

Twin-width can be exponential in treewidth

Édouard Bonnet   

Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Hugues Déprés  

Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Abstract

For any small positive real ε and integer $t > \frac{1}{\varepsilon}$, we build a graph with a vertex deletion set of size t to a tree, and twin-width greater than $2^{(1-\varepsilon)t}$. In particular, this shows that the twin-width is sometimes exponential in the treewidth, in the so-called oriented twin-width and grid number, and that adding an apex may multiply the twin-width by at least $2 - \varepsilon$. Except for the one in oriented twin-width, these lower bounds are essentially tight.

1 Introduction

Twin-width is a graph parameter introduced by Bonnet, Kim, Thomassé, and Watrigant [12]. It is defined by means of trigraphs. A *trigraph* is a graph with some edges colored black, and some colored red. A (vertex) *contraction* consists of merging two (non-necessarily adjacent) vertices, say, u, v into a vertex w , and keeping every edge wz black if and only if uz and vz were previously black edges. The other edges incident to w become red (if not already), and the rest of the trigraph remains the same. A *contraction sequence* of an n -vertex graph G is a sequence of trigraphs $G = G_n, \dots, G_1 = K_1$ such that G_i is obtained from G_{i+1} by performing one contraction. A *d-sequence* is a contraction sequence in which every vertex of every trigraph has at most d red edges incident to it. The *twin-width* of G , denoted by $\text{tw}(G)$, is then the minimum integer d such that G admits a d -sequence. Figure 1 gives an example of a graph with a 2-sequence, i.e., of twin-width at most 2. Twin-width can be naturally




Figure 1 A 2-sequence witnessing that the initial graph has twin-width at most 2.

extended to matrices (unordered [12] or ordered [9]) over a finite alphabet, and hence to any binary structures. Classes of binary structures with bounded twin-width include graphs with bounded treewidth, bounded clique-width, K_t -minor free graphs, posets with antichains of bounded size, strict subclasses of permutation graphs, map graphs, bounded-degree string graphs [12], segment graphs with no $K_{t,t}$ subgraph, visibility graphs of 1.5D terrains without large half-graphs, visibility graphs of simple polygons without large independent sets [6], as well as $\Omega(\log n)$ -subdivisions of n -vertex graphs, classes with bounded queue number or bounded stack number, and some classes of cubic expanders [7].

Despite their apparent generality, classes of bounded twin-width are small [7], χ -bounded [8], even quasi-polynomially χ -bounded [17], preserved (albeit with a higher upper bound) by first-order transductions [12], and by the usual graph products when one graph has bounded degree [16, 7], have VC density 1 [11, 19], admit, when $O(1)$ -sequences are given, a fixed-parameter tractable first-order model checking [12], an (almost) single-exponential

parameterized algorithm for various problems that are W[1]-hard in general [8], as well as a parameterized fully-polynomial linear algorithm for counting triangles [15], an (almost) linear representation [18], a stronger regularity lemma [19], etc.

In all these applications, the upper bound on twin-width, although somewhat hidden in the previous paragraph, plays a role. There is then an incentive to obtain as low as possible upper bounds on particular classes of bounded twin-width. To give one concrete algorithmic example, an independent set of size k can be found in time $O(k^2 d^{2k} n)$ in an n -vertex graph given with a d -sequence [8]. This is relatively practical for moderate values of k , with the guarantee that d is below 10, but not when d is merely upperbounded by 10^{10} . Another motivating example: triangle-free graphs of twin-width at most d are $d + 2$ -colorable [8], a stronger fact in the former case than in the latter.

In that line of work, Balabán and Hliněný show that posets of width k (i.e., with antichains of size at most k) have twin-width at most $9k$ [2]. Unit interval graphs have twin-width at most 2 [8], and proper k -mixed-thin graphs (a recently proposed generalization of unit interval graphs) have twin-width $O(k)$ [3]. Every graph obtained by subdividing at least $2 \log n$ (throughout the paper, all logs are in base 2) times each edge of an n -vertex graph has twin-width at most 4 [4]. Schidler and Szeider report the (exact) twin-width of a collection of graphs [20], obtained via SAT encodings. Jacob and Pilipczuk [14] give the current best upper bound of 183 on the twin-width of planar graphs, while graphs with genus g have twin-width $O(g)$ [13]. Most relevant to our paper, for every graph G , $\text{tw}(G) \leq 3 \cdot 2^{\text{tw}(G)-1}$ [14], where $\text{tw}(G)$ denotes the treewidth of G .

Conversely, one may ask the following.

► **Question 1.** *What is the largest twin-width a graph of treewidth k can have?*

A lower bound of $\Omega(k)$ comes from the existence of n -vertex graphs with twin-width $\Omega(n)$ (since the treewidth is trivially upperbounded by $n - 1$). This is almost surely the case of graphs drawn from $G(n, 1/2)$. Alternatively, the n -vertex Paley graph (for a prime n such that $n \equiv 1 \pmod{4}$) has precisely twin-width $(n - 1)/2$ [1]. Another example to derive the linear lower bound is the power set graph [14]. Improving on this lower bound is not obvious, and $\Theta(k)$ is indeed the answer to Question 1 within the class of planar graphs [14], or when replacing 'treewidth' by 'cliquewidth' or 'pathwidth'.

When switching 'twin-width' and 'treewidth' in Question 1, the gap is basically as large as possible: There are n -vertex graphs with treewidth $\Omega(n)$ and twin-width at most 6, in the iterated 2-lifts of K_4 [7, 5].

An important characterization of bounded twin-width is via the absence of complex divisions of an adjacency matrix. A matrix has a k -mixed minor if its row (resp. column) set can be partitioned into k sets of consecutive rows (resp. columns), such that each of the k^2 cells defined by this k -division contains at least two distinct rows and at least two distinct columns. The *mixed number of a matrix* M is the largest integer k such that M admits a k -mixed minor. The *mixed number of a graph* G , denoted by $\text{mxn}(G)$, is the minimum, taken among all the adjacency matrices M of G , of the mixed number of M . The following was shown.

► **Theorem 1** ([12]). *For every graph G , $(\text{mxn}(G) - 1)/2 \leq \text{tw}(G) \leq 2^{2^{O(\text{mxn}(G))}}$.*

In sparse graphs (here, excluding a fixed $K_{t,t}$ as a subgraph), the previous theorem is both simpler to formulate and has a better dependency. A matrix has a k -grid minor if it has a k -division with at least one 1-entry in each of its k^2 cells. The *grid number* of a matrix and of a graph G , denoted by $\text{gn}(G)$, are defined analogously to the previous paragraph. We

only state the inequality that is useful to bound the twin-width of a sparse class, but is valid in general.

► **Theorem 2** (follows from [12]). *For every graph G , $\text{tw}(G) \leq 2^{O(\text{gn}(G))}$.*

Theorems 1 and 2 allow to bound the twin-width of a class $\mathcal{C}$ by exhibiting, for every $G \in \mathcal{C}$, an adjacency matrix of G without large mixed or grid minor. Therefore one merely has to order $V(G)$ (the vertex set of G) in an appropriate way. The double (resp. simple) exponential dependency in mixed number (resp. grid number) implies relatively weak twin-width upper bounds. For several classes whose twin-width was originally upperbounded via Theorem 1, better bounds were later given by avoiding this theorem (see [7, 2, 14, 13, 4]). Still for some geometric graph classes, bypassing Theorem 1 seems complicated (see [6]). And in general (since this theorem is at the basis of several other applications, see for instance [7, 8, 9]) it would help to have an improved upper bound of $\text{tw}(G)$; in particular a negative answer to the following question.

► **Question 2.** *Is twin-width sometimes exponential in mixed and grid number?*

A variant of twin-width, called *oriented twin-width*, adds an orientation to the red edges (see [10]). The red edge (arc) is oriented away from the contracted vertex. The *oriented twin-width* d of a graph G , denoted by $\text{otw}(G)$, is then defined similarly as twin-width by tolerating more than d red arcs incident to a vertex, as long as at most d of them are out-going. Rather surprisingly twin-width and oriented twin-width are tied.

► **Theorem 3** ([10]). *For every graph G , $\text{otw}(G) \leq \text{tw}(G) \leq 2^{O(\text{otw}(G))}$.*

Classic results show that planar graphs have oriented twin-width at most 9 [10]. Thus it would be appreciable to lower the dependency of $\text{tw}(G)$ in $\text{otw}(G)$.

► **Question 3.** *Is twin-width sometimes exponential in oriented twin-width?*

An elementary argument shows that when adding an apex (i.e., an additional vertex with an arbitrary neighborhood) to a graph G , the twin-width of the obtained graph is at most $2 \cdot \text{tw}(G) + 1$. Again it is not clear whether this increase could be made smaller.

► **Question 4.** *Does twin-width sometimes essentially double when an apex is added?*

Note that Question 1 is asked by Jacob and Pilipczuk [14], and Question 3 is posed by Bonnet et al. [10], and is closely related to Question 2.

Our contribution.

With a single construction, we answer all these questions. The answer to Questions 2, 3, and 4 is affirmative, while the answer to Question 1 is $2^{\Theta(k)}$, which confirms the intuition of the authors of [14]. More precisely, we show the following.

► **Theorem 4.** *For every real $0 < \varepsilon \leq 1/2$ and integer $t > 1/\varepsilon$, there is a graph $G_{t,\varepsilon}$ with a feedback vertex set of size t and such that $\text{tw}(G_{t,\varepsilon}) > 2^{(1-\varepsilon)t}$.*

The graph $G_{t,\varepsilon}$ has in particular treewidth at most $t + 1$, grid number at most $t + 2$, and oriented twin-width at most $t + 1$. Thus

- $\text{tw}(G_{t,\varepsilon}) > 2^{(1-\varepsilon)(\text{tw}(G_{t,\varepsilon})-1)}$,
- $\text{tw}(G_{t,\varepsilon}) > 2^{(1-\varepsilon)(\text{gn}(G_{t,\varepsilon})-2)}$, and
- $\text{tw}(G_{t,\varepsilon}) > 2^{(1-\varepsilon)(\text{otw}(G_{t,\varepsilon})-1)}$.

Hence Theorem 4 has the following consequences.

► **Corollary 5.** *For every small $\varepsilon > 0$, there is a family $\mathcal{F}$ of graphs with unbounded twin-width such that for every $G \in \mathcal{F}$: $\text{tw}(G) > 2^{(1-\varepsilon)(\text{tw}(G)-1)}$.*

Up to multiplicative factors, it matches the known upper bound [14, 12], and essentially settles Question 1. The following answers Question 2.

► **Corollary 6.** *For every small $\varepsilon > 0$, there is a family $\mathcal{F}$ of graphs with unbounded twin-width such that for every $G \in \mathcal{F}$: $\text{tw}(G) > 2^{(1-\varepsilon)(\text{gn}(G)-2)}$.*

The following answers Question 3.

► **Corollary 7.** *For every small $\varepsilon > 0$, there is a family $\mathcal{F}$ of graphs with unbounded twin-width such that for every $G \in \mathcal{F}$: $\text{tw}(G) > 2^{(1-\varepsilon)(\text{otw}(G)-1)}$.*

The following answers Question 4.

► **Corollary 8.** *For every small $\varepsilon > 0$, there is a family $\mathcal{F}$ of graphs with unbounded twin-width such that for every $G \in \mathcal{F}$: $\text{tw}(G) > (2 - \varepsilon)\text{tw}(G - \{v\})$, where v is a single vertex of G .*

We leave as an open question if the twin-width upper bound in oriented twin-width and mixed number can be made single-exponential.

2 Preliminaries

For i and j two integers, we denote by $[i, j]$ the set of integers that are at least i and at most j . For every integer i , $[i]$ is a shorthand for $[1, i]$. We use the standard graph-theoretic notations: $V(G)$ denotes the vertex set of a graph G , $E(G)$ denotes its edge set, $G[S]$ denotes the subgraph of G induced by S , etc.

We give an alternative approach to contraction sequences. The *twin-width* of a graph, introduced in [12], can be defined in the following way (complementary to the one given in introduction). A *partition sequence* of an n -vertex graph G , is a sequence $\mathcal{P}_n, \dots, \mathcal{P}_1$ of partitions of its vertex set $V(G)$, such that $\mathcal{P}_n$ is the set of singletons $\{\{v\} : v \in V(G)\}$, $\mathcal{P}_1$ is the singleton set $\{V(G)\}$, and for every $2 \leq i \leq n$, $\mathcal{P}_{i-1}$ is obtained from $\mathcal{P}_i$ by merging two of its parts into one. Two parts P, P' of a same partition $\mathcal{P}$ of $V(G)$ are said *homogeneous* if either every pair of vertices $u \in P, v \in P'$ are non-adjacent, or every pair of vertices $u \in P, v \in P'$ are adjacent. Two non-homogeneous parts are also said *red-adjacent*. The *red degree* of a part $P \in \mathcal{P}$ is the number of other parts of $\mathcal{P}$ which are red-adjacent to P . Finally the twin-width of G , denoted by $\text{tw}(G)$, is the least integer d such that there is a partition sequence $\mathcal{P}_n, \dots, \mathcal{P}_1$ of G with every part of every $\mathcal{P}_i$ ($1 \leq i \leq n$) having red degree at most d .

The definition of the previous paragraph is equivalent to the one given in introduction, via contraction sequences. Indeed the trigraph G_i is obtained from partition $\mathcal{P}_i$, by having one vertex per part of $\mathcal{P}_i$, a black edge between any fully adjacent pair of parts, and a red edge between red-adjacent parts. A *partial contraction sequence* is a sequence of trigraphs $G_n, \dots, G_i$, for some $i \in [n]$. A (full) *contraction sequence* is one such that $i = 1$. We naturally consider the trigraph G_j to come *after* (resp. *before*) $G_{j'}$ if $j < j'$ (resp. $j > j'$). Thus when we write *the first trigraph of the sequence $\mathcal{S}$ to satisfy X* (or *the first time a trigraph of $\mathcal{S}$ satisfies X*) we mean the trigraph G_j with largest index j among those satisfying X . The same goes for partition sequences.

If u is a vertex of a trigraph H , then $u(G)$ denotes the set of vertices of G eventually contracted into u in H . We denote by $\mathcal{P}_G(H)$ (and $\mathcal{P}(H)$ when G is clear from the context) the partition $\{u(G) : u \in V(H)\}$ of $V(G)$. We may refer to a *part* of H as any set in $\{u(G) : u \in V(H)\}$. We may also refer to a *part* of a contraction/partition sequence as any part of one its trigraphs/partitions. A contraction *involves* a vertex v if it produces a new part (of size at least 2) containing v . In general, we use trigraphs and partitioned graphs somewhat interchangeably, when one notion appears more convenient than the other.

3 Proof of Theorem 4

We fix once and for all, $0 < \varepsilon \leq 1/2$, a possibly arbitrarily small positive real. We build for every integer $t > 1/\varepsilon$, a graph $G_{t,\varepsilon}$, that we shorten to G_t . We set

$$f(t) = \left\lceil 2 + C_t 2^{(1-\varepsilon)t(2+C_t(2^{(1-\varepsilon)t}+1))} \right\rceil$$

where $C_t = 2^{(1-\varepsilon)t}/\varepsilon$.

Construction of G_t . Let T be the full 2^t -ary tree of depth $f(t)$, i.e., with root-to-leaf paths on $f(t)$ edges. Let X be a set of t vertices, that we may identify to $[t]$. The vertex set of G_t is $X \uplus V(T)$. The edges of G_t are such that $G[X]$ is an independent set, and $G[V(T)] = T$. The edges between $V(T)$ and X are such that

- the root of T has no neighbor in X , and
 - the 2^t children (in T) of every internal node of T each have a distinct neighborhood in X .
- Note that this defines a single graph up to isomorphism. By a slight abuse of language, we may utilize the usual vocabulary on trees directly on G_t . By *root*, *internal node*, *child*, *parent*, *leaf* of G_t , we mean the equivalent in T .

We start with this straightforward observation.

► **Lemma 9.** G_t has treewidth at most $t + 1$.

Proof. The set X is a feedback vertex set of G_t of size t , thus $\text{tw}(G_t) \leq \text{fvs}(G_t) + 1 \leq t + 1$. ◀

The following is the core lemma, which occupies us for the remainder of the section.

► **Lemma 10.** G_t has twin-width greater than $2^{(1-\varepsilon)t}$.

Proof. We assume, by way of contradiction, that G_t admits a d -sequence with $d \leq 2^{(1-\varepsilon)t}$. We consider the partial d -sequence $\mathcal{S}$, starting at G_t , and ending right before the first contraction involving a child of the root. We first show that no vertex of X can be involved in a contraction of $\mathcal{S}$. Note that it implies, in particular, that the root cannot be involved in a contraction of $\mathcal{S}$.

▷ **Claim 11.** No part of $\mathcal{S}$ contains more than one vertex of X .

PROOF OF THE CLAIM: Observe that, for every $i \neq j \in [t]$, there are 2^{t-1} sets of $2^{[t]}$ containing exactly one of i, j : 2^{t-2} only contain i , and 2^{t-2} only contain j . Recall now that by assumption, in every trigraph of $\mathcal{S}$, every child of the root is alone in its part. Thus a part P of $\mathcal{S}$ such that $|P \cap X| \geq 2$ would have red degree at least $2^{t-1} > 2^{(1-\varepsilon)t} \geq d$. ◊

▷ **Claim 12.** No part of $\mathcal{S}$ intersects both X and $V(T)$.

PROOF OF THE CLAIM: For the sake of contradiction, consider the first occurrence of a part $P \supseteq \{x, v\}$ with $x \in X$ and $v \in V(T)$. Vertex x is adjacent to half of the children of the root, whereas v is adjacent to at most one of them, or all of them (if v is itself the root). In both cases, this entails at least $2^{t-1} - 1$ red edges for P towards children of the root. If v is not a grandchild of the root, the red degree of P is at least 2^{t-1} . We thus assume that v is a grandchild of the root.

As $t \geq 2$, there is a $y \in X \setminus \{x\}$. Let v' be the child of v whose neighborhood in X is exactly $\{y\}$. This vertex exists since $f(t) \geq 3$. If P contains v' , P is also red-adjacent to $\{y\}$ (indeed a part, by Claim 11). If instead, P does not contain v' , then P is also red-adjacent to the part containing v' .

Thus, in any case, the red degree of P is at least $2^{t-1} > 2^{(1-\varepsilon)t} \geq d$. $\diamond$

From Claims 11 and 12, we immediately obtain:

▷ **Claim 13.** Every part of $\mathcal{S}$ intersecting X is a singleton.

Crucial to the proof, we introduce two properties $\mathcal{P}$, and later $\mathcal{Q}$, on internal nodes $v \in V(T)$ in trigraphs $H \in \mathcal{S}$. Property $\mathcal{P}$ is defined by

$$\mathcal{P}(v, H) = \text{“At least } 2^{\varepsilon t} \text{ children of } v \text{ are in the same part of } \mathcal{P}(H)\text{.”}$$

We first remark that any internal node in a non-singleton part verifies $\mathcal{P}$.

▷ **Claim 14.** Let H be any trigraph of $\mathcal{S}$ and v be any internal node of T whose part in $\mathcal{P}(H)$ is not a singleton. Then $\mathcal{P}(v, H)$ holds.

PROOF OF THE CLAIM: Let P be the part of v (i.e., the one containing v) in $\mathcal{P}(H)$, and $u \in P \setminus \{v\}$. At least $2^t - 1$ children of v are not adjacent to u . Thus these $2^t - 1$ vertices have to be in at most $d + 1 \leq 2^{(1-\varepsilon)t} + 1$ parts. These parts are part P , plus at most d parts linked to P by a red edge. Since $(2^{\varepsilon t} - 1)(2^{(1-\varepsilon)t} + 1) < 2^t - 1$ (recall that $\varepsilon < 1/2$), one of these parts (possibly P) contains at least $2^{\varepsilon t}$ children of v . $\diamond$

As the merge of a singleton part $\{v\}$ with any other part does not change the intersections of parts with the set of children of v , we get a slightly stronger claim.

▷ **Claim 15.** Let v be an internal node of T , and H be the last trigraph of $\mathcal{S}$ for which v is in a singleton part of $\mathcal{P}(H)$. Then $\mathcal{P}(v, H)$ holds.

A *preleaf* is an internal node of T adjacent to a leaf, i.e., the parent of some leaves. We obtain the following as a direct consequence of Claim 14.

▷ **Claim 16.** In any trigraph $H \in \mathcal{S}$, any non-preleaf internal node $v \in V(T)$ that verifies $\mathcal{P}(v, H)$ has at least $2^{\varepsilon t}$ children u verifying $\mathcal{P}(u, H)$.

We define the property $\mathcal{Q}$ on internal nodes v of T and trigraphs $H \in \mathcal{S}$ by induction:

$$\mathcal{Q}(v, H) = \begin{cases} \mathcal{P}(v, H) & \text{if } v \text{ is a preleaf, and otherwise} \\ \mathcal{Q}(u_1, H) \wedge \mathcal{Q}(u_2, H) & \text{for some pair } u_1 \neq u_2 \text{ of children of } v. \end{cases}$$

That is, $\mathcal{Q}$ is defined as $\mathcal{P}$ for preleaves, and otherwise, $\mathcal{Q}$ holds when it holds for at least two of its children. Observe that $\mathcal{P}$ and $\mathcal{Q}$ are monotone in the following sense: If $\mathcal{P}(v, H)$ (resp. $\mathcal{Q}(v, H)$) holds, then $\mathcal{P}(v, H')$ (resp. $\mathcal{Q}(v, H')$) holds for every subsequent trigraph H' of the partial d -sequence $\mathcal{S}$. We may write that v *satisfies* $\mathcal{P}$ (resp. $\mathcal{Q}$) *in* H when $\mathcal{P}(v, H)$ (resp. $\mathcal{Q}(v, H)$) holds, and may add *for the first time* if no trigraph $H' \in \mathcal{S}$ before H is such that $\mathcal{P}(v, H')$ (resp. $\mathcal{Q}(v, H')$) holds.

Let B the vertices $w \in V(F)$ such that $w(G)$ contains at least $2^{\varepsilon t}$ children of the same node of T . Each vertex of B is red-adjacent to at least $\log(2^{\varepsilon t}) = \varepsilon t$ (singleton) parts of X . Therefore, since the red degree of (singleton) parts of X is at most $2^{(1-\varepsilon)t}$:

$$|B| \leq \frac{2^{(1-\varepsilon)t}}{\varepsilon}.$$

Next we show that there is relatively large set of vertices of F each corresponding to a non-singleton part that contains an internal node of T .

▷ **Claim 19.** There is a set $B' \subseteq V(F)$ of size at least

$$\frac{1}{(1-\varepsilon)t} \log \left(\frac{f(t)-2}{|B|} \right) - 1$$

such that for every $b \in B'$ there is an internal node v of T with $v \in b(G_t)$ and $|b(G_t)| \geq 2$.

PROOF OF THE CLAIM: Let $s := \frac{1}{(1-\varepsilon)t} \log \left(\frac{f(t)-2}{|B|} \right) - 1$. Our goal is to construct a sequence $(b_i)_{i \in [0, s]}$ of distinct vertices of F such that for every $i \in [s]$,

part $b_i(G_t)$ is *not* a singleton and contains an internal node of T . (1)

We first focus on finding b_0 . Note that b_0 need not satisfy Invariant (1), but will be chosen to force the existence of b_1 itself satisfying (1) and starting the induction.

Let $Q := \{q_j : 0 \leq j \leq f(t) - 3\} \subset V(T)$ be as described in Claim 18. Every $q_j \in Q$ has (at least) one descendant q'_j that is a preleaf and satisfies $\mathcal{Q}$, hence $\mathcal{P}$, in F . The q'_j s are pairwise distinct because no two nodes of Q are in an ancestor-descendant relationship. We set $Q' := \{q'_j : 0 \leq j \leq f(t) - 3\}$.

Now for every q'_j , at least $2^{\varepsilon t}$ of its children are in the same part of $\mathcal{P}(F)$; hence, this part corresponds to a vertex in B . By the pigeonhole principle, there is a $b_0 \in B$ that contains at least $2^{\varepsilon t}$ children of at least $(f(t) - 2)/|B|$ nodes of Q' .

For each b_i , we define $Q_i \subset Q$ as the set of vertices q_j such that

- $b_i(G_t)$ contains a (not necessarily strict) descendant z of q_j , and
- no part $b_{i'}(G_t)$ with $i' < i$ contains a node on the path between q_j and z in T .

Thus $|Q_0| \geq (f(t) - 2)/|B|$.

We now assume that $b_i \in V(F)$, for some $0 \leq i < s$, has been found with

$$|Q_i| \geq \frac{f(t) - 2}{|B| \cdot 2^{i(1-\varepsilon)t}}. \quad (2)$$

Observe that Q_0 satisfies (2). We construct b_{i+1}, Q_{i+1} satisfying the invariants (1) and (2).

For each $q_j \in Q_i$, consider the highest descendant z_j of q_j in $b_i(G_t)$, and z'_j the parent of z_j in T . By construction, the part P_j of $\mathcal{P}(F)$ containing z'_j is not a $b_k(G_t)$ for any $k \leq i$. Part P_j is linked to $b_i(G_t)$ by a red edge. Therefore there are at most $2^{(1-\varepsilon)t}$ such parts P_j . In particular, there is a $b_{i+1} \in V(F)$ such that $b_{i+1}(G_t)$ contains at least

$$\frac{|Q_i|}{d} \geq \frac{f(t) - 2}{|B| \cdot 2^{i(1-\varepsilon)t}} \cdot \frac{1}{2^{(1-\varepsilon)t}} = \frac{f(t) - 2}{|B| \cdot 2^{(i+1)(1-\varepsilon)t}}$$

parents z'_j of highest descendants z_j .

Remark that $b_{i+1}(G_t)$ has size at least two while $(f(t) - 2)/(|B| \cdot 2^{(i+1)(1-\varepsilon)t}) > 1$, which holds since $i < s$. Thus $b_{i+1}(G_t)$ does not contain any parent v_j of a q_j (since the v_j s are in singleton parts). In particular, $|Q_{i+1}| \geq (f(t) - 2)/(|B| \cdot 2^{(i+1)(1-\varepsilon)t})$, and b_{i+1}, Q_{i+1} satisfy (1) and (2).

Finally, the set $B' := \{b_i : 1 \leq i \leq s\}$ has the required properties. $\diamond$

We can now finish the proof of the lemma.

For every $b_i \in B'$, let $u_i \in b_i(G_t)$ be an internal node of T . As $b_i(G_t) \geq 2$, u_i satisfies $\mathcal{P}$ in F . This implies that b_i or a red neighbor of b_i is in B . Therefore, the total number of red edges incident to a vertex of B is at least $|B'| - |B|$. Thus there is a vertex in B with red degree at least $(|B'| - |B|)/|B|$. This is a contradiction since

$$\begin{aligned} \frac{|B'| - |B|}{|B|} &= \frac{|B'|}{|B|} - 1 \geq \left(\frac{1}{(1-\varepsilon)t} \log \left(\frac{f(t) - 2}{|B|} \right) - 1 \right) \cdot \frac{1}{|B|} - 1 \\ &\geq \left(\frac{1}{(1-\varepsilon)t} \log \left(2^{(1-\varepsilon)t(2+C_t \cdot (2^{(1-\varepsilon)t} + 1))} \right) - 1 \right) \cdot \frac{1}{|B|} - 1 \\ &= \left((2 + C_t \cdot (2^{(1-\varepsilon)t} + 1)) - 1 \right) \cdot \frac{1}{|B|} - 1 > 2^{(1-\varepsilon)t} + 1 - 1 = 2^{(1-\varepsilon)t} \geq d. \end{aligned}$$

since, we recall, $f(t) = \left\lceil 2 + C_t \cdot 2^{(1-\varepsilon)t(2+C_t \cdot (2^{(1-\varepsilon)t} + 1))} \right\rceil$ and $C_t = \frac{2^{(1-\varepsilon)t}}{\varepsilon} \geq |B|$. $\blacktriangleleft$

Since X is a feedback vertex set of size t of G_t , Lemma 10 implies Theorem 4, and hence Corollary 5.

As the twin-width of T is 2, adding the t apices in X , multiplies the twin-width by at least $2^{t(1-\varepsilon-\frac{1}{t})}$. Thus one apex in X multiplies the twin-width by at least $2^{1-\varepsilon-\frac{1}{t}}$, which can be made arbitrarily close to 2. This establishes Corollary 8.

4 Oriented twin-width and grid number

In this section, we check that G_t has oriented twin-width at most $t + 1$, and grid number at most $t + 2$.

A *(partial) oriented contraction sequence* is defined similarly as a *(partial) contraction sequence* with every red edge replaced by a red arc leaving the newly contracted vertex. Then a *(partial) oriented d -sequence* is such that all the vertices of all its *ditrigraphs* have at most d out-going red arcs. The *oriented twin-width* of a graph G , denoted by $\text{otww}(G)$, is the minimum integer d such that G admits an oriented d -sequence.

► **Lemma 20.** *The oriented twin-width of G_t is at most $t + 1$.*

Proof. We observe that the 2-sequence for trees [12] is an oriented 1-sequence. We contract T to a single vertex (without touching X) in that manner. This yields a partial oriented $t + 1$ -sequence for G_t ending on a $t + 1$ -vertex ditrigraph, which can be contracted in any way. This contraction sequence witnesses that $\text{otww}(G_t) \leq t + 1$. $\blacktriangleleft$

Thus Corollary 7 holds.

We finish by establishing Corollary 6.

► **Lemma 21.** *The grid number of G_t is at most $t + 2$.*

Proof. Recall that $V(G_t) = X \uplus V(T)$. Let $\prec$ be the total order on $V(G_t)$ that puts first all the vertices of X in any order, then from left to right, all the leaves of T , followed by the preleaves, the nodes at depth $f(t) - 2$, the nodes at depth $f(t) - 3$, and so on, up to the root. We denote by M the adjacency matrix of G_t ordered by $\prec$.

Let M_T be the submatrix of M obtained by deleting the t rows and t columns corresponding to X . Note that the grid number of M is at most $\text{gn}(M_T) + t$. We claim that there is no 3-grid minor in M_T .

Indeed, in the order $\prec$, above the diagonal of M_T there is no pair of 1-entries in strictly decreasing positions. Thus overall there is no triple of 1-entries in strictly decreasing positions. Thus no 3-grid minor is possible in M_T . $\blacktriangleleft$

References

- 1 Jungho Ahn, Kevin Hendrey, Donggyu Kim, and Sang-il Oum. Bounds for the twin-width of graphs. *CoRR*, abs/2110.03957, 2021. URL: <https://arxiv.org/abs/2110.03957>, arXiv:2110.03957.
- 2 Jakub Balabán and Petr Hlinený. Twin-width is linear in the poset width. In Petr A. Golovach and Meirav Zehavi, editors, *16th International Symposium on Parameterized and Exact Computation, IPEC 2021, September 8-10, 2021, Lisbon, Portugal*, volume 214 of *LIPIcs*, pages 6:1–6:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.IPEC.2021.6.
- 3 Jakub Balabán, Petr Hlinený, and Jan Jedelský. Twin-width and transductions of proper k-mixed-thin graphs. *CoRR*, abs/2202.12536, 2022. URL: <https://arxiv.org/abs/2202.12536>, arXiv:2202.12536.
- 4 Pierre Bergé, Édouard Bonnet, and Hugues Déprés. Deciding twin-width at most 4 is NP-complete. *CoRR*, abs/2112.08953, 2021. URL: <https://arxiv.org/abs/2112.08953>, arXiv:2112.08953.
- 5 Yonatan Bilu and Nathan Linial. Lifts, discrepancy and nearly optimal spectral gap*. *Combinatorica*, 26(5):495–519, 2006. doi:10.1007/s00493-006-0029-7.
- 6 Édouard Bonnet, Dibyayan Chakraborty, Eun Jung Kim, Noleen Köhler, Raul Lopes, and Stéphan Thomassé. Twin-width VIII: delineation and win-wins. *CoRR*, abs/2204.00722, 2022. arXiv:2204.00722, doi:10.48550/arXiv.2204.00722.
- 7 Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width II: small classes. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1977–1996, 2021. doi:10.1137/1.9781611976465.118.
- 8 Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width III: max independent set, min dominating set, and coloring. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, pages 35:1–35:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.ICALP.2021.35.
- 9 Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, Pierre Simon, Stéphan Thomassé, and Szymon Toruńczyk. Twin-width IV: ordered graphs and matrices. *CoRR*, abs/2102.03117, 2021, accepted at STOC 2022. URL: <https://arxiv.org/abs/2102.03117>, arXiv:2102.03117.
- 10 Édouard Bonnet, Eun Jung Kim, Amadeus Reinald, and Stéphan Thomassé. Twin-width VI: the lens of contraction sequences. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1036–1056. SIAM, 2022.
- 11 Édouard Bonnet, Eun Jung Kim, Amadeus Reinald, Stéphan Thomassé, and Rémi Watrigant. Twin-width and polynomial kernels. In Petr A. Golovach and Meirav Zehavi, editors, *16th International Symposium on Parameterized and Exact Computation, IPEC 2021, September 8-10, 2021, Lisbon, Portugal*, volume 214 of *LIPIcs*, pages 10:1–10:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.IPEC.2021.10.
- 12 Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. *J. ACM*, 69(1):3:1–3:46, 2022. doi:10.1145/3486655.

- 13 Édouard Bonnet, O-joung Kwon, and David R. Wood. Reduced bandwidth: a qualitative strengthening of twin-width in minor-closed classes (and beyond). *CoRR*, abs/2202.11858, 2022. URL: <https://arxiv.org/abs/2202.11858>, arXiv:2202.11858.
- 14 Hugo Jacob and Marcin Pilipczuk. Bounding twin-width for bounded-treewidth graphs, planar graphs, and bipartite graphs. *CoRR*, abs/2201.09749, 2022. URL: <https://arxiv.org/abs/2201.09749>.
- 15 Stefan Kratsch, Florian Nelles, and Alexandre Simon. On triangle counting parameterized by twin-width. *CoRR*, abs/2202.06708, 2022. URL: <https://arxiv.org/abs/2202.06708>, arXiv:2202.06708.
- 16 William Pettersson and John Sylvester. Bounds on the twin-width of product graphs. *CoRR*, abs/2202.11556, 2022. URL: <https://arxiv.org/abs/2202.11556>, arXiv:2202.11556.
- 17 Michal Pilipczuk and Marek Sokolowski. Graphs of bounded twin-width are quasi-polynomially χ -bounded. *CoRR*, abs/2202.07608, 2022. URL: <https://arxiv.org/abs/2202.07608>, arXiv:2202.07608.
- 18 Michal Pilipczuk, Marek Sokolowski, and Anna Zych-Pawlewicz. Compact representation for matrices of bounded twin-width. In Petra Berenbrink and Benjamin Monmege, editors, *39th International Symposium on Theoretical Aspects of Computer Science, STACS 2022, March 15-18, 2022, Marseille, France (Virtual Conference)*, volume 219 of *LIPICs*, pages 52:1–52:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.STACS.2022.52.
- 19 Wojciech Przybyszewski. VC-density and abstract cell decomposition for edge relation in graphs of bounded twin-width, 2022. URL: <https://arxiv.org/abs/2202.04006>, doi:10.48550/ARXIV.2202.04006.
- 20 André Schidler and Stefan Szeider. A SAT approach to twin-width. *CoRR*, abs/2110.06146, 2021, accepted at ALENEX 2022. URL: <https://arxiv.org/abs/2110.06146>, arXiv:2110.06146.