded expansion translated to a direct characterisation of nowhere denseness, but this has not yet been the case for neighbourhood complexity. It would be remarkable if one could only characterise the property of bounded expansion through neighbourhood complexity and not that of nowhere denseness. So far, it is only known that ν_1 is bounded by $O(|G|^{o(1)})$ in nowhere dense classes [18]. We pose as an interesting open question whether this holds true for ν_r for all r, or whether nowhere dense classes can indeed have a neighbourhood complexity that cannot be bounded by such a function.

References

[1] J. Alber, M. R. Fellows, and R. Niedermeier. Polynomial-time data reduction for Dominating Set. *Journal of the ACM*, 51:363–384, 2004.

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

[3] A. Dawar. Homomorphism preservation on quasi-wide classes. *Journal of Computer and System Sciences*, 76(5):324–332, 2010.

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

[5] E. D. Demaine, M. Hajiaghayi, and K. Kawarabayashi. Decomposition, approximation, and coloring of odd-minor-free graphs. In *Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 329–344. SIAM, 2010.

[6] E. D. Demaine, F. Reidl, P. Rossmanith, F. Sánchez Villaamil, S. Sikdar, and B. D. Sullivan. Structural sparsity of complex networks: Bounded expansion in random models and real-world graphs. 2014, abs/1406.2587.

[7] R. Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*. Springer, 2012.

[8] R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Monographs in Computer Science. Springer, 1999.

[9] P. G. Drange, M. Dregi, F. V. Fomin, S. Kreutzer, S. Lokshtanov, M. Pilipczuk, M. Pilipczuk, F. Reidl, F. Sánchez Villaamil, S. Saurabh, S. Siebertz, and S. Sikdar. Kernelization and Sparseness: the Case of Dominating Set. In *33rd Symposium on Theoretical Aspects of Computer*

Science (STACS), volume 47, pages 31:1–31:14, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.

[10] Z. Dvořák. *Asymptotical Structure of Combinatorial Objects*. PhD thesis, Charles University, Faculty of Mathematics and Physics, 2007.

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

[12] G. Fertin, A. Raspaud, and B. A. Reed. Star coloring of graphs. *Journal of Graph Theory*, 47(3):163–182, 2004.

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

[14] F. V. Fomin, D. Lokshtanov, S. Saurabh, and D. M. Thilikos. Bidimensionality and kernels. In *Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 503–510. SIAM, 2010.

[15] F. V. Fomin, D. Lokshtanov, S. Saurabh, and D. M. Thilikos. Linear kernels for (connected) dominating set on H-minor-free graphs. In *Proceedings of the 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 82–93. SIAM, 2012.

[16] F. V. Fomin, D. Lokshtanov, S. Saurabh, and D. M. Thilikos. Linear kernels for (connected) dominating set on graphs with excluded topological subgraphs. In *Proceedings of the 30th Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 20 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 92–103. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2013.

[17] M. Frick and M. Grohe. Deciding first-order properties of locally tree-decomposable structures. *Journal of the ACM*, 48(6):1184–1206, 2001.

[18] J. Gajarský, P. Hliněný, J. Obdržálek, S. Ordyniak, F. Reidl, P. Rossmanith, F. Sánchez Villaamil, and S. Sikdar. Kernelization using structural parameters on sparse graph classes. *To appear in Journal of Computer and System Sciences*, 2015.

[19] M. Grohe, S. Kreutzer, and S. Siebertz. Characterisations of nowhere dense graphs. In *Proceedings of the 32nd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*,

volume 24 of *LIPIcs*, pages 21–40. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2013.

[20] M. Grohe, S. Kreutzer, and S. Siebertz. Deciding first-order properties of nowhere dense graphs. In *Proceedings of the 46th ACM Symposium on Theory of Computing*, pages 89–98, 2014.

[21] M. Grohe and D. Marx. Structure theorem and isomorphism test for graphs with excluded topological subgraphs. In *Proceedings of the 44th ACM Symposium on Theory of Computing (STOC)*, pages 173–192, 2012.

[22] E. J. Kim, A. Langer, C. Paul, F. Reidl, P. Rossmanith, I. Sau, and S. Sikdar. Linear kernels and single-exponential algorithms via protrusion-decompositions. *To appear in Transaction on Algorithms*, 2015.

[23] J. Nešetřil and P. Ossona de Mendez. Grad and classes with bounded expansion I. Decompositions. *European Journal of Combinatorics*, 29(3):760–776, 2008.

[24] J. Nešetřil and P. Ossona de Mendez. Grad and classes with bounded expansion II. Algorithmic aspects. *European Journal of Combinatorics*, 29(3):777–791, 2008.

[25] J. Nešetřil and P. Ossona de Mendez. *Sparsity: Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and Combinatorics*. Springer, 2012.

[26] J. Nešetřil, P. Ossona de Mendez, and D. R. Wood. Characterisations and examples of graph classes with bounded expansion. *European Journal of Combinatorics*, 33(3):350–373, 2012.

[27] F. Reidl. *Structural sparseness and complex networks*. Dr., Aachen, Techn. Hochsch., Aachen, 2015. Aachen, Techn. Hochsch., Diss., 2015.

[28] N. Robertson and P. Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43–76, 2003.

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

[30] P. Wollan. The structure of graphs not admitting a fixed immersion. *Journal of Combinatorial Theory, Series B*, 110:47–66, 2015.

[31] X. Zhu. Colouring graphs with bounded generalized colouring number. *Discrete Mathematics*, 309(18):5562–5568, 2009.