Coloring Hardness on Low Twin-Width Graphs

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

Abstract

As the class $\mathcal{T}_4$ of graphs of twin-width at most 4 contains every finite subgraph of the infinite grid and every graph obtained by subdividing each edge of an n-vertex graph at least $2 \log n$ times, most NP-hard graph problems, like Max Independent Set, Dominating Set, Hamiltonian Cycle, remain so on $\mathcal{T}_4$. However, Min Coloring and k-Coloring are easy on both families because they are 2-colorable and 3-colorable, respectively.

We show that Min Coloring is NP-hard on the class $\mathcal{T}_3$ of graphs of twin-width at most 3. This is the first hardness result on $\mathcal{T}_3$ for a problem that is easy on cographs (twin-width 0), on trees (whose twin-width is at most 2), and on unit circular-arc graphs (whose twin-width is at most 3). We also show that for every $k \geqslant 3$, k-Coloring is NP-hard on $\mathcal{T}_4$. We finally make two observations: (1) there are currently very few problems known to be in P on $\mathcal{T}_d$ (graphs of twin-width at most d) and NP-hard on $\mathcal{T}_{d+1}$ for some nonnegative integer d, and (2) unlike $\mathcal{T}_4$, which contains every graph as an induced minor, the class $\mathcal{T}_3$ excludes a fixed *planar* graph as an induced minor; thus it may be viewed as a special case (or potential counterexample) for conjectures about classes excluding a (planar) induced minor. These observations are accompanied by several open questions.

Funding This work has been supported by the French National Research Agency through the project TWIN-WIDTH with reference number ANR-21-CE48-0014.

1 Introduction

The graph parameter *twin-width* was introduced in 2020 [15]. The family of graph classes of bounded twin-width is broad and diverse: it includes classes of bounded clique-width, d-dimensional grids, classes excluding a fixed minor [15], some cubic expanders [11]. Moreover, there are algorithms that work on any class of effectively[1] bounded twin-width: a fixed-parameter tractable first-order model checking algorithm [15], single-exponential parameterized algorithms [12], improved approximation algorithms [12, 6], fast shortest-path algorithms [12, 4], etc. Most NP-hard graph problems remain intractable on graphs of twin-width at most 4. This paper tackles the hardness of coloring graphs of low twin-width.

For any nonnegative integer d, we denote by $\mathcal{T}_d$ the class of graphs of twin-width at most d. Finite subgraphs of the infinite planar grid [15] and ($\geqslant 2 \log n$)-subdivisions of n-vertex graphs [5] are in $\mathcal{T}_4$. Problems such as Max Independent Set, Vertex Cover, Max Clique, Feedback Vertex Set, Dominating Set, and Max Induced Matching are NP-hard on the latter family, while Hamiltonian Path and Hamiltonian Cycle are NP-hard on the former. Therefore all these problems are NP-hard on $\mathcal{T}_4$. Notably, Min Coloring is easy on subgraphs of the grid, since they are bipartite. It is also easy on *strict* subdivisions (i.e., when every edge is subdivided at least once), since such graphs are always 3-colorable and bipartiteness can be checked in linear time. We show that Min Coloring is already NP-hard on $\mathcal{T}_3$. This is the first problem that is shown NP-hard on $\mathcal{T}_3$ while being in P on classes known to have twin-width at most 3 (cographs, trees, unit circular-arc graphs).

[1] A class $\mathcal{C}$ has *effectively bounded twin-width* if there is an integer d and a polynomial-time algorithm that outputs a d-sequence (see Section 2 for all relevant definitions) of any input graph from $\mathcal{C}$.

▶ **Theorem 1.** *Min Coloring is* NP*-hard and, unless the ETH fails, requires* $2^{\Omega(\sqrt{n})}$ *time on* n*-vertex graphs of twin-width at most 3 (even if a 3-sequence is provided).*

In the previous theorem, the ETH (for Exponential-Time Hypothesis) [28] asserts that there exists a constant $\lambda > 1$ such that n-variable 3-Sat cannot be solved in time $O(\lambda^n)$. The reduction requires the number of colors to grow with n.

We also show that 3-Coloring is NP-hard on $\mathcal{T}_4$. Note that planar graphs have twin-width at least 7 [33] (and at most 8 [27]). So we cannot just invoke the hardness of 3-Coloring in planar graphs. One would need to provide 4-sequences for the graphs produced by the reduction. Another option would be to tune long subdivisions and turn them into hard 3-Coloring instances. Instead, we give an *ad hoc* construction for which both the correctness and membership in $\mathcal{T}_4$ are easy to verify.

▶ **Theorem 2.** *3-Coloring is* NP*-hard on graphs of twin-width at most 4.*

This again holds even if a 4-sequence is provided and directly implies that k-Coloring is also hard on $\mathcal{T}_4$, for any fixed $k \geqslant 3$.

On the algorithmic side, most graph problems are in P on $\mathcal{T}_1$ since this class has bounded clique-width [13] and is a subclass of permutation graphs [3]. Although 1-sequences can be computed in polynomial time in $\mathcal{T}_1$ [14], and even in linear time [3], the corresponding algorithms do not require a sequence to be provided as part of the input. Max Independent Set (and consequently, Vertex Cover and Max Clique) can be solved in polynomial time in $\mathcal{T}_2$ if 2-sequences are provided as part of the input [12].

While $\mathcal{T}_1$ can be recognized in polynomial time, it is NP-hard to decide if a graph (of $\mathcal{T}_5$) is in $\mathcal{T}_4$. The complexity of the recognition of $\mathcal{T}_2$ and of $\mathcal{T}_3$ is open.

	$\mathcal{T}_1$	$\mathcal{T}_2$	$\mathcal{T}_3$	$\mathcal{T}_4$
Min Coloring	P [3, 21]	?	NP-c (Theorem 1)	
3-Coloring	P [13, 16]	?	?	NP-c (Theorem 2)
k-Coloring, $k \geqslant 3$	P [13, 16]	?	?	NP-c (Corollary 3)
Max Independent Set	P [13, 16]	P [12]	?	NP-c [5]
Vertex Cover	P [13, 16]	P [12]	?	NP-c [5]
Max Clique	P [13, 16]	P [12]	?	NP-c [5]
Feedback Vertex Set	P [13, 16]	?	?	NP-c [5]
Dominating Set	P [13, 16]	?	?	NP-c [5]
Max Induced Matching	P [13, 16]	?	?	NP-c [5]
Hamiltonian Path	P [3, 19]	?	?	NP-c [15, 30]
Hamiltonian Cycle	P [3, 20]	?	?	NP-c [15, 30]
Recognition	P [14, 3]	?	?	NP-c [5]

Table 1 Complexity of some of the main NP-complete graph problems in $\mathcal{T}_d$ for $d = 1, 2, 3, 4$. The P in gray in the $\mathcal{T}_2$ column means that 2-sequences are required by the known algorithm.

Table 1 suggests the task of closing these gaps.

▶ **Question 1.** *Establish, for a problem* Π *of Table 1, a nonnegative integer* d *such that* Π *on* $\mathcal{T}_d$ *is in* P *and* Π *on* $\mathcal{T}_{d+1}$ *is* NP*-hard.*

To our knowledge, the only problems for which Question 1 is settled are Firefighter and Restricted Vertex Multicut, which are in P on $\mathcal{T}_1$ (as they are polynomial-time solvable on permutation graphs [24, 34]) and NP-hard on $\mathcal{T}_2$ (as they are hard on trees [23, 17]).

As the current paper is on graph coloring, we ask the following question.

▶ **Question 2.** *Is* Min Coloring *on* $\mathcal{T}_2$ *in* P*?*

As we mentioned, Min Coloring is the first problem to be shown NP-hard on $\mathcal{T}_3$ while being tractable on classes known to have twin-width at most 3: cographs (which coincide with $\mathcal{T}_0$ [15]), trees (which are in $\mathcal{T}_2$ [15]), and unit circular-arc graphs (which are in $\mathcal{T}_3$ [9, below Theorem 5.5]). Without this requirement, other examples would include Achromatic Number, which is NP-hard on cographs [8], Bandwidth, which is NP-hard on trees [25], and Partial Representation Extension, which is NP-hard on unit circular-arc graphs [22]. Theorem 1 thus involved a novel encoding that incurred a quadratic blow-up. We wonder whether this blow-up can be avoided.

▶ **Question 3.** *Can* Min Coloring *be solved in time* $2^{O(\sqrt{n})}$ *on* n*-vertex graphs of* $\mathcal{T}_3$*?*

The class $\mathcal{T}_1$ is now well understood [3]. In particular, it is a subclass of permutation graphs with bounded clique-width. We believe that the classes $\mathcal{T}_2$ and $\mathcal{T}_3$ have some still hidden structure. While the membership problem in $\mathcal{T}_4$ is NP-hard [5], that of $\mathcal{T}_2$ and of $\mathcal{T}_3$ are open. It was proven that *weakly sparse* subclasses (i.e., excluding a biclique $K_{t,t}$ as a subgraph) of $\mathcal{T}_2$ have bounded treewidth, whereas this is not the case for $\mathcal{T}_3$ [7].

However, bounded-degree subclasses of $\mathcal{T}_3$ have bounded treewidth. This is because bounded-degree graphs of large treewidth admit subdivisions of large walls or their line graphs as induced subgraphs [32] and those graphs have twin-width (exactly) 4 [2]. Furthermore, the latter result combined with [1, 10] (extending [32]) implies that every subclass of $\mathcal{T}_3$ without large subdivided cliques as subgraphs has bounded treewidth. As 3-Coloring is in P on any class of bounded treewidth, if 3-Coloring is NP-hard on $\mathcal{T}_3$, then the hard instances must contain arbitrarily large clique subdivisions as subgraphs. The hard 3-Coloring instances that the proof of Theorem 2 builds do not: they have a single vertex of degree more than 4.

▶ **Question 4.** *Is 3-*Coloring *on* $\mathcal{T}_3$ *in* P*?*

The same question holds for Max Independent Set.

▶ **Question 5.** *Is* Max Independent Set *on* $\mathcal{T}_3$ *in* P*?*

Another consequence of the abovementioned result of [2] is that $\mathcal{T}_3$ excludes a fixed planar graph as an induced minor. Thus Questions 4 and 5 are special cases of the same questions on any class excluding a planar induced minor, which was previously raised for Max Independent Set [18]. The class $\mathcal{T}_3$ is a good candidate for a negative answer to the latter question (on the pessimistic side), or Question 5 could serve as a preliminary step in positively answering it (on the optimistic side). There are other questions and conjectures on classes excluding a planar induced minor. They can be revisited on the particular class $\mathcal{T}_3$, like for instance the following conjecture appearing in [26].

▶ **Question 6.** *Is there a universal constant* c *such that every graph of* $\mathcal{T}_3$ *admits a balanced separator included in the neighborhood of at most* c *vertices?*

We recall the relevant definitions and notation in Section 2, show Theorem 1 in Section 3, and Theorem 2 in Section 4.

2 Preliminaries

For two integers i and j, 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 $[1, i]$.

2.1 Standard graph-theoretic definitions and notation

We denote by $V(G)$ and $E(G)$ the set of vertices and edges of a graph G, respectively. For $S \subseteq V(G)$, the *subgraph of G induced by S*, denoted $G[S]$, is obtained by removing from G all the vertices that are not in S. Then $G - S$ is a shorthand for $G[V(G) \setminus S]$. We denote by $N_G(v)$ the set of neighbors of v in G. A *subdivision* of a graph G is any graph H obtained from G by replacing edges e of G by paths with at least one edge whose extremities are the endpoints of e. A $(\geqslant s)$*-subdivision* is a subdivision where every edge is replaced by a path with at least s internal vertices.

In this paper, a coloring of a graph is implicitly assumed to be a proper vertex-coloring.

2.2 Trigraphs, partition sequences, and twin-width

A *trigraph* G has vertex set $V(G)$, black edge set $E(G)$, red edge set $R(G)$ such that $E(G) \cap R(G) = \emptyset$ (and $E(G), R(G) \subseteq \binom{V(G)}{2}$). Two vertices u, v such that $uv \in R(G)$ are called *red neighbors*. The *red degree* of u is its number of red neighbors. The *maximum red degree* of G is the maximum red degree among all its vertices.

Given a (tri)graph G and a partition $\mathcal{P}$ of $V(G)$, the *quotient trigraph $G/\mathcal{P}$* is the trigraph with vertex set $\mathcal{P}$, where PP' is a black edge if these two parts are fully adjacent via black edges (i.e., for every $u \in P$ and every $v \in P'$, $uv \in E(G)$), and a red edge if there is $u \in P$ and $v \in P'$ such that $uv \in R(G)$ or $u_1, u_2 \in P$ and $v_1, v_2 \in P'$ such that $u_1v_1 \in E(G)$ and $u_2v_2 \notin E(G)$.

A *partition sequence* of an n-vertex (tri)graph G is a sequence $\mathcal{P}_n, \mathcal{P}_{n-1}, \ldots, \mathcal{P}_1$ of partitions of $V(G)$ such that $\mathcal{P}_n = \{\{v\} : v \in V(G)\}$, $\mathcal{P}_1 = \{V(G)\}$, and for every $i \in [n-1]$, $\mathcal{P}_i$ is obtained from $\mathcal{P}_{i+1}$ by merging $P, P' \in \mathcal{P}_{i+1}$ into $P \cup P'$. A *d-sequence* of G is a partition sequence $\mathcal{P}_n, \ldots, \mathcal{P}_1$ such that for every $i \in [n]$, the maximum red degree of $G/\mathcal{P}_i$ is at most d. The twin-width of a (tri)graph is the least integer d such that it admits a d-sequence.

A partition sequence is called *partial* when we relax the condition that the last partition has a single part. We say that a trigraph is *fully red* if it does not have any black edges. We will use the simple fact that turning some (or all) black edges red cannot decrease the twin-width of a trigraph. Therefore, when showing twin-width upper bounds, we may sometimes assume that all the edges (black and red) are in fact red, if this simplifies a later argument.

3 Hardness of Coloring Graphs of Twin-Width 3

As a decision problem, Min Coloring inputs a graph G and an integer k and asks whether the chromatic number of G, $\chi(G)$, is at most k. As a function problem, one is given the mere graph G and has to output a (proper) $\chi(G)$-coloring of G. We show that Min Coloring is NP-hard on $\mathcal{T}_3$, even in its decision form and when a 3-sequence of G is given in input. As Min Coloring is polynomial-time solvable on cographs, on trees, and on unit circular-arc graphs (all of which have twin-width at most 3), and planar graphs are not included in $\mathcal{T}_3$, this requires finding a novel kind of encoding of rich structures (here, 3-Sat instances) onto graphs of twin-width at most 3.

▶ **Theorem 1.** *Min Coloring is NP-hard and, unless the ETH fails, requires $2^{\Omega(\sqrt{N})}$ time on N-vertex graphs of twin-width at most 3 (even if a 3-sequence is provided).*

Proof. We reduce from 3-Sat, which is NP-complete [31], and unless the ETH fails, its n-variable m-clause instances cannot be solved in time $2^{o(n+m)}$ by the so-called Sparsification

Lemma [29]; it is indeed shown that, under the ETH, n-variable 3-SAT cannot be solved in time $2^{o(n)}$ even on instances with $O(n)$ clauses. Let $x_1, \ldots, x_n$ be the variables of a 3-SAT instance φ, and $c_1, \ldots, c_m$ be its clauses.

We describe the construction of a graph $G := G(\varphi)$ with chromatic number (at most) $2n$ if and only if φ is satisfiable. We start with two paths P, P' each on $2np$ vertices with $p := 2n + m$, which we partition evenly into p subpaths on $2n$ vertices. Let $A_1, A_2, \ldots, A_p$ (resp. $B_1, B_2, \ldots, B_p$) be the vertex sets of the subpaths of P (resp. P') from left to right. We set $A := \bigcup_{i \in [p]} A_i = V(P)$ and $B := \bigcup_{i \in [p]} B_i = V(P')$. We denote by $a_{i,1}, a_{i,2}, \ldots, a_{i,2n}$ the $2n$ vertices of the subpath $P[A_i]$ from left to right. Similarly we denote by $b_{i,1}, b_{i,2}, \ldots, b_{i,2n}$ the $2n$ vertices of the subpath $P'[B_i]$ from left to right. We now define $G[A]$ (resp. $G[B]$) as $P^{\leqslant 2n-1}$ (resp. $P'^{\leqslant 2n-1}$), that is, for every $u, v \in A$ (resp. $u, v \in B$), $uv \in E(G)$ whenever u and v are at distance at most $2n-1$ in P (resp. in P'). In particular, each A_i (and each B_i) is a clique of size $2n$, and for all $i, i' \in [p]$ with $i < i'$, it holds that $a_{i,j}$ and $a_{i',j'}$ are adjacent if and only if $i' = i+1$ and $j' < j$. There is no edge between A and B.

For every $i \in [p]$, we add one vertex v_i adjacent to a subset of $A_i \cup B_i$, as follows.

- For every $i \in [n]$, $N_G(v_{2i-1}) = (A_{2i-1} \cup B_{2i-1}) \setminus \{a_{2i-1,2i-1}, a_{2i-1,2i}, b_{2i-1,2i-1}\}$ and $N_G(v_{2i}) = (A_{2i} \cup B_{2i}) \setminus \{a_{2i,2i-1}, a_{2i,2i}, b_{2i,2i}\}$.
- For every $i \in [2n+1, 2n+m]$, say c_{i-2n} is the clause $s_1 x_{j_1} \vee s_2 x_{j_2} \vee s_3 x_{j_3}$ (repeat a literal if it had only two), where s_1, s_2, s_3 are signs in $\{\neg, \varepsilon\}$ (ε is the positive sign). Then, $N_G(v_i) = (B_i \setminus \{b_{i,2j_1}, b_{i,2j_2}, b_{i,2j_3}\}) \cup \{a_{i,2j_1-f(s_1)}, a_{i,2j_2-f(s_2)}, a_{i,2j_3-f(s_3)}\}$ with $f(\neg) = 0$ and $f(\varepsilon) = 1$.

This finishes the construction of G; see Figure 1 for an illustration. Note that G has $N := (4n+1)p = O(n(n+m))$ vertices, and that it can be constructed in polynomial time from φ. Thus, provided the reduction is correct, which we next prove, a $2^{o(\sqrt{N})}$-time algorithm for MIN COLORING on the produced instances would imply a $2^{o(n+m)}$-time algorithm for n-variable m-clause 3-SAT, hence refute the ETH.

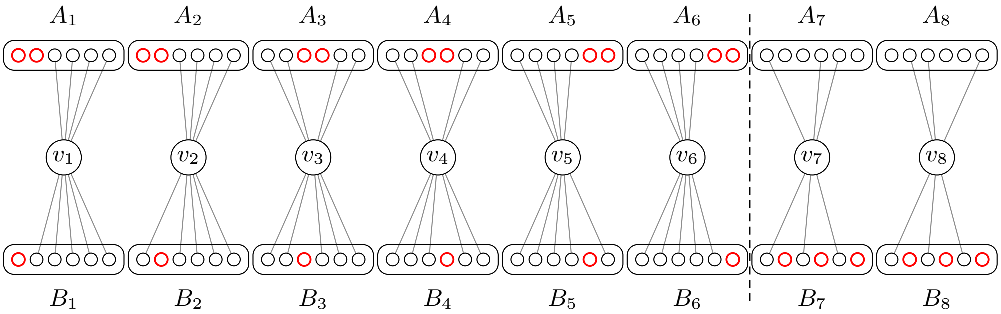

Figure 1 The graph $G(\varphi)$ for φ consisting of two clauses $c_1 = x_1 \vee \neg x_2 \vee x_3$ and $c_2 = \neg x_1 \vee x_2 \vee \neg x_3$. The variable gadgets and clause gadgets are delimited by the vertical dashed line. Not to clutter the picture, we have not drawn the edges in $G[A \cup B]$.

If φ is satisfiable, then G is $2n$-colorable. Let $\mathcal{A}$ be a satisfying assignment for φ. We then describe a (proper) $2n$-coloring c of G.

- For every $i \in [p]$ and every $j \in [2n]$, we set $c(b_{i,j}) = j$.
- For every $i \in [p]$ and every $j \in [n]$, we set $c(a_{i,2j-1}) = 2j-1$ and $c(a_{i,2j}) = 2j$ if $\mathcal{A}$ assigns x_j to true, or $c(a_{i,2j}) = 2j-1$ and $c(a_{i,2j-1}) = 2j$, otherwise.

As the $2n$ colors repeat with period $2n$ along each path of $\{P, P'\}$, no edge of $G[A \cup B]$ is monochromatic. We now simply need to argue that for every vertex v_i, there is at least one color c' of $[2n]$ not present in $N_G(v_i)$; and set $c(v_i) = c'$.

For every $i \in [2n]$, we can set $c(v_i) = i$, as this color does not appear in $N_G(v_i)$. For every $i \in [2n+1, 2n+m]$, with $c_{i-2n} = s_1 x_{j_1} \lor s_2 x_{j_2} \lor s_3 x_{j_3}$ and $s_1, s_2, s_3 \in \{\neg, \varepsilon\}$, there is an $h \in \{1,2,3\}$ such that $\mathcal{A}$ satisfies $s_h x_{j_h}$. Then we can set $c(v_i) = 2j_h$. Indeed, if $\mathcal{A}$ sets x_{j_h} to true, then $f(s_h) = 1$ and $a_{i,2j_h-1}$ is colored $2j_h - 1$, whereas if $\mathcal{A}$ sets x_{j_h} to false, then $f(s_h) = 0$ and $a_{i,2j_h}$ is colored $2j_h - 1$; thus, in both cases, color $2j_h$ is still available for v_i. See Figure 2 for an example.

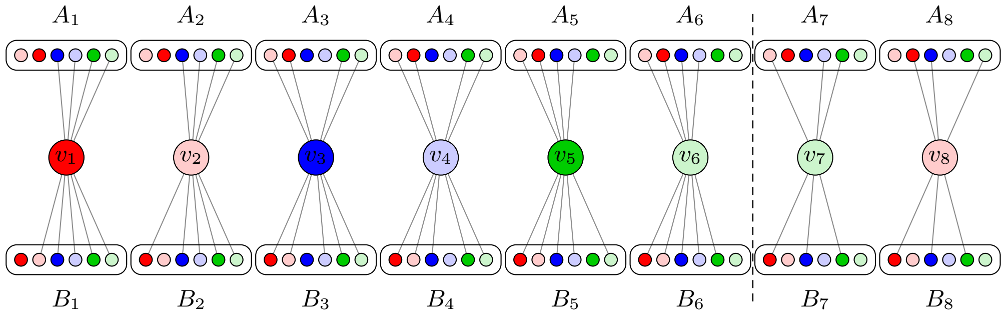

Figure 2 A (proper) $2n$-coloring of the graph $G(\varphi)$ of Figure 1 corresponding to the satisfying assignment $\mathcal{A} = \{x_1 \mapsto \text{false}, x_2 \mapsto \text{true}, x_3 \mapsto \text{true}\}$. The coloring of v_7 (resp. v_8) witnesses that c_1 (resp. c_2) is satisfied by literal x_3 (resp. $\neg x_1$).

If G is $2n$-colorable, then φ is satisfiable. Fix a (proper) coloring $c \colon V(G) \to [2n]$ of G, with $c(S) := \{c(v) : v \in S\}$ for any $S \subseteq V(G)$. As B_1 is a clique, we can further assume that $c(b_{1,j}) = j$ for every $j \in [2n]$. By a straightforward induction, for every $i \in [p-1]$ and $j \in [2n]$, $c(a_{i,j}) = c(a_{i+1,j})$ and $c(b_{i,j}) = c(b_{i+1,j}) = j$, and for every $i \in [p]$, $c(A_i) = [2n] = c(B_i)$. Indeed both A_1 and B_1 are $2n$-vertex cliques, thus $c(A_1) = c(B_1) = [2n]$, and when $2n$-coloring $P^{\leqslant 2n-1}$ (resp. $P'^{\leqslant 2n-1}$) from left to right, the only available color for $a_{i+1,j}$ (resp. $b_{i+1,j}$) is $c(a_{i,j})$ (resp. $c(b_{i,j})$).

We first show that, for every $i \in [p]$ and $j \in [n]$, $\{c(a_{i,2j-1}), c(a_{i,2j})\} = \{2j-1, 2j\}$. Assume for contradiction that there is an $h \in \{c(a_{i,2j-1}), c(a_{i,2j})\} \setminus \{2j-1, 2j\}$. By the previous paragraph, observe that h does not depend on i. This implies that $2j - 1 \notin \{c(a_{i,2j-1}), c(a_{i,2j})\}$ or $2j \notin \{c(a_{i,2j-1}), c(a_{i,2j})\}$. In the former case, v_{2j-1} would need a $(2n+1)$-st color, whereas in the latter case, the same would happen to v_{2j}.

Let $\mathcal{A}$ be the truth assignment that sets, for every $j \in [n]$, x_j to true if $c(a_{i,2j-1}) = 2j-1$ and $c(a_{i,2j}) = 2j$, and to false, if $c(a_{i,2j-1}) = 2j$ and $c(a_{i,2j}) = 2j-1$ (again, note that this does not depend on i). Let us show that $\mathcal{A}$ satisfies all the clauses of φ. For each $i \in [2n+1, 2n+m]$, say $c_{i-2n} = s_1 x_{j_1} \lor s_2 x_{j_2} \lor s_3 x_{j_3}$ with $s_1, s_2, s_3 \in \{\neg, \varepsilon\}$. Since v_i is adjacent to $B_i \setminus \{b_{i,2j_1}, b_{i,2j_2}, b_{i,2j_3}\}$, it holds that $c(v_i) \in \{2j_1, 2j_2, 2j_3\}$; say, $c(v_i) = 2j_h$ with $h \in \{1,2,3\}$. This implies that $c(a_{i,2j_h - f(s_h)}) \neq 2j_h$, thus $c(a_{i,2j_h - f(s_h)}) = 2j_h - 1$. Therefore, $s_h = \varepsilon$ and $\mathcal{A}$ sets x_{j_h} to true, or $s_h = \neg$ and $\mathcal{A}$ sets x_{j_h} to false. In both cases, the literal $s_h x_{j_h}$ of c_{i-2n} is satisfied by $\mathcal{A}$.

G admits a 3-sequence. We describe a partition sequence for G of width at most 3. The sequence has two stages. In the first stage, each A_i and each B_i is merged into a single part. Throughout this stage, we denote by P_i (resp. P'_i) the current part containing $a_{i,1}$ (resp. $b_{i,1}$). For j going from 2 to $2n$, for i going from 1 to p, merge P_i and singleton $\{a_{i,j}\}$, and merge P_i' and singleton $\{b_{i,j}\}$.

Let us argue that this partial sequence has width at most 3. At any point, the singleton parts $\{a_{i,j+1}\}, \dots, \{a_{i,2n}\}$ *in between* $P_i = \{a_{i,1}, \dots, a_{i,j}\}$ and P_{i+1} are all fully adjacent to $P_i \cup P_{i+1}$, some of these singletons are adjacent to v_i, and they have no other adjacencies. Thus these parts have red degree 0. Note that the singleton parts $\{a_{p,j+1}\}, \dots, \{a_{p,2n}\}$ do not have incident red edges either. Each singleton part $\{v_i\}$ has red degree at most 2, without any red neighbor outside $\{P_i, P_i'\}$. Finally, at any point, each P_i has at most three red neighbors: P_{i-1} (if it exists), v_i, and P_{i+1} (if it exists). Indeed, all the other parts are either fully nonadjacent to P_i or have red degree 0. Symmetrically, P_i' has always at most three red neighbors.

After the first stage is completed, the resulting trigraph is fully red (we may well assume that $n \geqslant 2$), has $3p$ vertices, and consists of two p-vertex paths whose ith vertices have an additional shared neighbor of degree 2, for every $i \in [p]$. In the second stage, we finish the 3-sequence, as follows. Let us call Q the part containing v_1. For i going from 2 to p, we merge P_1 and P_i, and keep denoting the resulting part by P_1, we merge P_1' and P_i', and denote the resulting part by P_1', and then we merge Q and $\{v_i\}$. Finally, only three parts remain that we merge to a single part in any way. One can observe that at any time of the second stage when at least three parts are left, every part P_i or P_i' (including P_1, P_1') has at most three red neighbors, and Q, as well as every part $\{v_i\}$, has two red neighbors. So this finishes the description of a 3-sequence of $G(\varphi)$.

As this sequence can be computed in polynomial time for any formula φ, the hardness of Min Coloring on $\mathcal{T}_3$ further holds when a 3-sequence is provided as part of the input. ◀

4 Hardness of 3-Coloring Graphs of Twin-Width 4

The hardness proof of the previous section requires an unbounded number of colors. In this section we show that deciding if a graph of twin-width at most 4 is 3-colorable is NP-hard. We also get a similar ETH lower bound, and readily extend this hardness result to k-Coloring.

▶ **Theorem 2.** *3-Coloring is* NP*-hard and, unless the ETH fails, requires $2^{\Omega(\sqrt{N})}$ time on N-vertex graphs of twin-width at most 4 (even if a 4-sequence is provided).*

Proof. We reduce from the NP-complete Not-All-Equal 3-Sat problem [35] (where each clause is on three distinct variables). Furthermore, by the Sparsification Lemma [29], unless the ETH fails, there is no $2^{o(n)}$-time algorithm for n-variable $O(n)$-clause instances. Indeed, the reduction from 3-Sat to Not-All-Equal 3-Sat in [35] is linear. Let $x_1, \dots, x_n$ be the variables of an instance φ and $c_1, \dots, c_m$ its clauses, with $m = O(n)$. We build a graph $G := G(\varphi)$ as follows.

For each clause $c_j = \ell_1 \vee \ell_2 \vee \ell_3$, add a triangle t_j on three vertices u_j, v_j, w_j, where u_j, v_j, w_j correspond to ℓ_1, ℓ_2, ℓ_3, respectively.

For each variable x_i, start with the m-vertex path $x_{i,1} x_{i,2} \cdots x_{i,m}$. Let $j_1 < j_2 < \cdots < j_q$ be the clause indices in which x_i appears, with $q \geqslant 1$. For each $h \in [q-1]$, let $s_h, s_{h+1} \in \{\neg, \varepsilon\}$ be the signs of x_i in clauses c_{j_h} and $c_{j_{h+1}}$. If either

- $|j_{h+1} - j_h|$ is even and $s_h \neq s_{h+1}$, or
- $|j_{h+1} - j_h|$ is odd and $s_h = s_{h+1}$,

subdivide once the edge $x_{i,j_{h+1}-1} x_{i,j_{h+1}}$ and denote the new vertex by $x'_{i,j_{h+1}}$. Let P_i be the resulting (still induced) path, which we refer to as the *variable path* (of x_i).

We add one vertex z adjacent to every vertex of $\bigcup_{i \in [n]} V(P_i)$. Finally, for each clause $c_j = s_1 x_{i_1} \vee s_2 x_{i_2} \vee s_3 x_{i_3}$, we add the edges $u_j x_{i_1,j}$, $v_j x_{i_2,j}$, and $w_j x_{i_3,j}$ (regardless of s_1, s_2, s_3). This concludes the construction; see Figure 3 for an illustration.

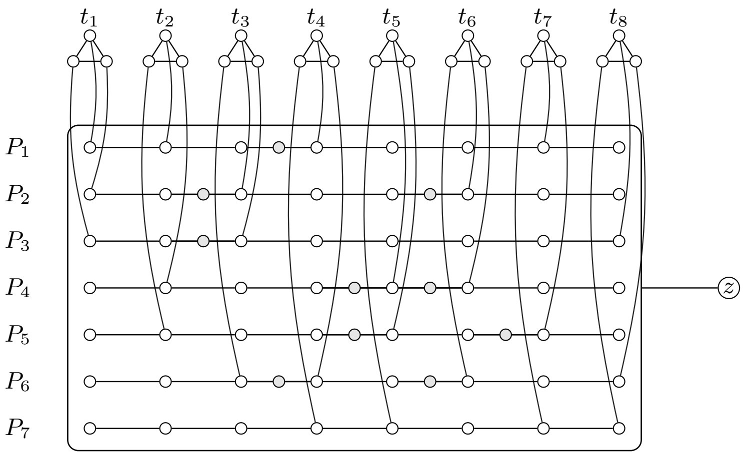

Figure 3 The graph $G(\varphi)$ for the 7-variable 8-clause formula φ with: $c_1 = x_1 \vee \neg x_2 \vee x_3$, $c_2 = \neg x_1 \vee x_4 \vee x_5$, $c_3 = x_2 \vee \neg x_3 \vee x_6$, $c_4 = x_1 \vee x_6 \vee \neg x_7$, $c_5 = x_4 \vee x_5 \vee x_7$, $c_6 = x_2 \vee x_4 \vee \neg x_6$, $c_7 = \neg x_1 \vee \neg x_5 \vee x_7$, $c_8 = x_3 \vee \neg x_6 \vee \neg x_7$.

Note that G has N vertices with $N \leqslant 3m + (2m-1)n + 1 = O(n^2)$. As we will next show that the reduction is correct and the twin-width of G is at most 4, this establishes the claimed ETH lower bound. Indeed, deciding if G is 3-colorable in $2^{o(\sqrt{N})}$ time would solve the Not-All-Equal 3-Sat instance φ in $2^{o(n)}$ time.

Parity property along the variable paths. Fix a variable x_i with occurrences at indices $j_1 < \cdots < j_q$ and signs $s_1, \ldots, s_q$, respectively. By construction, for every $h \in [q-1]$, the distance in P_i between x_{i,j_h} and $x_{i,j_{h+1}}$ is even if and only if $s_h = s_{h+1}$. Consequently, in any (proper) 2-coloring of P_i,

$$c(x_{i,j_h}) = c(x_{i,j_{h'}}) \iff s_h = s_{h'}, \quad \text{for all } h, h' \in [q],$$

i.e., all positive occurrences receive the same color and all negative occurrences receive the other color. Indeed, an even distance between x_{i,j_h} and $x_{i,j_{h+1}}$ (thus with equal color in a 2-coloring of P_i) does not change the sign of x_i, whereas an odd distance flips the sign of x_i.

If φ is satisfiable, then G is 3-colorable. Let $\mathcal{A}$ be a truth assignment such that every clause of φ has at least one literal set to true and at least one literal set to false. We define a 3-coloring c of G. First, set $c(z) = 3$.

For each variable x_i, pick (arbitrarily) one clause index j^* such that x_i appears in c_{j^*}, and let s^* be the sign of x_i in c_{j^*}. Set $c(x_{i,j^*}) = 1$ if $\mathcal{A}$ sets literal $s^* x_i$ to true, and $c(x_{i,j^*}) = 2$ otherwise. Since z has color 3, the path P_i must use only colors in $\{1,2\}$, and the coloring of P_i is uniquely determined by the color of x_{i,j^*}. By the parity property above, for every occurrence of x_i in a clause c_j with sign s, the vertex $x_{i,j}$ has color 1 if and only if $\mathcal{A}$ sets $s x_i$ to true (and color 2 otherwise).

Now fix a clause $c_j = \ell_1 \vee \ell_2 \vee \ell_3$ with triangle $t_j = \{u_j, v_j, w_j\}$. The three neighbors $x_{i_1,j}, x_{i_2,j}, x_{i_3,j}$ of t_j are colored in $\{1,2\}$ and, since $\mathcal{A}$ satisfies φ, these three colors are not all equal. Then, let $a, b \in [3]$ be such that $c(x_{i_a,j}) \neq c(x_{i_b,j})$. Give the neighbor of $x_{i_a,j}$ (resp. $x_{i_b,j}$) in t_j the color of $\{1,2\}$ opposite to $c(x_{i_a,j})$ (resp. $c(x_{i_b,j})$). Give the third vertex of t_j color 3. This yields a proper 3-coloring of t_j extending the already fixed colors on the variable paths. Doing this independently for all clauses completes a (proper) 3-coloring of G; see Figure 4.

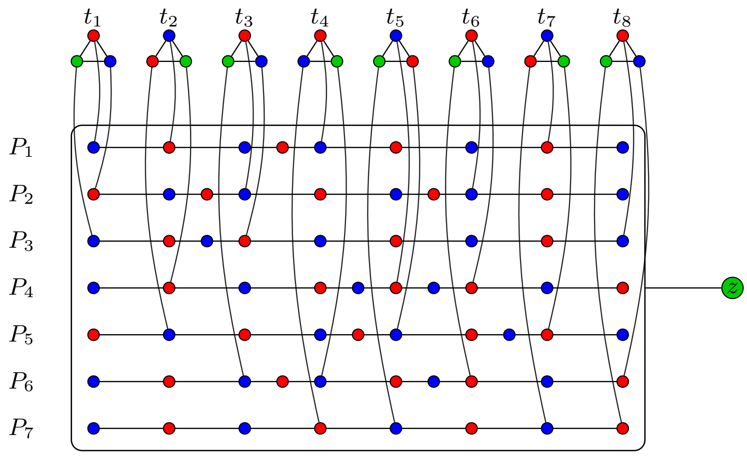

Figure 4 A 3-coloring of $G(\varphi)$ for the formula φ of Figure 3 corresponding to the assignment where all variables are set to true except x_4, which is set to false.

If G is 3-colorable, then φ is satisfiable. Let c be a (proper) 3-coloring of G. Up to color permutation, we can assume that $c(z) = 3$. Since z is adjacent to every vertex of every P_i, each path P_i uses only colors in $\{1, 2\}$.

We define an assignment $\mathcal{A}$ of the variables of φ as follows. For each variable x_i, if x_i appears positively in some clause c_j, $\mathcal{A}$ sets x_i to true if $c(x_{i,j}) = 1$, and to false otherwise; if x_i appears only negatively, $\mathcal{A}$ sets x_i to false if $c(x_{i,j}) = 1$ for some clause c_j containing literal $\neg x_i$, and to true otherwise. This is well-defined by the parity property: all positive occurrences have the same color and all negative occurrences have the other color (in $\{1, 2\}$).

Consider any clause $c_j = \ell_1 \vee \ell_2 \vee \ell_3$ represented by triangle t_j. If the three literals ℓ_1, ℓ_2, ℓ_3 had the same truth value under $\mathcal{A}$, then the three neighbors of t_j on the variable paths would all receive the same color in $\{1, 2\}$. In that case, each of u_j, v_j, w_j would be forbidden to use that color, leaving only two colors for the triangle t_j, which would not result in a 3-coloring of G. Hence, in every clause, not all three literals have the same truth value under $\mathcal{A}$. Therefore, $\mathcal{A}$ satisfies φ (as a Not-All-Equal 3-Sat instance).

G admits a 4-sequence. We describe a partition sequence for G of width at most 4. It is computable in polynomial time and does not depend on φ beyond knowing n and m. The sequence has four stages.

In the first stage, we merge each clause gadget into a single part. For each $j \in [m]$, merge $\{u_j\}$ with $\{v_j\}$, then merge the result with $\{w_j\}$. We denote the resulting part by $T_j := \{u_j, v_j, w_j\}$. Throughout this stage, every non-singleton part has red degree at most 3, and every singleton part has at most one red neighbor.

In the second stage, we merge every existing $x'_{i,j}$ with $x_{i,j}$, and set $X_{i,j} := \{x'_{i,j}, x_{i,j}\}$ if $x'_{i,j}$ exists, and $X_{i,j} := \{x_{i,j}\}$ otherwise. This does not create any part of red degree at least 4. The rest of the partition sequence has width at most 4, even if every variable path (which has now exactly m vertices) is fully red. We thus assume that this is the case. Henceforth, for each $j \in [m]$, we keep denoting by $X_{1,j}$ the part containing $x_{1,j}$.

In the third stage, for i going from 2 to n, for j going from 1 to m, we merge $X_{1,j}$ with $X_{i,j}$. This merge cannot increase the red degree of a part that is not the resulting part (recall that we assumed that the variable paths are fully red), and the latter has at most four red neighbors: T_j (if it is adjacent to), $X_{1,j-1}$ (if it exists), $X_{1,j+1}$ and $X_{i,j+1}$ (if they exist). Indeed note that at this point, $X_{i,j-1}$ has already been merged with $X_{1,j-1}$.

In the last stage, for j going from 2 to m, we merge $X_{1,1}$ and $X_{1,j}$, and keep denoting the resulting part by $X_{1,1}$, and we merge T_1 and T_j, and keep denoting the resulting part by T_1. This never creates a part with red degree larger than 3. Finally, there are three parts left: $T_1, X_{1,1}, \{z\}$, which we merge in any order. The concatenation of the four stages is a 4-sequence for G, which concludes the proof. ◀

▶ **Corollary 3.** *For every $k \geqslant 3$, k-COLORING is NP-hard and, unless the ETH fails, requires $2^{\Omega(\sqrt{n})}$ time on n-vertex graphs of twin-width at most 4.*

Proof. Adding $k-3$ universal vertices to the graph $G(\varphi)$ built in the proof of Theorem 2 does not change its twin-width and makes the resulting graph k-colorable if and only if $G(\varphi)$ is 3-colorable. ◀

References

1 Tara Abrishami, Bogdan Alecu, Maria Chudnovsky, Sepehr Hajebi, and Sophie Spirkl. Induced subgraphs and tree decompositions VII. basic obstructions in H-free graphs. *J. Comb. Theory B*, 164:443–472, 2024. URL: https://doi.org/10.1016/j.jctb.2023.10.008, doi:10.1016/J.JCTB.2023.10.008.

2 Jungho Ahn, Debsoumya Chakraborti, Kevin Hendrey, and Sang-il Oum. Twin-width of subdivisions of multigraphs. *SIAM J. Discret. Math.*, 39(2):783–833, 2025. URL: https://doi.org/10.1137/23m1621514, doi:10.1137/23M1621514.

3 Jungho Ahn, Hugo Jacob, Noleen Köhler, Christophe Paul, Amadeus Reinald, and Sebastian Wiederrecht. Twin-width one. In Olaf Beyersdorff, Michal Pilipczuk, Elaine Pimentel, and Kim Thang Nguyen, editors, *42nd International Symposium on Theoretical Aspects of Computer Science, STACS 2025, March 4-7, 2025, Jena, Germany*, volume 327 of *LIPIcs*, pages 6:1–6:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025. URL: https://doi.org/10.4230/LIPIcs.STACS.2025.6, doi:10.4230/LIPICS.STACS.2025.6.

4 Max Bannach, Florian Andreas Marwitz, and Till Tantau. Faster graph algorithms through DAG compression. In Olaf Beyersdorff, Mamadou Moustapha Kanté, Orna Kupferman, and Daniel Lokshtanov, editors, *41st International Symposium on Theoretical Aspects of Computer Science, STACS 2024, Clermont-Ferrand, France, March 12-14, 2024*, volume 289 of *LIPIcs*, pages 8:1–8:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. URL: https://doi.org/10.4230/LIPIcs.STACS.2024.8, doi:10.4230/LIPICS.STACS.2024.8.

5 Pierre Bergé, Édouard Bonnet, and Hugues Déprés. Deciding twin-width at most 4 is NP-complete. In *Proc. ICALP 2022*, volume 229 of *LIPIcs*, pages 18:1–18:20. Schloss Dagstuhl, 2022. arXiv:2112.08953, doi:10.4230/LIPIcs.ICALP.2022.18.

6 Pierre Bergé, Édouard Bonnet, Hugues Déprés, and Rémi Watrigant. Approximating highly inapproximable problems on graphs of bounded twin-width. In Petra Berenbrink, Patricia Bouyer, Anuj Dawar, and Mamadou Moustapha Kanté, editors, *40th International Symposium on Theoretical Aspects of Computer Science, STACS 2023, Hamburg, Germany, March 7-9, 2023*, volume 254 of *LIPIcs*, pages 10:1–10:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. URL: https://doi.org/10.4230/LIPIcs.STACS.2023.10, doi:10.4230/LIPICS.STACS.2023.10.

7 Benjamin Bergougnoux, Jakub Gajarský, Grzegorz Guspiel, Petr Hlinený, Filip Pokrývka, and Marek Sokolowski. Sparse graphs of twin-width 2 have bounded tree-width. In Satoru Iwata and Naonori Kakimura, editors, *34th International Symposium on Algorithms and Computation, ISAAC 2023, Kyoto, Japan, December 3-6, 2023*, volume 283 of *LIPIcs*, pages 11:1–11:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. URL: https://doi.org/10.4230/LIPIcs.ISAAC.2023.11, doi:10.4230/LIPICS.ISAAC.2023.11.

8 Hans L. Bodlaender. Achromatic number is np-complete for cographs and interval graphs. *Inf. Process. Lett.*, 31(3):135–138, 1989. doi:10.1016/0020-0190(89)90221-4.

9 Édouard Bonnet. *Twin-width and contraction sequences.* Habilitation à diriger des recherches, École Normale Supérieure de Lyon, 2024.

10 Édouard Bonnet. Sparse induced subgraphs of large treewidth. *J. Comb. Theory B*, 173:184–203, 2025. URL: `https://doi.org/10.1016/j.jctb.2025.03.002`, `doi:10.1016/J.JCTB.2025.03.002`.

11 Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width II: small classes. *Comb. Theory*, 2(2), 2022. URL: `https://doi.org/10.5070/c62257876`, `doi:10.5070/C62257876`.

12 Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width III: max independent set, min dominating set, and coloring. *SIAM J. Comput.*, 53(5):1602–1640, 2024. URL: `https://doi.org/10.1137/21m142188x`, `doi:10.1137/21M142188X`.

13 Édouard Bonnet, Eun Jung Kim, Amadeus Reinald, and Stéphan Thomassé. Twin-width VI: the lens of contraction sequences. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1036–1056. SIAM, 2022.

14 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`.

15 Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. *J. ACM*, 69(1):3:1–3:46, 2022. `doi:10.1145/3486655`.

16 Bruno Courcelle, János A. Makowsky, and Udi Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory of Computing Systems*, 33(2):125–150, 2000. `doi:10.1007/s002249910009`.

17 Gruia Cualinescu, Cristina G. Fernandes, and Bruce A. Reed. Multicuts in unweighted graphs and digraphs with bounded degree and bounded tree-width. *J. Algorithms*, 48(2):333–359, 2003. `doi:10.1016/S0196-6774(03)00073-7`.

18 Clément Dallard, Martin Milanic, and Kenny Storgel. Treewidth versus clique number. III. tree-independence number of graphs with a forbidden structure. *J. Comb. Theory B*, 167:338–391, 2024. URL: `https://doi.org/10.1016/j.jctb.2024.03.005`, `doi:10.1016/J.JCTB.2024.03.005`.

19 Peter Damaschke, Jitender S Deogun, Dieter Kratsch, and George Steiner. Finding hamiltonian paths in cocomparability graphs using the bump number algorithm. *Order*, 8(4):383–391, 1991.

20 Jitender S. Deogun and George Steiner. Polynomial algorithms for hamiltonian cycle in cocomparability graphs. *SIAM J. Comput.*, 23(3):520–552, 1994. `doi:10.1137/S0097539791200375`.

21 Shimon Even, Amir Pnueli, and Abraham Lempel. Permutation graphs and transitive graphs. *J. ACM*, 19(3):400–410, 1972. `doi:10.1145/321707.321710`.

22 Jirí Fiala, Ignaz Rutter, Peter Stumpf, and Peter Zeman. Extending partial representations of circular-arc graphs. In Michael A. Bekos and Michael Kaufmann, editors, *Graph-Theoretic Concepts in Computer Science - 48th International Workshop, WG 2022, Tübingen, Germany, June 22-24, 2022, Revised Selected Papers*, volume 13453 of *Lecture Notes in Computer Science*, pages 230–243. Springer, 2022. `doi:10.1007/978-3-031-15914-5\_17`.

23 Stephen Finbow, Andrew D. King, Gary MacGillivray, and Romeo Rizzi. The firefighter problem for graphs of maximum degree three. *Discret. Math.*, 307(16):2094–2105, 2007. URL: `https://doi.org/10.1016/j.disc.2005.12.053`, `doi:10.1016/J.DISC.2005.12.053`.

24 Fedor V. Fomin, Pinar Heggernes, and Erik Jan van Leeuwen. The firefighter problem on graph classes. *Theor. Comput. Sci.*, 613:38–50, 2016. URL: `https://doi.org/10.1016/j.tcs.2015.11.024`, `doi:10.1016/J.TCS.2015.11.024`.

25 Michael R Garey, Ronald L Graham, David S Johnson, and Donald Ervin Knuth. Complexity results for bandwidth minimization. *SIAM Journal on Applied Mathematics*, 34(3):477–495, 1978.

26 Peter Gartland. *Quasi-Polynomial Time Techniques for Independent Set and Beyond in Hereditary Graph Classes*. PhD thesis, UC Santa Barbara, 2023.

27 Petr Hlinený and Jan Jedelský. Twin-width of planar graphs is at most 8, and some related bounds. *SIAM J. Discret. Math.*, 39(4):2003–2048, 2025. URL: `https://doi.org/10.1137/23m1623823`, `doi:10.1137/23M1623823`.

28 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001. URL: `https://doi.org/10.1006/jcss.2000.1727`, `doi:10.1006/JCSS.2000.1727`.

29 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001. URL: `https://doi.org/10.1006/jcss.2001.1774`, `doi:10.1006/JCSS.2001.1774`.

30 Alon Itai, Christos H. Papadimitriou, and Jayme L. Szwarcfiter. Hamilton paths in grid graphs. *SIAM Journal on Computing*, 11(4):676–686, 1982. `doi:10.1137/0211056`.

31 Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, *Proceedings of a symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, USA*, The IBM Research Symposia Series, pages 85–103. Plenum Press, New York, 1972. `doi:10.1007/978-1-4684-2001-2\_9`.

32 Tuukka Korhonen. Grid induced minor theorem for graphs of small degree. *J. Comb. Theory B*, 160:206–214, 2023. URL: `https://doi.org/10.1016/j.jctb.2023.01.002`, `doi:10.1016/J.JCTB.2023.01.002`.

33 Daniel Král' and Ander Lamaison. Planar graph with twin-width seven. *Eur. J. Comb.*, 123:103749, 2025. URL: `https://doi.org/10.1016/j.ejc.2023.103749`, `doi:10.1016/J.EJC.2023.103749`.

34 Charis Papadopoulos. Restricted vertex multicut on permutation graphs. *Discret. Appl. Math.*, 160(12):1791–1797, 2012. URL: `https://doi.org/10.1016/j.dam.2012.03.021`, `doi:10.1016/J.DAM.2012.03.021`.

35 Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing*, STOC '78, page 216–226, New York, NY, USA, 1978. Association for Computing Machinery. `doi:10.1145/800133.804350`.