



Twin-width II: small classes

Edouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, Rémi Watrigant

► To cite this version:

Edouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, Rémi Watrigant. Twin-width II: small classes. ACM-SIAM Symposium on Discrete Algorithms (SODA21), Jan 2021, Alexandria, United States. ⟨hal-03107577⟩

HAL Id: hal-03107577

<https://hal.science/hal-03107577v1>

Submitted on 13 Aug 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Twin-width II: small classes*

Édouard Bonnet[†] Colin Geniet[‡] Eun Jung Kim[§] Stéphan Thomassé[¶]
Rémi Watrigant^{||}

Abstract

The recently introduced *twin-width* of a graph G is the minimum integer d such that G has a d -contraction sequence, that is, a sequence of $|V(G)| - 1$ iterated vertex identifications for which the overall maximum number of red edges incident to a single vertex is at most d , where a red edge appears between two sets of identified vertices if they are not homogeneous in G (not fully adjacent nor fully non-adjacent). We show that if a graph admits a d -contraction sequence, then it also has a linear-arity tree of $f(d)$ -contractions, for some function f . Informally if we accept to worsen the twin-width bound, we can choose the next contraction from a set of $\Theta(|V(G)|)$ pairwise disjoint pairs of vertices. This has two main consequences. First it permits to show that every bounded twin-width class is *small*, i.e., has at most $n!c^n$ graphs labeled by $[n]$, for some constant c . This unifies and extends the same result for bounded treewidth graphs [Beineke and Pippert, JCT '69], proper subclasses of permutations graphs [Marcus and Tardos, JCTA '04], and proper minor-free classes [Norine et al., JCTB '06]. It implies in turn that bounded-degree graphs, interval graphs, and unit disk graphs have unbounded twin-width. The second consequence is an $O(\log n)$ -adjacency labeling scheme for bounded twin-width graphs, confirming several cases of the implicit graph conjecture.

We then explore the *small conjecture* that, con-

versely, every small hereditary class has bounded twin-width. The conjecture passes many tests. Inspired by sorting networks of logarithmic depth, we show that $\log_{\Theta(\log \log d)} n$ -subdivisions of K_n (a small class when d is constant) have twin-width at most d . We obtain a rather sharp converse with a surprisingly direct proof: the $\log_{d+1} n$ -subdivision of K_n has twin-width at least d . Secondly graphs with bounded stack or queue number (also small classes) have bounded twin-width. These sparse classes are surprisingly rich since they contain certain (small) classes of expanders. Thirdly we show that cubic expanders obtained by iterated random 2-lifts from K_4 [Bilu and Linial, Combinatorica '06] also have bounded twin-width. These graphs are related to so-called separable permutations and also form a small class. We suggest a promising connection between the small conjecture and group theory.

Finally we define a robust notion of sparse twin-width. We show that for a hereditary class \mathcal{C} of bounded twin-width the five following conditions are equivalent: every graph in \mathcal{C} (1) is $K_{t,t}$ -free for some fixed t , (2) has an adjacency matrix without a d -by- d division with a 1 entry in each d^2 cells for some fixed d , (3) has at most linearly many edges, (4) the subgraph closure of \mathcal{C} has bounded twin-width, and (5) \mathcal{C} has bounded expansion. We discuss how sparse classes with similar behavior with respect to clique subdivisions compare to bounded sparse twin-width.

1 Introduction

We continue to develop the theory of twin-width, a novel graph and matrix invariant introduced in the first paper of the series [7]. We start with a bird's eye view of our results. The exact definitions of some objects and concepts will be deferred to the next section, but this introduction can be read by taking them as black boxes. Furthermore Section 2 includes a summary of the first paper, so that the current paper is self-contained.

A trigraph is a graph with two disjoint edge sets: black edges (regular edges) and red edges (error edges). The graph induced by the red edges (resp. black edges) is called the red graph (resp. black graph). A d -trigraph has a red graph with maximum degree at

*This work was supported by the grants from French National Agency under PRC program (project Digraphs, ANR-19-CE48-0013-01), under JCJC program (project ASSK, ANR-18-CE40-0025-01), and by the LABEX MILYON (ANR-10-LABX-0070) of Université de Lyon, within the program "Investissements d'Avenir" (ANR-11-IDEX-0007) operated by the French National Research Agency (ANR).

[†]Univ Lyon, CNRS, ENS de Lyon, UCBL 1, edouard.bonnet@ens-lyon.fr

[‡]University of Warsaw, colin.geniet@ens-paris-saclay.fr

[§]Université Paris-Dauphine, PSL University, eun-jung.kim@dauphine.fr

[¶]Univ Lyon, CNRS, ENS de Lyon, UCBL 1, and Institut Universitaire de France, stephan.thomasse@ens-lyon.fr

^{||}Univ Lyon, CNRS, ENS de Lyon, UCBL 1, remi.watrigant@univ-lyon1.fr

most d . A contraction in a trigraph identifies two (non-necessarily adjacent) vertices, and puts black edges towards shared neighbors in the black graph, and red edges towards the other (non-necessarily shared) neighbors (see Fig. 1). A d -contraction sequence, or d -sequence, of an n -vertex graph G is a sequence of d -trigraphs $G = G_n, G_{n-1}, \dots, G_2, G_1$ such that G_i is obtained by performing a single contraction in G_{i+1} . In particular G_1 is the one-vertex graph K_1 . The twin-width of G is the minimum d such that it admits a d -sequence.

A contraction sequence of G may be seen as a path with at the left end, G , at the right end, K_1 , and the current trigraph gets smaller and smaller when we walk from left to right. We show that this path can be made a tree of large arity. Now G is at the root of the tree, all the leaves contain the graph K_1 , and every child is obtained by performing a single contraction in the parent node. A d -contraction tree is such a tree with a d -trigraph at every node. More precisely, we show that if a graph G has a d -contraction sequence, then it has a D_d -contraction tree with linear arity. By linear arity, we mean that every non-leaf node H has $\Theta(|V(H)|)$ distinct children.

Denoting the class of graphs with twin-width at most d by \mathcal{C}_d , the first consequence is that the number of graphs in \mathcal{C}_d on the vertex set $[n]$ is at most $n!f(d)^n$. Intuitively the large-arity tree tells us that many $n-1$ -vertex graphs of \mathcal{C}_d can be obtained from the same n -vertex graph of \mathcal{C}_d . By inverting the process, there are not so many *distinct* n -vertex graphs in \mathcal{C}_d , obtained by splitting a vertex in $n-1$ -vertex graphs of \mathcal{C}_d . This crucial fact makes the inductive proof works. Our result generalizes several similar theorems in enumerative combinatorics.

The first one is an over 50-year old result that bounded treewidth graphs on vertex set $[n]$ have a similar growth in $n!c^n$ [2]. Graph classes with such a growth are called small. The second one is comparatively much more recent, it is the celebrated answer to the Stanley-Wilf conjecture, now the Marcus-Tardos theorem. Marcus and Tardos [20] showed that there are at most c_σ^n permutations over $[n]$ avoiding a fixed permutation pattern σ . In other words, every proper subclass of permutations (where a class of permutations is closed under taking subpermutations) has at most single-exponential growth, much below $n!$, the growth of the full class. Expressed in the language of graph classes, proper subclasses of permutation graphs are small. The third one, due to Norine et al. [24], is that the number of graphs on vertex set $[n]$ not containing a fixed minor H is at most $n!c_H^n$. Thus proper minor-closed classes are small.

We previously showed [7] that bounded treewidth

(even rank-width) graphs, proper subclasses of permutation graphs, and proper minor-closed classes have bounded twin-width. Thus the fact that bounded twin-width classes are small unifies and extends all the above-mentioned theorems. We then explore the converse statement. Could it be that every small hereditary class has bounded twin-width? We do not answer this question, dubbed the small conjecture, but instead we give some evidences it may be true. This comes in the form of showing that many potential counterexamples, that is, seemingly complex small hereditary classes, actually have bounded twin-width. If the conjecture is true, it gives a universal explanation for the single-exponential growth (up to isomorphism) of combinatorial classes: Translate the objects into graphs or matrices, a bound or lack thereof in the twin-width of the class decides the existence of such a bound in the growth.

Another by-product of the contraction tree is that we can always contract in parallel a linear number of disjoint pairs of vertices. This gives rise to so-called parallel d -sequences of logarithmic length. This will be instrumental in showing that bounded twin-width classes admit an $O(\log n)$ -adjacency labeling scheme. This verifies a variety of particular cases of the implicit graph conjecture which posits that such labeling schemes exist for every factorial hereditary class, i.e., hereditary class with growth $n!^{O(1)}$.

Finally we show that five different ways of restricting twin-width to sparse classes actually lead to the same notion. For example, bounded sparse twin-width classes can be equivalently defined as hereditary classes with bounded twin-width that are $K_{t,t}$ -free or where every graph has at most linearly many edges. A first but challenging step towards the small conjecture is to show that small sparse classes have bounded (sparse) twin-width. For instance, do classes with polynomial expansion have bounded twin-width? We discuss (possible) containments and strict containments of established sparse classes with respect to bounded sparse twin-width.

2 Preliminaries and outline

In this section we recall the relevant notations and definitions, summarize the important bits of the first paper, and outline our new results.

2.1 Notations and definitions. We denote by $[i, j]$ the set of integers $\{i, i+1, \dots, j-1, j\}$, and by $[i]$ the set of integers $[1, i]$. If \mathcal{X} is a set of sets, we denote by $\cup \mathcal{X}$ their union. Unless stated otherwise, all graphs are assumed undirected and simple, that is, they do not have parallel edges or self-loops. We denote by $V(G)$ and $E(G)$, the set of vertices and edges,

respectively, of a graph G . For $S \subseteq V(G)$, we denote the *open neighborhood* (or simply *neighborhood*) of S by $N_G(S)$, i.e., the set of neighbors of S deprived of S , and the *closed neighborhood* of S by $N_G[S]$, i.e., the set $N_G(S) \cup S$. We simplify $N_G(\{v\})$ into $N_G(v)$, and $N_G[\{v\}]$ into $N_G[v]$. We denote by $G[S]$ the subgraph of G induced by S , and $G - S := G[V(G) \setminus S]$. For two disjoint sets $A, B \subseteq V(G)$, $E(A, B)$ denotes the set of edges in $E(G)$ with one endpoint in A and the other one in B . Two distinct vertices u, v such that $N(u) = N(v)$ are called *false twins*, and *true twins* if $N[u] = N[v]$. Two vertices are *twins* if they are false twins or true twins. For two vertices $u, v \in V(G)$, the *distance* $d_G(u, v)$ is the number of edges in a shortest path from u to v , and ∞ if u and v are in two distinct connected components of G . In all the notations with a graph subscript, we may omit it if the graph is clear from the context.

A *graph class* is a family of graphs closed under isomorphism (i.e., under renaming the vertices). Since we will be interested in the “size” of a class, we will further impose that the vertex set of n -vertex graphs is precisely¹ $[n]$. With that requirement the number of n -vertex graphs in a class \mathcal{C} is a well-defined (finite) number. Observe that every single n -vertex graph in a class \mathcal{C} implies that at least $n!$ graphs are in \mathcal{C} , namely all its relabelings. A graph class is said *hereditary* if it is closed under taking induced subgraphs. It is said *monotone* or *subgraph-closed* if it is even closed under taking subgraphs.

A graph is *H-free* if it does not contain H as an induced subgraph. However we make an exception for $H = K_{t,t}$. A $K_{t,t}$ -free graph is a graph with no biclique $K_{t,t}$ as a subgraph. A class is *H-free* if all its graphs are H -free. When t is not yet defined, we may say that a class \mathcal{C} is K_t -free (resp. $K_{t,t}$ -free) to mean that there exists a finite integer t such that \mathcal{C} is K_t -free (resp. $K_{t,t}$ -free).

We denote by $\Delta(G)$ the maximum degree of a vertex in G , and $\Delta(\mathcal{C}) := \sup_{G \in \mathcal{C}} \Delta(G)$. A class \mathcal{C} has *bounded degree* if $\Delta(\mathcal{C}) < \infty$. More generally, for any graph invariant ι , we say that \mathcal{C} has *bounded* ι if $\iota(\mathcal{C}) := \sup_{G \in \mathcal{C}} \iota(G) < \infty$. The strong product $G \boxtimes H$ of two graphs G and H has vertex set $V(G) \times V(H)$ and $(u, v)(u', v') \in E(G \boxtimes H)$ if and only if $[u = u' \text{ or } uu' \in E(G)]$ and $[v = v' \text{ or } vv' \in E(H)]$. We denote by $\mathcal{G} \boxtimes \mathcal{H}$ the class $\{G \boxtimes H \mid G \in \mathcal{G}, H \in \mathcal{H}\}$, where \mathcal{G} and \mathcal{H} are two sets of graphs. Given a class \mathcal{C} , we denote by $\text{Sub}(\mathcal{C})$ the class of all subgraphs of members of \mathcal{C} . The class $\text{Sub}(\mathcal{C})$ is by definition subgraph-closed,

and is called the *subgraph closure* of \mathcal{C} . Similarly the *hereditary closure* of a class \mathcal{C} consists of all the induced subgraphs of members of \mathcal{C} , and is hereditary by design.

An *edge contraction* of two adjacent vertices u, v consists of merging u and v into a single vertex adjacent to $N(\{u, v\})$ (and deleting u and v). A graph H is a *minor* of a graph G if H can be obtained from G by a sequence of vertex and edge deletions, and edge contractions. Equivalently a minor H with vertex set say, $\{v_1, \dots, v_{|V(H)|}\}$, of G can be defined as a vertex partition $B_1, \dots, B_{|V(H)|}$ of a subgraph of G , such that every $G[B_i]$ is connected and $E_G(B_i, B_j) \neq \emptyset$ whenever $v_i v_j \in E(H)$. Indeed after contracting each B_i into a single vertex (which is possible since they induce connected subgraphs), H appears as a subgraph. The set B_i is called the *branch set* of $v_i \in V(H)$. A graph G is said *H-minor free* if H is not a minor of G . A class is said *minor-closed* if every minor of a member of the class is in the class, and *proper minor-closed* if further the class is *not* the set of all graphs.

The *radius* $\text{rad}(G)$ of a graph G is defined as $\min_{u \in V(G)} \max_{v \in V(G)} d_G(u, v)$. The *radius* $\text{rad}_G(S)$ of a subset of vertices $S \subseteq V(G)$ is simply defined as $\text{rad}(G[S])$. Note that two vertices can be further away in $G[S]$ than in G . An *r-shallow minor* H of G is a minor of G with branch sets $B_1, \dots, B_{|V(H)|}$ satisfying $\text{rad}_G(B_i) \leq r$ for every $i \in [|V(H)|]$. We denote that by $H \preceq_r G$. In particular 0-shallow minors correspond to subgraphs. The theory of graph sparsity pioneered by Ossona de Mendez and Nešetřil [23] introduces the following invariants for a graph G and a class \mathcal{C} :

$$\nabla_r(G) := \sup_{H \preceq_r G} \frac{|E(H)|}{|V(H)|}, \text{ and } \nabla_r(\mathcal{C}) := \sup_{G \in \mathcal{C}} \nabla_r(G).$$

Note that $\nabla_0(G)$ is tied to the maximum average degree of G .

A class \mathcal{C} of graphs is said to have *bounded expansion* if $\nabla_r(\mathcal{C}) < \infty$ for every $r \in \mathbb{N}$. More generally \mathcal{C} has *expansion* f if $\nabla_r(\mathcal{C}) \leq f(r)$ for every $r \in \mathbb{N}$. A class has *polynomial expansion* if it has expansion f for a polynomial function f . Proper minor-closed classes even have constant expansion, i.e., expansion f for a constant function f .

2.2 Summary of the previous paper. In the previous paper of the series [7], we introduced a new graph and matrix invariant dubbed *twin-width*, inspired by the work of Guillemot and Marx on permutations [17]. We proved that many classes such as, bounded rank-width graphs, proper minor-free classes, proper subclasses of permutation graphs, and posets with antichains of bounded size have bounded twin-width. For all these classes, we showed how to find in polynomial-

¹If it is sometimes more convenient to use a different vertex set for the class definition, this will implicitly come with a canonical mapping from this vertex set to $[n]$.

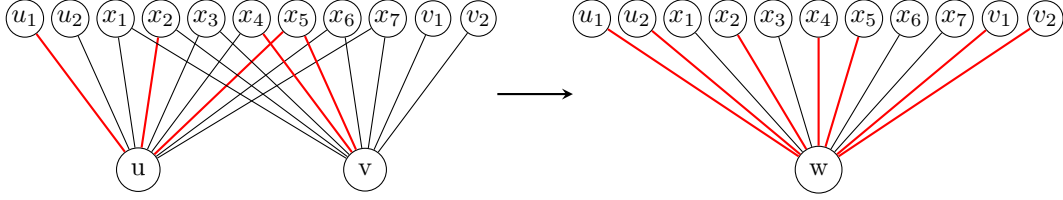


Figure 1: Contraction of vertices u and v , and how the edges of the trigraph are updated.

time a so-called d -sequence, witnessing that the twin-width is at most a constant d . Finally given a d -sequence of a binary structure G on n elements and a first-order (FO) formula φ of quantifier-depth ℓ , we provided an FO model checking algorithm deciding $G \models \varphi$ in time $f(d, \ell)n$.

We start by recalling the definition of twin-width, and then we summarize the milestones of [7] that will also be useful in the current paper.

2.2.1 Trigraphs, contraction sequences, and twin-width of a graph. A *trigraph* G has vertex set $V(G)$, (black) edge set $E(G)$, and red edge set $R(G)$ (the error edges), with $E(G)$ and $R(G)$ being disjoint. The *set of neighbors* $N_G(v)$ of a vertex v in a trigraph G consists of all the vertices adjacent to v by a black or red edge. A d -trigraph is a trigraph G such that the *red graph* $(V(G), R(G))$ has degree at most d . In that case, we also say that the trigraph has *red degree* at most d . In the context of trigraphs and twin-width, we will somewhat overload the term “contraction”. A *contraction* or *identification* in a trigraph G consists of merging two (non-necessarily adjacent) vertices u and v into a single vertex w , and updating the edges of G in the following way. Every vertex of the symmetric difference $N_G(u) \Delta N_G(v)$ is linked to w by a red edge. Every vertex x of the intersection $N_G(u) \cap N_G(v)$ is linked to w by a black edge if both $ux \in E(G)$ and $vx \in E(G)$, and by a red edge otherwise. The rest of the edges (not incident to u or v) remain unchanged. We insist that the vertices u and v (together with the edges incident to these vertices) are removed from the trigraph. See Fig. 1 for an illustration.

A *sequence of d -contractions* or d -sequence is a sequence of d -trigraphs G_n, G_{n-1}, \dots, G_1 , where $G_n = G$, $G_1 = K_1$ is the graph on a single vertex, and G_{i-1} is obtained from G_i by performing a single contraction of two (non-necessarily adjacent) vertices. We observe that G_i has precisely i vertices, for every $i \in [n]$. The twin-width of G , denoted by $\text{tw}(G)$, is the minimum integer d such that G admits a d -sequence. Going back to the overload of the word “contraction”, in case we actually refer to the classical (edge) contraction, either

we will use the term “edge contraction”, or it will be clear from the context what is meant.

2.2.2 Partitions, divisions, red number, and twin-width of a matrix. We now give two equivalent definitions for the twin-width of a matrix. The first is based on a *contraction sequence* where we progressively reduce the size of the matrix, and introduce error symbols r . The second (equivalent) definition is based on a *coarsening sequence* where we progressively coarsen a partition of the rows and columns of the matrix.

The red number of a matrix is the maximum number of r entries (error entry, the r stands for red) in a single row or column. Given an $n \times m$ matrix M and two columns C_i and C_j (resp. two rows R_i and R_j), the *contraction* of C_i and C_j (resp. R_i and R_j) is obtained by deleting C_j (resp. R_j) and replacing every entry $m_{k,i}$ of C_i (resp. every entry $m_{i,k}$ of R_i) by r whenever $m_{k,i} \neq m_{k,j}$ (resp. $m_{i,k} \neq m_{i,j}$). A d -contraction sequence of matrix M is sequence of successive contractions starting at M , ending at some 1×1 matrix, such that all matrices of the sequence have red number at most d . The *twin-width* of a matrix M is the smallest integer d such that M admits a d -contraction sequence.

We observe that when M has twin-width at most d , one can reorder its rows and columns such that every contraction is on two consecutive rows or two consecutive columns. The reordered matrix is then called *d -twin-ordered*. The *symmetric twin-width* of an $n \times n$ matrix M is defined similarly, except that the contraction of rows i and j (resp. columns i and j) is immediately followed by the contraction of columns i and j (resp. rows i and j). The symmetric twin-width of the adjacency matrix of a graph G corresponds to the twin-width of G .

For the second definition of the twin-width of a matrix, we need to introduce a bit of vocabulary on partitions. We say that a partition \mathcal{P} of a set S *refines* a partition \mathcal{P}' of S if every part of \mathcal{P} is contained in a part of \mathcal{P}' . Conversely we say that \mathcal{P}' is a *coarsening* of \mathcal{P} . We will further assume that a coarsening is proper, that is, \mathcal{P}' and \mathcal{P} are distinct. Given a partition \mathcal{P} and two distinct parts P, P' of \mathcal{P} , the *elementary coarsening*

of P and P' yields the coarsening $\mathcal{P} \setminus \{P, P'\} \cup \{P \cup P'\}$. Informally an elementary coarsening is the merge of two parts.

Given an $n \times m$ matrix M , we call *row-partition* (resp. *column-partition*) a partition of the rows (resp. columns) of M . A (k, ℓ) -*partition*, or simply *partition*, of a matrix M is a pair $(\mathcal{R} = \{R_1, \dots, R_k\}, \mathcal{C} = \{C_1, \dots, C_\ell\})$ where \mathcal{R} is a row-partition and \mathcal{C} is a column-partition. In a matrix partition $(\mathcal{R}, \mathcal{C})$, each part $R \in \mathcal{R}$ is called a *row-part*, and each part $C \in \mathcal{C}$ is called a *column-part*. An *elementary coarsening* of a partition $(\mathcal{R}, \mathcal{C})$ of a matrix M is obtained by performing one elementary coarsening in \mathcal{R} or in \mathcal{C} . We distinguish two canonical partitions of an $n \times m$ matrix M : the *finest partition* where $(\mathcal{R}, \mathcal{C})$ have size n and m , respectively, and the *coarsest partition* where $|\mathcal{R}| = |\mathcal{C}| = 1$.

A *coarsening sequence* of an $n \times m$ matrix M is a sequence of partitions $(\mathcal{R}^1, \mathcal{C}^1), \dots, (\mathcal{R}^{n+m-1}, \mathcal{C}^{n+m-1})$ where

- $(\mathcal{R}^1, \mathcal{C}^1)$ is the finest partition,
- $(\mathcal{R}^{n+m-1}, \mathcal{C}^{n+m-1})$ is the coarsest partition, and
- for every $i \in [n + m - 2]$, $(\mathcal{R}^{i+1}, \mathcal{C}^{i+1})$ is an elementary coarsening of $(\mathcal{R}^i, \mathcal{C}^i)$.

Given a subset R of rows and a subset C of columns in a matrix M , the *zone* $R \cap C$ denotes the submatrix of all entries of M at the intersection between a row of R and a column of C . A *zone* of a matrix partitioned by $(\mathcal{R}, \mathcal{C}) = (\{R_1, \dots, R_k\}, \{C_1, \dots, C_\ell\})$ is any $R_i \cap C_j$ for $i \in [k]$ and $j \in [\ell]$. A zone is *constant* if all its entries are identical. The *error value* of a row-part R_i (resp. a column-part C_j) is the number of non-constant zones among all zones in $\{R_i \cap C_1, \dots, R_i \cap C_\ell\}$ (resp. $\{R_1 \cap C_j, \dots, R_k \cap C_j\}$). The *error value* of $(\mathcal{R}, \mathcal{C})$ is the maximum error value of a part, taken over all parts R_i and C_j . Now the twin-width of a matrix M can be equivalently defined as the minimum d for which M admits a coarsening sequence in which all partitions have error value at most d .

We will work with particular partitions, called *divisions*, where every part consists of a set of consecutive rows, or a set of consecutive columns. If the matrix is d -twin-ordered, there is a coarsening sequence with error value at most d , in which all the partitions are divisions. We call *division sequence* such a coarsening sequence.

2.2.3 Grid minor theorem for twin-width.

A (t, t) -*division* is a division $(\mathcal{R}, \mathcal{C})$ such that $|\mathcal{R}| = |\mathcal{C}| = t$. A t -*grid minor* is a (t, t) -division whose t^2 zones contains a non-zero entry. As for the PERMUTATION PATTERN breakthrough algorithm of Guillemot and Marx [17], a crucial engine of twin-width is the following celebrated theorem by Marcus and Tardos.

THEOREM 2.1. ([20]) *For every integer t , there is some c_t such that every $n \times m$ 0,1-matrix M with at least $c_t \max(n, m)$ entries 1 has a t -grid minor.*

Informally, if a matrix has sufficiently many entries 1, then there is a large grid structure where each cell is “complicated”. The current best bound for c_t , due to Cibulka and Kynčl [8], is $8/3(t+1)^2 2^{4t}$.

To leverage Marcus-Tardos theorem in the dense regime, too, we modify the definition of “complicated” from “containing a 1” to “being mixed”. A zone is *horizontal* if all its columns are equal (restricted to the zone), and *vertical* if all the rows are equal. Equivalently each row (resp. column) within a horizontal zone (resp. vertical zone) consists of a repeated same entry. Note that a zone is constant (consists of a same entry repeated) if it is horizontal and vertical. A zone is *mixed* if it is not horizontal nor vertical.

We can now introduce the notions of t -mixed minors and t -mixed freeness. A t -*mixed minor* of a matrix M is a (t, t) -division of M such that every zone is mixed. A matrix is t -*mixed free* if it does not admit a t -mixed minor. We showed that having small twin-width and admitting no large mixed minors are equivalent in the following sense.

THEOREM 2.2. ([7]) *Let α be the alphabet size for the matrix entries, and $c_t := 8/3(t+1)^2 2^{4t}$.*

- *Every t -twin-ordered matrix is $2t + 2$ -mixed free.*
- *Every t -mixed free matrix has twin-width at most $4c_t \alpha^{4c_t+2} = 2^{2^{O(t)}}$.*

The first item is a relatively simple observation. The difficulty lies in the second item. In a nutshell, if the matrix is t -mixed free, we find, using Marcus-Tardos theorem, a sequence of divisions with small number of mixed zones per column and per row. From this favorable sequence of divisions, we are able to extract an $f(t)$ -contraction sequence.

One simple but important ingredient is a local characterization of mixedness by means of a *corner*. A *corner* in a matrix $M = (m_{i,j})_{i,j}$ is a mixed zone made by four contiguous entries $m_{i,j}, m_{i+1,j}, m_{i,j+1}, m_{i+1,j+1}$. A 0,1-corner is a corner where each entry is in $\{0, 1\}$.

LEMMA 2.1. ([7]) *A matrix is mixed if and only if it contains a corner.*

In Section 3 we will work with specifically divided 0, 1, r -matrices, respecting the following invariants. Every zone is filled with r entries, or is *non-mixed* (that is, horizontal or vertical) and has only 0 and 1 entries. In this context, we will redefine the mixed zones as those filled with r entries. The coarsenings will be followed by updating the entries of the matrix to keep the invariants.

Namely every zone with a 0, 1-corner is filled with r entries. This new viewpoint mixes contraction sequence and coarsening sequence. It will turn out useful to find, in a t -mixed free matrix, not just one “good contraction” (as in Theorem 2.2) but a linear number of disjoint pairs of “good contractions”. This will have two main consequences. It will enable us to show that bounded twin-width classes are *small* (see Section 2.3 for a formal definition). This will also be used to find $O(\log n)$ -bits adjacency labeling schemes (see Section 2.4) for n -vertex graphs in classes of bounded twin-width.

2.2.4 Closure by FO transduction. Bounded twin-width behaves surprisingly well with respect to first-order logic. In addition to the fixed-parameter tractable algorithm running in time $f(d, |\phi|)n$ for model checking a first-order sentence ϕ on an n -vertex graph given with a d -contraction sequence, we show that bounded twin-width is preserved by first-order (FO) transductions.

THEOREM 2.3. ([7]) *Every transduction of a bounded twin-width class has bounded twin-width.*

A formal definition of FO transductions can be found in several papers (see for instance [4, 15]). As this definition is somewhat technical and we will only use Theorem 2.3 in a black-box fashion, we refer the interested reader to these paper. Informally an FO transduction of a graph G defines several new graphs. It consists of² a non-deterministic “coloring” of $V(G)$ by a constant number of unary relations, followed by a redefinition of the edges by means of a fixed FO formula using the former edge predicate as well as these new unary relations. The unary relations are then discarded, and we here further allow to take any induced subgraph of the obtained graph (to preserve the class heredity). An FO transduction of a class \mathcal{C} is simply the union of the graphs obtained by FO transduction of G , for every $G \in \mathcal{C}$.

2.3 Small classes and the small conjecture. We recall that a hereditary class is a class closed under taking induced subgraphs. Formally if G is in a hereditary class \mathcal{C} , then for every induced subgraph H of G , it also holds that H is in \mathcal{C} . The overwhelming majority of the usually considered classes of graphs are hereditary.³

²The original definition allows first to duplicate the graph G into constantly many copies of G . As any disjoint union of bounded twin-width graphs has bounded twin-width, we omit this step here.

³Notable exceptions include regular graphs, connected graphs, and visibility graphs of a point set.

A class of graph \mathcal{C} is said *small* (resp. *factorial*), if there exists a constant c , such that the number of n -vertex graphs of \mathcal{C} is at most $n!c^n$ (resp. $n!^c = 2^{O(n \log n)}$), for every $n \in \mathbb{N}$. Recall that our n -vertex graphs are all assumed to be on the vertex set $[n]$, and that we count up to equality and *not* up to isomorphism. Norine et al. [24] show that the number of K_t -minor free graphs on $[n]$ is at most $n!c^n$, for some integer c depending only on t . In other words, proper minor-closed classes are small. Marcus-Tardos theorem [20], combined with an argument due to Klazar [19], implies that the number of $n \times n$ 0, 1-matrices avoiding a fixed permutation submatrix is at most c^n , for some constant c . In particular the number of permutations on n elements avoiding a fixed permutation grows in $2^{O(n)}$. A translation of this result to graphs is that proper subclasses of permutation graphs are small.

We say that a class \mathcal{C} has bounded twin-width if there exists an integer $d_{\mathcal{C}}$ such that every member of \mathcal{C} has twin-width at most $d_{\mathcal{C}}$. Thus $\text{tw}(\mathcal{C}) := \sup_{G \in \mathcal{C}} \text{tw}(G) < \infty$.

One of the main contributions of the paper is the following.

THEOREM 2.4. *Every class with bounded twin-width is small.*

This generalizes the smallness of proper minor-closed classes [24], proper subclasses of permutation graphs [20, 19], and graphs with bounded treewidth [2], as we previously showed that all these classes have bounded twin-width [7]. We then explore a possible converse for Theorem 2.4. Of course it is easy to artificially build an unbounded twin-width class with only $n!$ graphs of size n . For example, by taking in the class a single (up to isomorphism) n -vertex graph among the n -vertex graphs with maximum twin-width, for every n . However this is not a satisfactory counterexample. In combinatorics, classes of objects are often required to be closed under substructures. For instance, a *class of permutations* is by definition closed under taking sub-permutations. Same goes for graphs: Hereditary classes have richer properties than non-hereditary ones. Many interesting questions on hereditary classes have trivial answers or are not even well-defined on general classes.

We provocatively conjecture the following converse of Theorem 2.4.

CONJECTURE 1. (SMALL CONJECTURE) *Every small hereditary class has bounded twin-width.*

It may seem ambitious to expect that the converse of Theorem 2.4 holds for hereditary classes. Why would the mere limited number of graphs guarantee anything close to a d -contraction sequence? A typical example of

a class with unbounded twin-width contains an infinite sequence of graphs G_1, G_2, \dots where every distinct pair $u, v \in V(G_i)$ satisfies $|N_{G_i}(u) \Delta N_{G_i}(v)| \geq i$. Indeed any first contraction in G_i creates a vertex with red degree at least i . A class is said to have *unbounded symmetric difference* if it contains such a sequence, and *bounded symmetric difference*, otherwise. So for every class \mathcal{C} with bounded symmetric difference, there is an integer d such that for every graph $G \in \mathcal{C}$, there exist two distinct vertices $u, v \in V(G)$ satisfying $|N(u) \Delta N(v)| \leq d$. For example, the $i \times i$ rook graphs (with vertex set $[i] \times [i]$ and an edge between (a, b) and (c, d) if $a = c$ or $b = d$), with $i \geq 3$, is a class with unbounded symmetric difference. However the hereditary closure of this class is not small.

Having bounded symmetric difference is a prerequisite to having bounded twin-width. A first step towards Conjecture 1 would be to show that small hereditary classes have bounded symmetric difference. Even that is unclear. For $K_{2,2}$ -free classes or classes with girth at least 5, bounded symmetric difference simply implies bounded minimum degree. Thus a very particular case of Conjecture 1 is that there every small $K_{2,2}$ -free hereditary class has bounded minimum degree.

Let us present some elements supporting the conjecture. First and foremost, bounded twin-width seems to “stop at the right place” in the sparse and dense realms. Unit interval graphs (a small class) have bounded twin-width while interval graphs (a non-small class) do not. Similarly among sparse classes, proper minor-closed classes (small) have bounded twin-width, whereas subcubic graphs (non-small) have unbounded twin-width. We will also see that some expander classes have bounded twin-width (and are small), unlike random cubic graphs.

An interesting test is the case of the s -subdivisions (where each edge of a graph is subdivided $s \geq 1$ times). Since the number of subcubic graphs on $[n]$ is $n^{3n/2 + O(n/\log n)}$, the $o(\log n)$ -subdivisions of subcubic graphs is still a non-small class. Thus by Theorem 2.4, they have unbounded twin-width. We show a more fine-grained version of that fact by a direct proof. We also build in polynomial time $O(1)$ -sequences for $\Omega(\log n)$ -subdivisions of K_n , which yields the following.

THEOREM 2.5. *The s -subdivision of K_n has bounded twin-width if and only if $s = \Omega(\log n)$. More precisely, for every integer d , there are $\ell_d < u_d$ such that the $\lfloor c \log n \rfloor$ -subdivision of K_n has twin-width at least d for every $1 \leq c \leq \ell_d$, and at most d for every $c \geq u_d$.*

The hereditary closure of $\Omega(\log n)$ -subdivisions of K_n is indeed a small class. But Theorem 2.5 in particular implies that this class does have bounded twin-width. Dvořák and Norine [14] show that, for any

constants $c, \varepsilon > 0$, classes with expansion $r \mapsto c^{r^{1/3-\varepsilon}}$ are small, while the class of all graphs with expansion $r \mapsto 6 \cdot 3^{\sqrt{r \log(r+\varepsilon)}}$ is not small. If the small conjecture is true, then bounded twin-width contains polynomial expansion (actually even expansion $r \mapsto 2^{r^{0.33}}$). Thus another possible first step to Conjecture 1 is to show that classes with polynomial expansion have bounded twin-width.

A supplementary motivation for the small conjecture appears if its proof is algorithmic, that is, yields on any small hereditary class a polytime algorithm which takes any graph of the class and outputs a (non-necessarily optimal) $O(1)$ -sequence. In light of Theorem 2.4 and considering that $\omega(1)$ -sequences are not as algorithmically useful, that would be almost as good as a constant approximation of twin-width in general graphs.

2.4 Implicit representations. A class \mathcal{C} has an $f(n)$ -bits adjacency labeling scheme (or simply *labeling scheme*, for short) if there is a decoding function $A : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}$ such that for every n -vertex graph $G \in \mathcal{C}$ there is a labeling function $\ell : V(G) \rightarrow \{0, 1\}^*$, satisfying $|\ell(u)| \leq f(n)$ for every $u \in V(G)$, and $A(\ell(u), \ell(v)) = 1$ if and only if $uv \in E(G)$. Here we will further impose that the labeling function ℓ is injective. For example trees now have $\log n + O(1)$ -bits adjacency labeling scheme [1], which up to the constant term, is optimal. It is known that a class \mathcal{C} has a $c \log n$ -bits adjacency labeling scheme if and only if, for every integer n , there is a *universal graph* graph U_n (not necessarily in \mathcal{C}) on at most n^c vertices such that every n -vertex graph of \mathcal{C} is an induced subgraph of U_n (see for instance [26]). This becomes apparent when one considers the possible labels as the vertex set of the universal graph.

Several classes, such as interval graphs and K_t -minor free graphs, are known to have $O(\log n)$ -bits labeling schemes. By a direct counting argument, only factorial classes can expect to admit $O(\log n)$ -bits labeling scheme. Indeed the number of distinct labels is $2^{O(\log n)} = n^{O(1)}$. Thus the number of n -vertex graphs that can be induced subgraphs of the universal graph is only $\binom{n^{O(1)}}{n} = n^{O(n)}$. The *implicit graph conjecture* asserts that every factorial hereditary class has an $O(\log n)$ -bits labeling scheme [18]. We show the conjecture in the particular case of bounded twin-width classes.

THEOREM 2.6. *Every bounded twin-width class admits an $O(\log n)$ -bits labeling scheme.*

This result is at the same time quite strong and quite weak. Its strength lies in its broad generality.

We produce a unified labeling scheme for very different sparse and dense classes. However there are two caveats, both linked to its generality. The first one is that we still do not know if the labeling function can be computed in polynomial time. Indeed it requires a d -sequence (even a so-called *parallel D -sequence* of logarithmic length). If we know how to compute this sequence in many bounded twin-width classes, we do not know in the full generality of all the graphs with twin-width at most d . In the latter case, we currently need exponential time to find the sequence, and then to compute the labeling. The decoding function, that is the adjacency test, runs in time $O(\log n)$ in the RAM model with unit-cost arithmetic operations over words of logarithmic length. The second caveat is that when restricted to particular classes, the multiplicative constant preceding $\log n$ given by our proof is much larger than in the shortest known labeling schemes. For instance, the current best labeling scheme for K_t -minor free graphs requires $2 \log n + o(\log n)$ bits per vertex [16], while our multiplicative constant is double-exponential in t .

Improving the constant c of existing $(c + o(1)) \log n$ -bits labeling schemes is topical in implicit representations. Recently planar graphs were shown to admit a $(1 + o(1)) \log n$ -bits adjacency labeling scheme [9]. It is optimal up to the second-order term. The labeling scheme is actually more general, and works for all subgraphs of strong products $H \boxtimes P$ where H is a bounded-treewidth chordal graph (or k -tree, for some fixed k), and P is a path. A class \mathcal{C} is said *flat* if there is an integer k such that $\mathcal{C} \subseteq \text{Sub}(\mathcal{H} \boxtimes \mathcal{P})$ where \mathcal{P} is the set of all paths, and \mathcal{H} is a set of graphs with treewidth at most k . An ongoing program (not specific to adjacency labeling schemes), dubbed *graph product structure theorem*, established that many small and sparse classes are flat. This was initiated by Dujmović et al. [11], building up on a paper by Pilipczuk and Siebertz [25] which shows a similar result for planar graphs. This property was extended to apex-minor free [11], bounded-degree minor-free [10], and k -planar classes [12]. Hence they all enjoy a $(1 + o(1)) \log n$ -bits adjacency labeling scheme. Interestingly all these classes have bounded twin-width (minor-free classes and k -planar graphs have bounded twin-width [7]). This is no coincidence. We will see that the strong product of two bounded twin-width graphs, one of which has bounded degree, has bounded twin-width.

THEOREM 2.7. *Let G and H be two graphs. Then $\text{tw}(G \boxtimes H) \leq \max\{\text{tw}(G)(\Delta(H) + 1) + 2\Delta(H), \text{tw}(H) + \Delta(H)\}$.*

As cliques have twin-width 0, taking subgraphs does in general not preserve twin-width at all. Nevertheless

on “sparse” classes, bounded twin-width is subgraph-closed. We show that if the strong product of a bounded twin-width class \mathcal{G} with a bounded-degree bounded twin-width class \mathcal{H} is $K_{t,t}$ -free, then the subgraphs of $\mathcal{G} \boxtimes \mathcal{H}$ have bounded twin-width.

THEOREM 2.8. *(\star) Let \mathcal{G} and \mathcal{H} two classes such that $\mathcal{G} \boxtimes \mathcal{H}$ is $K_{t,t}$ -free. Then $\text{tw}(\text{Sub}(\mathcal{G} \boxtimes \mathcal{H})) \leq f(\text{tw}(\mathcal{G}), \text{tw}(\mathcal{H}), \Delta(\mathcal{H}), t)$.*

In particular flat classes have bounded twin-width (since graphs with bounded treewidth have bounded twin-width, and flat classes are $K_{t,t}$ -free). By essence, the “flat class” approach to $(1 + o(1)) \log n$ -bits labeling scheme is limited to classes that are K_t -free. Another interesting limit case is minor-free classes which are not apex-minor free, like all the K_6 -minor free graphs for example. Dujmović et al. [11] show that these classes are not flat.

We hope that the *versatile tree of contractions* (see Lemma 3.7) or the *short parallel contraction sequence* (see Lemma 4.2) may help for small dense classes and K_t -minor free graphs. We optimistically conjecture that our Theorem 2.6 can be improved to an optimal labeling scheme up to the second-order term.

CONJECTURE 2. *Every bounded twin-width class has a $(1 + o(1)) \log n$ -bits labeling scheme.*

2.5 Sparse twin-width. The trace of bounded twin-width on sparse classes is also an interesting and potentially new class. There are five natural ways of forcing a bounded twin-width class to be “sparse”: forbidding $K_{t,t}$ as a subgraph, forbidding a d -grid minor in its adjacency matrix (and not a mere d -mixed minor), requiring that every graph has bounded average degree, requiring that the subgraphs also have bounded twin-width, and requiring that the class has bounded expansion. Let $A_\sigma(G)$ denote the adjacency matrix of G when $V(G)$ is ordered by σ . We say that a class \mathcal{C} is *d -grid free* if for every $G \in \mathcal{C}$ there is an ordering σ of $V(G)$ such that $A_\sigma(G)$ is d -grid free.

We show that all five definitions are actually equivalent.

THEOREM 2.9. *(\star) If \mathcal{C} is a hereditary class of bounded twin-width, the following are equivalent.*

- (i) *There is an integer t such that no graph of \mathcal{C} contains $K_{t,t}$ as a subgraph.*
- (ii) *There is an integer d such that \mathcal{C} is d -grid free.*
- (iii) *There is an integer g such that every n -vertex graph $G \in \mathcal{C}$ has at most gn edges.*
- (iv) *The subgraph closure $\text{Sub}(\mathcal{C})$ has bounded twin-width.*
- (v) *There is a function f such that $\nabla_r(\mathcal{C}) \leq f(r)$ for every r .*

Ignoring item (iv), a compact version of this theorem reads: For a hereditary class of bounded twin-width having bounded grid minors, bicliques, average degree, or expansion are all equivalent.

Thus we say that a hereditary class has *bounded sparse twin-width* if it has bounded twin-width and satisfies any of the five items (that is, satisfies all five). One may wonder whether bounded sparse twin-width coincides with some existing sparse class. More generally it is interesting to see how bounded sparse twin-width compares to the established sparse classes. A few candidates come to mind: polynomial expansion, bounded expansion, bounded queue number, bounded stack number, bounded nonrepetitive coloring classes. Although we do not prove it for bounded queue or stack number, we argue that these classes do not coincide with bounded sparse twin-width.

As cubic graphs have unbounded twin-width, bounded expansion is strictly more general than bounded sparse twin-width. For the same reason, bounded nonrepetitive coloring does not imply bounded sparse twin-width. It is possible however that bounded sparse twin-width classes have bounded nonrepetitive coloring. The existence of an infinite family of cubic expanders with bounded twin-width implies that bounded sparse twin-width classes do not necessarily have polynomial expansion. If the small conjecture is true, polynomial expansion would be a strict subset of bounded sparse twin-width. We will show that classes with bounded queue number or bounded stack number have bounded (sparse) twin-width. We believe that this inclusion is strict and that the expanders based on random 2-lifts have unbounded queue and stack numbers.

2.6 Organization of the rest of paper. In Section 3 we show Theorem 2.4, that every class of bounded twin-width is small. From this we conclude that non-small classes such as subcubic graphs, interval graphs, and triangle-free unit segment graphs have unbounded twin-width. This can be respectively put in perspective with the fact that some cubic expanders (as we see in Section 5), unit interval graphs, and K_t -free unit d -dimensional ball graphs, have bounded twin-width [7]. In Section 4 we leverage the results from the previous section to present $O(\log n)$ -bits adjacency labeling schemes on bounded twin-width classes. We then explore the converse of Theorem 2.4 for hereditary classes. In Section 5 we show that the small class of cubic expanders obtained by iterated 2-lifts from K_4 has indeed bounded twin-width. In Section 6 we prove that the s -subdivision of the clique K_n , with $s > 0$, has bounded twin-width if and only if $s = \Omega(\log n)$. In Section 7 we prove Theorem 2.9, the list of characterizations of

bounded sparse twin-width. We then show that flat classes, and classes with bounded queue or stack number have bounded (sparse) twin-width. In Section 8 we investigate the twin-width of the finite induced subgraphs of a fixed Cayley graph. We show that such classes are small for every finitely generated group. This is a rare example of a small class for which we still do not know if the twin-width is bounded.

The proofs of statements marked with a \star were removed to fulfill the space constraints and can be found in the full version of the paper [5].

3 Bounded twin-width classes are small

In this section we show that graphs of bounded twin-width have bounded *versatile twin-width*. Informally it says that whenever we can find a sequence (or path) of d -contractions, we can even find a tree of D -contractions with linear arity, for some D bounded by a function of d . This result is fairly technical but shares some ideas and arguments with Section 5 of our previous paper [7]. We made the current section self-contained. We nevertheless mention some frequent parallels with [7]. Finally we can follow the end of the proof of Norine et al. [24] –that proper minor-closed classes are small– to extend the result to bounded twin-width.

3.1 The proof for proper minor-closed classes and how (not) to tune it.

Let us first give a brief sketch of Norine et al.’s proof, which works by induction on n . They say that a vertex is d -good if it has degree at most d and either has a twin or has a neighbor with degree at most d . They show the following technical lemma: K_t -minor free n -vertex graphs have at least n/d d -good vertices, for some d function of t only. Let $\mathcal{I}_{n,t}$ be the set of K_t -minor free graphs on $[n]$, and $\mathcal{K}_{n,t}$ be the subset of all those graphs of $\mathcal{I}_{n,t}$ where vertex n is d -good. By their lemma $n/d \cdot |\mathcal{I}_{n,t}| \leq n|\mathcal{K}_{n,t}|$, hence $|\mathcal{I}_{n,t}| \leq d|\mathcal{K}_{n,t}|$. Furthermore, any graph of $\mathcal{K}_{n,t}$ admits an index $i \in [n-1]$ such that either i and n are false twins, or i and n are adjacent and have at most $d-1$ other neighbors each. Therefore any $G \in \mathcal{K}_{n,t}$ can be obtained from a $G' \in \mathcal{I}_{n-1,t}$ and $i \in [n-1]$ by either introducing a new vertex n false twin of i (one graph), or by splitting i into i and a new vertex n adjacent to i , and by distributing in G the at most $2(d-1)$ neighbors of i in G' into: neighbors of i only, neighbors of n only, and common neighbors (at most $3^{2(d-1)}$ graphs). Hence $|\mathcal{I}_{n,t}| \leq d(1 + 3^{2(d-1)})(n-1)|\mathcal{I}_{n-1,t}| \leq (1 + 3^{2(d-1)})(n-1) \cdot (n-1)!c^{n-1} \leq n!c^n$, by taking $c := d(1 + 3^{2(d-1)})$.

We need to redefine the notion of being d -good for bounded twin-width classes. A very natural candidate

for that would be to say that a vertex is d -good if it admits a d -contraction with another vertex. After all, there is always such a vertex (or such a pair of vertices) in a d -trigraph. However, we cannot expect d -trigraphs to have linearly many such vertices. Think for instance of a path on n vertices. It has twin-width 1, but only four vertices (the two endpoints and their neighbor) that can be contracted to yield a 1-sequence. Surely we could allow mere D -contractions, for some $D \gg d$, but then we would leave the class of d -trigraphs. So it would be unclear which class we are bounding the size of. It is indeed noteworthy in the above sketch that by deleting a vertex or contracting adjacent vertices, one remains in the class of K_t -minor free graphs.

To overcome that issue, we introduce a more robust notion of bounded twin-width. A *tree of d -contractions* of a d -trigraph G is a rooted tree, whose root is labeled by G , and whose leaves are all labeled by 1-vertex graphs K_1 , and such that one can go from any parent to any child by a d -contraction. With this new definition, d -sequences coincide with trees of d -contractions which are in fact paths. We say that a trigraph G has *versatile twin-width d* if there exists some p , function of d only, such that G admits a tree of d -contractions in which every internal node has at least $|V(\cdot)|/p$ children with distinct labels (where $|V(\cdot)|$ denotes the number of vertices of the corresponding node label). Such a tree is then called a *versatile tree of d -contractions*.

Let us say that a contraction is d -correct (or simply *correct* when we precise that it is a d -contraction) if the obtained graph has twin-width at most d . The inductive nature of versatile twin-width provides us the desired stability. Not only G admits linearly many correct d -contractions, but it admits linearly many d -contractions towards graphs of versatile twin-width d . This is indeed witnessed by the subtrees rooted at each child of the root labeled by G . We now focus on proving that every trigraph with twin-width d has a versatile tree of D -contractions, for a larger D function of d only. This is a bit technical, but once it is done, we will be able to mimic the end of Norine et al.'s proof.

3.2 Neatly divided symmetric $0,1,r$ -matrices. Recall that r (for red) is the error symbol. It will now be convenient to tune some of the notions developed in our previous paper specifically for $0,1,r$ -matrices with particular divisions. The notions introduced without a definition are all formalized in Section 2 of the present paper, as well as in [7, Section 5]. Reading first [7, Section 5] does not harm, but it is not necessary to understand the current section.

We will manipulate divisions of $0,1,r$ -matrices such that every zone either contains only r entries or contains

no r entry and is horizontal or vertical (or both). Let us call *neat* such a division. Zones filled with r entries are now called *mixed*. A *neatly divided matrix* is a pair $(M, (\mathcal{R}, \mathcal{C}))$ where M is a $0,1,r$ -matrix and $(\mathcal{R}, \mathcal{C})$ is a neat division of M . A *t -mixed minor* in a neatly divided matrix is a (t,t) -division which coarsens the neat subdivision, and contains in each of its t^2 zones at least one mixed zone (filled with r entries) or a $0,1$ -corner. See Fig. 2 for an illustration. A neatly divided matrix is said *t -mixed free* if it does not admit a t -mixed minor.

A *mixed cut of a row-part $R \in \mathcal{R}$* of a neat division $(\mathcal{R}, \mathcal{C} = \{C_1, C_2, \dots\})$ is an index i such that both $R \cap C_i$ and $R \cap C_{i+1}$ are *non-mixed*, and there is a $0,1$ -corner in the 2 -by- $|R|$ zone defined by the last column of C_i , the first column of C_{i+1} , and R . Importantly, a mixed cut cannot border a mixed zone. (This is a difference with the definition of [7, Section 5].) The *mixed value of a row-part $R \in \mathcal{R}$* of a neat division $(\mathcal{R}, \mathcal{C} = \{C_1, C_2, \dots\})$ is the number of mixed zones $R \cap C_j$ plus the number of mixed cuts between two (adjacent non-mixed) zones $R \cap C_j$ and $R \cap C_{j+1}$. Note that a mixed cut counts for one unit in the mixed value, regardless of the number of corners overlapping the two adjacent zones. We similarly define the mixed value of a column-part $C \in \mathcal{C}$. The *mixed value of a neat division* of a $0,1,r$ -matrix is the maximum of the mixed values taken over every part. The *part size* of a division (resp. partition) $(\mathcal{R}, \mathcal{C})$ is defined as $\max(\max_{R \in \mathcal{R}} |R|, \max_{C \in \mathcal{C}} |C|)$. A division is *symmetric* if the largest row index of each row-part and the largest column index of each column-part define the same set of integers, that is informally, if the horizontal separations are symmetric of the vertical separations about the main diagonal. For instance the division depicted on Fig. 2 is symmetric since both the largest row indices of the row-parts and the largest column indices of the column-parts define the set $\{2, 3, 4, 6\}$. We call *symmetric fusion* of a symmetric division the fusion of two consecutive parts in \mathcal{C} and of the two corresponding parts in \mathcal{R} . A symmetric fusion on a symmetric division yields another symmetric division. A matrix $A := (a_{i,j})_{i,j}$ is said *symmetric* in the usual sense, namely, for every entry $a_{i,j}$ of A , $a_{i,j} = a_{j,i}$.

The following definition is crucial. It lists the invariants that we want to keep in our neatly divided matrices in order to build a versatile tree of contractions.

DEFINITION 1. Let $\mathcal{M}_{n,d}$ be the class of the neatly divided $n \times n$ symmetric $0,1,r$ -matrices $(M, (\mathcal{R}, \mathcal{C}))$, such that $(\mathcal{R}, \mathcal{C})$ is symmetric and has:

- *mixed value at most $4c_d$,*
- *part size at most 2^{4c_d+2} , and*

| | | | | | | | |
|-----|-----|---|---|-----|-----|-----|-----|
| 1 | 1 | 0 | 1 | 0 | 0 | r | r |
| 1 | 1 | 1 | 1 | 1 | 1 | r | r |
| 0 | 1 | 1 | 0 | 0 | 1 | 0 | 1 |
| 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 |
| r | r | 0 | 0 | 1 | 0 | 1 | 0 |
| r | r | 0 | 1 | 1 | 0 | 1 | 0 |
| 1 | 0 | 1 | 1 | r | r | 0 | 1 |
| 1 | 0 | 1 | 1 | r | r | 0 | 1 |

| | | | | | | | |
|-----|-----|---|---|-----|-----|-----|-----|
| 1 | 1 | 0 | 1 | 0 | 0 | r | r |
| 1 | 1 | 1 | 1 | 1 | 1 | r | r |
| 0 | 1 | 1 | 0 | 0 | 1 | 0 | 1 |
| 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 |
| r | r | 0 | 0 | 1 | 0 | 1 | 0 |
| r | r | 0 | 1 | 1 | 0 | 1 | 0 |
| 1 | 0 | 1 | 1 | r | r | 0 | 1 |
| 1 | 0 | 1 | 1 | r | r | 0 | 1 |

Figure 2: To the left, a neat division: each zone is horizontal, or vertical, or full with r entries (mixed zone). Note that the division is symmetric but not the matrix. To the right, in bold, a 3-mixed minor of the neat division. Observe that it coarsens the neat division and contains in each of its 9 zones either a 0,1-corner or a mixed zone (framed by red dashed boxes).

- *no d -mixed minor.*

In the previous definition, $c_d := 8/3(d+1)^2 2^{4d}$ as defined in the improvement of Marcus-Tardos bound [8]. The conditions of the first and second bullets are enough to bound the red number of a neatly divided matrix of $\mathcal{M}_{n,d}$.

LEMMA 3.1. *Let $(M, (\mathcal{R}, \mathcal{C}))$ be in $\mathcal{M}_{n,d}$. The red number of M is at most $4c_d \cdot 2^{4c_d+2}$.*

Proof. Any row or column intersects at most $4c_d$ mixed zones (filled with r entries). Each mixed zone has width and length bounded by the part size 2^{4c_d+2} . Hence the maximum total number of r entries on a single row or column is at most $4c_d \cdot 2^{4c_d+2}$. \square

3.3 Finding invariant-preserving coarsenings.

A *coarsening* of a neatly divided matrix $(M, (\mathcal{R}, \mathcal{C}))$ is a neatly divided matrix $(M', (\mathcal{R}', \mathcal{C}'))$ such that $(\mathcal{R}', \mathcal{C}')$ is a coarsening of $(\mathcal{R}, \mathcal{C})$, and M' is obtained from M by setting to r all entries that lie, in M divided by $(\mathcal{R}', \mathcal{C}')$, in a zone with at least one r entry or a 0,1-corner. We also refer to the process of going from $(M, (\mathcal{R}, \mathcal{C}))$ to $(M', (\mathcal{R}', \mathcal{C}'))$ as *coarsening operation* (or simply *coarsening*). A coarsening operation from $(M, (\mathcal{R}, \mathcal{C})) \in \mathcal{M}_{n,d}$ to $(M', (\mathcal{R}', \mathcal{C}'))$ is said *invariant-preserving* if $(M', (\mathcal{R}', \mathcal{C}')) \in \mathcal{M}_{n,d}$, and *elementary* if it consists of a single symmetric fusion. The following lemma shows that not having a t -mixed minor is preserved for free in coarsenings of neatly divided matrices.

LEMMA 3.2. *Every coarsening $(M', (\mathcal{R}', \mathcal{C}'))$ of a t -mixed neatly divided free matrix $(M, (\mathcal{R}, \mathcal{C}))$ is t -mixed free itself.*

Proof. Assume there is a t -mixed minor $(\mathcal{R}^*, \mathcal{C}^*)$ of $(M', (\mathcal{R}', \mathcal{C}'))$. Let us consider the (t, t) -division $(\mathcal{R}^*, \mathcal{C}^*)$ in $(M, (\mathcal{R}, \mathcal{C}))$. By transitivity, $(\mathcal{R}^*, \mathcal{C}^*)$ coarsens $(\mathcal{R}, \mathcal{C})$.

There are two possibilities for a zone Z of $(\mathcal{R}^*, \mathcal{C}^*)$ in $(M', (\mathcal{R}', \mathcal{C}'))$. Either it contains a 0,1-corner, but then, Z contains the same 0,1-corner in $(M, (\mathcal{R}, \mathcal{C}))$. This is because the coarsening operation of a neatly divided matrix never replaces entries by 0 or 1 entries (we may only add r entries). Or Z contains an r entry, or more precisely a zone $Z' \subseteq Z$ of $(M', (\mathcal{R}', \mathcal{C}'))$ filled with r entries. Either one of these r entries was already present in $(M, (\mathcal{R}, \mathcal{C}))$, or the r entries of Z' appear after the fusion of a zone $Z_1 \subseteq Z'$ adjacent to a zone $Z_2 \subseteq Z'$ such that $Z_1 \cup Z_2$ contained a 0,1-corner (and $Z_1 \cup Z_2 \subseteq Z' \subseteq Z$).

Therefore, in any case, Z in $(M, (\mathcal{R}, \mathcal{C}))$ contains an 0,1-corner or an r entry. We conclude that $(\mathcal{R}^*, \mathcal{C}^*)$ is a t -mixed minor of $(M, (\mathcal{R}, \mathcal{C}))$. \square

The previous lemma will in particular give us some control on the average mixed value among the parts of a coarsening of a neatly divided matrix in $\mathcal{M}_{n,d}$. This turns out crucial to find a coarsening which preserves the imposed upper bound on the overall mixed value.

LEMMA 3.3. (\star) *Let $(M, (\mathcal{R}, \mathcal{C}))$ be in $\mathcal{M}_{n,d}$, and $(M', (\mathcal{R}', \mathcal{C}'))$ be a coarsening of $(M, (\mathcal{R}, \mathcal{C}))$ with $|\mathcal{C}'| \geq \lceil |\mathcal{C}|/2 \rceil$. Then the average mixed value among all the parts of \mathcal{C}' on $(M', (\mathcal{R}', \mathcal{C}'))$ is at most $2c_d$.*

Finally we check again, with our slightly different definition of mixed value (compared to that of [7, Section 5]), that the column-part fusions can only decrease the mixed value of row-parts (and vice versa).

LEMMA 3.4. (\star) *Let $(M', (\mathcal{R}', \mathcal{C}'))$ be the coarsening of a neatly divided matrix $(M, (\mathcal{R}, \mathcal{C}))$ resulting from the fusion of a single pair of consecutive parts $C, C' \in \mathcal{C}$, with $C \cup C' = C^*$. Then for every part $R \in \mathcal{R}$, the mixed value of R on $(M', (\mathcal{R}', \mathcal{C}'))$ is at most the mixed value of R on $(M, (\mathcal{R}, \mathcal{C}))$.*

We are now equipped to find invariant-preserving coarsenings.

LEMMA 3.5. (\star) *We set $\ell := 2^{4c_d+1}$ and $s := 8\ell$. Every neatly divided matrix $(M, (\mathcal{R}, \mathcal{C})) \in \mathcal{M}_{n,d}$ has an invariant-preserving coarsening $(M', (\mathcal{R}', \mathcal{C}')) \in \mathcal{M}_{n,d}$ with at least $\lfloor n/s \rfloor$ disjoint pairs of identical columns.*

Now it will become apparent why we are filling the mixed zones with r entries. This allows to simulate a contraction as a simple deletion of an equal row (and a symmetric equal column). The following lemma is straightforward and states that this operation is invariant-preserving in $\mathcal{M}_{n,d}$.

LEMMA 3.6. *Let $(M, (\mathcal{R}, \mathcal{C})) \in \mathcal{M}_{n,d}$ be a neatly divided matrix with two equal rows ρ, ρ' in a part $R \in \mathcal{R}$, hence symmetrically two equal columns γ, γ' in a part $C \in \mathcal{C}$. Then removing row ρ' and the symmetric column γ' yields a neatly divided matrix of $\mathcal{M}_{n-1,d}$.*

Proof. By design the new matrix and division are symmetric. The new neatly divided matrix remains d -mixed free. The part size can only decrease, as well as the mixed value. \square

3.4 Bounded twin-width classes have bounded versatile twin-width. We can now use Lemma 3.5 to find linearly many pairs of vertices that can be contracted, and Lemma 3.6 to recurse. This will be our scheme to find a versatile tree of contractions.

LEMMA 3.7. *(\star) Every trigraph of twin-width d has versatile twin-width at most $4c_{2d+2}2^{4c_{2d+2}+2}$.*

3.5 Finishing the proof. Lemma 3.7 is all we need to mimic Norine et al.'s proof for K_t -minor free graphs [24], as described in Section 3.1.

THEOREM 3.1. *(\star) There is a triple-exponential function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that the number of n -vertex trigraphs with twin-width at most d is at most $n!f(d)^n$.*

3.6 Showing that a class has unbounded twin-width by counting. We have shown that bounded twin-width classes are small. This may be used to establish that the twin-width of some graphs is unbounded, namely if these graphs do *not* form a small class. It is not so easy to show that cubic graphs have unbounded twin-width by direct arguments. Theorem 3.1 implies this fact by a simple counting argument. A bipartite cubic graph is the disjoint union of three perfect matchings. Each matching can be defined in $(n/2)!$ different ways, leading to at least $(n/2)!^3/3^{3n/2} = n^{3n/2+o(n)}$ graphs on vertex set $[n]$, well above $n!c^n = n^{n+o(n)}$. Similarly, two arbitrary total orders on $[n]$ can be defined in $(n!)^2$ ways, hence cannot have bounded twin-width.

We will now define a simple class of graphs capturing two arbitrary orders. Then we will show that these graphs are representable by intervals and by unit disks, and conclude that interval graphs and unit disk graphs have unbounded twin-width. Of course we did not expect these classes to have bounded twin-width⁴, since FO model checking is W[1]-hard on interval graphs [22],

while the mere MAXIMUM INDEPENDENT SET is W[1]-hard on unit disk graphs [21]. We give a more satisfactory proof of that fact, not using the complexity-theoretic assumption $\text{FPT} \neq \text{W}[1]$.

We define the (non-hereditary) class \mathcal{B} by its slices \mathcal{B}_n of graphs on vertex set $[3n]$. Each graph of \mathcal{B}_n has its vertex set partitioned into three cliques of size n , say, (A, B, C) . There is no edge between A and C . There are two arbitrary *half-graphs* between A and B , and between B and C . To build a half-graph between A and B , we first choose an order for the vertices of A , say, a_1, a_2, \dots, a_n , and an order for B , b_1, b_2, \dots, b_n . Then we put an edge between a_i and b_j if and only if $i < j$. The half-graph between B and C is built similarly. We choose another order for the vertices of B , say, b'_1, b'_2, \dots, b'_n , and an order for C , c_1, c_2, \dots, c_n . Then we put an edge between b'_i and c_j if and only if $i < j$. It is important that the choice of the orders b_1, \dots, b_n and b'_1, \dots, b'_n are independent.

Let us estimate the number of graphs in \mathcal{B}_n , ignoring the single-exponential factors such as the one required to fix the partition (A, B, C) . The half-graph between A and B is defined by choosing a total order for A and a total order for B . There are $n!^2$ such pairs of orders. Defining the half-graph between B and C requires an additional total order for B (recall that this second ordering of B is independent of its order for the half-graph on $A \cup B$) and a total order for C . Again this amounts to $n!^2$. Overall there are more than $n!^4$ graphs in \mathcal{B}_n . Thus $|\mathcal{B}_n|$ grows like $n^{4n+o(n)}$, while the number of bounded twin-width graphs with vertices labeled by $[3n]$ is only at most $(3n)!c^{3n} = n^{3n+o(n)}$.

One can describe an unlabeled graph of \mathcal{B}_n with a single permutation σ over $[n]$ such that $b'_{\sigma(i)} = b_i$. Fig. 3 shows how to realize a graph of \mathcal{B}_n as the intersection graph of intervals or as the intersection graph of unit disks, for any given permutation σ .

Unit d -dimensional ball intersection graphs with bounded clique number have bounded twin-width [7]. One could wonder if K_t -free string graphs have bounded twin-width. Fig. 4 shows that even triangle-free unit segment graphs have unbounded twin-width. Indeed it shows how to represent any graph of \mathcal{B}'_n with axis-parallel triangle-free unit segments, where \mathcal{B}'_n is defined analogously to \mathcal{B}_n but the sets A, B, C induce now independent sets, and not cliques. The same argument establishes that the growth of \mathcal{B}'_n is not the one of a small class.

Let us say that a class \mathcal{C} is t -bounded if there is a function $f_{\mathcal{C}}$ such that every K_t -free graph G of \mathcal{C} have twin-width at most $f_{\mathcal{C}}(t)$. The previous remark shows that there are classes that are χ -bounded but not t -bounded, since unit segment graphs are χ -bounded

⁴In [7] we show that FO model checking is FPT on bounded twin-width graphs given with a d -sequence.

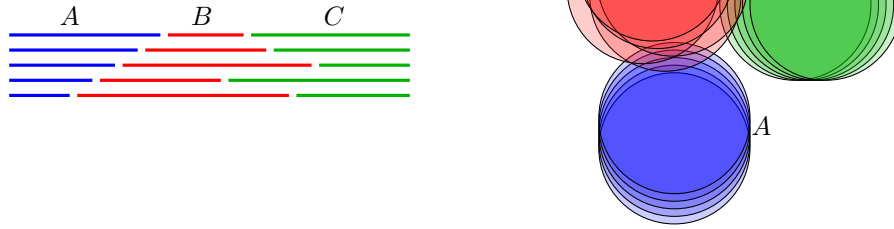


Figure 3: To the left, a representation of a graph of \mathcal{B}_5 by intervals. All intervals are obviously stacked up on a single real line, by projection on the x -axis. To the right, the same graph represented with unit disks. The permutation σ associated to the graph is 41532. In both representations, one can read out the permutation matrix of 41532, where the first row is the bottom one, not the top one. For the intervals, this permutation matrix appears in the small gaps between the intervals of B and C , while for the unit disks the matrix appears in the centers of the disks of B .

[27]. In a subsequent paper [6], we show that classes of bounded twin-width are χ -bounded. This implies in particular that every t -bounded class is χ -bounded, hence the set of t -bounded classes is a proper subset of the set of χ -bounded classes.

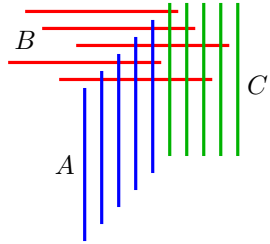


Figure 4: A representation of the graph of Fig. 3, where the cliques induced by A, B, C are replaced by independent sets, with axis-parallel triangle-free unit segments. The upside-down permutation matrix of $\sigma = 41532$ is still visible as the right endpoints of the red segments.

4 Short parallel d -sequences and adjacency labeling schemes

Every d -contraction sequence of an n -vertex graph has length exactly $n - 1$, since each of its steps contracts exactly one pair of vertices. What if we allow *parallel contractions* where disjoint pairs of vertices may be contracted in a single step? In this section we adapt the results of Section 3 on versatile twin-width to prove the existence of parallel contraction sequences of logarithmic length. We then use them to provide an $f(d) \log n$ -adjacency labeling scheme for graphs of twin-width at most d .

A *parallel contraction* in a trigraph G consists of the successive contractions of any number of pairs of vertices $\{a_1, b_1\}, \dots, \{a_\ell, b_\ell\}$, where $a_1, \dots, a_\ell, b_1, \dots, b_\ell$ are all distinct. One can check that the resulting trigraph does not depend on the order in which the pairs are contracted. Thus instead of the contraction of a *sequence* of pairs, we may as well speak of the parallel contraction of a *set* of disjoint pairs. A *sequence of parallel d -contractions*, or *parallel d -sequence* of a trigraph G is a sequence of d -trigraphs G_k, \dots, G_1 where $G_k = G$, $G_1 = K_1$ is the one-vertex (tri)graph, and G_{i-1} is obtained from G_i by a parallel contraction (of disjoint pairs of vertices). It is noteworthy that the existence of a parallel contraction sequence is equivalent to the existence of a (regular) contraction sequence, up to a multiplicative factor in the red degree.

LEMMA 4.1. (\star) *Let G be a trigraph, and $d \in \mathbb{N}$.*

- *If G admits a d -sequence, then G also admits a parallel d -sequence.*
- *If G admits a parallel d -sequence, then G also admits a $(2d + 1)$ -sequence.*

Our main result on parallel contraction sequences is that one can always find a parallel sequence of logarithmic length, at the cost of an increase in the red degree. This is a variant of the versatile twin-width theorem presented in Section 3 (Lemma 3.7).

LEMMA 4.2. (\star) *Any n -vertex graph G with twin-width at most d admits a parallel D -sequence of length $O(s \cdot \log n)$ where s, D are double exponential functions of d .*

From Lemmas 4.1 and 4.2, we immediately conclude the following.

THEOREM 4.1. *The class of graphs with twin-width at most d admits a $g(d) \log n$ -bits adjacency labeling scheme, where n is the number of vertices and g is a double-exponential function.*

The labeling scheme can in particular be used to encode an n -vertex graph of twin-width at most d on $2^{2^{\gamma(d+1)}} n \log n$ bits, for some constant γ . This offers a significant compression over adjacency lists, since cliques for instance have twin-width 0. Now if the aim is only to globally compress the whole graph, and not to balance the lengths of the vertex labels, there is a simpler encoding with a better dependency in d . It basically consists of “reading” the d -sequence $G = G_n, \dots, G_1 = K_1$ backwards. The encoding of K_1 is an identifier on $\lceil \log n \rceil$ bits. Then to go from G_i to G_{i+1} , we write $3 \lceil \log n \rceil + 2$ bits corresponding to the “split vertex” w , in which two vertices u, v vertex w is split, and whether there is a non-edge, a black edge, or a red edge between u and v , followed by $d(\lceil \log n \rceil + 4)$ bits corresponding to the edges between u, v and the at most d vertices adjacent to w in the red graph of G_i . The latter part is carried by writing down the identifier of each red neighbor z of w followed by two pairs of bits encoding if there is a non-edge, a black edge, or a red edge between u and z , and between v and z . This permits to reconstruct G , and store it on only $(d+3)n \lceil \log n \rceil + (4d+2)n$ bits.

5 Expanders with bounded twin-width

A 2 -lift of a graph G is a graph G' on twice as many vertices, built by duplicating every vertex $v \in V(G)$ into two copies, say, v_1 and v_2 , and for every edge $vw \in E(G)$, adding to $E(G')$ either the edges v_1w_1 and v_2w_2 (parallel) or the edges v_1w_2 and v_2w_1 (crossing). The choice, for each edge of G , of having two *parallel* edges or two *crossing* edges is called the *signing* of the edges. See Fig. 5 for an example of a 2 -lift. Observe that G has $2^{|E(G)|}$ possible 2 -lifts or signings. For instance, the all-parallel signing gives two disjoint copies of G , while the all-crossing signing gives the bipartite adjacency graph of G .

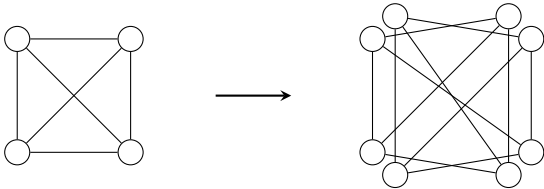


Figure 5: An example of a 2 -lift of K_4 .

For n a power of 2, performing a sequence of $\log n - 2$ randomly-signed 2 -lifts starting on K_4 yields

an n -vertex expander almost surely [3]. Observe that the obtained graph is necessary cubic since the 2 -lift operation preserves the degree. Bilu and Linial [3] even exhibit a deterministic polytime procedure to actually find the signings leading from K_4 to a cubic expander. The next result shows that cubic expanders can have bounded twin-width.

LEMMA 5.1. *Every graph obtained from K_4 by performing a sequence of 2 -lifts has twin-width at most 6.*

Proof. We show that if G is a cubic graph and G' is a 2 -lift of G , then G can be obtained from G' by a sequence of contractions in which the maximum degree never goes above 6. It is enough to conclude since K_4 is obviously 6-collapsible, and we can assume that the cubic trigraph we start from has all its edges red.

Let v^1, v^2, \dots, v^n be the vertices of G , and v_1^i, v_2^i be the duplicates of v^i in G' . For each i running from 1 to n , we contract v_1^i and v_2^i . By definition of a 2 -lift, after these n contractions, the graph obtained is G . We contracted disjoint pairs of vertices of degree 3, so we could not create vertices of degree more than 6. \square

6 Subdivisions of cliques

For any non-negative integer k , the k -subdivision of a graph G , denoted by $G^{(k)}$, is the graph obtained by subdividing every edge of G exactly k times. For any $f : \mathbb{N} \rightarrow \mathbb{N}$, let \mathcal{G}_f be the class formed by the $f(|V(G)|)$ -subdivision of every graph G .

THEOREM 6.1. *For every positive and non-decreasing f , \mathcal{G}_f has bounded twin-width if and only if $f(n) = \Omega(\log n)$.*

Let us first observe that for any integer $k > 0$ and n -vertex graph G , $G^{(k)}$ is an induced subgraph of $K_n^{(k)}$. Thus the class \mathcal{G}_f is contained in the hereditary closure of the graphs $K_n^{(f(n))}$ for $n \geq 0$. Since twin-width never increases when taking induced subgraphs, it suffices to consider graphs of the form $K_n^{(f(n))}$. As hinted at in Section 2.3, the forward implication of Theorem 6.1 could be derived from Theorem 3.1 and the fact that $o(\log n)$ -subdivisions does *not* form a small class. We give a direct proof of a stronger statement.

PROPOSITION 6.1. *For $d \geq 0$ and $k > 0$ integers, if $K_n^{(k)}$ has twin-width at most d , then $k \geq \log_{d+1}(n-1) - 1$.*

Proof. Let G be $K_n^{(k)}$, for some positive integer k . Assuming that G has twin-width at most d , we show that $k \geq \log_{d+1}(n-1) - 1$. Note that the assumption $k > 0$ is required because $K_n^{(0)} = K_n$ has twin-width 0.

In a d -contraction sequence of G , let us consider the first step in which two vertices x, y of the original K_n are contracted. Let \mathcal{P} the partition of $V(G)$ at this step, and $P_0 \in \mathcal{P}$ the part containing x and y . In G , consider the $n - 1$ paths, on $k + 1$ edges each, resulting from the subdivided edges starting at x . We partition the vertices of these paths as V_1, \dots, V_{k+1} , where V_i contains all the vertices at distance i of x . Then V_{k+1} contains all the vertices of the original K_n except x . In particular, no two vertices of V_{k+1} are in the same part of \mathcal{P} .

All the vertices of V_1 are neighbors of x but not of y , thus for any part $P \in \mathcal{P} \setminus \{P_0\}$ intersecting V_1 , PP_0 is a red edge in $G_{\mathcal{P}}$. Thus at most $d + 1$ parts of \mathcal{P} intersect V_1 , and there exists $P_1 \in \mathcal{P}$ such that $|P_1 \cap V_1| \geq \frac{n-1}{d+1}$. Observe that P_1 may well be equal to P_0 . Similarly the vertices in $P_1 \cap V_1$ have pairwise-disjoint neighborhoods in V_2 , hence $V_2 \cap N_G(P_1 \cap V_1)$, of size at least $\frac{n-1}{d+1}$, is split in at most $d + 1$ parts in \mathcal{P} . Thus there is a part $P_2 \in \mathcal{P}$ (that may be P_0 or P_1) which contains at least $\frac{n-1}{(d+1)^2}$ vertices of V_2 . It follows by induction that for every $i \in [k + 1]$, there exists a part of \mathcal{P} containing at least $\frac{n-1}{(d+1)^i}$ vertices of V_i . However no part of \mathcal{P} contains more than one vertex of V_{k+1} . Hence $\frac{n-1}{(d+1)^{k+1}} \leq 1$, and $k \geq \log_{d+1}(n - 1) - 1$. \square

The converse relies on some results on decompositions of permutations. We now encode a permutation σ in the usual way, as the sparse matrix with entry 1 at position $(i, \sigma(i))$, and 0 elsewhere. (This is unlike the more cumbersome but technically-motivated dense encodings used in Section 3.6 and [7, Section 6.1].)

A permutation σ is a t -merge if its domain can be partitioned into t possibly-empty discrete intervals I_1, \dots, I_t such that the restriction of σ to I_i is increasing. Merging t sorted lists can be expressed as the application of some well chosen t -merge to the concatenation of the lists. A permutation σ is a *parallel t -merge* if its domain can be partitioned into an arbitrary number of intervals J_1, \dots, J_r such that σ operates independently on each J_i (i.e., $\sigma(J_i) = J_i$), and the restriction $\sigma|_{J_i}$ is a t -merge. See Fig. 6 for an example of a parallel 2-merge.

LEMMA 6.1. *For any $t, \ell \in \mathbb{N}$, any permutation on t^ℓ elements can be decomposed as a product of at most ℓ parallel t -merges.*

Proof. The case $t = 2$ corresponds to a merge sort. In the recursion tree of a merge sort, each level of inductive calls can be expressed as a single parallel 2-merge. To sort up to 2^ℓ elements, a merge sort with recursion depth limited to ℓ suffices, and this can be expressed as the composition of ℓ parallel 2-merges. This generalizes easily to t -merges, and composing ℓ parallel t -merges allows to sort up to t^ℓ elements. \square

| | | | | | | |
|---|---|---|---|---|---|---|
| | | | | | 1 | 0 |
| | | | | | 0 | 1 |
| | | | | | | |
| | | | | | 1 | |
| 0 | 0 | 0 | 1 | 0 | | |
| 1 | 0 | 0 | 0 | 0 | | |
| | | | | | | |
| 0 | 0 | 0 | 0 | 1 | | |
| 0 | 0 | 1 | 0 | 0 | | |
| 0 | 1 | 0 | 0 | 0 | | |

Figure 6: A parallel 2-merge matrix, corresponding to the permutation 23514687. Note that the first row is at the bottom, as is common with permutation matrices. It is composed of three blocks, each of which can be partitioned in two increasing subsequences, indicated by the dashes. Empty areas are filled with 0.

The previous lemma is reminiscent of the theory of sorting networks, in that we decompose arbitrary permutations as a product of few base permutations—in our case parallel t -merges. However, sorting networks consider more restricted base permutations (e.g., separable permutations), whereas we merely need the base permutations to have bounded twin-width.

LEMMA 6.2. *Matrices of parallel t -merges are $(t + 1)$ -grid free.*

Proof. Let σ a parallel t -merge, with its domain partitioned into intervals J_1, \dots, J_r such that $\sigma(J_i) = J_i$, and every $\sigma|_{J_i}$ is a t -merge. Assume for a contradiction that σ contains a $(t + 1)$ -grid. Then it contains a decreasing subsequence of length $t + 1$.

For any $i < j$, $x \in J_i$ and $y \in J_j$, one has $x < y$ and $\sigma(x) < \sigma(y)$ because J_i, J_j are disjoint intervals, with $\sigma(J_i) = J_i$ and $\sigma(J_j) = J_j$. It follows that any decreasing subsequence is contained entirely in one of the J_k . Thus, there exist a t -merge $\sigma|_{J_k}$ which contains a decreasing subsequence of length $t + 1$.

Since $\sigma|_{J_k}$ is a t -merge, J_k is itself partitioned into intervals I_1, \dots, I_t such that σ is increasing on I_i . Hence each I_i can contain at most one element of a decreasing subsequence, and $\sigma|_{J_k}$ contains no decreasing subsequence of length more than t , a contradiction. \square

PROPOSITION 6.2. *For any $c > 0$, the class of cliques K_n subdivided at least $\frac{\log n}{c}$ times has twin-width at most $f(c)$ for some triple-exponential function f .*

Proof. Let $k \geq \frac{\log n}{c}$, and let G be $K_n^{(k)}$. We want to order $V(G)$ such that the adjacency matrix of G in that order is r -grid free, for some r depending only on c . This implies the desired twin-width bound by Theorem 2.2.

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix} \quad
\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad
\begin{bmatrix} & & & 0 & 0 & 1 & 0 \\ & & & 0 & 0 & 0 & 1 \\ & & & 1 & 0 & 0 & 0 \\ & & & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & & & \\ 0 & 0 & 0 & 1 & & & \\ 0 & 1 & 0 & 0 & & & \\ 0 & 0 & 1 & 0 & & & \end{bmatrix}$$

Figure 7: Left: the permutation τ to sort. Center: the 2-merge permutation σ to use on τ . Right: the composition $\sigma^{-1} \circ \tau$, one may inductively sort the two blocks by applying further parallel 2-merges.

Choose an arbitrary orientation of the edges of K_n . In G , the edges of K_n become directed paths on $k+1$ edges. Then, for $0 \leq i \leq k$, let $V_i \subset V(G)$ contain every i -th vertex along these directed paths. In particular, V_0 corresponds to the vertices of K_n , while V_1, \dots, V_k are all the vertices created by the subdivision. Thus, V_0, \dots, V_k is a partition of $V(G)$.

Let us now define an order within each V_i . Choose x_1, \dots, x_n an arbitrary order on V_0 . The extremal set V_1 is ordered according to the neighbors in V_0 , i.e., with first the neighbors of x_1 in any order, then the neighbors of x_2 , etc. We proceed similarly for V_k . The disjoint paths in $G - V_0$ define a bijection between V_1 and V_k , which can be interpreted as a permutation σ on $\frac{n(n-1)}{2}$ elements according to the previous orderings. Then, choosing orderings for V_2, \dots, V_{k-1} is equivalent to decomposing σ as a product $\sigma_1 \circ \dots \circ \sigma_{k-1}$. By Lemma 6.1, we may choose $\sigma_1, \dots, \sigma_{k-1}$ to be parallel t -merges for any t such that $t^{\log(n)/c} \geq \frac{n(n-1)}{2}$. This is satisfied by $t = \lceil 2^{2^c} \rceil$, which crucially is independent of n . With this choice of decomposition for σ , we have ordered V_2, \dots, V_{k-1} . Finally, $V(G)$ is ordered as $V_0 < \dots < V_k$, where V_i is ordered as previously defined.

Let M be the adjacency matrix of G respecting this ordering. Let R_0, \dots, R_k (resp. C_0, \dots, C_k) the partition of the rows (resp. columns) of M induced by the partition V_0, \dots, V_k of $V(G)$. Then $(\mathcal{R}, \mathcal{C}) = (\{R_0, \dots, R_k\}, \{C_0, \dots, C_k\})$ is a division of M . For $i, j \in [0, k]$, let $M_{i,j}$ be the zone $R_i \cap C_j$, which corresponds to the adjacency matrix between V_i and V_j . The zone $M_{i,j}$ is non-zero if and only if $i = j \pm 1$ modulo $k+1$. Thus, there are $2k+2$ non-zero zones, forming a double diagonal with corners (see Fig. 8).

CLAIM 1. *Every zone of the division $(\mathcal{R}, \mathcal{C})$ of M is $(t+1)$ -grid free.*

PROOF OF THE CLAIM: For $1 \leq i < k$, the zones $M_{i,i+1}$ and $M_{i+1,i}$ are parallel t -merges or transposes thereof, hence are $t+1$ -grid free by Lemma 6.2. The zones $M_{0,1}$,

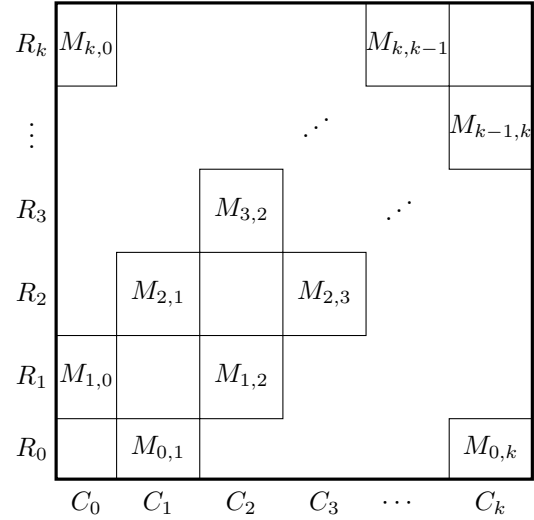


Figure 8: The adjacency matrix M of G , with the appropriate ordering of the vertices.

$M_{1,0}$, $M_{0,k}$, and $M_{k,0}$ are composed of a single monotone sequence, hence are 2-grid free. \diamond

Let us now consider an ℓ -grid minor $(\mathcal{R}', \mathcal{C}') = (\{R'_1, \dots, R'_\ell\}, \{C'_1, \dots, C'_\ell\})$ of M , i.e., every zone $R'_i \cap C'_j$ contains at least one entry 1.

CLAIM 2. *There is a set $A \subset \mathcal{C}$ of at most 5 column-parts such that every $C' \in \mathcal{C}'$ satisfies $C' \cap \bigcup A \neq \emptyset$.*

PROOF OF THE CLAIM: Let i be minimal such that $R'_1 \subseteq R_0 \cup \dots \cup R_i$. Then for $i+1 < j < k$, one may verify from the structure of M that $R'_1 \cap C'_j$ is full 0. Thus any $C' \in \mathcal{C}'$ must intersect one of C_0, \dots, C_{i+1} or C_k . Symmetrically, let i' be maximal such that $R'_\ell \subseteq R_{i'} \cup \dots \cup R_k$. Then any $C' \in \mathcal{C}'$ must intersect one of C_0 or $C_{i'-1}, \dots, C_k$. Define $A := \{C_{i'-1}, \dots, C_{i+1}\} \cup \{C_0, C_k\}$. The above implies that any $C' \in \mathcal{C}'$ must intersect some $C \in A$. Finally we have $i \leq i'$, which implies $|A| \leq 5$. \diamond

CLAIM 3. *There exists $C \in \mathcal{C}$ such that at least $\frac{\ell-10}{5}$ parts of \mathcal{C}' are subsets of C .*

PROOF OF THE CLAIM: Consider an arbitrary $C' \in \mathcal{C}'$. By Claim 2, there is some $C \in A$ such that $C' \cap C \neq \emptyset$. We consider two cases, depending on whether or not $C' \subseteq C$:

- If $C' \not\subseteq C$, it means that C' contains one of the two boundaries of C . For a given $C \in A$, there can only be two $C' \in \mathcal{C}'$ for which it is the case. Thus this case applies to at most $2|A| \leq 10$ elements of \mathcal{C}' .
- Otherwise—and this applies to at least $\ell - 10$ elements of \mathcal{C}' —we have $C' \subseteq C$ for some $C \in A$. Since $|A| \leq 5$, by pigeonhole principle, there exist $C \in A$ such that at least $\frac{\ell-10}{5}$ elements of \mathcal{C}' are subsets of C .

◇

Of course, Claims 2 and 3 still hold when inverting the roles of rows and columns. Thus, there are $R \in \mathcal{R}, C \in \mathcal{C}$ such that R (resp. C) contains at least $\frac{\ell-10}{5}$ parts of \mathcal{R}' (resp. \mathcal{C}') as subsets. Hence the zone $R \cap C$ contains an $\frac{\ell-10}{5}$ -grid induced by the corresponding parts of \mathcal{R}' and \mathcal{C}' . By Claim 1, it follows that $\frac{\ell-10}{5} \leq t$, or $\ell \leq 5t + 10$. Recall that t was chosen as $t = \lceil 2^{2c} \rceil$. Hence we have proved that M is $g(c)$ -grid free for $g(c) = 5\lceil 2^{2c} \rceil + 11$.

A fortiori M is $g(c)$ -mixed free, and by Theorem 2.2 the twin-width of G is at most $f(c)$ for some $f(c)$ double-exponential in $g(c)$, hence triple-exponential in c . □

In the next section, we will show that graphs with queue number t have twin-width $2^{2^{O(t)}}$ (see Theorem 7.4). This can be used to get an alternative proof to Proposition 6.2, albeit not self-contained. Indeed it was shown that the $2\lceil \log_d \lfloor n/2 \rfloor \rceil + 1$ -subdivision of K_n (see [13, Theorem 4]) has queue number at most d .

7 Sparse twin-width

We start this section by showing the list of equivalences of Theorem 2.9.

7.1 Characterizations. We recall that $A_\sigma(G)$ is the adjacency matrix of G when $V(G)$ is ordered by σ , and that a class \mathcal{C} is said d -grid free if for every $G \in \mathcal{C}$ there is an ordering σ of $V(G)$ such that $A_\sigma(G)$ is d -grid free.

THEOREM 7.1. *(*) If \mathcal{C} is a hereditary class of bounded twin-width, the following are equivalent.*

- (i) *There is an integer t such that no graph of \mathcal{C} contains $K_{t,t}$ as a subgraph.*
- (ii) *There is an integer d such that \mathcal{C} is d -grid free.*

- (iii) *There is an integer g such that every n -vertex graph $G \in \mathcal{C}$ has at most gn edges.*
- (iv) *The subgraph closure $\text{Sub}(\mathcal{C})$ has bounded twin-width.*
- (v) *There is a function f such that $\nabla_r(\mathcal{C}) \leq f(r)$ for every r .*

The equivalences (i) \Leftrightarrow (ii) \Leftrightarrow (iv) \Leftrightarrow (v) hold for every (possibly non-hereditary) class of bounded twin-width. Bounded sparse twin-width classes remain surprisingly diverse. They for instance contain K_t -minor free graphs and bounded-degree bounded twin-width graphs, which in turn contain some expander classes. In particular bounded sparse twin-width graphs do not have polynomial expansion.

7.2 Flat classes. For any graph invariant ι , we say that a class \mathcal{C} is ι flat if it is included in $\text{Sub}(\mathcal{G} \boxtimes \mathcal{H})$ with \mathcal{G} and \mathcal{H} two classes of bounded ι , and \mathcal{H} also has bounded degree. Recalling the definition in Section 2, a class is flat if it is treewidth flat. We will see that twin-width flat classes have bounded twin-width. It will imply that (treewidth) flat classes are other examples of bounded sparse twin-width classes.

We say that G is a *trigraph over a graph H* if $(V(G), E(G) \cup R(G))$ is isomorphic to H . Thus G is obtained from the graph H by coloring red some of its edges. More generally G is a *trigraph over a trigraph H* if there is a graph isomorphism from $(V(G), E(G) \cup R(G))$ to $(V(H), E(H) \cup R(H))$ such that every black edge of G is mapped to a black edge of H . Again G is obtained from the trigraph H by coloring red some of its black edges. We start by bounding the twin-width of trigraphs over graphs with bounded degree and bounded twin-width.

LEMMA 7.1. *Every trigraph over a graph H has twin-width at most $\text{tw}(H) + \Delta(H)$.*

Proof. Consider a $\text{tw}(H)$ -sequence of H . A simple but important observation is that the black degree of a vertex never increases in a contraction sequence. Thus each trigraph of the sequence has total degree at most $\Delta(H) + \text{tw}(H)$. Therefore, when the same sequence is applied to any trigraph over H , the overall maximum (red) degree is also bounded by $\Delta(H) + \text{tw}(H)$. □

We can now show the following.

THEOREM 7.2. *Let G and H be two graphs. Then $\text{tw}(G \boxtimes H) \leq \max\{\text{tw}(G)(\Delta(H) + 1) + 2\Delta(H), \text{tw}(H) + \Delta(H)\}$.*

Proof. We set $d_G := \text{tw}(G)$, $d_H := \text{tw}(H)$, and $\Delta := \Delta(H)$, i.e., the maximum degree of H . Let

$G = G_n, \dots, G_1 = K_1$ be a sequence of d_G -contraction, and let $[h]$ be the vertex set of H , hence $h = |V(H)|$. We set $d := \max\{(d_G + 2)\Delta, d_H + \Delta\}$, and present a d -sequence for $G \boxtimes H$. For a fixed $j \in [h]$, we call j -th copy of G , the vertices (v, j) of $G \boxtimes H$ for every $v \in V(G)$.

First we contract $G \boxtimes H$ to a trigraph over H by a sequence containing as intermediate steps trigraphs over $G_n \boxtimes H, G_{n-1} \boxtimes H, \dots, G_1 \boxtimes H$. Say G_i is obtained from G_{i+1} by contracting $u, v \in V(G_{i+1})$, into vertex w , then the part of the d -sequence from a trigraph over $G_{i+1} \boxtimes H$ to one over $G_i \boxtimes H$ consists of contracting, in any order, the vertices (u, j) and (v, j) , into vertex (w, j) , for every $j \in [h]$. As the red degree of $w \in V(G_i)$ is at most d_G , vertex (w, j) has red degree at most $d_G(\Delta + 1) + 2\Delta$. This is because the j -th copy of G is linked to the j' -th copy only if $j' \in N_H[j]$. This explains the $d_G(\Delta + 1)$ term. The additional 2Δ accounts for possible red edges between (w, j) and (\star, j') , where $\star \in \{u, v, w\}$ and $j' \neq j$.

We can now finish the d -sequence from the obtained trigraph over $K_1 \boxtimes H$, which is isomorphic to H , using the d_H -sequence of H . Indeed by Lemma 7.1 this trigraph admits a $d_H + \Delta$ -sequence. \square

THEOREM 7.3. *(\star) Let \mathcal{G} and \mathcal{H} two classes such that $\mathcal{G} \boxtimes \mathcal{H}$ is $K_{t,t}$ -free. Then $\text{tww}(\text{Sub}(\mathcal{G} \boxtimes \mathcal{H})) \leq f(\text{tww}(\mathcal{G}), \text{tww}(\mathcal{H}), \Delta(\mathcal{H}), t)$.*

LEMMA 7.2. *(\star) If G is $K_{t,t}$ -free, then $G \boxtimes H$ is $K_{s,s}$ -free where $s = 2t(\Delta(H) + 1)$.*

Theorem 2.8 and Lemma 7.2 imply that flat classes have bounded twin-width, since bounded treewidth classes have sparse bounded twin-width (they are $K_{t,t}$ -free and have bounded twin-width). In particular, it provides an alternative proof that planar graphs have bounded twin-width (see [7, Section 6]). The obtained bound remains bad since we still need to use Theorem 2.2 to justify that the subgraph closure of a $K_{t,t}$ -free bounded twin-width class has bounded twin-width.

7.3 Classes with bounded queue or stack number. A pair of edges uv and xy is said *independent* if u, v, x, y are four distinct vertices. An independent pair of edges uv and xy is *nested* with respect to a linear ordering σ of the vertex set, if $u \prec_\sigma x \prec_\sigma y \prec_\sigma v$, and *overlaps* if $u \prec_\sigma x \prec_\sigma v \prec_\sigma y$. A *queue* (resp. *stack*) *layout* of a graph G is a linear ordering σ of $V(G)$ and a partition of $E(G)$ into t parts, called queues (resp. stacks), such that no independent pair of edges within the same part is nested (resp. overlaps) with respect to σ . The *queue number* (resp. *stack number*) is defined as the minimum integer t such that such a queue layout (resp. stack layout) exists.

LEMMA 7.3. *Let σ be a linear ordering on the vertex set of a graph G . If G admits an edge partition into t parts such that each part forms a queue (resp. a stack) with respect to σ , then the adjacency matrix $A_\sigma(G)$ is $2(t + 1)$ -grid free.*

Proof. Assume, for the sake of contradiction, that $A_\sigma(G)$ has a $2(t + 1)$ -grid minor $(\mathcal{R}, \mathcal{C}) := (\{R_1, \dots, R_{2t+2}\}, \{C_1, \dots, C_{2t+2}\})$, i.e., each zone $R_i \cap C_j$ contains an entry 1. Let us consider the 2×2 coarsening of $(\mathcal{R}, \mathcal{C})$ where each part contains the first/last $t + 1$ row/column parts of $(\mathcal{R}, \mathcal{C})$. At least one of the two off-diagonal zones of this coarsening does *not* cross the main diagonal of $A_\sigma(G)$. Without loss of generality, let us assume that all the vertices of R_1, \dots, R_{t+1} precedes all the vertices of C_{t+2}, \dots, C_{2t+2} in the order σ . Now for each $i \in [t + 1]$, choose one edge $u_i v_i$ from the zone $R_{t+2-i} \cap C_{t+1+i}$. From the previous observation, we know that $u_i \prec_\sigma v_i$ for each i . With respect to σ , the vertices u_i (for i going from 1 to $t + 1$) form a decreasing sequence in σ while the vertices v_i form an increasing sequence. Therefore the chosen $t + 1$ edges are pairwise nested, contradicting that G admits an edge partition into t queues with respect to σ .

For the stack number, we choose $t + 1$ edges $u_i v_i$ from the zones $R_i \cap C_{t+1+i}$ for $i \in [t + 1]$. They pairwise overlap, and thus contradict that there is a partition into t stacks with that vertex ordering. \square

The following is a direct consequence of Theorem 2.2 and Lemma 7.3.

THEOREM 7.4. *Classes with queue or stack number t have twin-width bounded by $2^{2^{O(t)}}$.*

8 Twin-width of finitely generated groups

We investigate here an algebraic approach to constructing small graph classes. Let Γ be a (multiplicative) countable group where the identity is denoted by ε . We assume that Γ is generated by a finite set S . We form the Cayley graph $\text{Cay}(\Gamma, S)$ which has vertex set Γ and edge set all pairs $\{x, x \cdot s\}$ where $x \in \Gamma$ and $s \in S$.

For example when Γ is the free group generated by $S = \{a, b\}$, the graph $\text{Cay}(\Gamma, S)$ is the infinite tree where all the vertices have degree 4. Furthermore, if we quotient Γ by the relation $aba^{-1}b^{-1} = \varepsilon$, we obtain the infinite two-dimensional grid. Both trees and grids are examples of classes with bounded twin-width. Thus a natural question is whether this could hold for all finitely generated groups. Let us denote by $F(\Gamma, S)$ the set of all finite induced subgraphs of $\text{Cay}(\Gamma, S)$. Our main question in this section is the following.

CONJECTURE 3. *For every group Γ generated by a finite set S , the class $F(\Gamma, S)$ has bounded twin-width.*

This is a far-reaching generalization of the case of trees and grids. It could provide some insights on both the structure of finite induced subgraphs of $\text{Cay}(\Gamma, S)$, but also in the global (infinite) structure of $\text{Cay}(\Gamma, S)$ as illustrated by the following result.

PROPOSITION 8.1. *If all the finite induced subgraphs of an infinite (possibly uncountable) graph G have twin-width at most t , then there is a linear order L on $V(G)$ such that the adjacency matrix of G , ordered by L , has no $f(t)$ -mixed minor.*

Proof. Let $F(G)$ the class of finite non-empty induced subgraphs of G . We assume that graphs in $F(G)$ have twin-width at most t , hence there exists an integer $f(t)$ such that any $H \in F(G)$ has a linear order L_H such that the adjacency matrix of H has no $f(t)$ -mixed minor. Let \mathcal{M}_H be the logical structure formed by H equipped with the order L_H . The proof proceeds in two parts. First, we will build an ultraproduct of $(\mathcal{M}_H)_{H \in F(G)}$, following a standard construction used, for example, to prove the compactness theorem. We will then show that G is an induced subgraph of this ultraproduct. The order L on $V(G)$ is then obtained by restriction of the order on the ultraproduct.

For $H \in F(G)$, let $\uparrow H = \{H' \in F(G) \mid V(H) \subseteq V(H')\}$ be its upward closure in $F(G)$. The family of all $\uparrow H$ generates a proper filter on $F(G)$, which is contained in some ultrafilter U . Let $\mathcal{M}' = \prod_{H \in F(G)} \mathcal{M}_H / U$ be the corresponding ultraproduct. By Łoś's theorem, any first-order formula satisfied by every \mathcal{M}_H is also satisfied by \mathcal{M}' . Being a linear order, and being $f(t)$ -mixed free with respect to that order can both be expressed in first-order logic, hence \mathcal{M}' is an infinite graph equipped with a linear order for which it is $f(t)$ -mixed free.

Let us show that G is an induced subgraph of \mathcal{M}' . For $v \in V(G)$, choose $\bar{v} \in \prod_{H \in F(G)} \mathcal{M}_H$ to be a tuple “equal to v when possible”, that is $\bar{v}(H) = v$ when $V(H) \ni v$ (and unconstrained if $V(H) \not\ni v$). We then map v to the equivalence class of \bar{v} , which is a vertex in \mathcal{M}' . This mapping is injective: If $u \neq v$, then $\bar{u}(H) \neq \bar{v}(H)$ for any H such that $u, v \in V(H)$, i.e., whenever $H \in \uparrow G[\{u, v\}]$. Since $\uparrow G[\{u, v\}]$ is an element of U , by Łoś's theorem, \bar{u} and \bar{v} are not equated in \mathcal{M}' . The same arguments show that this mapping preserves edges and non-edges. Hence G is an induced subgraph of \mathcal{M}' , and it follows that G is $f(t)$ -mixed free for the linear order on \mathcal{M}' restricted to G . \square

We suspect that bounded-degree Cayley graphs have bounded twin-width since they form a small class.

LEMMA 8.1. (\star) *The class $F(\Gamma, S)$ is small.*

Should the small conjecture be true, $F(\Gamma, S)$ would have bounded twin-width. We finally observe that having bounded twin-width is a group invariant, i.e., does not depend on the choice of the finite generating set S .

LEMMA 8.2. (\star) *If S and S' are two finite generating sets of the group Γ , then $F(\Gamma, S)$ has bounded twin-width if and only if $F(\Gamma, S')$ has bounded twin-width.*

Therefore, if the small conjecture does not hold, the class of finitely generated groups splits into bounded twin-width groups and unbounded twin-width groups. This could reflect a known dichotomy for groups. A natural candidate for a finitely generated group of unbounded twin-width, would be a group with no finite presentation. For instance the lamplighter group is an interesting test case, but its associated class of graphs has indeed bounded twin-width. A first step towards Conjecture 3 is to show that finitely presented groups have bounded twin-width.

References

- [1] Stephen Alstrup, Søren Dahlgaard, and Mathias Bæk Tejs Knudsen. Optimal induced universal graphs and adjacency labeling for trees. *J. ACM*, 64(4):27:1–27:22, 2017. URL: <https://doi.org/10.1145/3088513>, doi:10.1145/3088513.
- [2] Lowell W Beineke and Raymond E Pippert. The number of labeled k -dimensional trees. *Journal of Combinatorial Theory*, 6(2):200–205, 1969.
- [3] Yonatan Bilu and Nathan Linial. Lifts, discrepancy and nearly optimal spectral gap*. *Combinatorica*, 26(5):495–519, 2006. URL: <https://doi.org/10.1007/s00493-006-0029-7>, doi:10.1007/s00493-006-0029-7.
- [4] Achim Blumensath and Bruno Courcelle. On the monadic second-order transduction hierarchy. *Logical Methods in Computer Science*, 6(2), 2010. URL: [https://doi.org/10.2168/LMCS-6\(2:2\)2010](https://doi.org/10.2168/LMCS-6(2:2)2010), doi:10.2168/LMCS-6(2:2)2010.
- [5] Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width II: small classes. *CoRR*, abs/2006.09877, 2020. URL: <https://arxiv.org/abs/2006.09877>, arXiv:2006.09877.
- [6] Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width III: Max Independent Set and Coloring. *Manuscript*, 2020.
- [7] Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. *CoRR*, abs/2004.14789, 2020. To appear at FOCS 2020. URL: <https://arxiv.org/abs/2004.14789>, arXiv:2004.14789.

- [8] Josef Cibulka and Jan Kyncl. Füredi-hajnal limits are typically subexponential. *CoRR*, abs/1607.07491, 2016. URL: <http://arxiv.org/abs/1607.07491>, arXiv:1607.07491.
- [9] Vida Dujmovic, Louis Esperet, Gwenaël Joret, Cyril Gavoille, Piotr Micek, and Pat Morin. Adjacency labelling for planar graphs (and beyond). *CoRR*, abs/2003.04280, 2020. URL: <https://arxiv.org/abs/2003.04280>, arXiv:2003.04280.
- [10] Vida Dujmović, Louis Esperet, Pat Morin, Bartosz Walczak, and David R Wood. Clustered 3-colouring graphs of bounded degree. *arXiv preprint arXiv:2002.11721*, 2020.
- [11] Vida Dujmovic, Gwenaël Joret, Piotr Micek, Pat Morin, Torsten Ueckerdt, and David R. Wood. Planar graphs have bounded queue-number. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 862–875. IEEE Computer Society, 2019. URL: <https://doi.org/10.1109/FOCS.2019.00056>, doi:10.1109/FOCS.2019.00056.
- [12] Vida Dujmovic, Pat Morin, and David R. Wood. The structure of k-planar graphs. *CoRR*, abs/1907.05168, 2019. URL: <http://arxiv.org/abs/1907.05168>, arXiv:1907.05168.
- [13] Vida Dujmovic and David R. Wood. Stacks, queues and tracks: Layouts of graph subdivisions. *Discret. Math. Theor. Comput. Sci.*, 7(1):155–202, 2005. URL: <http://dmtcs.episciences.org/346>.
- [14] Zdenek Dvorák and Serguei Norine. Small graph classes and bounded expansion. *J. Comb. Theory, Ser. B*, 100(2):171–175, 2010. URL: <https://doi.org/10.1016/j.jctb.2009.06.001>, doi:10.1016/j.jctb.2009.06.001.
- [15] Jakub Gajarský, Stephan Kreutzer, Jaroslav Nesetril, Patrice Ossona de Mendez, Michal Pilipczuk, Sebastian Siebertz, and Szymon Toruńczyk. First-order interpretations of bounded expansion classes. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, pages 126:1–126:14, 2018. URL: <https://doi.org/10.4230/LIPIcs.ICALP.2018.126>, doi:10.4230/LIPIcs.ICALP.2018.126.
- [16] Cyril Gavoille and Arnaud Labourel. Shorter implicit representation for planar graphs and bounded treewidth graphs. In Lars Arge, Michael Hoffmann, and Emo Welzl, editors, *Algorithms - ESA 2007, 15th Annual European Symposium, Eilat, Israel, October 8-10, 2007, Proceedings*, volume 4698 of *Lecture Notes in Computer Science*, pages 582–593. Springer, 2007. URL: https://doi.org/10.1007/978-3-540-75520-3_52, doi:10.1007/978-3-540-75520-3_52.
- [17] Sylvain Guillemot and Dániel Marx. Finding small patterns in permutations in linear time. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 82–101, 2014. URL: <https://doi.org/10.1137/1.9781611973402.7>, doi:10.1137/1.9781611973402.7.
- [18] Sampath Kannan, Moni Naor, and Steven Rudich. Implicit representation of graphs. *SIAM J. Discret. Math.*, 5(4):596–603, 1992. URL: <https://doi.org/10.1137/0405049>, doi:10.1137/0405049.
- [19] Martin Klazar. The füredi-hajnal conjecture implies the stanley-wilf conjecture. In *Formal power series and algebraic combinatorics*, pages 250–255. Springer, 2000.
- [20] Adam Marcus and Gábor Tardos. Excluded permutation matrices and the stanley-wilf conjecture. *J. Comb. Theory, Ser. A*, 107(1):153–160, 2004. URL: <https://doi.org/10.1016/j.jcta.2004.04.002>, doi:10.1016/j.jcta.2004.04.002.
- [21] Dániel Marx. Efficient approximation schemes for geometric problems? In *Algorithms - ESA 2005, 13th Annual European Symposium, Palma de Mallorca, Spain, October 3-6, 2005, Proceedings*, pages 448–459, 2005. URL: https://doi.org/10.1007/11561071_41, doi:10.1007/11561071_41.
- [22] Dániel Marx and Ildikó Schlotter. Cleaning interval graphs. *Algorithmica*, 65(2):275–316, 2013. URL: <https://doi.org/10.1007/s00453-011-9588-0>, doi:10.1007/s00453-011-9588-0.
- [23] Jaroslav Nesetril and Patrice Ossona de Mendez. *Sparcity - Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and combinatorics*. Springer, 2012. URL: <https://doi.org/10.1007/978-3-642-27875-4>, doi:10.1007/978-3-642-27875-4.
- [24] Serguei Norine, Paul D. Seymour, Robin Thomas, and Paul Wollan. Proper minor-closed families are small. *J. Comb. Theory, Ser. B*, 96(5):754–757, 2006. URL: <https://doi.org/10.1016/j.jctb.2006.01.006>, doi:10.1016/j.jctb.2006.01.006.
- [25] Michal Pilipczuk and Sebastian Siebertz. Polynomial bounds for centered colorings on proper minor-closed graph classes. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1501–1520. SIAM, 2019. URL: <https://doi.org/10.1137/1.9781611975482.91>, doi:10.1137/1.9781611975482.91.
- [26] Jeremy P. Spinrad. *Efficient graph representations*, volume 19 of *Fields Institute monographs*. American Mathematical Society, 2003. URL: <http://www.ams.org/bookstore-getitem/item=fim-19>.
- [27] Andrew Suk. Coloring intersection graphs of x-monotone curves in the plane. *Combinatorica*, 34(4):487–505, 2014. URL: <https://doi.org/10.1007/s00493-014-2942-5>, doi:10.1007/s00493-014-2942-5.