

Deciding twin-width at most 4 is NP-complete

Pierre Bergé

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

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

We show that determining if an n -vertex graph has twin-width at most 4 is NP-complete, and requires time $2^{\Omega(n/\log n)}$ unless the Exponential-Time Hypothesis fails. Along the way, we give an elementary proof that n -vertex graphs subdivided at least $2\log n$ times have twin-width at most 4. We also show how to encode trigraphs H (2-edge colored graphs involved in the definition of twin-width) into graphs G , in the sense that every d -sequence (sequence of vertex contractions witnessing that the twin-width is at most d) of G inevitably creates H as an induced subtrigraph, whereas there exists a partial d -sequence that actually goes from G to H . We believe that these facts and their proofs can be of independent interest.


2012 ACM Subject Classification Theory of computation \rightarrow Graph algorithms analysis; Theory of computation \rightarrow Fixed parameter tractability

Keywords and phrases Twin-width, lower bounds

1 Introduction

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 stays 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 where all the trigraphs have red degree at most d . The *twin-width* of G , denoted by $tw(G)$, is then the minimum integer d such that G admits a d -sequence. See Figure 1 for an example of a graph admitting a 2-sequence. The *red graph* of a trigraph is obtained by simply deleting its black edges. A *partial d -sequence* is similar to a d -sequence but ends on any trigraph G_i , instead of on the 1-vertex (tri)graph G_1 . Twin-width can be naturally extended to matrices over a finite alphabet (in an unordered [6], or an ordered setting [4]), and hence to any binary structure.

Surprisingly many classes turn out to be of bounded twin-width. Such is the case of graphs with bounded clique-width, H -minor free graphs for any fixed H , posets with antichains of bounded size, strict subclasses of permutation graphs, map graphs, bounded-degree string graphs [6], as well as $\Omega(\log n)$ -subdivisions of n -vertex graphs, and some classes of cubic



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

expanders [3]. One of the main algorithmic interests with twin-width is that first-order (FO) model checking, that is, deciding if a sentence φ holds in a graph G , can be decided in fixed-parameter time (FPT) $f(|\varphi|, d) \cdot |V(G)|$ for some computable function f , when given a d -sequence of G [6]. As for most classes known to have bounded twin-width, one can compute $O(1)$ -sequences in polynomial time for members of the class, the latter result unifies and extends several known results [14, 17, 18, 16, 21] for hereditary (but not necessarily monotone) classes.

For monotone (i.e., subgraph-closed) classes, the FPT algorithm of Grohe, Kreutzer, and Siebertz [20] for FO model checking on nowhere dense classes, is complemented by W[1]-hardness on classes that are somewhere dense (i.e., *not* nowhere dense) [11], and even AW[*]-hardness on classes that are *effectively* somewhere dense [9]. The latter results mean that, for monotone classes, FO model checking is unlikely to be FPT beyond nowhere dense classes.

The missing piece for an FO model-checking algorithm in FPT time on any class of bounded twin-width is a polynomial-time algorithm and a computable function f , that given a constant integer bound c and a graph G , either finds an $f(c)$ -sequence for G , or correctly reports that the twin-width of G is greater than c . The running time of the algorithm could be $n^{g(c)}$, for some function g . However to get an FPT algorithm in the combined parameter *size of the sentence + bound on the twin-width*, one would further require that the approximation algorithm takes FPT time in c (now seen as a parameter), i.e., $g(c)n^{O(1)}$. We know such an algorithm for instance on ordered graphs (more generally, ordered binary structures) [4], graphs of bounded clique-width, proper minor-closed classes [6], but not in general graphs.

On the other hand, prior to this paper, no algorithmic lower bound was known for computing the twin-width. Our main result rules out an (exact) XP algorithm to decide $\text{tww}(G) \leq k$, that is, an algorithm running in time $n^{f(k)}$ for some computable function f . Indeed we show that deciding if the twin-width of a graph is at most 4 is intractable. We refer the reader to Section 2 for some context on the Exponential-Time Hypothesis (ETH), which implies that n -variable 3-SAT cannot be solved in time $2^{o(n)}$.

► **Theorem 1.** *Deciding if a graph has twin-width at most 4 is NP-complete. Furthermore, no algorithm running in time $2^{o(n/\log n)}$ can decide if an n -vertex graph has twin-width at most 4, unless the ETH fails.*

As far as approximation algorithms are concerned, our result only rules out a ratio better than $5/4$ for determining the twin-width. This still leaves plenty of room for an $f(\text{OPT})$ -approximation, which would be good enough for most of the (theoretical) algorithmic applications. Note that such algorithms exist for treewidth in polytime [12] and FPT time [27], for pathwidth [12], and for clique-width via rank-width [30].

Is Theorem 1 surprising? On the one hand, it had to be expected that deciding, given a graph G and an integer k , whether $\text{tww}(G) \leq k$ would be NP-complete. This is the case for example of treewidth [1], pathwidth [29, 26, 28], clique-width [13], rank-width [23], mim-width [32], and bandwidth [19]. On the other hand, the parameterized complexity of these width parameters is more diverse and harder to predict. Famously, Bodlaender’s algorithm is a linear FPT algorithm to exactly compute treewidth [2] (and a non-uniform FPT algorithm came from the Graph Minor series [31]). In contrast, while there is an XP algorithm to compute bandwidth [33], an FPT algorithm is highly unlikely [10]. It is a long-standing open whether an FPT or a mere XP algorithm exist for computing clique-width exactly, or even simply if one can recognize graphs of clique-width at most 4 in polynomial time (deciding clique-width at most 3 is indeed tractable [7]).

Theorem 1 almost completely resolves the parameterized complexity of exactly computing twin-width on general graphs. Two questions remain: can graphs of twin-width at most 2, respectively at most 3, be recognized in polynomial time. Graphs of twin-width 0 are cographs, which can be recognized in linear time [22], while it was recently shown that graphs of twin-width at most 1 can be recognized in polynomial time [5]. In the course of establishing Theorem 1 we show and generalize the following, where an $(\geq s)$ -subdivision of a graph is obtained by subdividing each of its edges at least s times.

► **Theorem 2.** *Any $(\geq 2 \log n)$ -subdivision of an n -vertex graph has twin-width at most 4.*

That those graphs have bounded twin-width was known [3], but not with the explicit bound of 4. Another family of graphs with twin-width at most 4 is the set of grids, walls, their subgraphs, and subdivisions. Even if there is no proof of that fact, sufficiently large grids, walls, or their subdivisions likely have twin-width *at least* 4; it is actually surprising that the 6×8 grid still has twin-width 3 [34]. We also believe that long subdivisions of “sufficiently complicated” graphs have twin-width *at least* 4. That would make graphs of twin-width at most 3 considerably simpler than those of twin-width at most 4, especially among sparse graphs.

Contrary to the hardness proof for treewidth [1], which involves some structural characterizations by chordal completions, and the intermediate problems MINIMUM CUT LINEAR ARRANGEMENT, MAX CUT, and MAX 2-SAT [19], our reduction is “direct” from 3-SAT. This makes the proven hardness of twin-width more robust, and easier to extend to restricted classes of graphs, especially sparse ones. Theorem 1 holds for bounded-degree input graphs. Furthermore performing our reduction from PLANAR 3-SAT produces graphs excluding a fixed minor, more specifically subgraphs of constant powers of planar graphs (while admittedly weakening the ETH lower bound from $2^{\Omega(n/\log n)}$ to $2^{\Omega(\sqrt{n}/\log n)}$). Hence, while the complexity status of computing treewidth on planar graphs is a famous long-standing open question, one can probably extend the NP-hardness of *twin-width at most 4* to planar graphs, by tuning and/or replacing the few non-planar gadgets of our reduction.

Let us point out that, in contrast with subset problems, there is no $2^{O(n)}$ -time algorithm known to compute twin-width. The exhaustive search takes time $n^{2n+O(1)}$ by considering all sequences of $n - 1$ pairs of vertices. We leave as an open question whether the ETH lower bound of computing twin-width can be brought from $2^{\Omega(n/\log n)}$ to $2^{\Omega(n)}$, or even $2^{\Omega(n \log n)}$. The latter lower bound is known to hold for SUBGRAPH ISOMORPHISM [8] (precisely, given a graph H and an n -vertex graph G , deciding if H is isomorphic to a subgraph of G requires time $2^{\Omega(n \log n)}$, unless the ETH fails), or computing the Hadwiger number [15] (i.e., the size of the largest clique minor).

1.1 Outline of the proof of Theorem 1

We propose a quasilinear reduction from 3-SAT. Given an n -variable instance I of 3-SAT, we shall construct an $O(n \log n)$ -vertex graph $G = G(I)$ which has twin-width at most 4 if and only if I is satisfiable.

Half of our task is to ensure that no 4-sequence will exist if I is unsatisfiable. This is challenging since many contraction strategies are to be considered and addressed. We make this task more tractable by attaching *fence gadgets* to some chosen vertex subsets. The effect of the fence *enclosing* S is that no contraction can involve vertices in S with vertices outside of S , while S is not contracted into a single vertex. The *maximal or outermost* fences (we may nest two or more fence gadgets) partition the rest of the vertices. This significantly tames the potential 4-sequences of G .

Our basic building block, the *vertical set*, consists of a pair of vertices (*vertical pair*) enclosed by a fence. It can be thought of as a bit set to 0 as long as the pair is not contracted, and to 1 when the pair gets contracted. It is easy to assemble vertical sets as prescribed by an auxiliary digraph D (of maximum degree 3), in such a way that, to contract (by a partial 4-sequence) the pair of a vertical set V , one first has to contract all the vertical sets that can reach V in D . This allows to propagate and duplicate a bit in a so-called *wire* (corresponding to an out-tree in D), and to perform the logical AND of two bits.

The bit propagation originates from a variable gadget (we naturally have one per variable appearing in I) that offers two alternatives. One can contract the “top half” of the gadget of variable x_i , which then lets one contract the vertical sets in the wire of literal x_i , or one can contract instead the “bottom half” of the gadget, as well as the vertical sets in the wire of literal $\neg x_i$. Concretely, these two contraction schemes represent the two possible assignments for variable x_i . A special “lock” on the variable gadget (called *half-guards*) prevents its complete contraction, and in particular, performing contractions in both the wires of a literal and its negation.

The leaves of the literal wires serve as inputs for 3-clause gadgets. One can contract the output (also a vertical set) of a clause gadget if and only if one of its input is previously contracted. We then progressively make the AND of the clauses via a “path” of binary AND gadgets fed by the clause outputs. We eventually get a vertical set, called *global output*, which can be contracted by a partial 4-sequence only if I is satisfiable. Indeed at this point, the variable gadgets are still locked so at most one of their literals can be propagated. This ticks one of our objective off. We should now ensure that a 4-sequence is possible from there, when I is satisfiable.

For that purpose, we add a wire from the global output back to the half-guards (or locks) of the variable gadgets. One can contract the vertical sets of that wire, and in particular the half-guards. Once the variable gadgets are “unlocked,” they can be fully contracted. As a consequence, one can next contract the wires of literals set to false, and *all* the remaining vertical sets involved in clause gadgets.

At this point, the current trigraph H roughly has one vertex per outermost fence with red edges linking two adjacent gadgets (and no black edge). We will guarantee that the (red) degree of H is at most 4, its number of vertices of degree at least 3 is at most βn , for some constant β . Besides we will separate gadgets by degree-2 wires of length $2 \log(\beta n)$ beforehand. This is crucial so that the red graph of H is a $(2 \log n')$ -subdivision of an n' -vertex graph. We indeed show that such trigraphs have twin-width at most 4. A complicated proof in [3] shows that $\Theta(\log n')$ -subdivisions of n' -vertex graphs have bounded twin-width. Here we give an elementary proof of a similar fact with an explicit upper bound of 4.

This finishes to describe our overall plan for the reduction and its correctness. It happens that fence gadgets are easier to build as trigraphs, while the rest of the gadgetry can be directly encoded by graphs. We thus show how to encode trigraphs by graphs, as follows. For any trigraph J whose red graph has degree at most d , and component size at most h , there is a graph G on at most $f(d, h) \cdot |V(J)|$ vertices such that J has twin-width at most $2d$ if and only if G has twin-width at most $2d$. This uses some local replacements and confluence properties of certain partial contraction sequences.

1.2 Organization of the paper

The rest of the paper is organized as follows. In Section 2 we formally introduce the relevant definitions and notations, and state some easy observations. In Section 3 we show a generalization to trigraphs of Theorem 2, crucial to the subsequent hardness proof. Section 4

is devoted to the hardness construction. In Section 4.1 we present how to encode trigraphs into graphs. In Section 4.2, we quickly recap the overall plan. The subsequent subsections go through the various gadgets. Finally Section 4.9 details the quasilinear reduction from 3-SAT to the problem of deciding if the twin-width is at most 4, and its correctness.

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. An ($\geq s$)-subdivision (resp. s -subdivision) of a graph G is obtained by subdividing every edge of G at least s times (resp. exactly s times).

2.1 Definitions and notations related to twin-width

We first define partition sequences, which is an alternative approach to contraction sequences.

Partition sequences. The *twin-width* of a graph, introduced in [6], 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. Finally the twin-width of G , denoted by $tw(G)$, is the least integer d such that there is partition sequence $\mathcal{P}_n, \dots, \mathcal{P}_1$ of G with every part of every \mathcal{P}_i ($1 \leq i \leq n$) being homogeneous to every other parts of \mathcal{P}_i but at most d .

Contraction sequences. A *trigraph* G has vertex set $V(G)$, black edge set $E(G)$, red edge set $R(G)$. Its *red graph* $(V(G), R(G))$ may be denoted $\mathcal{R}(G)$, and *total graph* $(V(G), E(G) \cup R(G))$, $\mathcal{T}(G)$. The *red degree* of a trigraph is the degree of its red graph. A trigraph G' is an *induced subtrigraph* of trigraph G if $V(G') \subseteq V(G)$, $E(G') = E(G) \cap \binom{V(G')}{2}$, and $R(G') = R(G) \cap \binom{V(G')}{2}$. Then we say that G is a *supertrigraph* of G' , and we may also denote G' by $G[V(G')]$. 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 any non-homogeneous pair of parts. A *partial contraction sequence* is then a sequence of trigraphs of $\mathcal{P}_n, \dots, \mathcal{P}_i$ (for some $i \in [n]$). A (non-partial) *contraction sequence* is one such that $i = 1$. Note that going from the trigraph of \mathcal{P}_{i+1} to the one of \mathcal{P}_i corresponds to the contraction operation described in Section 1. A (*partial*) *d -sequence* is a (partial) contraction sequence where all the trigraphs have red degree at most d . In the above definitions of twin-width, nothing prevents the starting structure G to be a trigraph itself. We may then talk about the twin-width of a trigraph.

Conversion partition/contraction, and convenient notations and definitions.

The reason we gave two (equivalent) definitions of twin-width is that both viewpoints are incomparably useful and convenient. It is always cumbersome to describe a sequence of partitions, so we will prefer the trigraph and contraction viewpoint to describe d -sequences. However the trigraph “loses” a part of the information that the partitioned graph preserves; namely the exact adjacencies between non-homogeneous parts, and the names of the original

vertices. To navigate between these two worlds, and keep the proofs compact, we use the following notations and vocabulary.

In this paragraph, we assume that there is a partial contraction sequence from (tri)graph G to trigraph H . If u is a vertex of H , then $u(G)$ denotes the set of vertices eventually contracted into u in H . We denote by $\mathcal{P}(H)$ the partition $\{u(G) : u \in V(H)\}$ of $V(G)$. If G is clear from the context, we may refer to a *part* of H as any set in $\{u(G) : u \in V(H)\}$. We may say that two parts $y(G), z(G)$ of $\mathcal{P}(H)$ are *in conflict* if $yz \in R(H)$. We say that a contraction of two vertices $u, u' \in V(H)$ *involves* a vertex $v \in V(G)$ if $v \in u(G)$ or $v \in u'(G)$. A contraction *involves* a pair of vertices v, v' if $v \in u(G)$ and $v' \in u'(G)$ (or $v \in u'(G)$ and $v' \in u(G)$). Observe that the two vertices should appear in two distinct parts. By extension, we may say that a contraction *involves* a set S , if it involves a vertex of S , or a pair of sets S, T if it involves a pair in $S \times T$.

2.2 Useful observations

The twin-width can only decrease when taking induced subtrigraph.

► **Observation 3.** *Let G' be an induced subtrigraph of trigraph G . Then $tww(G') \leq tww(G)$.*

It can be observed that adding red edges can only increase the twin-width, since the same contraction sequence for the resulting trigraph works at least as well for the initial trigraph.

► **Observation 4.** *Let G be a trigraph and G' another trigraph obtained from G by turning some non-edges and some edges into red edges. Then $tww(G') \geq tww(G)$.*

Trees admit a simple 2-sequence, that gives a d -sequence on red trees of degree at most d .

► **Lemma 5** ([6]). *Every (black) tree has twin-width at most 2. Every red tree has twin-width at most its maximum degree.*

Proof. Root the tree arbitrarily. Contract two leaves with the same parent whenever possible. If not possible, contract a deepest leaf with its parent. Observe that this cannot result in red degree larger than the maximum of 2 and the initial degree of the tree. ◀

2.3 The Exponential-Time Hypothesis

The Exponential-Time Hypothesis (ETH, for short) was proposed by Impagliazzo and Paturi [24] and asserts that there is no subexponential-time algorithm solving 3-SAT. More precisely, there is an $\varepsilon > 0$ such that n -variable 3-SAT cannot be solved in time $2^{\varepsilon n}$. Impagliazzo et al. [25] present a subexponential-time Turing-reduction parameterized by a positive real $\varepsilon > 0$ that, given an n -variable m -clause k -SAT-instance I , produces at most $2^{\varepsilon n}$ k -SAT-instances I_1, \dots, I_t such that I is satisfiable if and only if at least one of I_1, \dots, I_t is satisfiable, each I_i having no more than n variables and $C_\varepsilon n$ clauses for some constant C_ε depending only on ε . This crucial reduction is known as the Sparsification Lemma. Due to it, there is an $\varepsilon' > 0$ such that there is no algorithm solving m -clause 3-SAT in time $2^{\varepsilon' m}$, unless the ETH fails.

A classic reduction from Tovey [35], linear in the number of clauses, then shows the following.

► **Theorem 6** ([35, 25]). *The n -variable 3-SAT problem where each variable appears at most twice positively, and at most twice negatively, is NP-complete, and cannot be solved in time $2^{o(n)}$, unless the ETH fails.*

3 Long subdivisions have twin-width at most four

In [3], it is proved that the $\Omega(\log n)$ -subdivision of any n -vertex graph has bounded twin-width. The proof is rather involved, relies on a characterization by *mixed minors* established in [6], and does not give an explicit constant bound. Here we give an elementary proof that any $(\geq 2\lceil \log n \rceil - 1)$ -subdivision of an n -vertex graph has twin-width at most 4.

► **Theorem 7.** *Let G be a trigraph obtained by subdividing each edge of an n -vertex graph H at least $2\lceil \log n \rceil - 1$ times, and by turning red any subset of its edges as long as the red degree of G remains at most 4, and no vertex with red degree 4 has a black neighbor. Then $\text{tw}(G) \leq 4$.*


Proof. By no more than doubling the number of vertices of H , we can assume that n is a power of 2. Indeed, padding H with isolated vertices up to the next power of 2 does not change the quantity $\lceil \log |V(H)| \rceil$.

Let G' be a supertrigraph of G obtained by arbitrarily arranging the vertices of H (in G) at the leaves of a “virtual” full binary tree of height $\log n$. So that the red degree does not exceed 4, we so far omit the edges of the tree incident to a leaf (i.e., a vertex of H), while we put in red all the other edges of the tree. (The missing edges of the tree will naturally appear in red.) The internal nodes of the tree are all fresh vertices, not present in G . See Figures 2 and 3 for an illustration. We show that $\text{tw}(G') \leq 4$, hence by Observation 3, $\text{tw}(G) \leq 4$ since G is an induced subtrigraph of G' .

We label $1, 2, \dots, n$ the vertices of H . If there is an edge $ij \in E(H)$, it is subdivided into a path, say, $i, s(ij, 1), s(ij, 2), \dots, s(ij, z), j$ in G with $z \geq 2 \log n - 1$. First, we repeatedly contract adjacent vertices in the middle of this path until it consists of exactly $2 \log n$ edges. If $z > 2 \log n - 1$, we had to contract at least one pair of adjacent vertices. Thus the vertex in the middle of the path necessarily has now two red edges incident to it. Note that the other edges of the path can be black or red indifferently. To avoid cumbersome notations, we rename the inner vertices of the path $s(ij, 1), s(ij, 2), \dots, s(ij, z)$ with now $z = 2 \log n - 1$.


► **Claim 8.** There is a partial 4-sequence from G' to $G' - \{s(ij, 1), \dots, s(ij, z)\}$.

Proof. Intuitively we “zip” the subdivision of ij with the walk made by the union of the path from leaf i to the root, and the path from the root to leaf j . Let $i, v_1, v_2, \dots, v_z, j$ be the concatenation of the simple path from i to the root of the tree, and the simple path from



■ **Figure 2** Contracting the pairs of vertices with the same color, from the greenest to the bluest, is a partial 4-sequence, which acts as a deletion of the subdivided edge $(3, 11)$.

the root to j . Its length is thus $2 \log n = z + 1$. For h going from 1 to $z = 2 \log n - 1$, we contract v_h and $s(ij, h)$ (see Figure 2). After each contraction, the newly formed vertex has red degree at most 4. The red degree of vertices that are neither the new vertex nor a leaf of the tree is either unchanged or at most 2. The red degree of a leaf ℓ of the tree may increase by 1. This may only happen the first time a neighbor of ℓ is involved in a contraction, and that contraction merges a black neighbor of ℓ with the parent of ℓ in the tree (like is the case for leaf 3 from Figure 2 to Figure 3). By assumption, this implies that ℓ had red degree at most 3, thus its red degree does not exceed 4. Thus, what we defined is indeed a partial 4-sequence. One can finally notice that after these z contractions, we indeed reach trigraph $G' - \{s(ij, 1), \dots, s(ij, z)\}$. \triangleleft



■ Figure 3 The picture after the first three contractions. The newly formed vertex has red degree 4.

We apply Claim 8 for each edge of H (or rather, subdivided edge in G). We are then left with a red full binary tree which admits a 3-sequence by Lemma 5. Hence there is a 4-sequence for G' , and in particular, for G . \blacktriangleleft

We will only use the following consequence.

► **Lemma 9.** Let G be a trigraph obtained by subdividing at least $2\lceil \log n \rceil - 1$ times each edge of an n -vertex graph H of degree at most 4, and by turning red all its edges. Then $tw(G) \leq 4$.

4 Hardness of determining if the twin-width is at most four

Here we show the main result of the paper.

► **Theorem 1.** Deciding if a graph has twin-width at most 4 is NP-complete. Furthermore, no algorithm running in time $2^{o(n/\log n)}$ can decide if an n -vertex graph has twin-width at most 4, unless the ETH fails.

The membership to NP is ensured by the d -sequence: a polynomial-sized certificate that a graph has twin-width at most d , checkable in polynomial time. We thus focus on the hardness part of the statement, and design a quasilinear reduction from 3-SAT.

4.1 Encoding a trigraph by a graph

In this subsection, we present a construction allowing to encode trigraphs into (plain) graphs. Our objective is, given a trigraph H with red degree at most d , to produce a graph G such that H admits a $2d$ -sequence iff G admits a $2d$ -sequence.

Formally, we build a graph G for which every $2d$ -sequence \mathcal{S} inevitably creates H as an induced subtrigraph of a trigraph of \mathcal{S} . By Observation 3, the existence of a $2d$ -sequence for G implies the existence of one for H . Moreover, this construction is such that there is a partial $2d$ -sequence which, from G , exactly reaches H . Therefore, any $2d$ -sequence for H can be augmented to become a $2d$ -sequence for G .

At first sight, our construction works only for trigraphs H with a connected red graph. We could show the following.

► **Lemma 10.** *For every trigraph H such that $\mathcal{R}(H)$ is connected and of degree at most d , there is a graph G such that:*

- every $2d$ -sequence of G goes through a supertrigraph of H , and
- there is a partial $2d$ -sequence from G to H .

However, we will need a stronger version, allowing to simulate trigraphs with disconnected red graphs. First we show how to turn the trigraph induced by one connected component of the red graph into a plain graph (this is Lemma 11). Later we loop through all connected components of the red graph with size at least 2, and apply Lemma 11 (that is done in Lemma 15).

► **Lemma 11.** *Let H be a trigraph, and $S \subseteq V(H)$ be the vertex set of a connected component of $\mathcal{R}(H)$ with degree at most d . There is a trigraph G and a set $T \subseteq V(G)$ satisfying the following statements:*

1. $|T|$ is bounded by a function of d and $|S|$ only,
2. no red edge touches T (in particular, $G[T]$ has no red edge),
3. $G - T$ is isomorphic to $H - S$,
4. there is a partial $2d$ -sequence from G to H , and
5. every $2d$ -sequence of G goes through a supertrigraph of some trigraph \tilde{H} , where \tilde{H} can be obtained from H by performing contractions not involving vertices of S , nor creating red edges incident to S .

Proof. We begin with the construction of G . The vertex set of G can be split into two sets: on one hand $V(H) \setminus S$, and on the other hand T , which will be defined in function of $H[S]$. Let $\{v_1, v_2, \dots, v_{|S|}\}$ be the set S . To build set T , we blow up every vertex v_i of S into a copy L_i of the biclique $K_{t,t}$ with $t = 2(2d + 2) \cdot (2d)^{|S|-1} + 1$. The two sides of each biclique L_i are denoted by $A_i = \{a_{i,j}\}_{j \in [t]}$ and $B_i = \{b_{i,j}\}_{j \in [t]}$, and are respectively called *A-side* and *B-side*. The vertex set $T \subseteq V(G)$ is:

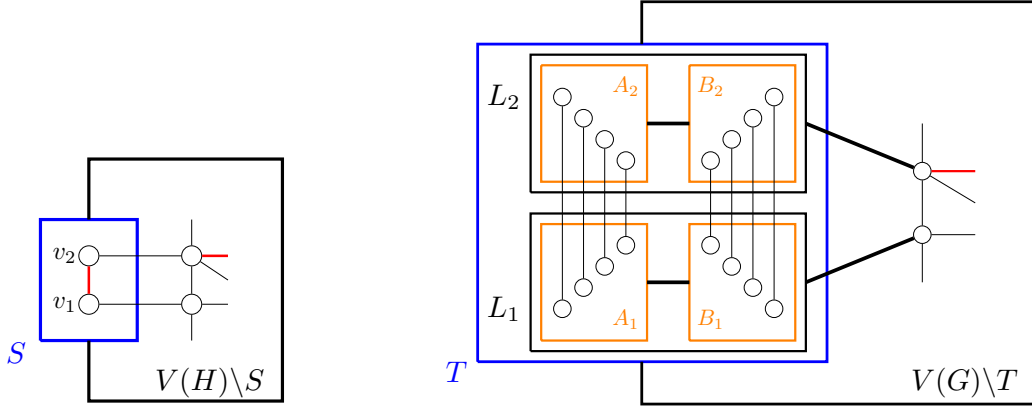
$$T = \bigcup_{v_i \in S} V(L_i) = \bigcup_{v_i \in S, j \in [t]} \{a_{i,j}, b_{i,j}\}.$$

In addition to the edges forming each biclique, we replace every black edge of $H[S]$ by a complete bipartite graph between the two corresponding bicliques. Furthermore, every red edge of $H[S]$ becomes a black canonical matching, with edges $a_{i,j}a_{i',j}$ and $b_{i,j}b_{i',j}$. Formally,

$$E(G[T]) = \bigcup_{v_i \in S} A_i \times B_i \cup \bigcup_{v_i, v_{i'} \in E(H)} L_i \times L_{i'} \cup \bigcup_{\substack{v_i, v_{i'} \in \mathcal{R}(H) \\ j \in [t]}} \{a_{i,j}a_{i',j}, b_{i,j}b_{i',j}\}.$$

We finish the construction by considering the black edges initially connecting S and $V(H) \setminus S$ in the trigraph H . For any edge $v_i z \in E(H)$ with $v_i \in S$ and $z \in V(H) \setminus S$, we add a black edge between any vertex of the biclique L_i and $z \in V(G) \setminus T$. In summary,

$$E(G) = E(G[T]) \cup E(H - S) \cup \bigcup_{\substack{v_i \in S \\ z \in V(H) \setminus S \\ v_i z \in E(H)}} L_i \times \{z\}.$$



(a) An example of trigraph H . Set S consists of two vertices v_1 and v_2 , linked by a red edge. Vertices of $H - S$ adjacent to S are also represented.

(b) Graph G obtained from H . Thick black edges link every vertex of one endpoint to every vertex of the other endpoint. Set T contains two bicliques L_1 and L_2 , inherited from $S = \{v_1, v_2\}$.

■ **Figure 4** An example of encoding: induced trigraph S (left) and the plain graph T (right).

Statements 1-3. The first three statements are satisfied by this construction. In particular the size of T is $2t|S|$, and t was defined as a function of d and $|S|$. Moreover, no vertex of $V(H) \setminus S$ was modified by this construction, so $G - T$ is isomorphic to $H - S$.

Statement 4. We focus now on the fourth statement: We exhibit a partial $2d$ -sequence \mathcal{S}^* from G to H . Let us begin with a short description of \mathcal{S}^* . This particular sequence never contracts two vertices of T lying in different bicliques. Moreover, after a contraction of two vertices belonging to the same biclique, it forces the same contraction in all other bicliques of T . We now describe \mathcal{S}^* in detail.

We start by contracting $a_{i,1}$ and $a_{i,2}$ for each $v_i \in S$. Then, we contract $\{a_{i,1}, a_{i,2}\}$ with $a_{i,3}$ for each $v_i \in S$, and so on. When the A -side of each biclique L_i is contracted into a single vertex $\{a_{i,1}, a_{i,2}, \dots, a_{i,t}\}$, we proceed similarly with the B -side. Finally, each biclique L_i contains exactly two contracted vertices, which are respectively A_i and B_i . We contract all these pairs in each L_i : every biclique is now contracted into a single vertex.

The obtained trigraph is isomorphic to H . Indeed, if $v_i v_{i'}$ is a black edge (resp. a non-edge) in $H[S]$, then L_i and $L_{i'}$ are homogeneous and, as contracted vertices, they are connected with a black edge (resp. a non-edge). However, if $v_i v_{i'}$ is a red edge in $H[S]$, then there is a semi-induced matching connecting bicliques L_i and $L_{i'}$ in $G[T]$, so the contracted vertices L_i and $L_{i'}$ are linked by a red edge. Finally, the black edges between S and $V(H) \setminus S$ are preserved because the contractions of \mathcal{S}^* occur inside bicliques of T which are homogeneous (i.e., fully adjacent or fully non-adjacent) with every vertex $z \in V(H) \setminus S$.

Let us check that the partial sequence \mathcal{S}^* only goes through trigraphs with red degree at most $2d$. As we only contract vertices within the same side of each biclique first, no red edge appears between two vertices of the same biclique. Moreover, if $v_i v_{i'} \notin R(H)$, then

bicliques L_i and $L_{i'}$ are homogeneous, so no red edge appears between two vertices belonging respectively to L_i and $L_{i'}$. Thus, the red edges may only bridge two different bicliques $L_i, L_{i'}$ such that $v_i v_{i'} \in R(H)$. The way we progressively contract the matching between L_i and $L_{i'}$, a contracted vertex within L_i is the endpoint of at most two red edges towards $L_{i'}$. Assume without loss of generality that $u(G) = \{a_{i,1}, a_{i,2}, \dots, a_{i,j+1}\}$ is contracted in A_i while only $u'(G) = \{a_{i',1}, a_{i',2}, \dots, a_{i',j}\}$ is contracted in $A_{i'}$ for some $j \in [t-1]$. There are two red edges between bicliques L_i and $L_{i'}$ which are uu' and uu'' , where $u''(G) = \{a_{i',j+1}\}$. As the red degree of $H[S]$ is at most d , we indeed have that \mathcal{S}^* is a partial $2d$ -sequence.

Statement 5. We terminate by showing that any $2d$ -sequence \mathcal{S} from graph G necessarily produces at some moment a supertrigraph of \tilde{H} , where \tilde{H} is obtained from H with contractions not involving set S . As this part of the proof is more intricate than the previous ones, two intermediate claims are stated. We fix $g(d, t) = (t-1)/(2d+2) = 2(2d)^{|S|^{-1}}$.

▷ **Claim 12.** Let G_h be any trigraph of $\mathcal{S} = G, \dots, G_h, \dots, K_1$. If a vertex u of G_h is such that $u(G)$ intersects two distinct bicliques L_i and $L_{i'}$, then there is a vertex w (possibly u) of G_h intersecting one side of L_i in at least $g(d, t)$ elements: $|w(G) \cap A_i| \geq g(d, t)$ or $|w(G) \cap B_i| \geq g(d, t)$.

Proof. Let $a \in u(G) \cap V(L_i)$ and $z \in u(G) \cap V(L_{i'})$. We assume w.l.o.g. that the partite set of the biclique L_i in which a lies is A_i . We identify a set $X \subset V(G)$ of size at least $t-1$ such that, for any vertex $u' \in V(G_h) \setminus \{u\}$ such that $u'(G) \cap X \neq \emptyset$, there is a red edge uu' . The composition of set X depends on the adjacency between v_i and $v_{i'}$ in H .

- If $v_i v_{i'}$ is a non-edge in H , then we fix $X = B_i$. Indeed, any element of B_i is adjacent to a but not to z , so there is a red edge between u and any contracted vertex u' intersecting B_i . Moreover, $|B_i| = t$.
- If $v_i v_{i'}$ is a black edge in H , then we fix $X = A_i \setminus \{a\}$. Any element of $A_i \setminus \{a\}$ is not adjacent to a but adjacent to z . Thus, any contracted vertex u intersecting $A_i \setminus \{a\}$ is connected to u in red. Moreover, $|A_i \setminus \{a\}| = t-1$.
- If $v_i v_{i'}$ is a red edge in H , then we fix $X = B_i \setminus \{z'\}$, where z' is the matching neighbor of z in L_i . Observe that z' may belong to A_i and, in this case, $X = B_i$. Any element of X is adjacent to a but not to z . Moreover, $|X| \in \{t-1, t\}$.

In summary, there is a set X of at least $t-1$ vertices such that any contracted vertex $u' \in G_h$ intersecting X is linked by a red edge to u . Consequently, there cannot be more than $2d+1$ contracted vertices $w \in G_h$ intersecting X , otherwise the red degree of u is necessarily larger than $2d$. For this reason, at least one vertex $w \in G_h$ which intersects X has size at least $\lceil (t-1)/(2d+1) \rceil \geq g(d, t)$. In each case, either $X \subsetneq A_i$ or $X \subseteq B_i$, so either $|w(G) \cap A_i| \geq g(d, t)$ or $|w(G) \cap B_i| \geq g(d, t)$. ◁

Let G' be the first trigraph of sequence \mathcal{S} having a vertex u such that $|u(G) \cap T| \geq g(d, t)$. We denote by G'' the trigraph preceding G' in \mathcal{S} . We show that every contracted vertex $w \in G'$ is either fully contained in some biclique L_i or contained in $V(G) \setminus T$.

▷ **Claim 13.** Every contracted vertex $u \in V(G')$ satisfies:

- either $u(G) \subseteq A_i$ or $u(G) \subseteq B_i$ for some L_i of T ,
- or $u(G) \subseteq V(G) \setminus T$.


Furthermore, there is no red edge between a contracted vertex inside a biclique of T and a contracted vertex in $G - T$.

Proof. We begin by proving that if a contracted vertex $w \in V(G')$ intersects T , then it intersects at most one of its bicliques, that is, $w(G) \cap T \subseteq L_i$ for some index i . According to Claim 12, all the contracted vertices of G'' intersect T within a single biclique. This holds

in particular for the two vertices, say $x, y \in V(G'')$, whose contraction yields u . Assume that x and y are both contained within different bicliques, say, L_i and $L_{i'}$, respectively. By Claim 12, there is a vertex of G' intersecting L_i in at least $g(d, t)$ elements. By definition of trigraph G' , this vertex is necessarily u . This yields a contradiction, as the cardinality of $x(G) = u(G) \cap V(L_i)$ would be at least $g(d, t)$. Therefore, vertex u intersects at most one biclique. In brief, from now on, every set $w(G)$ with $w \in V(G')$ can be written as the union of elements of some biclique L_i with vertices of $G - T$.

Each side, A_i or B_i , of a biclique L_i satisfies the following property: There are at least $2d + 2$ contracted vertices of G' intersecting every A_i (or B_i). Any contracted vertex in G'' covers less than $g(d, t)$ elements of T , by definition. So, at least $2d + 3$ contracted vertices of G'' must intersect A_i as $|A_i| = t$. In G' , this property holds except for vertex u , which is the contraction of two vertices $x, y \in V(G'')$. Therefore, as announced, at least $2d + 2$ contracted vertices w of G' satisfy $w(G) \cap A_i \neq \emptyset$ (resp. $w(G) \cap B_i \neq \emptyset$) for every side A_i (resp. B_i).

Based on this property, we show that, for any $w \in V(G')$, set $w(G)$ cannot contain both vertices of T and of $V(G) \setminus T$. Assume it is the case and w.l.o.g. that $w(G) \cap A_i \neq \emptyset$. All vertices of L_i play the same role regarding the adjacencies between T and $G - T$: either they are all connected to some $z \in G - T$ (if $v_i z \in E(H)$) or none of them is connected to z (if $v_i z \notin E(H)$). Consequently, we can distinguish only two cases. First, assume that at least one vertex of $w(G) \setminus A_i$ is adjacent (black or red) to L_i in G (Figure 5a). In that case, w admits at least $2d + 1$ red neighbors in G' which are the other contracted vertices intersecting A_i . Indeed, these vertices contain elements of A_i which are adjacent to $w(G) \setminus A_i$ but not to $w(G) \cap A_i$. Second, assume that no edge of G connects $w(G) \setminus A_i$ with L_i (Figure 5b). We can identify $2d + 1$ red neighbors of w : among the at least $2d + 2$ contracted vertices of G' intersecting B_i , at least $2d + 1$ of them are different from w . These contracted vertices contain elements of B_i which are adjacent to $w(G) \cap A_i$ but not to $w(G) \setminus A_i$. In summary, for any $w \in G'$, either $w(G) \subseteq L_i$ or $w(G) \subseteq V(G) \setminus T$.



(a) Case when v_i is adjacent to some vertex of $w(G) \setminus A_i$ in H . (b) Case when v_i is not adjacent to any vertex of $w(G) \setminus A_i$ in H

■ **Figure 5** How contractions intersecting both T and $V(G) \setminus T$ would imply large red degree in G' . The green boxes represent contracted vertices in G' .

We now prove that there is no red edge $ww' \in R(G')$ such that $w(G) \subseteq L_i$ and $w'(G) \subseteq V(G) \setminus T$. Suppose by way of contradiction that such an edge exists. As all vertices of L_i play the same role in the adjacencies between T and $G - T$, there is a red edge between w' and any vertex of G' intersecting L_i . As at least $2d + 2$ of these vertices intersect A_i , the red degree of w' is at least $2d + 2$, a contradiction.

To end the proof, we show that any contracted vertex $w(G) \subseteq L_i$ verifies either $w(G) \subseteq A_i$

or $w(G) \subseteq B_i$. Consider the at least $2d + 1$ vertices of G' intersecting A_i different from w . They contain elements of A_i which are adjacent to $w(G) \cap A_i$ but not to $w(G) \cap B_i$. Therefore, the red degree of w is at least $2d + 1$, a contradiction. \triangleleft

In brief, the partial sequence from G to G' does not use contractions mixing vertices of T and vertices of $V(G) \setminus T$. As a consequence, G' can be split in two parts: its induced subtrigraph G'_T obtained from contractions on the bicliques and its induced subtrigraph G'_{G-T} obtained from the contractions on $G - T$. Our objective is to prove that H appears in the first part.

We now focus more specifically on the vertex u . W.l.o.g., we assume that $u(G)$ intersects $A_i \subset L_i$. We denote by J_0 a set of exactly $g(d, t)$ distinct indices j such that $a_{i,j} \in u(G)$. For every $J \subseteq [t]$, we denote by $A_{i,J}$ the set $\{a_{i,j} : j \in J\}$. It holds that $|J_0| = |A_{i,J_0}| = g(d, t)$.

▷ **Claim 14.** $H[S]$ is an induced subtrigraph of G'_T .

Proof. We initialize a subset $Y \subseteq S$ to $\{v_i\}$, an injective mapping $\rho : Y \hookrightarrow V(G')$ to $v_i \mapsto u$, and a subset $J \subseteq [t]$ to J_0 . We maintain the following invariant:

$$(*) \ A_{h,J} \subseteq \rho(v_h)(G) \text{ for every } v_h \in Y, \text{ and } 2 \leq |J| = \frac{t-1}{(2d+2) \cdot (2d)^{|Y|-1}}.$$

The invariant initially holds by construction of J_0 . While $Y \neq S$, we pick a pair $v_h \in Y, v_{h'} \in S \setminus Y$ such that $v_h v_{h'} \in R(H)$. Since $v_h(G) \supseteq A_{h,J}$ and there is an induced matching between $A_{h,J}$ and $A_{h',J}$, the set $A_{h',J}$ can be spanned by at most $2d$ parts of $\mathcal{P}(G')$. Thus we select of vertex $w_{h'} \in V(G')$ and a subset $J' \subseteq J$ of size $\frac{|J|}{2d}$ such that $w_{h'}(G) \supseteq A_{h',J'}$. We set the new J to J' , and we augment ρ with $v_{h'} \mapsto w_{h'}$. Finally, we add $v_{h'}$ to Y .

The invariant $(*)$ is preserved by construction. Since the graph $\mathcal{R}(H)$ is connected, this process ends when $Y = S$ and $|J| = \frac{g(d,t)}{(2d)^{|S|-1}} = 2$. We claim that trigraph $G'[\rho(S)]$, which is an induced subtrigraph of G'_T , is isomorphic to $H[S]$. Indeed, let $v_h, v_{h'} \in S$ be any pair of vertices, and let $w_h = \rho(v_h)$ and $w_{h'} = \rho(v_{h'})$. Assume first that $v_h v_{h'} \in R(H)$. As $w_h(G) \supseteq A_{h,J}, w_{h'}(G) \supseteq A_{h',J}, |J| = 2$, and there is an induced matching between $A_{h,J}$ and $A_{h',J}$, there is a red edge in G' between w_h and $w_{h'}$. If $v_h v_{h'}$ is a black edge in $H[S]$, then any pair of vertices in $w_h(G) \times w_{h'}(G)$ is a black edge in G . Hence, sets $w_h(G)$ and $w_{h'}(G)$ are homogeneous and $w_h w_{h'}$ is a black edge in G' . Eventually, if $v_h v_{h'}$ is a non-edge in $H[S]$, both $w_h(G)$ and $w_{h'}(G)$ are homogeneous in the sense that they are not adjacent at all. So, $w_h w_{h'}$ is a non-edge in G' . In brief, $G'[\rho(S)]$ is isomorphic to $H[S]$. \triangleleft

This concludes the proof of Lemma 11. Indeed, consider the subtrigraph of G' induced on vertices of both sets $\rho(S)$ and G'_{G-T} . The subtrigraph $G'[\rho(S)]$ is isomorphic to $H[S]$ (Claim 14) and the second part G'_{G-T} is obtained from contractions on $G - T$ which do not make red edges appear towards S (Claim 13). \blacktriangleleft

If the red graph $\mathcal{R}(H)$ is not connected, one can apply Lemma 11 for each of its connected components.

► **Lemma 15.** *Given any trigraph H whose red graph has degree at most d and connected components of size at most h , one can compute in time $O_{d,h}(|V(H)|)$ a graph G on $O_{d,h}(|V(H)|)$ vertices such that H has a $2d$ -sequence if and only if G has a $2d$ -sequence.*

Proof. Let S_1, \dots, S_r be the connected components of $\mathcal{R}(H)$ of size at least 2. We define a collection of trigraphs $H_0 = H, H_1, \dots, H_r$, where H_r is a plain graph isomorphic to G . Each trigraph $H_{\ell+1}$ will be built from H_ℓ by applying Lemma 11 on $S_{\ell+1}$. We proceed by induction and prove, for any $0 \leq \ell \leq r$, that:

- $|V(H_\ell)|$ is linear in $|V(H)|$ if $d, h = O(1)$: $|V(H_\ell)| = f_\ell(d, h)|V(H)|$,
- the connected components of $\mathcal{R}(H_\ell)$ of size at least 2 have vertex sets $S_{\ell+1}, \dots, S_r$,
- H_ℓ has a $2d$ -sequence iff H has a $2d$ -sequence.

The base case is trivial, as $H_0 = H$. Let us assume, for the induction step, that H_ℓ satisfies these three properties. We consider the red component $S_{\ell+1}$ in H_ℓ to build trigraph $H_{\ell+1}$. We apply Lemma 11 with trigraph H_ℓ and vertex set $S_{\ell+1}$. The trigraph $H_{\ell+1}$ substitutes $H_\ell[S_{\ell+1}]$ with a plain graph $H_{\ell+1}[T]$ not touched by a red edge. Thus the vertex sets of connected components of $\mathcal{R}(H_{\ell+1})$ of size at least 2 are those of $\mathcal{R}(H_\ell)$ minus $S_{\ell+1}$.

Moreover, the size of T is bounded by d and $|S_{\ell+1}|$ only. As a consequence,

$$|V(H_{\ell+1})| = |T| + |V(H_\ell) \setminus S_{\ell+1}| \leq h(d, |S_{\ell+1}|) + f_\ell(d, h)|V(H)| \leq f_{\ell+1}(d, h)|V(H)|,$$

where h is the function defined accordingly to Lemma 11.

We finally show that $H_{\ell+1}$ has a $2d$ -sequence iff H_ℓ has a $2d$ -sequence. This is a direct consequence of Lemma 11. On one hand, there is a $2d$ -partial sequence from $H_{\ell+1}$ to H_ℓ , so if H_ℓ has a $2d$ -sequence, then $H_{\ell+1}$ has one, too. On the other hand, if $H_{\ell+1}$ has a $2d$ -sequence, then one trigraph of the sequence is a supertrigraph of \tilde{H}_ℓ , where \tilde{H}_ℓ is a trigraph which can be obtained from H_ℓ by performing contractions preserving S_ℓ as a red connected component of the trigraph. One can deduce from this statement a $2d$ -sequence of H_ℓ : first contract H_ℓ to obtain \tilde{H}_ℓ , then conclude by Observation 3.

As a conclusion, trigraph $G = H_r$ satisfies the induction hypotheses: its red graph is edgeless and it has a $2d$ -sequence if and only if H admits one, too. ◀

In the sequel, we will use Lemma 15 with $d = 2$ and $h = 12$. In particular, the size of the encoding graph will be linear in the trigraph. As paths are connected graphs of degree at most 2, we will invoke this scaled-down version.

► **Lemma 16.** *Given any trigraph H whose red graph is a disjoint union of 12-vertex paths and isolated vertices, one can compute in polynomial time a graph G on $O(|V(H)|)$ vertices such that H has twin-width at most 4 if and only if G has twin-width at most 4.*

4.2 Foreword to the reduction

Our task is now slightly simpler. Given a 3-SAT instance I , we may design a *trigraph* satisfying the requirements of Lemma 16 with twin-width at most 4 if and only if I is satisfiable.

As already mentioned, given an instance I of 3-SAT we will create an equivalent instance of the problem: *is trigraph G of twin-width at most 4?* We now present the various gadgets used in our reduction. The only building block featuring red edges is the fence gadget. Its red graph is a 12-vertex path that is a connected component in the red graph of the overall construction G . Hence G can be turned into a graph with only a constant multiplicative blow-up, by Lemma 16.

The correctness of the reduction naturally splits into two implications:

- (i) If G admits a 4-sequence, then I is satisfiable.
- (ii) If I is satisfiable, then G admits a 4-sequence.

We will motivate all the gadgets along the way, by exhibiting key properties that they impose on a potential 4-sequence. These properties readily leads to a satisfying assignment for I . So the proof of (i) will mostly consist of aggregating lemmas specific to individual gadgets.

We also describe partial 4-sequences to reduce most of the gadgets. However some preconditions (specifying the context in which a particular gadget stands) tend to be technical,


and make more sense after the construction of G . In those cases, to avoid unnecessarily lengthy lemmas, we only give an informal strategy, and postpone the adequate contraction sequence to the final proof of (ii).

4.3 Fence gadget

We now design a gadget F partitioned into two vertex sets A, B . The gadget is *attached* to a non-empty subset $S \subseteq V(G) \setminus (A \cup B)$ by making A and S fully adjacent, and B and S fully non-adjacent. Our intent is that, in a 4-sequence, a vertex of F can be contracted with another vertex of the graph only when S has been contracted into a single vertex. We will ensure that every vertex outside $S \cup V(F)$ is fully adjacent to B or fully non-adjacent to A . As $|A| = |B| = 6 > 4$ will hold, the effect is that vertices of S have to be contracted to a single vertex before any vertex of S (or subset of vertices within S) can be contracted with a vertex or subset of vertices in $G - S$. To summarize, F encloses S into an “unbreakable unit:” the inside of S cannot be contracted with the outside as long as S is not a single vertex. Hence we call F a *fence gadget*.

The fence gadget is defined as a trigraph whose red graph is a simple path on 12 vertices, in particular, a connected graph of maximum degree 2. In the overall trigraph G , no red edge will link a vertex in the gadget to a vertex outside of it. Thus the fence gadget can be replaced by a graph by Lemma 16.

We now define the fence gadget. Its vertex set is $A \cup B$ with $A = \{a_1, a_2, a_3, a_4, a_5, a_6\}$ and $B = \{b_1, b_2, b_3, b_4, b_5, b_6\}$. Its black edge set consists of 13 edges: the cycles $a_1 a_2 a_3 a_4 a_5 a_6 a_1$ and $b_1 b_2 b_3 b_4 b_5 b_6 b_1$, plus the edge $b_1 a_6$. Its red edge set consists of 11 edges: $a_i b_i$ for each $i \in [6]$, and $a_i b_{i+1}$ for each $i \in [5]$. Finally A is made fully adjacent to S . See Figure 6 for an illustration.



■ **Figure 6** The fence gadget F , with $A = \{a_i \mid 1 \leq i \leq 6\}$ and $B = \{B_i \mid 1 \leq i \leq 6\}$.

We will later nest fence gadgets. Thus we have to tolerate that F has other neighbors than S in G . Actually we even allow $V(F)$ to have neighbors outside of S and the fence gadgets surrounding F . We however always observe the following rule.

► **Definition 17 (Attachment rule).** *A fence gadget F with vertex bipartition (A, B) , and attached to S , satisfies the attachment rule in a trigraph H if F is a connected component of $\mathcal{R}(H)$, and there is a set $X \subseteq V(H) \setminus (A \cup B \cup S)$ such that:*

- $\forall x \in A, N(x) \setminus V(F) = X \cup S,$
- $\forall x \in B, N(x) \setminus V(F) = X,$ and
- $\forall x \in X, S \subset N(x).$


Initially in G , we make sure that all the fence gadgets satisfy the attachment rule. This will remain so until we decide to contract them.

We denote by Y the set $V(G) \setminus (V(F) \cup S \cup X)$, when X is defined in the attachment rule. We make the following observations on the three possible neighborhoods that vertices outside of F have within $V(F)$.


► **Observation 18.** *The fence gadget definition and the attachment rule implies:*

- $\forall x \in S$, it holds $N(x) \cap V(F) = A$,
- $\forall x \in X$, it holds $N(x) \cap V(F) = V(F) = A \cup B$, and
- $\forall x \in Y$, it holds $N(x) \cap V(F) = \emptyset$.

Henceforth we will represent every fence gadget as a brown rectangle surrounding the set S it is attached to. The vertices of X are linked to the brown rectangle, as they are fully adjacent to $S \cup V(F)$. See Figure 7 for an illustration of the attachment rule, and a compact representation of fence gadgets, and Figure 8 for two nested fence gadgets and how the attachment rule is satisfied for both.



■ **Figure 7** Left: The forced adjacencies (solid lines, all edges between the two sets) and non-adjacencies (dashed lines, no edge between the two sets), as specified by the attachment rule. Right: Symbolic representation of the fence gadget attached to S by a brown rectangle; the vertices in X are linked to the brown box, while the potential neighbors of S in Y are only linked individually to their neighbors in S , and are fully non-adjacent to the vertices of the fence $V(F)$. Possible edges between X and Y are *not* represented.



■ **Figure 8** Left: Two nested fences, where the set X of the innermost fence is $\{x, x_1, x_2\} \cup V(F')$, while the set X' of the outermost fence is $\{x_1, x_2\}$. Right: Its compact representation.

Constraints of the fence gadget on a 4-sequence. The following lemmas are preparatory steps for the milestone that no part in a 4-sequence of G can overlap S (that is, intersects S without containing it).

► **Lemma 19.** *The first contraction involving two vertices of F results in a vertex of red degree at least 5, except if it is a contraction of some $a_i \in A$ with some $b_j \in B$.*

Proof. We consider a pair of distinct vertices u, v of $A \times A$ or $B \times B$. If u and v are non-adjacent in F , they both have at least three private neighbors in the total graph $\mathcal{T}(F)$. Thus their contraction results in a vertex of red degree at least 6, since by assumption these at least 6 vertices lie in distinct parts.

If u and v are adjacent in F , they both have at least two private neighbors in $\mathcal{T}(F)$, and a common neighbor w such that at least one of uw and vw is red in F . Hence their contraction results in a vertex with red degree at least 5. ◀

► **Lemma 20.** *If the first contraction involving two vertices of F is of some $a_i \in A$ with some $b_j \in B$, the red degree within F of the created vertex is at least 3.*

Proof. Let $u \in A$ and $v \in B$. Then, in the total graph $\mathcal{T}(F)$, either u and v both have at least two private neighbors, or they both have at least one private neighbor and a common neighbor linked to u or v in red in the trigraph F . By assumption, these at least 4 neighbors or at least 3 neighbors are in distinct parts. Hence, in both cases, the contraction of u and v results in a vertex of red degree at least 3 within F . ◀

The last preparatory step is this easy lemma.

► **Lemma 21.** *Before a contraction involves two vertices of $V(F)$, the following holds in a partial 4-sequence:*

- no part intersects both X and S ,
- no part intersects both Y and S , and
- no part intersects both X and Y .

Proof. By Observation 18, such a part would have red degree $|B| = 6$, $|A| = 6$, and $|A \cup B| = 12$, respectively. ◀

As a consequence we obtain the following.

► **Lemma 22.** *In a partial 4-sequence of G , the first contraction involving a vertex in $V(F)$ and a vertex in $V(F) \cup S$ has to be done after S is contracted into a single vertex.*

Proof. We consider the first time a vertex $u \in V(F)$ is involved in a contraction with a vertex of $V(F) \cup S$. Either (case 1) the part of u , P_u , is contracted with a part P_v containing $v \in V(F)$, or (case 2) P_u is contracted with a part P intersecting S but not $V(F)$.

In case 1, by Lemma 19, u and v hit both A and B . Thus, by Lemma 20, the red degree within F of the resulting vertex z is at least 3. Moreover z is linked by a red edge to every part within S , since S is fully adjacent to A , and fully non-adjacent to B . Thus S should at this point consist of a single part.

We now argue that case 2 is impossible in a partial 4-sequence. By Lemma 21, part P cannot intersect $X \cup Y$ (nor $V(F)$, by construction). Thus $P \subseteq S$. If $u \in A$, then the contraction of P_u and P has incident red edges toward at least 5 vertices: three vertices of A non-adjacent to u and two private neighbors of u (in the total graph) within B . If instead $u \in B$, the red degree of the contracted part is at least 6, as witnessed by two neighbors of u in B , and four non-neighbors of u in A . ◀

We can now establish the main lemma on how a fence gadget constrains a 4-sequence. Lemmas 21 and 22 have the following announced consequence: While S is not contracted into a single vertex, no part within S can be contracted with a part outside of S , and similarly vertices of X cannot be contracted with vertices of Y .

► **Lemma 23.** *In a partial 4-sequence, while S is not contracted to a single vertex,*

- (i) no part intersects both S and $V(G) \setminus S$, nor
- (ii) both X and Y .

Proof. Let H be a trigraph obtained by a partial 4-sequence from G , such that S is not contained in a part of $\mathcal{P}(H)$. By Lemma 22, no pair of vertices in $V(F) \times (S \cup V(F))$ are in the same part of $\mathcal{P}(H)$ (since S is not contracted to a single vertex). Thus we conclude by Lemma 21. ◀

Contracting the fence gadget. The previous lemmas establish some constraints that the fence gadget imposes on a supposed (partial) 4-sequence. We now see how a partial 4-sequence actually contracts a fence gadget.

Every time we are about to contract a fence gadget F attached to S , we will ensure that the following properties hold:

- no prior contraction has involved a vertex of $V(F)$,
- no red edge has one endpoint in $V(F)$ and one endpoint outside $V(F)$, and
- S is contracted into a single vertex with red degree at most 3.


In particular, the fence gadget F still satisfies the attachment rule.

► **Lemma 24.** *Let H be a trigraph containing a fence gadget F attached to a single vertex s of red degree at most 3. We assume that F respects the attachment rule in H .*

Then there is a partial 4-sequence from H to H' , where H' is the trigraph obtained from H by contracting $V(F)$ into a single vertex.

Proof. As F respects the attachment rule in H , every vertex of F has the same (fully black) neighborhood in $V(H) \setminus (V(F) \cup \{s\})$. Thus, contractions within $V(F)$ will only create red edges within $V(F) \cup \{s\}$. We can therefore focus on the trigraph induced by $V(F) \cup \{s\}$.

Recall the vertex labels of Figure 6. We first contract a_1 and b_1 . This creates a vertex c_1 of red degree 4, and in particular it adds one to the red degree of s , which has now red degree at most 4 (see left-hand side of Figure 9). Now we will contract within A , and within B , not to increase more the red degree of s .



■ **Figure 9** The first three contractions of the fence gadget.

Vertices b_2 and b_3 now have only 4 neighbors in total in the fence gadget. Thus we can contract them into b_{23} and keep red degree at most 4 (see middle of Figure 9). In turn, a_2 and a_3 can be contracted into a_{23} for the same reason (see right of Figure 9). Then we contract b_4 with b_{23} , forming b_{234} , and a_4 with a_{23} , forming a_{234} . We contract b_5 and b_{234} into b_{2345} , and a_5 and a_{234} into a_{2345} . We finally contract b_6 and b_{2345} into b , and a_6 and a_{2345} into a . We contract a and c_1 , then we contract the resulting vertex and b .


Crucially the edge e stays black until the contraction of a_6 and a_{2345} , so c_1 maintains a red degree of at most 4. Also importantly, the only contractions involving a vertex of A and a vertex of B are the first and last two contractions. Thus s is incident in red to at most one vertex of the fence gadget. ◀

4.4 Propagation, wire, and long chain

A *vertical set* V consists of two vertices x, y combined with a fence gadget F attached to $\{x, y\}$. Thus $V = \{x, y\} \cup V(F)$. We call *vertical pair* the vertices x and y . We will usually add a superscript to identify the different copies of vertical sets: Every vertex whose label is of the form u^j belongs to the vertical set V^j .

The *propagation gadget* from V^j to $V^{j'}$ puts all the edges between V^j and $x^{j'}$ (and no other edge). We call these edges an *arc from V^j to $V^{j'}$* . We also say that the vertical set $V^{j'}$ is *guarded* by V^j . The pair $V^j, V^{j'}$ is said *adjacent*. Here, singleton $\{x_j\}$ plays the role of X (Definition 17) for the attachment rule of vertical set V_j .

The *propagation digraph* of G , denoted by $\mathcal{D}(G)$, has one vertex per vertical set and an arc between two vertical sets linked by an arc (in the previous sense). A (maximal) *wire* W is an induced subgraph of G corresponding in $\mathcal{D}(G)$ to a (maximal) out-tree on at least two vertices. See Figure 10 for an illustration of a wire made by simply concatenating propagation gadgets. Eventually $\mathcal{D}(G)$ will have out-degree at most 2, in-degree at most 2, and total degree at most 3.



■ **Figure 10** Left: A non-branching wire made of 4 vertical sets and 3 propagation gadgets. Every vertical set is guarded by the vertical set just to its left. Center: A more compact representation, which corresponds to the propagation digraph. Right: Symbolic representation of the long chain, that is, of the represented wire if $L = 4$.


The *children* of a vertical set V are the vertical sets that V guards. The *root of wire* W is the unique vertical set of in-degree 0 in $\mathcal{D}(W)$. The *leaves of wire* W are the vertical sets of out-degree 0 in $\mathcal{D}(W)$. A wire is said *primed* when the vertical pair of its root has been contracted.

A wire W is *non-branching* if every vertex of $\mathcal{D}(W)$ has out-degree at most 1; hence, $\mathcal{D}(W)$ is a directed path. A *long chain* is a wire W such that $\mathcal{D}(W)$ is a directed path on L vertices, where integer L will be specified later (and can be thought as logarithmic in the total number of fences which do not belong to vertical sets). Otherwise, if $\mathcal{D}(W)$ has at least one vertex with out-degree at least 2, wire W is said *branching*. A vertical set with two children is also said *branching*. See Figure 11 for an example of a branching wire with exactly one branching vertical set.

Constraints of the propagation gadget on a 4-sequence. We provide the proof that a contraction in a vertical set V is only possible when the vertical pair of all the vertical sets V' with a directed path to V in $\mathcal{D}(G)$ has been contracted.

► **Lemma 25.** *Let V^j and $V^{j'}$ be two vertical sets with an arc from V^j to $V^{j'}$. In a partial 4-sequence from G , any contraction involving two vertices of $V^{j'}$ has to be preceded by the contraction of x^j and y^j .*

Proof. We recall the notations of Section 4.3. Let F be the fence gadget attached to $S = \{x^j, y^j\}$, X the neighborhood of $V(F)$ outside of $S \cup V(F)$, and Y the vertices that are not in $S \cup V(F) \cup X$. (We always assume that the attachment rule is satisfied.) We have $x^{j'} \in X$ and $y^{j'} \in Y$, therefore by the second item of Lemma 23, their contraction has to be preceded by the contraction of x^j and y^j . Now applying the first item of Lemma 23 to the fence gadget F' attached to $S' = \{x^{j'}, y^{j'}\}$, and Lemma 22, any contraction involving a pair of $V^{j'}$ distinct from S' has to be preceded by the contraction of $x^{j'}$ and $y^{j'}$. ◀



■ **Figure 11** An example of a branching wire with its corresponding propagation digraph.

Henceforth, when we say that a contraction is *preceded* by another contraction, it includes the case that the two contractions are in fact the same. By a straightforward induction, we obtain the following from Lemma 25 (and Lemma 22).

► **Lemma 26.** *In a partial 4-sequence from G , any contraction involving a pair of vertices in a vertical set V has to be preceded by the contraction of the vertical pair of every vertical set V' such that there is a directed path from V' to V in $\mathcal{D}(G)$.*

Contracting wires. As the roots and leaves of wires will be connected to other gadgets, we postpone the description of how to contract wires until after building the overall construction G . Intuitively though, contracting a wire (in the vacuum) consists of contracting the vertical pair of its root, then its fence gadget by applying Lemma 24, and finally recursively contracting the subtrees rooted at its children. Since $\mathcal{D}(G)$ has total degree at most 3, every vertex has red degree at most 4 (for the at most 3 adjacent vertical sets, plus the pendant vertex of the fence gadget).


4.5 Binary AND gate

The *binary AND gate* (AND gadget, for short) simply consists of three vertical sets V^1, V^2, V^3 with an arc from V^1 to V^3 , and an arc from V^2 to V^3 . As usual, the vertical pairs of V^1, V^2 , and V^3 , are $\{x^1, y^1\}$, $\{x^2, y^2\}$, and $\{x^3, y^3\}$, respectively. We call the vertical sets V^1, V^2 the *inputs* of the AND gadget, and the vertical set V^3 the *output* of the AND gadget. See Figure 12 for an illustration.

Constraint of the AND gadget on a 4-sequence. By Lemma 25, we readily derive:

► **Lemma 27.** *Assume G contains an AND gadget with inputs V^1, V^2 , and output V^3 . In a partial 4-sequence from G , any contraction involving two vertices of V^3 has to be preceded by the contraction of x^1 and y^1 , and the contraction of x^2 and y^2 .*

Contraction of an AND gadget. Once V^1 and V^2 are contracted into single vertices, one can contract the vertical pair x^3, y^3 . This results in a vertex of red degree 2. Thus one can contract the fence gadget of V^3 by applying Lemma 24.




■ **Figure 12** An AND gadget, and the corresponding propagation digraph.

4.6 Binary OR gate

The *binary OR gate* (OR gadget) is connected to three vertical sets: two *inputs* V^1, V^2 , and one *output* V^3 . We start by building two vertical sets V, V' whose vertical pairs are $\{a, b\}$ and $\{c, d\}$, respectively. The edges ac and bd are added, as well as a vertex e adjacent to a and to c . Finally a fence is attached to $\{e\} \cup V \cup V'$.

The OR gadget is connected to its inputs and output, in the following way. Vertex a is made adjacent to x^1 and to y^1 (but not to their fence gadget). Similarly vertex b is linked to x^2 and y^2 . Finally x^3 is adjacent to all the vertices of the OR gadget, that is, $\{e\} \cup V \cup V'$ plus the vertices of the outermost fence. See Figure 13 for a representation of the OR gadget.



■ **Figure 13** An OR gadget attached to inputs V^1, V^2 and output V^3 , and its symbolic representation. These vertical sets are technically not part of the OR gate, so we represent them slightly dimmer. The outermost fence of the OR gadget can only be contracted when at least one the pairs x^1, y^1 and x^2, y^2 have been contracted. Only after that, can x^3 and y^3 be contracted together.

Constraint of the OR gadget on a 4-sequence. By design, one can only start contracting the OR gadget after the vertical pair of at least one of its two inputs has been contracted. This implies that no contraction can involve V^3 before at least one the vertical pairs $\{x^1, y^1\}$ and $\{x^2, y^2\}$ is contracted.

► **Lemma 28.** *Assume G contains an OR gadget attached to inputs V^1 and V^2 . In a partial 4-sequence from G , the contractions of a, b and of c, d have to be preceded by the contraction of x^1, y^1 or the contraction of x^2, y^2 .*

Proof. Assume none of the pairs $\{a, b\}, \{c, d\}, \{x^1, y^1\}, \{x^2, y^2\}$ have been contracted. Because of the fences, by Lemma 23, all the vertices $x^1, y^1, a, b, c, d, x^2, y^2$ and e are in distinct parts. Therefore contracting a and b would create a vertex of red degree at least 5, considering the (singleton) parts of x^1, y^1, c, d, e . Symmetrically contracting c and d yields at least five red neighbors, considering the (singleton) parts of x^2, y^2, a, b, e . ◀

From Lemma 28 we get the following.

► **Lemma 29.** *Assume G contains an OR gadget attached to inputs V^1 and V^2 . In a partial 4-sequence from G , no contraction involving a vertex of V^3 can happen before either the pair x^1, y^1 or the pair x^2, y^2 is contracted.*

Proof. Suppose neither x^1, y^1 nor x^2, y^2 is contracted. By Lemma 28, the pairs $\{a, b\}$ and $\{c, d\}$ cannot be contracted. By the first item of Lemma 23, no contraction can involve a vertex of $\{a, b, c, d, e\}$. As x^3 is adjacent to all these vertices but not y^3 , one cannot contract the vertical pair $\{x^3, y^3\}$. Hence by Lemma 22 no contraction can involve a vertex of V^3 . ◀

Contraction of the OR gadget. We now show how to contract the OR gate when the vertical pair of one of its inputs has been contracted.

► **Lemma 30.** *Assume that x^1 and y^1 have been contracted into z^1 , and that z^1, x^2 , and y^2 all have red degree at most 3. Then there is a partial 4-sequence that contracts the whole OR gadget to a single vertex with only three red neighbors: z^1, x^2 , and y^2 . (The same holds symmetrically if x^2 and y^2 have been contracted into a single vertex.)*

Proof. First contract a and b into vertex α of red degree 4. At this point the fence of $\{a, b\}$ cannot be contracted yet, as this would make the red degree of α go above 4. Hence we next contract c and d into γ , decreasing the red degree of α to 3. Note that γ has only 4 red neighbors: α, e, x^2, y^2 .

By Lemma 24, we can now contract the fence gadget of $\{a, b\}$ to a single vertex. Next we contract this latter vertex with α , and the resulting vertex with t ; we call α' the obtained vertex. Now γ has only red degree 3, so we can contract the fence gadget of $\{c, d\}$ to a single vertex that we further contract with γ ; we call γ' the obtained vertex. We contract α' and γ' in a vertex ε of red degree 3; its three red neighbors are z^1, x^2 , and y^2 . Again by Lemma 24, the outermost fence of the OR gadget can be contracted into a single vertex, that we finally contract with ε . This results in a vertex with three red neighbors: z^1, x^2 , and y^2 .

Throughout this process the red degree of z^1, x^2 , and y^2 never goes above 4. Indeed the red degree of these vertices is initially at most 3, while they have exactly one black neighbor in the entire OR gadget (so at most one part to be in conflict with). ◀


4.7 Variable gadget

We first describe the variable gadget for, say, a variable x . We start by attaching a fence on a set formed by three vertices: x, \top, \perp . We add two disjoint copies of an OR gadget, with vertices $a^\top, b^\top, c^\top, d^\top, e^\top$ (resp. $\{a^\perp, b^\perp, c^\perp, d^\perp, e^\perp\}$) plus the vertices in the fence gadgets; see Section 4.6 and Figure 13. We call T' , respectively U' , the vertex sets of the two OR gadgets. We link a^\top to x and to \top , and we link a^\perp to x and to \perp . We then add a vertex f^\top (resp. f^\perp), make it adjacent to c^\top (resp. c^\perp), and add a fence F'^\top (resp. F'^\perp) attached to $T' \cup \{f^\top\}$ (resp. $U' \cup \{f^\perp\}$). We add another vertex g^\top (resp. g^\perp), make it adjacent to c^\top (resp. c^\perp), and add a fence F^\top (resp. F^\perp) attached to $T = T' \cup V(F'^\top) \cup \{f^\top, g^\top\}$ (resp. $U = U' \cup V(F'^\perp) \cup \{f^\perp, g^\perp\}$).

The variable gadget is *half-guarded* by two vertical sets V^1, V^2 (with vertical pairs x^1, y^1 and x^2, y^2): We make \top adjacent to x^1 and y^1 , and we make \perp adjacent to x^2 and y^2 . Finally T guards a vertical set V^3 (with vertical pair x^3, y^3), and U guards a vertical set V^4 (with vertical pair x^4, y^4), that is, x^3 is fully adjacent to T , and x^4 is fully adjacent to U . Vertical set V^3 is the root of what we call the *wire of literal* $+x$, which is the maximal wire

containing V^3 . We add the $+$ to differentiate the *literal* from the *variable* x . Similarly, V^4 is the root of the *wire of literal* $\neg x$, that is, the maximal wire containing V^4 .

This finishes the construction of the variable gadget; see Figure 14 for an illustration.



■ **Figure 14** Top: A variable gadget half-guarded by V^1, V^2 , and with outputs V^3, V^4 . Bottom: Its compact representation.

Constraints of the variable gadget on a 4-sequence. Because of the fence gadget attached to $\{\top, x, \perp\}$, one has at some point to contract \top and \perp (be it with or without x). A first observation is that this has to wait that the vertical pair of V^1 or V^2 is contracted.

► **Lemma 31.** *Assume G has a variable gadget half-guarded by vertical sets V^1 and V^2 . In a partial 4-sequence from G , the contraction of \top and \perp has to be preceded by the contraction of the pair x^1, y^1 or of the pair x^2, y^2 .*

Proof. The pairs $\{x^1, y^1\}$, $\{x^2, y^2\}$, and $\{\top, \perp\}$ are contained in three disjoint sets S^1, S^2, S , respectively, to which a fence is attached. Thus before any of these pairs are contracted, by Lemma 23, a vertex outside $S^1 \cup S^2 \cup S$, like a^\perp , is in a different part than the six vertices $x^1, y^1, x^2, y^2, \top, \perp$. Therefore contracting \top and \perp would create a vertex with five red neighbors, considering the parts of $x^1, y^1, x^2, y^2, a^\perp$. ◀

We next show that no contraction is possible in U (resp. in T), while x and \perp (resp. x and \top) are not contracted.

► **Lemma 32.** *In a partial 4-sequence, the contractions of a^\perp and b^\perp (resp. a^\top and b^\top) and of c^\perp and d^\perp (resp. c^\top and d^\top) have to be preceded by the contraction of x and \perp (resp. x and \top). Therefore no contraction is possible in U (resp. T) before x and \perp (resp. x and \top) are contracted.*

Proof. Since the two statements are symmetric, we only show it with \perp . Assume none of the pairs $\{x, \perp\}$, $\{a^\perp, b^\perp\}$, $\{c^\perp, d^\perp\}$ are contracted. Because of the fence gadgets, by the first item of Lemma 23, the vertices $x, \perp, a^\perp, b^\perp, c^\perp, d^\perp, e^\perp, f^\perp$ are pairwise in distinct parts. Therefore contracting a^\perp and b^\perp or c^\perp and d^\perp would create a vertex of red degree at least 5. The structure of the fence gadgets in U thus prevents any contraction. \blacktriangleleft

We deduce that priming the wire of $\neg x$ (resp. $+x$) can only be done after x and \perp (resp. x and \top) are contracted.

► **Lemma 33.** *Assume that G has a variable gadget with outputs the vertical sets V^3 (root of the wire of $+x$) and V_4 (root of the wire of $\neg x$). In a partial 4-sequence from G , the contraction of x^4 and y^4 (resp. x^3 and y^3) has to be preceded by the contraction of x and \perp (resp. x and \top).*

Proof. By the second item of Lemma 23 applied to the fence attached to U , the pair x^4, y^4 cannot be contracted until U is not contracted to a single vertex. Thus by Lemma 32, the pair x^4, y^4 can only be contracted after the pair x, \perp is contracted. The other statement is obtained symmetrically. \blacktriangleleft

Contraction of the variable gadget. We show two options to contract a “half” of the variable gadget, either T and its fence, or U and its fence, into a single vertex.

► **Lemma 34.** *There is a partial 4-sequence that contracts x and \top together, and $T \cup F^\top$ into a single vertex. Symmetrically there is a partial 4-sequence that contracts x and \perp together, and $U \cup F^\perp$ into a single vertex.*

Proof. We first contract x and \top into a vertex $+x$. Observe that $+x$ has exactly three red neighbors: x^1, y^1 , and a^\perp . Thus $\{a^\top, b^\top, c^\top, d^\top, e^\top\}$ and their three fences can be contracted exactly like an OR gadget. So by Lemma 30, there is a partial 4-sequence that contracts all these vertices to a single vertex u , with three red neighbors ($+x, f^\top$, and g^\top). We can now contract u and f^\top into u' , followed by contracting their fence gadget F'^\top into a single vertex, by Lemma 24. That pendant vertex can be contracted to u' , and the result to g^\top , forming vertex v . Finally, again by Lemma 24, the fence F^\top attached to T can be contracted into a single vertex, which can be contracted with v .

The other sequence is the symmetric. \blacktriangleleft

We now see how the second “half” of the variable gadget can be contracted once the vertical pairs of the *half-guards* V^1, V^2 have been contracted.


► **Lemma 35.** *Assume $T \cup F^\top$ (resp. $U \cup F^\perp$) has been contracted into a single vertex u , and that the pairs $\{\top, x\}$ (resp. $\{\perp, x\}$), $\{x^1, y^1\}$, and $\{x^2, y^2\}$ have been contracted into $+x$ (resp. $\neg x$), z^1 , and z^2 , respectively. We further assume that the red degree of z^2 (resp. z^1) is at most 3. Then there is partial 4-sequence that contracts \top, x, \perp and their fence into a single vertex, and $U \cup F^\perp$ (resp. $T \cup F^\top$) into a single vertex.*

Proof. We contract \perp with $+x$ into v of red degree 4. This increases the red degree of z^2 by one, which remains at most 4. We then contract $U \cup F^\perp$ into a single vertex w , like in Lemma 34. We contract u and w into y , a vertex of red degree at most 3. Now v has degree 3. So we can contract its fence gadget by Lemma 24. We further contract v with its pendant neighbor, and finally with y . What results is a unique vertex with four red neighbors.

The other partial sequence is symmetric. \blacktriangleleft

4.8 Clause gadget

A *3-clause gadget* (or simply *clause gadget*) has for *inputs* three vertical sets V^1, V^2, V^3 , and for *output* one vertical set V^4 . It consists of combining two OR gadgets, and using long chains to make the OR gadgets *distant enough*. We add an OR gadget with input V^1 and V^2 , and output V . We then add a long chain from V to V' , and an OR gadget with input V' and V^3 , and output V^4 . The clause gadget is depicted in Figure 15. We call *first OR gadget* of the clause gadget, the one with output V , and *second OR gadget*, the one with output V^4 .



■ **Figure 15** Left: A 3-clause gadget. Right: A shorthand for the gadget.

Constraint of the clause gadget on a 4-sequence. As a consequence of Lemmas 26 and 29, we get that once a contraction involves the output, at least one of the vertical pairs of the inputs has been contracted.

► **Lemma 36.** *Assume G contains a clause gadget with inputs V^1, V^2, V^3 and output V^4 . In a partial 4-sequence from G , any contraction involving a vertex of V^4 is preceded by the contraction of the vertical pair of one of V^1, V^2 , or V^3 .*

Proof. Assume a contraction involves a vertex of V^4 . By Lemma 29 applied to the second OR gadget, the vertical pair of V^3 or of V' has to be contracted beforehand. If the vertical pair V^3 has been contracted, we conclude. Otherwise, the vertical pair of V' has been contracted, and by Lemma 26, this implies that the vertical pair of V has been contracted. By Lemma 29 applied to the first OR gadget, this in turn implies that the vertical pair of V^1 or of V^2 has been contracted. ◀


Contraction of the clause gadget. The OR gates of the clause gadgets will be contracted as specified in Lemma 30, while we wait the overall construction to describe the contraction of the wires.

4.9 Overall construction and correctness

Let $I = (C_1, \dots, C_m)$ be an instance of 3-SAT, that is, a collection of m 3-clauses over the variables x_1, \dots, x_n . We further assume that each variable appears at most twice positively, and at most twice negatively in I . The 3-SAT problem remains NP-complete with that restriction, and without $2^{o(n)}$ -time algorithm unless the ETH fails; see Theorem 6. We build a trigraph G that has twin-width at most 4 if and only if I is satisfiable. As trigraph G will satisfy the condition of Lemma 16, it can be replaced by a graph on $O(|V(G)|)$ vertices. We set L the length of the long chain to $2\lceil \log(5n + 3m) \rceil = O(\log n)$.

4.9.1 Construction

We now piece the gadgets described in the previous sections together.



■ **Figure 16** An example of the overall construction G on a 3-SAT instance with 6 variables and 8 clauses. The first two clauses are $\neg x_1 \vee x_3 \vee x_4$ and $x_1 \vee x_2 \vee \neg x_5$.

Variable to clause gadgets. For every variable x_i , we add a variable gadget half-guarded by V_i^1, V_i^2 , and with outputs V_i^3, V_i^4 . We add a long chain starting at vertical set V_i^3 (resp. V_i^4), and ending at a branching vertical set from which starts two long chains stopping at vertical sets $V_i^{\top,1}$ and $V_i^{\top,2}$ (resp. $V_i^{\perp,1}$ and $V_i^{\perp,2}$). Vertical set $V_i^{\top,1}$ (resp. $V_i^{\perp,1}$) serves as the input of the first clause gadget in which x_i appears positively (resp. negatively), while $V_i^{\top,2}$ (resp. $V_i^{\perp,2}$) becomes the input of the second clause gadget in which x_i appears positively (resp. negatively). If a literal has only one occurrence, then we omit the corresponding vertical set, and the long chain leading to it. We nevertheless assume that each literal has at least one occurrence, otherwise the corresponding variable could be safely assigned.

Clause gadgets to global output. For every $j \in [m]$, we add a long chain from the output of the clause gadget of C_j , to a vertical set, denoted by V_j^c . For every $j \in [2, m]$, we then add a long chain starting at V_j^c and ending at V_{j-1}^c . We add an arc from V_1^c to a new vertical set V^o , which we call the *global output*.

Global output back to half-guarding the variable gadgets. For every $i \in [n]$, we add two vertical sets $V_i^{1,r}, V_i^{2,r}$, and puts a long chain starting at $V_i^{1,r}$ and ending at $V_i^{2,r}$. We add a long chain from V^o to $V_1^{1,r}$. We also add a long chain from $V_i^{2,r}$ to $V_{i+1}^{1,r}$, for every $i \in [n-1]$. Finally we add a long chain from $V_i^{a,r}$ to V_i^a for every $a \in \{1, 2\}$ and every $i \in [n]$. Recall that V_i^1 and V_i^2 are half-guarding the variable gadget of x_i .

This finishes the construction of the graph $G = G(I)$. See Figure 16 for an illustration.

4.9.2 Correctness

If G has twin-width at most 4, then I is satisfiable. Let us consider the trigraph H obtained after the vertical pair of the global output V^o is contracted (see top-right picture of Figure 17). This has to happen in a 4-sequence by the first item of Lemma 23 applied to the fence of V^o . By Lemma 26, no contraction involving vertices of the vertical sets V_i^1, V_i^2 can have happened (for any $i \in [n]$). This is because there is a directed path in the propagation digraph $\mathcal{D}(G)$ from V^o to V_i^1 and V_i^2 .

Thus by Lemmas 31 and 34, for every variable x_i ($i \in [n]$), at most one of the wire of $+x_i$ and the wire of $\neg x_i$ has been primed. We can define the corresponding truth assignment \mathcal{A} : x_i is set to true if the wire of $\neg x_i$ is *not* primed, and to false if instead the wire of $+x_i$ is *not* primed. Besides, by Lemma 26 applied to the contraction in V^o , every vertical pair of a clause-gadget output has been contracted. Then Lemma 36 implies that the vertical pair of at least one input of each clause gadget has been contracted. But such a vertical pair can be contracted only if it corresponds to a literal set to true by \mathcal{A} . For otherwise, the root of the wire of that literal cannot be contracted. This implies that \mathcal{A} is a satisfying assignment.

If I is satisfiable, then G has twin-width at most 4. In what follows, when we write that we *can* contract a vertex, a set, or make a sequence of contractions, we mean that there is a partial 4-sequence that does the job. Let \mathcal{A} be a satisfying assignment of I . We start by contracting “half” of the variable gadget of each x_i . We add a subscript matching the variable index to the vertex and set labels of Figure 14. For every $i \in [n]$, we contract vertices x_i and \top_i together, and $T_i \cup F_i^\top$ to a single vertex v_i , if \mathcal{A} sets variable x_i to true, and x_i and \perp_i together, and $U_i \cup F_i^\perp$ to a single vertex v_i , if instead \mathcal{A} sets x_i to false. By Lemma 34, this can be done by a partial 4-sequence.


Next we contract the wire of $+x_i$ if $\mathcal{A}(x_i)$ is true, or the wire of $\neg x_i$ if $\mathcal{A}(x_i)$ is false. This is done as described in the end of Section 4.4. We contract the vertical pair of the root of the wire into z_i (the red degree of v_i goes from 1 to 2). We then contract its fence by Lemma 24. We can now contract the resulting vertex with z_i . Inductively, we contract the children of the current vertical set to single vertices, in a similar fashion. As the propagation digraph has degree at most 3, this never creates a vertex of red degree more than 4. At this point the trigraph corresponds to the top-left drawing of Figure 17.

At the leaves of the wire (the vertical sets $V_i^{\top,1}$, $V_i^{\top,2}$ or $V_i^{\perp,1}$, $V_i^{\perp,2}$), we make an exception, and only contract the vertical pair. We then contract, by Lemma 30, all the (non-already reduced) OR gadgets (involved in clause gadgets) at least one input of which has its vertical pair contracted. After that, applying Lemma 24, we finish the contraction of the OR inputs whose vertical pairs was contracted. Next we contract each output of a contracted OR gadget into a single vertex.

In those clauses where the third literal is not satisfied by \mathcal{A} , the output of the clause gadget is not contracted at that point. However, as \mathcal{A} is a satisfying assignment, the output of the first OR gate of the gadget is contracted. We then contract each vertical set of the long chain leading to the input of the second OR gate. We only contract the vertical pair of that input, and contract the incident OR gadget, by Lemma 30. We finally proceed by contracting the input vertical set into a single vertex, and the output vertical set into a single vertex.

At this point, each output of the clause gadgets is contracted into a single vertex. The (*not* strongly) connected component C of the propagation digraph $\mathcal{D}(G)$ containing the global output V_o is acyclic and has total degree at most 3. All the vertical sets of C with in-degree 0 (ant out-degree 1) in $\mathcal{D}(G)$ are the clause outputs, which have been contracted to single vertices. Thus each vertical set of C can be contracted to single vertices, by repeated use of Lemma 24, followed by contracting the pendant vertex (resulting from the fence contraction) with its unique (red) neighbor.

Note that this process terminates by contracting the half-guards V_i^1, V_i^2 (for every $i \in [n]$). The conditions of Lemma 35 are now satisfied, so we can finish contracting each variable gadget into two vertices that we further contract together. This results in a vertex of red degree 4. See top-right and bottom-left drawings of Figure 17 for the additional contractions leading to that point.



■ **Figure 17** The different stages of the contraction sequence. Gadgets in black are contracted to single vertices, while gadgets in gray are only partially contracted.

We can then contract the wire of the literal that was set to false by \mathcal{A} . Again, we only contract the vertical pair of the inputs of OR gates that are not contracted yet. Then we contract those OR gadgets, and finish by fully contracting the vertical set of those inputs. We next contract each output of the newly contracted OR gates. We eventually contract into a single vertex each vertical set of the long chain in clause gadgets where this was not already done.

At this point, the current trigraph H has only red edges (and corresponds to the bottom-right drawing of Figure 17). Thus it can be interpreted as a graph, and we write *degree* instead of *red degree*. H has $4n + 3m$ vertices of degree 3, n vertices of degree 4, and the rest of its vertices have degree 2. Because we added long chains to separate what now corresponds to vertices of degree at least 3, H is an $(\geq L)$ -subdivision of a graph on $5n + 3m$ vertices. Since $L = 2\lceil \log(5n + 3m) \rceil$, by Lemma 9, H finally admits a 4-sequence.

4.9.3 Wrapping up

The initial trigraph G comprises $O((n + m)L) = O(n \log n)$ gadgets and vertical sets, each consisting of $O(1)$ vertices. Hence $|V(G)| = O(n \log n)$. It is immediate that the construction of G can be made in polynomial time.

The red graph of G is a disjoint union of paths of length 12 (fence gadgets), and isolated vertices (the rest of the trigraph). Thus, by Lemma 16, there is a graph G' on $O(|V(G)|) = O(n \log n)$ vertices such that G' has twin-width at most 4 if and only if G has twin-width at most 4 if and only if I is satisfiable. This concludes the proof of Theorem 1 since a $2^{o(N/\log N)}$ -time algorithm deciding if an N -vertex graph has twin-width at most 4, would allow to decide 3-SAT where each variable appears at most twice positively and at most twice negatively, in time $2^{o(\frac{n \log n}{\log n + \log \log n})} = 2^{o(n)}$, contradicting the ETH.

References

- 1 Stefan Arnborg, Derek G Corneil, and Andrzej Proskurowski. Complexity of finding embeddings in a k -tree. *SIAM Journal on Algebraic Discrete Methods*, 8(2):277–284, 1987.
- 2 Hans L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996. doi:10.1137/S0097539793251219.
- 3 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.
- 4 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. URL: <https://arxiv.org/abs/2102.03117>, arXiv:2102.03117.
- 5 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.
- 6 Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 601–612. IEEE, 2020. doi:10.1109/FOCS46700.2020.00062.
- 7 Derek G. Corneil, Michel Habib, Jean-Marc Lanlignel, Bruce A. Reed, and Udi Rotics. Polynomial-time recognition of clique-width ≤ 3 graphs. *Discret. Appl. Math.*, 160(6):834–865, 2012. doi:10.1016/j.dam.2011.03.020.
- 8 Marek Cygan, Fedor V. Fomin, Alexander Golovnev, Alexander S. Kulikov, Ivan Mihajlin, Jakub Pachocki, and Arkadiusz Socala. Tight lower bounds on graph embedding problems. *J. ACM*, 64(3):18:1–18:22, 2017. doi:10.1145/3051094.
- 9 Anuj Dawar and Stephan Kreutzer. Parameterized complexity of first-order logic. *Electronic Colloquium on Computational Complexity (ECCC)*, 16:131, 2009. URL: <http://eccc.hpi-web.de/report/2009/131>.
- 10 Markus Sortland Dregi and Daniel Lokshtanov. Parameterized complexity of bandwidth on trees. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, volume 8572 of *Lecture Notes in Computer Science*, pages 405–416. Springer, 2014. doi:10.1007/978-3-662-43948-7_34.
- 11 Zdenek Dvorák, Daniel Král, and Robin Thomas. Testing first-order properties for subclasses of sparse graphs. *J. ACM*, 60(5):36:1–36:24, 2013. doi:10.1145/2499483.
- 12 Uriel Feige, MohammadTaghi Hajiaghayi, and James R. Lee. Improved approximation algorithms for minimum weight vertex separators. *SIAM J. Comput.*, 38(2):629–657, 2008. doi:10.1137/05064299X.
- 13 Michael R. Fellows, Frances A. Rosamond, Udi Rotics, and Stefan Szeider. Clique-width is NP-complete. *SIAM J. Discret. Math.*, 23(2):909–939, 2009. doi:10.1137/070687256.
- 14 Jörg Flum and Martin Grohe. Fixed-parameter tractability, definability, and model-checking. *SIAM J. Comput.*, 31(1):113–145, 2001. doi:10.1137/S0097539799360768.
- 15 Fedor V. Fomin, Daniel Lokshtanov, Ivan Mihajlin, Saket Saurabh, and Meirav Zehavi. Computation of hadwiger number and related contraction problems: Tight lower bounds. *ACM Trans. Comput. Theory*, 13(2):10:1–10:25, 2021. doi:10.1145/3448639.
- 16 Jakub Gajarský, Petr Hlinený, Daniel Lokshtanov, Jan Obdržálek, Sebastian Ordyniak, M. S. Ramanujan, and Saket Saurabh. FO model checking on posets of bounded width. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 963–974, 2015. doi:10.1109/FOCS.2015.63.

- 17 Jakub Gajarský, Petr Hlinený, Jan Obdržálek, and Sebastian Ordyniak. Faster existential FO model checking on posets. *Logical Methods in Computer Science*, 11(4), 2015. doi:10.2168/LMCS-11(4:8)2015.
- 18 Robert Ganian, Petr Hlinený, Daniel Král, Jan Obdržálek, Jarett Schwartz, and Jakub Teska. FO model checking of interval graphs. *Logical Methods in Computer Science*, 11(4), 2015. doi:10.2168/LMCS-11(4:11)2015.
- 19 Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- 20 Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. *J. ACM*, 64(3):17:1–17:32, 2017. doi:10.1145/3051095.
- 21 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. doi:10.1137/1.9781611973402.7.
- 22 Michel Habib and Christophe Paul. A simple linear time algorithm for cograph recognition. *Discret. Appl. Math.*, 145(2):183–197, 2005. doi:10.1016/j.dam.2004.01.011.
- 23 Petr Hlinený and Sang-il Oum. Finding branch-decompositions and rank-decompositions. *SIAM J. Comput.*, 38(3):1012–1032, 2008. doi:10.1137/070685920.
- 24 Russell Impagliazzo and Ramamohan Paturi. On the Complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.
- 25 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001. doi:10.1006/jcss.2001.1774.
- 26 Toshinobu Kashiwabara. NP-completeness of the problem of finding a minimal-clique number interval graph containing a given graph as a subgraph. In *Proc. 1979 Int. Symp. Circuit Syst.*, pages 657–660, 1979.
- 27 Tuukka Korhonen. Single-exponential time 2-approximation algorithm for treewidth. *CoRR*, abs/2104.07463, 2021. URL: <https://arxiv.org/abs/2104.07463>, arXiv:2104.07463.
- 28 Thomas Lengauer. Black-white pebbles and graph separation. *Acta Informatica*, 16:465–475, 1981. doi:10.1007/BF00264496.
- 29 Tatsuo Ohtsuki, Hajimu Mori, Ernest S. Kuh, Toshinobu Kashiwabara, and Toshio Fujisawa. One-dimensional logic gate assignment and interval graphs. In *The IEEE Computer Society's Third International Computer Software and Applications Conference, COMPSAC 1979, 6-8 November, 1979, Chicago, Illinois, USA*, pages 101–106. IEEE, 1979. doi:10.1109/COMPSAC.1979.762474.
- 30 Sang-il Oum. Approximating rank-width and clique-width quickly. *ACM Trans. Algorithms*, 5(1):10:1–10:20, 2008. doi:10.1145/1435375.1435385.
- 31 Neil Robertson and Paul D. Seymour. Graph minors. XIII. The disjoint paths problem. *J. Comb. Theory, Ser. B*, 63(1):65–110, 1995. doi:10.1006/jctb.1995.1006.
- 32 Sigve Hortemo Sæther and Martin Vatshelle. Hardness of computing width parameters based on branch decompositions over the vertex set. *Theor. Comput. Sci.*, 615:120–125, 2016. doi:10.1016/j.tcs.2015.11.039.
- 33 James B. Saxe. Dynamic-programming algorithms for recognizing small-bandwidth graphs in polynomial time. *SIAM J. Algebraic Discret. Methods*, 1(4):363–369, 1980. doi:10.1137/0601042.
- 34 André Schidler and Stefan Szeider. A SAT approach to twin-width. *CoRR (accepted at ALENEX '22)*, abs/2110.06146, 2021. URL: <https://arxiv.org/abs/2110.06146>, arXiv:2110.06146.
- 35 Craig A. Tovey. A simplified NP-complete satisfiability problem. *Discret. Appl. Math.*, 8(1):85–89, 1984. doi:10.1016/0166-218X(84)90081-7.