



Contents lists available at ScienceDirect

European Journal of Combinatorics

journal homepage: www.elsevier.com/locate/ejcNeighbourhood complexity of graphs of bounded twin-width[☆]Édouard Bonnet^a, Florent Foucaud^{b,c}, Tuomo Lehtilä^{d,e},
Aline Parreau^f^a Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France^b Université Clermont Auvergne, CNRS, Clermont Auvergne INP, Mines Saint-Etienne, LIMOS, 63000 Clermont-Ferrand, France^c Univ. Orléans, INSA Centre Val de Loire, LIFO EA 4022, F-45067 Orléans Cedex 2, France^d Univ Lyon, UCBL, CNRS, LIRIS - UMR 5205, F69622, France^e University of Turku, Department of Mathematics and Statistics, Turku, Finland^f Univ Lyon, CNRS, INSA Lyon, UCBL, Centrale Lyon, Univ Lyon 2, LIRIS, UMR5205, F-69622 Villeurbanne, France

ARTICLE INFO

Article history:

Received 12 January 2023

Accepted 15 June 2023

Available online xxxxx

ABSTRACT

We give essentially tight bounds for, $\nu(d, k)$, the maximum number of distinct neighbourhoods on a set X of k vertices in a graph with twin-width at most d . Using the celebrated Marcus–Tardos theorem, two independent works (Bonnet et al., 2022; Przybyszewski, 2022) have shown the upper bound $\nu(d, k) \leq \exp(\exp(O(d)))k$, with a double-exponential dependence in the twin-width. The work of Gajarsky et al. (2022), using the framework of local types, implies the existence of a single-exponential bound (without explicitly stating such a bound). We give such an explicit bound, and prove that it is essentially tight. Indeed, we give a short self-contained proof that for every d and k

$$\nu(d, k) \leq (d + 2)2^{d+1}k = 2^{d+\log d+\Theta(1)}k,$$

and build a bipartite graph implying $\nu(d, k) \geq 2^{d+\log d+\Theta(1)}k$, in the regime when k is large enough compared to d .

© 2023 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license

(<http://creativecommons.org/licenses/by/4.0/>).

[☆] Florent Foucaud was financed by the French government IDEX-ISITE initiative 16-IDEX-0001 (CAP 20–25) and by the ANR project GRALMECO (ANR-21-CE48-0004). Tuomo Lehtilä's research was supported by the Finnish Cultural Foundation and by the Academy of Finland grant 338797.

E-mail address: tualeh@utu.fi (T. Lehtilä).

1. Introduction

The aim of this paper is to refine our understanding of how complex the neighbourhoods of graphs of bounded twin-width can be. We provide an improved bound on the neighbourhood complexity of such graphs, complemented by a construction showing that our bound is essentially tight. The improvements in the bounds for neighbourhood complexities translate directly to better structural bounds and algorithms, in some contexts which are explained below.

Twin-width. Twin-width is a recently introduced graph invariant [10]; see Section 2 for a definition. It can be naturally extended to matrices over finite alphabets and binary structures [7,10,12]. Although classes of bounded twin-width are broad and diverse, they allow (most of the time, provided a witness is given as an input) improved algorithms, compared to what is possible on general graphs or binary structures.

Most prominently, it was shown [10] that, on n -vertex graphs given with a d -sequence (a witness that their twin-width is at most d), deciding if a first-order sentence φ holds can be solved in time $f(d, \varphi)n$, for some computable function f . In some special cases, such as for k -INDEPENDENT SET or k -DOMINATING SET,¹ single-exponential parameterised algorithms running in time $2^{O_d(k)}n$ are possible [5]. In the same setting, the triangles of an n -vertex m -edge graph can be counted in time $O(d^2n + m)$ [21]. See [8,19,26] for more applications of twin-width with an algorithmic flavour.

Classes of binary structures with bounded twin-width include bounded treewidth, and more generally, bounded clique-width classes, proper minor-closed classes, posets of bounded width (that is, whose antichains are of bounded size), hereditary subclasses of permutations, as well as $\Omega(\log n)$ -subdivisions of n -vertex graphs [10], and particular classes of (bounded-degree) expanders [6]. A rich range of geometric graph classes have bounded twin-width such as map graphs, bounded-degree string graphs [10], classes with bounded queue number or bounded stack number [6], segment graphs with no $K_{t,t}$ subgraph, and visibility graphs of simple polygons without large independent sets [4], to give a few examples.

If efficiently approximating the twin-width is a challenging open question in general, this is known to be possible for the above-mentioned classes (albeit a representation may be needed for the geometric classes) and for ordered graphs [7]. By that, we mean that there are two computable functions f, g and an algorithm that, for an input n -vertex graph G from the class and an integer k , in time $g(k)n^{O(1)}$, either outputs an $f(k)$ -sequence (again, witnessing that the twin-width is at most $f(k)$) or correctly reports that the twin-width of G is larger than k .

Structural properties of graph classes of bounded twin-width include χ -boundedness [5], with a polynomial binding function [13], smallness (i.e., containing up to isomorphism $2^{O(n)}$ n -vertex graphs) [6,12], and Vapnik–Chervonenkis (VC) density at most 1 [9,19,27]. The latter property is the topic of the current article.

VC density and neighbourhood complexity. VC density is related to the celebrated VC dimension [30]. Given a set-system (or hypergraph) \mathcal{S} on a domain X , the *shatter function* $\pi_{\mathcal{S}} : \mathbb{N} \rightarrow \mathbb{N}$ is defined as

$$\pi_{\mathcal{S}}(n) = \max_{A \in \binom{X}{n}} |\{Y \subseteq A \mid \exists S \in \mathcal{S}, Y = A \cap S\}|.$$

The Perles–Sauer–Shelah lemma states that $\pi_{\mathcal{S}}(n) = O(n^d)$ if the VC dimension of \mathcal{S} (i.e., the supremum of $\{n \mid \pi_{\mathcal{S}}(n) = 2^n\}$) is a finite integer d . Then the VC density of \mathcal{S} is defined as $\inf\{c \in \mathbb{R} \mid \pi_{\mathcal{S}}(n) = O(n^c)\}$, and as $+\infty$ if the VC dimension is unbounded.

We define the *VC density* of an infinite class \mathcal{C} of finite graphs as the VC density of the infinite set-system formed by the neighbourhood hypergraph of the disjoint union of the graphs of \mathcal{C} , that is, $\{N_G(v) \mid v \in V(\uplus_{G \in \mathcal{C}} G)\}$, where $N_G(v)$ denotes the set of neighbours of v in G . The VC density is an important measure in finite model theory, often more tractable than the VC dimension (see for instance [1,2]). Tight bounds have been obtained for the VC density of (logically) definable

¹ That is, the problems of deciding whether in an input graph, there are k vertices that are pairwise non-adjacent or whose closed neighbourhood is the entire vertex set, respectively.

hypergraphs from graph classes of bounded clique-width [25] (with monadic second-order logic), and more recently, of bounded twin-width [19] (with first-order logic).

In structural graph theory and kernelisation [17] (a subarea of parameterised complexity [15]) the function $\pi_{\mathcal{N}(G)}$, where $\mathcal{N}(G)$ is the neighbourhood hypergraph of G , is often² called *neighbourhood complexity*. (See [3] for an algorithmic study of the computation of this notion.) In these contexts, obtaining the best possible upper bound for $\pi_{\mathcal{N}(G)}$ (and not just the exponent matching the VC density) translates to qualitatively better structural bounds and algorithms; see for instance [9,11,16,29].

The *r-neighbourhood complexity* of G is the neighbourhood complexity of G^r , with same vertex set as G , and an edge between two vertices at distance at most r in G . Reidl et al. [29] showed that among subgraph-closed classes, bounded expansion³ is equivalent to linear r -neighbourhood complexity. Indeed, the more general nowhere dense classes [23]⁴ have almost linear r -neighbourhood complexity [16]: there is a function $f : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ such that for every $\varepsilon > 0$, $\pi_{\mathcal{N}(G^r)}(n) \leq f(r, \varepsilon)n^{1+\varepsilon}$ for all n . On hereditary classes, i.e., closed under taking *induced* subgraphs, there is no known characterisation of linear neighbourhood complexity.

As we already mentioned in a different language, bounded twin-width classes have been proven to have linear neighbourhood complexity. See [9, Lemma 3] or [27, Section 3] for two independent proofs, both using the Marcus–Tardos theorem [22]. However, the dependence in the twin-width is doubly exponential in both papers. Setting $\nu(d, k)$ as the maximum number of distinct neighbourhoods on a set of size k within a graph of twin-width at most d , i.e., $\max\{\pi_{\mathcal{N}(G)}(k) \mid G \text{ has twin-width at most } d\}$, they show that $\nu(d, k) \leq \exp(\exp(O(d)))k$. There is a recent third proof not using the Marcus–Tardos theorem [19]. The authors tackle the more general problem of bounding the number of distinct first-order definable subsets within a fixed set. In the particular case of neighbourhoods, even though this is not made explicit in [19], their proof gives a similar upper bound of $\nu(d, k)$ to ours.

Our results. In this note, we give in Section 3 a short and self-contained proof (also not using the Marcus–Tardos theorem) that $\nu(d, k) \leq 2^{d+\log d+\Theta(1)}k$. In Section 4, we complement that proof with a construction of a bipartite graph witnessing that $\nu(d, k) \geq 2^{d+\log d+\Theta(1)}k$, which makes our single-exponential upper bound in twin-width essentially tight.

2. Preliminaries

We use the standard graph-theoretic notations: $V(G)$, $E(G)$, $G[S]$, $G - S$ respectively denote the vertex set, edge set, subgraph of G induced by S , and subgraph of G induced by $V(G) \setminus S$. If $v \in V(G)$, then the *open neighbourhood of vertex v in G* denoted by $N_G(v)$ (or $N(v)$ if G is clear from the context) is the set of neighbours of v in G . If $X \subseteq V(G)$, then an *X -neighbourhood* is a set $N(v) \cap X$ for some $v \in V(G)$.

We now define the twin-width of a graph, following the definition of [10].

A *trigraph* is a triple $G = (V(G), E(G), R(G))$ where $E(G)$ and $R(G)$ are two disjoint sets of edges on $V(G)$: the usual edges (also called *black edges*) and the *red edges*. Informally, a red edge between two vertices u and v means that some errors have been made between u and v . The *red degree* of a trigraph is the maximum degree of the graph $(V(G), R(G))$. Any graph G can be interpreted as a trigraph $G = (V(G), E(G), \emptyset)$. Given a trigraph and two vertices $u, v \in V(G)$ (not necessarily adjacent), the trigraph $G/u, v = G'$ is obtained by *contracting* u and v in a new vertex w such that:

- $V(G') = \{w\} \cup V(G) \setminus \{u, v\}$;
- the edges between vertices of $V(G) \setminus \{u, v\}$ are the same in G' ;
- we set the edges incident to w in the following way:

² Some authors define the neighbourhood complexity as $n \mapsto \frac{\pi_{\mathcal{N}(G)}(n)}{n}$.

³ A notion from the Sparsity theory of Nešetřil and Ossona de Mendez [24] extending bounded degree and proper minor-free classes.

⁴ Another invention of the Sparsity program [24].

- $wx \in E(G')$ if $xu \in E(G)$ and $xv \in E(G)$;
- $wx \notin E(G') \cup R(G')$ if $xu \notin E(G) \cup R(G)$ and $xv \notin E(G) \cup R(G)$;
- $wx \in R(G')$ otherwise.

In other words, the common black neighbours of u and v are black neighbours of w . All the other neighbours of u or v are red neighbours of w . Red edges stay red, black edges stay black, red and black edges become red. Moreover, non-edges stay as non-edges, non-edges and red edges become red edges, and non-edges and black edges become red edges. We say that $G/u, v$ is a contraction of G . A d -sequence of an n -vertex graph G is a sequence of n trigraphs $G = G_n, G_{n-1}, \dots, G_1$ such that each trigraph G_i is obtained from G_{i+1} by a contraction and has red degree at most d . The twin-width of G , denoted by $\text{tw}(G)$, is the minimum integer d such that G admits a d -sequence. Note that an induced subgraph of G has a twin-width smaller or equal to the twin-width of G [10].

If $u \in G_i$, then $u(G)$ denotes the set of vertices of G eventually contracted to u in G_i . Instead of considering the trigraphs G_i , we might prefer to deal with the partitions of $V(G)$ induced by the sets $u(G)$ for u in G_i : $\mathcal{P}_i = \{u(G) \mid u \in V(G_i)\}$. In this setting, we say that $u(G)$ is a part of \mathcal{P}_i . We say that there is a red edge, a black edge or a non-edge between two parts $u(G)$ and $v(G)$ of \mathcal{P}_i if uv is a red edge, a black edge or a non-edge in G_i .

3. Upper bound on the number of distinct neighbourhoods

We state and prove our upper bound on the maximum number of distinct X -neighbourhoods in bounded twin-width graphs.

Theorem 1. *Let G be an n -vertex graph of twin-width d , and $X \subseteq V(G)$ where $X \neq \emptyset$. Then the number of distinct X -neighbourhoods in G is at most $(d + 2)2^{d+1}|X| = 2^{d+\log d + \Theta(1)}|X|$.*

Proof. Fix non-empty $X \subseteq V(G)$. First of all, for all vertices of $V(G) \setminus X$ with the same X -neighbourhood, we keep only one representative. Note that the new graph G'' is an induced subgraph of G , thus its twin-width is at most d . We further modify graph G'' by adding for each $v \in X$ a new vertex u to G'' so that $N(u) = N(v)$ if such vertex does not exist in $V(G'') \setminus X$. We do this one vertex at a time. The new graph is called G' and it has the same twin-width as G'' .

Let $M = (d + 2)2^{d+1} + 1$. We prove by induction on n that an n -vertex graph of twin-width at most d with a set X of $k \geq 1$ vertices, where all vertices outside X have a distinct X -neighbourhood, satisfies $n \leq kM$. This will prove that G' has at most kM vertices, and thus that in G , there are at most $(M - 1)k$ distinct X -neighbourhoods.

The statement is trivially true for $n \leq 5$ since $M \geq 5$, for all $d \geq 0$.

Thus, assume $n \geq 6$. In particular, we have $k > 1$. Let $x \in X$. Let $X' = X \setminus \{x\}$ and let T_x be the set of pairs of vertices outside X that are twins with respect to X' , i.e.

$$T_x = \left\{ \{u, v\} \in \binom{V(G') \setminus X}{2} \mid N(u) \cap X' = N(v) \cap X' \right\}.$$

Since every vertex of $V(G') \setminus X$ has a distinct neighbourhood in X , there are at most two vertices of $V(G') \setminus X$ with the same (possibly empty) neighbourhood N in X' ; namely the vertices $u, v \in V(G') \setminus X$ with $N(u) \cap X = N$ and $N(v) \cap X = N \cup \{x\}$ (if they exist). Hence, T_x consists of pairwise-disjoint pairs of vertices.

We prove the following claim.

Claim A. *There exists a vertex x of X such that T_x comprises at most $M - 1$ pairs, in G' .*

Proof of claim. Consider a d -sequence of contractions G'_n, \dots, G'_1 of G' . Consider the last step G'_i of the sequence where all the parts of \mathcal{P}_i contain at most one vertex of X (that is, contrary to \mathcal{P}_i , some part of \mathcal{P}_{i-1} contains two vertices of X).

Let P be a part of \mathcal{P}_i . Let x be the unique (if there exists one) element of $P \cap X$. Then we claim that $|P \setminus X| \leq 2^{d+1}$. Indeed, any two vertices of $P \setminus X$ have some vertex in the symmetric difference of

their X -neighbourhoods, either it is x , or some vertex x' of X outside P . If that distinguishing vertex is some x' that is not in P , then there has to be a red edge between P and the part that contains x' . There are at most d red edges with P as an extremity. Since all the elements of X are in distinct parts in G'_i , it means that $d + 1$ vertices of X are enough to distinguish all the X -neighbourhoods of vertices of $P \setminus X$, and thus $|P \setminus X| \leq 2^{d+1}$.

We now consider the next contraction in the sequence, which leads to G'_{i-1} . By definition of G'_i , it must contract two vertices corresponding to two parts of \mathcal{P}_i that both contain an element of X . Let x_1 and x_2 be these two elements of X . Let Q be the part of \mathcal{P}_{i-1} that contains both x_1 and x_2 . Let $\{u, v\}$ be a pair of T_{x_1} and let T_{x_1} contain M' pairs. Since u and v have the same neighbourhood in $X \setminus \{x_1\}$, it means that they are either both adjacent or both non-adjacent to x_2 , and exactly one of them is adjacent to x_1 . Thus, necessarily, one vertex among the pair $\{u, v\}$ is adjacent to exactly one vertex among $\{x_1, x_2\}$. In particular, if this vertex is not in Q , then there has to be a red edge between the part containing this vertex and the part Q in G'_{i-1} . Since T_{x_1} contains M' pairs (which are disjoint) and Q has at most 2^{d+2} vertices not in X , there are at least $M' - 2^{d+2}$ vertices not in X whose part in G'_{i-1} has a red edge to Q . Since each other part has at most 2^{d+1} vertices not in X , it makes at least $\frac{M' - 2^{d+2}}{2^{d+1}}$ red edges incident to Q . Thus, we must have $\frac{M' - 2^{d+2}}{2^{d+1}} \leq d$, leading to $M' \leq 2^{d+1}(d + 2) = M - 1$, which proves the claim. \diamond

By **Claim A**, there exists a vertex $x \in X$ such that $|T_x| \leq M - 1$. Let Y be a set of $|T_x|$ vertices that intersects each pair of T_x exactly once. Let $G_Y = G' - (Y \cup \{x\})$. Then, $X' = X \setminus \{x\}$ is a vertex set of size $k - 1$ such that all X' -neighbourhoods of vertices outside X' are distinct. The graph G_Y has at least $n - M$ vertices, and twin-width at most d . By induction, we have $n - M \leq |V(G_Y)| \leq (k - 1)M$ and thus, $n \leq kM$. Hence, once we recall that no vertex in X has unique X -neighbourhood, there are at most $(M - 1)k$ distinct X -neighbourhoods, which completes the proof. \square

4. Lower bound on the number of distinct neighbourhoods

Notice that when $|X|$ and $\text{tw}(G)$ are roughly the same, the bound from **Theorem 1** cannot be sharp, since G' has at most $2^{|X|} + |X|$ vertices. However, when $|X|$ is large enough compared to $\text{tw}(G)$, we next show that the bound is sharp up to a constant factor.

Proposition 2. *There is a positive constant c , such that for any integer d , there is a bipartite graph G of twin-width at most d , and a large enough set $X \subseteq V(G)$, with at least $c \cdot d2^d |X| = 2^{d+\log d+\Theta(1)} |X|$ distinct X -neighbourhoods in G .*

Proof. Observe that the claim is clearly true for any small d . Thus, we do not need to consider separately graphs with small twin-width upper bounded by a constant. Hence, we assume from now on that $d \geq d'$ where d' is some positive constant (at least 3).

We construct the graph G as follows. Let $A, B, C \in \mathbb{Z}$ be three constants that will be given later (A and B will be roughly equal to \sqrt{d} and C will be roughly equal to d). Let $X = \{x_1, \dots, x_k\}$ be an independent set of $k \geq d + 2\sqrt{d-2} + 1$ vertices. Our goal is to construct G so that each vertex in $V(G) \setminus X$ has a unique X -neighbourhood. For any integers i, j, t with $1 \leq i \leq j \leq i + A - 1$, $j + 2 \leq t \leq j + 1 + B$ and $t \leq k - C$, we create a set $V_{i,j,t}$ of vertices as follows. Consider the set $X_t = \{x_{t+1}, \dots, x_{t+C}\}$. For every subset Y of X_t , let $Y' = \{x_i, \dots, x_j, x_t\} \cup Y$ and add a vertex $v_{Y'}$ to $V_{i,j,t}$, making it adjacent to the vertices of Y' . Each set $V_{i,j,t}$ has size 2^C and there are $\Theta(kAB)$ (for fixed A, B and C , and growing k) such sets. Thus there are $\Theta(kAB2^C)$ vertices in the graph.

Any two vertices not in X have distinct X -neighbourhoods. Indeed, by considering the natural ordering of X induced by the indices, any vertex not in X is first adjacent to a consecutive interval of vertices from x_i to x_j , then is not adjacent to vertices from x_{j+1} to x_{t-1} (which is not empty since $t \geq j + 2$), and then adjacent to x_t . Thus, if two vertices have the same X -neighbourhood, they must be in the same set $V_{i,j,t}$. But then, they have a distinct neighbourhood in $\{x_{t+1}, \dots, x_{t+C}\}$.

We now prove that the twin-width of G is at most $M = \max\{AB, C\} + 2$. For that, we give a sequence of contractions with red degree at most M .

The contraction sequence is split into $k - C$ steps. During these steps we first consider vertices of X one by one and then in the last one we deal with the remaining vertices of X . Let $0 \leq \ell \leq k - C - 1$. Step 0 corresponds to the starting point, where each vertex is alone. Let $\ell \geq 1$. After Step ℓ , there will be the following parts in the corresponding partition (vertices not in any of the mentioned parts are in corresponding singleton parts containing only the vertices themselves):

- Let $i = \ell$. For each j, t such that $i \leq j \leq i + A - 1$ and $j + 2 \leq t \leq j + 1 + B$, there is a part $B_{j,t}$. The parts $B_{i,t}$ (parts with $j = i$), contain all the vertices of the sets $V_{i',j',t}$ such that $j' \leq i$. The parts $B_{j,t}$ with $j > i$ contain all the vertices of the sets $V_{i',j',t}$ such that $i' \leq i$ and $j' = j$. Note that there are AB non-empty $B_{j,t}$ parts in total.
- There is a part X_0 that contains vertices from x_1 to x_ℓ of X .
- There is a part T (for “trash”) that contains all the vertices of the sets $V_{i',j,t}$ with $t \leq \ell + 1$.

All the other vertices are not yet contracted. This corresponds to the vertices from $x_{\ell+1}$ to x_k of X and to the vertices of the sets $V_{i',j,t}$ with $i' > i = \ell$. Indeed, if $i' \leq i$ and $t \leq i + 1$, then the vertices of $V_{i',j,t}$ are in T . If $t \geq i + 2$ but $j \leq i$, then they are in the part $B_{i,t}$. If $j > i$, then they are in the part $B_{j,t}$.

We first prove that the red degree after Step ℓ is at most M . Then, we explain how to get from Step ℓ to Step $\ell + 1$ by keeping the red degree at most M .

Consider the part $B_{j,t}$ at the end of Step ℓ . A vertex in this part belongs to some set $V_{i',j',t}$ with $i' \leq i = \ell$ and $j' = j$ if $j > i$ or $j' \leq i$ otherwise. In particular, two vertices of $B_{j,t}$ are adjacent to all the vertices between x_{i+1} and x_j , to no vertex between x_{j+1} and x_{t-1} , to x_t , and to no vertex after x_{t+C} . Thus, there is a red edge between the parts $B_{j,t}$ and X_0 , and C red edges between the part $B_{j,t}$ and the vertices $\{x_{t+1}, \dots, x_{t+C}\}$. Therefore, the number of red edges incident with $B_{j,t}$ is at most $C + 1$.

Consider now the part T . Vertices in T are adjacent only to vertices of X up to $x_{\ell+C+1}$. Since vertices x_1 to x_ℓ are all in the part X_0 , the red degree of T is at most $C + 2$.

Single vertices not in X have no incident red edges: indeed, they are all in some sets $V_{i',j,t}$ for $i' > i = \ell$ and thus are not adjacent to any vertex of X_0 . For the same reason, there are red edges incident to X_0 only to T and to the parts $B_{j,t}$. Hence, the red degree of X_0 is at most $AB + 1$. Similarly, the red degree of $x_{i'}$, $i' > i + 1$ is at most $AB + 1$. Moreover, the red degree of x_{i+1} is at most one. Indeed, the only red edge is between x_{i+1} and T .

Finally, the red degree after step ℓ is at most $\max\{AB + 1, C + 2\} \leq M$.

Let $\ell \geq 0$. We now explain how we perform the contractions to go from step ℓ to step $\ell + 1$.

1. (only if $\ell \geq 1$) Let $i = \ell$. For any $i + 3 \leq t \leq i + 2 + B$, merge the part $B_{i,t}$ with the part $B_{i+1,t}$ resulting in part $B_{i+1,t}$. The only new red edge this merging may lead to, when $B_{i,t}$ is non-empty, is between $B_{i+1,t}$ and x_{i+1} . Thus, we add only one red edge between x_{i+1} and $B_{i+1,t}$. Thus, the red degree of $B_{i+1,t}$ is at most $C + 2$ and the red degree of x_{i+1} is at most 2.
2. Add all the vertices of $V_{i+1,j,t}$ for some j, t to the part (that might be empty at this point) $B_{j,t}$. The red degree of $B_{j,t}$ is at most $C + 2$ since we might have a red edge between $B_{j,t}$ and x_{i+1} . The number of nonempty parts $B_{j,t}$ at this point is at most $AB + 1$ (there is still the part $B_{i,i+2}$). Adding T , this gives $AB + 2$ red edges incident to a vertex in X (or from part X_0).
3. Add $x_{\ell+1}$ to X_0 . The part X_0 can have red edges only to non-empty parts $B_{j,t}$ and to T , but no red edges to the single vertices. Thus, it has red degree at most $AB + 2$.
4. Put the part $B_{i,i+2}$ into T . This part is only adjacent to vertices up to $x_{\ell+2+C}$, and thus has at most $C + 2$ red edges.

Thus, at each point, the red degree is always at most $M = \max\{AB, C\} + 2$.

The process ends at step $\ell = k - C - 1$. Then, all the vertices not in X are in some parts, and there are at most $AB + 1$ such parts. On the other side of the bipartition, we have part X_0 and $C + 1$ single vertices. Thus, the graph is bipartite with both sides of size at most M . One can contract each part independently to finish the contraction sequence.

To conclude, taking $C = d - 2$ and $A = B = \lfloor \sqrt{d - 2} \rfloor$, we have $M \leq d$ and $kAB2^C = \Theta(kd2^d)$. Notice that we may assume that A, B and C are positive since $d \geq d'$ where d' was some well chosen positive constant. This concludes the proof. \square

5. Conclusion

We have given an essentially tight upper bound for the neighbourhood complexity of graphs of bounded twin-width together with a construction almost attaining this upper bound. Moreover, our method is simple and self-contained. A similar upper bound was implied by the techniques in [19] (though not stated explicitly).

It is known that the twin-width of G^r can be upper-bounded by a function of the twin-width of G and r [10]. Thus, graphs of twin-width at most d have linear r -neighbourhood complexity. Recently, improved bounds were given for planar graphs and proper minor-closed graph classes in [20] (such graphs also have bounded twin-width). We leave as an interesting open problem to obtain an essentially tight twin-width dependence for the r -neighbourhood complexity.

We remark that the neighbourhood complexity is also related to *identification problems* on graphs such as *identifying codes* or *locating-dominating sets*, where one seeks a (small) set A of vertices of a graph such that all other vertices have a distinct neighbourhood in A [18]. Some works in this area about specific graph classes, are equivalent to the study of the neighbourhood complexity of these graph classes: see for example [14,18,28]. Moreover, we note that for graph classes with VC density 1, since any solution has linear size, the natural minimisation versions of the above identification problems have a polynomial-time constant-factor approximation algorithm (trivially select the whole vertex set), while such an algorithm is unlikely to exist in the general case [14]. Thus, the bounds given in the current work imply a better approximation ratio for these problems, when restricted to input graph classes of bounded twin-width.

Acknowledgement

We thank an anonymous reviewer for pointing out the implicit upper bound of [19] mentioned in the introduction.

References

- [1] Matthias Aschenbrenner, Alf Dolich, Deirdre Haskell, Dugald Macpherson, Sergei Starchenko, Vapnik–Chervonenkis density in some theories without the independence property, II, *Notre Dame J. Formal Logic* 54 (3–4) (2013) 311–363.
- [2] Matthias Aschenbrenner, Alf Dolich, Deirdre Haskell, Dugald Macpherson, Sergei Starchenko, Vapnik–Chervonenkis density in some theories without the independence property, I, *Trans. Amer. Math. Soc.* 368 (8) (2016) 5889–5949.
- [3] Cristina Bazgan, Florent Foucaud, Florian Sikora, Parameterized and approximation complexity of Partial VC Dimension, *Theoret. Comput. Sci.* 766 (2019) 1–15.
- [4] Édouard Bonnet, Dibyayan Chakraborty, Eun Jung Kim, Noleen Köhler, Raul Lopes, Stéphan Thomassé, Twin-width VIII: delineation and win-wins, in: Holger Dell, Jesper Nederlof (Eds.), 17th International Symposium on Parameterized and Exact Computation, IPEC 2022, September 7–9, 2022, Potsdam, Germany, in: LIPIcs, Vol. 249, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, pp. 9:1–9:18.
- [5] Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, Rémi Watrigant, Twin-width III: max independent set, min dominating set, and coloring, in: Nikhil Bansal, Emanuela Merelli, James Worrell (Eds.), 48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12–16, 2021, Glasgow, Scotland (Virtual Conference), in: LIPIcs, Vol. 198, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, pp. 35:1–35:20.
- [6] Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, Rémi Watrigant, Twin-width II: small classes, *Combin. Theory 2* (2) (2022).
- [7] Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, Pierre Simon, Stéphan Thomassé, Szymon Toruńczyk, Twin-width IV: Ordered graphs and matrices, in: Stefano Leonardi, Anupam Gupta (Eds.), STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022, ACM, 2022, pp. 924–937.
- [8] Édouard Bonnet, Eun Jung Kim, Amadeus Reinald, Stéphan Thomassé, Twin-width VI: The lens of contraction sequences, in: Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, SIAM, 2022, pp. 1036–1056.
- [9] Édouard Bonnet, Eun Jung Kim, Amadeus Reinald, Stéphan Thomassé, Rémi Watrigant, Twin-width and polynomial kernels, *Algorithmica* 84 (2022) 3300–3337.
- [10] Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, Rémi Watrigant, Twin-width I: Tractable FO model checking, *J. ACM* 69 (1) (2022) 3:1–3:46.
- [11] Édouard Bonnet, O-joung Kwon, David R. Wood, Reduced bandwidth: A qualitative strengthening of twin-width in minor-closed classes (and beyond), 2022, CoRR, [abs/2202.11858](https://arxiv.org/abs/2202.11858).
- [12] Édouard Bonnet, Jaroslav Nešetřil, Patrice Ossona de Mendez, Sebastian Siebertz, Stéphan Thomassé, Twin-width and permutations, 2021, CoRR, [abs/2102.06880](https://arxiv.org/abs/2102.06880).

- [13] Romain Bourneuf, Stéphan Thomassé, Bounded twin-width graphs are polynomially χ -bounded, 2023, arXiv preprint arXiv:2303.11231.
- [14] Nicolas Bousquet, Aurélie Lagoutte, Zhentao Li, Aline Parreau, Stéphan Thomassé, Identifying codes in hereditary classes of graphs and VC-dimension, *SIAM J. Discrete Math.* 29 (4) (2015) 2047–2064.
- [15] Marek Cygan, Fedor V Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, Saket Saurabh, *Parameterized Algorithms*, Vol. 4, Springer, 2015.
- [16] Kord Eickmeyer, Archontia C. Giannopoulou, Stephan Kreutzer, O-joung Kwon, Michał Pilipczuk, Roman Rabinovich, Sebastian Siebertz, Neighborhood complexity and kernelization for nowhere dense classes of graphs, in: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, Anca Muscholl (Eds.), 44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10–14, 2017, Warsaw, Poland, in: LIPIcs, Vol. 80, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, pp. 63:1–63:14.
- [17] Fedor V Fomin, Daniel Lokshtanov, Saket Saurabh, Meirav Zehavi, *Kernelization: Theory of Parameterized Preprocessing*, Cambridge University Press, 2019.
- [18] Florent Foucaud, George B. Mertzios, Reza Naserasr, Aline Parreau, Petru Valicov, Identification, location-domination and metric dimension on interval and permutation graphs. I. Bounds, *Theoret. Comput. Sci.* 668 (2017) 43–58.
- [19] Jakub Gajarský, Michał Pilipczuk, Wojciech Przybyszewski, Szymon Toruńczyk, Twin-width and types, in: Mikołaj Bojanczyk, Emanuela Merelli, David P. Woodruff (Eds.), 49th International Colloquium on Automata, Languages, and Programming, ICALP 2022, July 4–8, 2022, Paris, France, in: LIPIcs, Vol. 229, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, pp. 123:1–123:21.
- [20] Gwenaél Joret, Clément Rambaud, Neighborhood complexity of planar graphs, 2023, arXiv preprint arXiv:2302.12633.
- [21] Stefan Kratsch, Florian Nelles, Alexandre Simon, On triangle counting parameterized by twin-width, 2022, CoRR, abs/2202.06708.
- [22] Adam Marcus, Gábor Tardos, Excluded permutation matrices and the Stanley-Wilf conjecture, *J. Comb. Theory Ser. A* 107 (1) (2004) 153–160.
- [23] Jaroslav Nešetřil, Patrice Ossona de Mendez, On nowhere dense graphs, *European J. Combin.* 32 (4) (2011) 600–617.
- [24] Jaroslav Nešetřil, Patrice Ossona de Mendez, Sparsity - Graphs, Structures, and Algorithms, in: *Algorithms and combinatorics*, Vol. 28, Springer, 2012.
- [25] Adam Paszke, Michał Pilipczuk, VC density of set systems definable in tree-like graphs, in: Javier Esparza, Daniel Král' (Eds.), 45th International Symposium on Mathematical Foundations of Computer Science, MFCS 2020, August 24–28, 2020, Prague, Czech Republic, in: LIPIcs, Vol. 170, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, pp. 78:1–78:13.
- [26] Michał Pilipczuk, Marek Sokołowski, Anna Zych-Pawlewicz, Compact representation for matrices of bounded twin-width, in: Petra Berenbrink, Benjamin Monmege (Eds.), 39th International Symposium on Theoretical Aspects of Computer Science, STACS 2022, March 15–18, 2022, Marseille, France (Virtual Conference), in: LIPIcs, Vol. 219, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, pp. 52:1–52:14.
- [27] Wojciech Przybyszewski, VC-density and abstract cell decomposition for edge relation in graphs of bounded twin-width, 2022, CoRR, abs/2202.04006.
- [28] Douglas Rall, Peter J. Slater, On location-domination numbers for certain classes of graphs, *Congr. Numer.* 45 (1984) 97–106.
- [29] Felix Reidl, Fernando Sánchez Villamil, Konstantinos S. Stavropoulos, Characterising bounded expansion by neighbourhood complexity, *European J. Combin.* 75 (2019) 152–168.
- [30] Vladimir N. Vapnik, Alexey Y. Chervonenkis, On the uniform convergence of relative frequencies of events to their probabilities, in: *Measures of Complexity*, Springer, 2015, pp. 11–30.