

Independent Set Hardness in Graphs of Bounded Twin-Width and Low-Radius Merge-Width

Édouard Bonnet   

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

Maël Dumas  

Institute of Informatics, University of Warsaw, Poland

Julien Duron 

Institute of Informatics, University of Warsaw, Poland

Abstract

For every $\varepsilon > 0$, MAX INDEPENDENT SET admits a polynomial-time n^ε -approximation algorithm on n -vertex graphs of effectively bounded twin-width [Bergé et al., STACS '23]. The approximation factor actually obtained is more precisely $n^{\mathcal{O}(1/\log \log n)}$. Prior to the current paper, no approximation hardness was known for this problem, and the existence of a polynomial-time approximation scheme (PTAS) was repeatedly raised as an open question. We answer this question in a strong sense: We show that there is a constant $\gamma > 0$ such that a polynomial-time $n^{\gamma/(\log \log n)^2}$ -approximation algorithm for MAX INDEPENDENT SET on graphs of twin-width at most 4 would refute the Exponential-Time Hypothesis (ETH). This lower bound further holds if a 4-sequence is provided as part of the input. We show the same hardness of approximation for MIN COLORING, which also has a nearly matching $n^{\mathcal{O}(1/\log \log n)}$ -approximation algorithm on graphs of effectively bounded twin-width.

We also clarify the parameterized complexity of k -INDEPENDENT SET on graphs of bounded radius- r merge-width when the range of r is limited. There is a fixed-parameter tractable algorithm for k -INDEPENDENT SET on graphs given with radius- $2^{\mathcal{O}(k^2)}$ merge sequences of bounded width [Dreier and Toruńczyk, STOC '25]. We complement this result by showing that k -INDEPENDENT SET is W[1]-hard on graphs given with radius- $o(k)$ merge sequences of bounded width. We further show that this result also holds for k -DOMINATING SET.

Funding JD received funding from ERC grant BUKA (No. 101126229). MD and JD were supported by the project BOBR that has received funding from ERC under the European Union's Horizon 2020 research and innovation programme, grant agreement No. 948057.

Acknowledgements We want to thank Karolina Drabik, Jakub Nowakowski, Nikolas Mählmann, and Szymon Toruńczyk for fruitful discussions at an early stage of this project.

1 Introduction

The problem of finding in a graph a largest subset of pairwise non-adjacent vertices, MAX INDEPENDENT SET (MIS for short), is NP-complete [20], and very hard to approximate: For any $\varepsilon > 0$, an $n^{1-\varepsilon}$ -approximation algorithm would imply that $P = NP$ [22, 26]. The best known polynomial-time approximation factor is $\frac{n(\log \log n)^2}{\log^3 n}$ [18]. In contrast, this central problem enjoys much better approximation algorithms on structured graph classes. A historical example is that of planar graphs where, due to Baker's shifting technique, a polynomial-time approximation scheme (PTAS) exists [2]. MIS also has a PTAS on minor-free classes [21] and on disk graphs [12]. If a geometric representation is provided, there is, for every $\varepsilon > 0$, an n^ε -approximation algorithm for the independence number of the intersection graph of n curves every two of which have at most a constant number of intersections [19], and a quasipolynomial-time approximation scheme (QPTAS) for general string graphs [1].

In this paper, our focus is on classes of bounded twin-width [11] and of bounded merge-width [17]. These graph parameters were introduced in 2020 and 2025, respectively. Classes of bounded twin-width include classes of bounded clique-width, classes excluding a fixed minor, unit interval graphs, and d -dimensional grids [11]. Classes of bounded merge-width are even more general and further include classes of bounded expansion [17]. The main algorithmic application of classes of bounded twin-width and of bounded merge-width, when bounded-width witnesses are given, is that first-order model checking is fixed-parameter tractable [11, 17].

As we detail in the next paragraph, some improved approximation algorithms exist on classes of bounded twin-width. In general, however, there is no precise understanding of the provably best algorithms achievable when twin-width or merge-width is bounded. This is reflected in large approximability gaps and the lack of tight lower bounds under the Exponential-Time Hypothesis (ETH). Our motivation is to fill those gaps.

In classes of effectively¹ bounded twin-width, n^ε -approximation algorithms for MIS, MIN COLORING, and other related problems were shown for every $\varepsilon > 0$ [4]. The proof in [4] reveals that the approximation factor for MIS is more specifically $n^{c/\log \log n}$, for some constant c . While it was known that MAX INDEPENDENT SET is NP-hard on graphs of bounded twin-width [11], even bounded by 4 [3] (also see [8] for the status of other graph problems on graphs of low twin-width), no conditional hardness of approximation was established on classes of bounded twin-width. Even a polynomial-time approximation scheme (PTAS) was not excluded. Indeed, the existence of a PTAS was raised several times [9, 10, 4, 14, 7]. It was nonetheless observed that any $\mathcal{O}(1)$ -approximation algorithm would imply the existence of a PTAS [10].

We rule this out under the ETH and show a much tighter hardness-of-approximation result that nearly matches the $n^{c/\log \log n}$ -approximation algorithm.

► **Theorem 1.** *Unless the ETH fails, there is a constant $\gamma > 0$ such that MAX INDEPENDENT SET does not admit a polynomial-time $n^{\gamma/(\log \log n)^2}$ -approximation algorithm on n -vertex graphs of twin-width at most 4, even when a 4-sequence is given as part of the input.*

The reduction uses, in order, three ingredients: (1) a classical reduction from 3-SAT to MIS, (2) a long even subdivision of the resulting instances, and (3) a lexicographic power of the resulting graphs.

“Long” subdivisions have twin-width at most 4 [3], which explains the role of (2). Even subdivisions preserve the NP-hardness of MIS [25]. However, “long” subdivisions do not preserve any hardness of approximation. There is a straightforward PTAS on these graphs: select every other vertex along each subdivided edge. So, there is a psychological barrier in using (2). Nevertheless, with the right exponent, (3) resurrects this seemingly defunct approximation gap, while keeping the reduction within the desired running-time budget.

Similar ingredients can be adapted to give the same result for MIN COLORING.

► **Theorem 2** (⊕). *Unless the ETH fails, there is a constant $\gamma > 0$ such that MIN COLORING does not admit a polynomial-time $n^{\gamma/(\log \log n)^2}$ -approximation algorithm on n -vertex graphs of twin-width at most 4, even when a 4-sequence is given as part of the input.*

Theorem 2 implies Theorem 1. Indeed, a polynomial-time r -approximation algorithm for MIS implies a polynomial-time $\mathcal{O}(r \log n)$ -approximation algorithm for MIN COLORING. Repeatedly removing an r -approximate maximum independent set and assigning it a new

¹ See this and all relevant definitions in Section 2.

color yields an $r(1 + \ln n)$ -approximate minimum coloring, by the classic analysis of SET COVER’s greedy algorithm. Moreover, $n^{\Theta(1)/(\log \log n)^2} \cdot \mathcal{O}(\log n) = n^{\Theta(1)/(\log \log n)^2}$.

We keep both proofs because the proof of Theorem 1 implements the above outline without the extra technicalities needed for MIN COLORING.

We then delineate the tractability boundary of k -INDEPENDENT SET (k -IS for short) on graphs of bounded merge-width. Strictly speaking, merge-width is a *family* of parameters. For every positive integer r , the radius- r merge-width, denoted by mw_r , must be bounded by a function of r , for the merge-width itself to be bounded. Low-radius merge-width, think mw_1 , mw_2 , or mw_c for some constant c , is of independent interest. There are some properties that only require these parameters to be bounded. Notably, classes of bounded merge-width were proven χ -bounded by showing the stronger statement that classes of bounded mw_2 are χ -bounded [6] (and the χ -boundedness of classes of bounded mw_1 remains open).

Given radius- 2^{ck^2} merge sequences of bounded width² for some appropriate constant $c > 0$, k -IS admits a fixed-parameter tractable (FPT) algorithm [17]. In fact, this is more generally the case for first-order model checking. Dreier and Toruńczyk ask whether a simpler and faster algorithm could exist for k -INDEPENDENT SET (and k -DOMINATING SET).

A related question is whether these problems admit FPT algorithms given radius- r merge sequences of bounded width already in the regime $r = \mathcal{O}(1)$. We rule this out for any radius $r = o(k)$, with a reduction from the W[1]-hard GRID TILING problem [24] to k -IS. We ensure that the produced graphs have bounded radius- r merge-width by spacing the GRID TILING cells out with “paths of co-matchings” of length $\Theta(r)$.

► **Theorem 3.** $\forall r = o(k)$, k -INDEPENDENT SET is W[1]-hard in graphs of bounded mw_r .

We then adapt the previous reduction to get the same statement for k -DOMINATING SET.

► **Theorem 4.** $\forall r = o(k)$, k -DOMINATING SET is W[1]-hard in graphs of bounded mw_r .

Future work. Our work naturally leads to the following questions. Combining Theorems 1 and 2 and [4], we now know that, unless the ETH fails, the best approximation factor for MAX INDEPENDENT SET and MIN COLORING lies somewhere between $n^{\Theta(1)/(\log \log n)^2}$ and $n^{\Theta(1)/\log \log n}$. Can one obtain matching bounds, up to constants in the exponent?

Ruling out a PTAS for MIS on graphs of bounded twin-width under the weaker assumption that $P \neq NP$ remains open. Let us also mention the question of obtaining nontrivial approximation algorithms for these problems when the contraction sequence is *not* provided.

Theorems 3 and 4 and [17] raise the question of the smallest radius $r(k)$ such that k -INDEPENDENT SET (and k -DOMINATING SET) is FPT on graphs given with a radius- $r(k)$ merge sequence of bounded width. We now know that this radius is between $\Omega(k)$ and $2^{\mathcal{O}(k^2)}$.

AI disclosure. After we proved Theorem 1, we suspected that the same result should hold for MIN COLORING, and could probably be proven in a similar fashion. We delegated this adaptation to ChatGPT Pro 5.5 Extended Thinking. We provided it with the proof of Theorem 1 with a concise prompt asking for a possible transfer to MIN COLORING. Its first response, produced after 17 minutes, pinpointed the places where the same proof would break. We then clarified that the task was the adaptation. After another 50 minutes, it provided a complete proof of Theorem 2. Although it follows the proof of Theorem 1 at a high level, this proof contains several new elements, notably the initial hardness result [5,

² witnessing that the radius- 2^{ck^2} merge-width is bounded; see Section 2.3 for the relevant definitions.

Thm. 14], a modified subdivision, and a reduction to leverage the chromatic number lower bound that comes from an upper bound on the independence number. The write-up of this result, and of the paper as a whole, is ours.

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 $[1, i]$. In this paper, the log function is in base 2.

The Exponential-Time Hypothesis (ETH) asserts that there is a real constant $\lambda > 1$ such that n -variable 3-SAT cannot be solved in time $\mathcal{O}(\lambda^n)$.

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. We denote by $\|G\|$ the total size of G , that is, $|V(G)| + |E(G)|$. 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 . We denote by $N_G(v)$ the set of neighbors of v in G . We denote by $\text{dist}_G(u, v)$ the length of a shortest path between u and v in G , and set $\text{Ball}_G^r(v) := \{w \in V(G) \mid \text{dist}_G(v, w) \leq r\}$.

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 $(\geq s)$ -*subdivision* is a subdivision where every edge is replaced by a path with at least s internal vertices.

2.2 Trigraphs, contraction 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.

For distinct vertices $u, v \in V(G)$, the *contraction* of u and v replaces them with a new vertex w . For every $x \in V(G) \setminus \{u, v\}$, the pair wx is a black edge if both ux and vx are black edges, it is a non-edge if both ux and vx are non-edges, and it is a red edge otherwise. All pairs not incident with u or v remain unchanged.

A *contraction sequence* of a graph G , viewed as a trigraph with no red edges, is a sequence $G = G_n, G_{n-1}, \dots, G_1$ such that $|V(G_i)| = i$ and G_{i-1} is obtained from G_i by contracting two vertices. The sequence is a *d -sequence* if, for every $i \in [n]$, the trigraph G_i has maximum red degree at most d . The *twin-width* of G , denoted by $\text{tw}(G)$, is the minimum integer d for which G admits a d -sequence.

To show Theorems 1 and 2, we will rely on the fact that long-enough subdivisions of general graphs have bounded twin-width.

► **Lemma 5** (Theorem 7 in [3]). *For every n -vertex graph G , the $2\lceil \log n \rceil$ -subdivision H of G has twin-width at most 4. Furthermore, a 4-sequence of H can be computed in $\mathcal{O}(\|H\|)$ time.*

The algorithmic statement of Lemma 5 is implicit in [3] and follows from the proof.

We say that a graph class $\mathcal{C}$ has *effectively bounded twin-width* if there is a polynomial-time algorithm $\mathcal{A}$ and a fixed integer d such that on every graph $G \in \mathcal{C}$, $\mathcal{A}$ outputs a d -sequence of G . We observe that most classes currently known to have bounded twin-width have effectively bounded twin-width.

2.3 Merge-width

For two sets X, Y , we write $XY := \{xy \mid x \in X, y \in Y, x \neq y\}$, where xy denotes the unordered pair $\{x, y\}$. Let $\mathcal{P}$ be a partition of the vertices V of a graph G . For a set $X \subseteq V$, we denote by $X/\mathcal{P}$ the set of parts of $\mathcal{P}$ intersecting X . Let $R \subseteq \binom{V}{2}$ be a set of pairs of vertices of G . The distance between x and y in (V, R) is denoted by $\text{dist}_R(x, y)$, and we define $\text{Ball}_R^r(x) := \{y \in V \mid \text{dist}_R(x, y) \leq r\}$. The *radius- r width* of $(\mathcal{P}, R)$ is defined as

$$\max_{v \in V} |\text{Ball}_R^r(v)/\mathcal{P}|.$$

We say that $\mathcal{P}$ is *homogeneous modulo R* (in G) if for any parts $X, Y \in \mathcal{P}$ (possibly $X = Y$) and pairs $xy, x'y' \in XY \setminus R$, we have that $xy \in E(G)$ if and only if $x'y' \in E(G)$. A partition $\mathcal{P}$ is *coarser* than a partition $\mathcal{P}'$, denoted by $\mathcal{P}' \preceq \mathcal{P}$, if every part of $\mathcal{P}'$ is contained in a part of $\mathcal{P}$.

► **Definition 6** (Definition 3.1 in [17]). *A merge sequence for a graph G is a sequence $(\mathcal{P}_1, R_1), \dots, (\mathcal{P}_m, R_m)$ such that:*

1. $\mathcal{P}_1 \preceq \mathcal{P}_2 \preceq \dots \preceq \mathcal{P}_m$ is a sequence of ever coarser partitions of $V(G)$, with $\mathcal{P}_1$ the partition into singletons and $\mathcal{P}_m$ the partition with one part,
2. $R_1 \subseteq \dots \subseteq R_m \subseteq \binom{V(G)}{2}$ is a monotone sequence of sets of pairs of vertices, and
3. $\mathcal{P}_t$ is homogeneous modulo R_t , for $t = 1, \dots, m$.

The radius- r width of this merge sequence is the maximum radius- r width of $(\mathcal{P}_t, R_{t+1})$, for $1 \leq t < m$. Finally, the radius- r merge-width of G , denoted by $\text{mw}_r(G)$, is the minimum radius- r width of a merge sequence for G .

The elements of R_t are called *resolved pairs* (at step t). Note that the partition $\mathcal{P}_{t+1}$ is obtained by merging some parts of $\mathcal{P}_t$. Note also that the index mismatch in $(\mathcal{P}_t, R_{t+1})$ is intentional, and prevents one from merging many parts and adding many resolved pairs all at once when going from step t to $t + 1$.

A graph class $\mathcal{C}$ has *bounded merge-width* if there is a function f such that for every $G \in \mathcal{C}$ and every positive integer r , we have $\text{mw}_r(G) \leq f(r)$.

Organization. We prove Theorem 1 in Section 3, Theorem 2 in Section 4, Theorem 3 in Section 5, and Theorem 4 in Section 6.

3 Approximation Hardness of MIS in Graphs of Bounded Twin-Width

In this paper, all 3-SAT instances have only 3-clauses. Equal literals within the same clause are allowed, but we disallow opposite literals of the same variable within the same clause. A 3-SAT- B formula has no variable occurring (positively or negatively) more than B times. The following theorem is a consequence of Dinur's almost-linear PCP [15] and some folklore reductions; see, for instance, [23, Thm. 7].

► **Theorem 7** (consequence of Theorem 8.3 in [15]). *Unless the ETH fails, there are positive constants $\beta < 1, B, c$ such that distinguishing n -variable m -clause 3-SAT- B satisfiable formulas from instances where at most βm clauses can be satisfied requires $2^{\Omega(n/\log^c n)}$ time.*

We denote by $e(\varphi)$ the total number of unordered pairs of opposing literals appearing in φ . For instance, if $\varphi = x_1 \vee \neg x_2 \vee x_4, \neg x_1 \vee \neg x_1 \vee x_2, \neg x_2 \vee x_3 \vee \neg x_4$, we have $e(\varphi) = 2 + 2 + 0 + 1 = 5$. We denote by $\text{val}(\varphi)$ the largest number of clauses satisfied by any truth assignment. In particular, if φ is satisfiable, then $\text{val}(\varphi)$ is equal to its number of clauses.

► **Theorem 8.** *There is an $\mathcal{O}((m + e(\varphi)) \log m)$ -time algorithm that, given an m -clause 3-SAT formula φ , outputs a graph G_φ of twin-width at most 4 such that*

- $|V(G_\varphi)| = 3m + (e(\varphi) + 3m) \cdot 2 \lceil \log(3m) \rceil$, and
- $\alpha(G_\varphi) = \text{val}(\varphi) + (e(\varphi) + 3m) \cdot \lceil \log(3m) \rceil$.

Furthermore, a 4-sequence of G_φ can be computed in $\mathcal{O}(\|G_\varphi\|)$ time.

Proof. Let J_φ be the graph with one vertex per literal occurrence in a clause of φ , and an edge between every two literals that are either in the same clause or that are of the same variable with opposite signs. In particular, J_φ has exactly $3m$ vertices and exactly $3m + e(\varphi)$ edges (recall our definition of 3-SAT).

Then, G_φ is the s -subdivision of J_φ with $s := 2 \lceil \log(3m) \rceil$. Note that G_φ can be built from φ in $\mathcal{O}((m + e(\varphi)) \log m)$ time. By Lemma 5, G_φ has twin-width at most 4, and a 4-sequence of G_φ can be computed in $\mathcal{O}(\|G_\varphi\|)$ time. Since G_φ is a s -subdivision of a $3m$ -vertex $(3m + e(\varphi))$ -edge graph, $|V(G_\varphi)| = 3m + (e(\varphi) + 3m) \cdot s$.

The second item dates back to Poljak [25]. We sketch this classic fact.

The inequality $\alpha(G_\varphi) \geq \text{val}(\varphi) + (e(\varphi) + 3m) \cdot \frac{s}{2}$ holds, as there is an independent set of G_φ with one vertex corresponding to a satisfied literal for each clause satisfied by a fixed truth assignment that satisfies $\text{val}(\varphi)$ clauses of φ , and $\frac{s}{2}$ vertices per subdivided edge of J_φ .

The inequality $\alpha(G_\varphi) \leq \text{val}(\varphi) + (e(\varphi) + 3m) \cdot \frac{s}{2}$ holds since a larger independent set of G_φ would contain more than $\text{val}(\varphi)$ original vertices (those of J_φ). This would imply that J_φ has an independent set of size greater than $\text{val}(\varphi)$. (We elaborate on this in the proof of Theorem 13.) The (consistent, partial) truth assignment setting the corresponding literals to true would contradict the definition of $\text{val}(\varphi)$. ◀

Given a graph G and a positive integer k , we denote by $G^{[k]}$ the k -th *lexicographic power* of G with $V(G^{[k]}) = V(G)^k$ and $(u_1, \dots, u_k)(v_1, \dots, v_k) \in E(G^{[k]})$ whenever there is some $i \in [k]$ such that $u_i v_i \in E(G)$ and for every $h \in [i - 1]$, $u_h = v_h$. The following is a folklore observation. We give a recent reference for completeness.

► **Lemma 9** (see Lemma 7.12 in [10]). *For every graph G and integer $k \geq 1$, $\alpha(G^{[k]}) = \alpha(G)^k$.*

Furthermore, there is an algorithm that, given an independent set S of $G^{[k]}$, outputs an independent set of G of size at least $|S|^{1/k}$ in $\mathcal{O}(\|G^{[k]}\|)$ time.

The twin-width of every lexicographic power of a graph G is equal to that of G .

► **Lemma 10** (Lemma 7.11 in [10]). *For every graph G and integer $k \geq 1$, $\text{tw}(G^{[k]}) = \text{tw}(G)$.*

Furthermore, given a d -sequence of G , one can compute a d -sequence of $G^{[k]}$ in time $\mathcal{O}(\|V(G^{[k]})\|)$.

We now have all the ingredients to show our main result, which we restate for convenience.

► **Theorem 1.** *Unless the ETH fails, there is a constant $\gamma > 0$ such that MAX INDEPENDENT SET does not admit a polynomial-time $n^{\gamma/(\log \log n)^2}$ -approximation algorithm on n -vertex graphs of twin-width at most 4, even when 4-sequences are given as part of the input.*

Proof. Fix the constants $\beta < 1, B, c$ so as to satisfy Theorem 7. We make the following reduction. Every N -variable m -clause 3-SAT- B formula φ is mapped to the graph $H := G_\varphi^{[k]}$ with $k := \lceil m / \log^{c+2} m \rceil$. We will assume that m is large enough. In particular, the following inequalities hold:

1. $\log(4(3B^2 + 6) \log(3m)) < \log(3m)$;
2. $\log m > 2 \log \log^{c+2} m + \log 3$;

$$3. \frac{(e(\varphi)+3m) \cdot \lceil \log(3m) \rceil}{(1-\beta)m} > 2.$$

Building H takes $\mathcal{O}(m \log m) + (m \log m)^{\mathcal{O}(k)}$ time. Indeed, note that $e(\varphi) \leq B^2 N \leq 3B^2 m$. So the first term in the running time corresponds to building G_φ , and the second term, to building its k -th lexicographic power H . By Theorem 8 and Lemma 10, H indeed has twin-width at most 4, and a 4-sequence of H can be computed in time polynomial in $\|H\|$.

Suppose there is an $n^{\mathcal{O}(1)}$ -time r -approximation algorithm $\mathcal{A}$ for MAX INDEPENDENT SET on n -vertex graphs of twin-width at most 4. Our proof will work even if $\mathcal{A}$ has access to a 4-sequence of H .

Let $\delta \leq \Delta := 6(B^2 + 1)$ be such that

$$(e(\varphi) + 3m) \cdot \lceil \log(3m) \rceil = \delta m \log(3m).$$

We set the constant

$$\gamma := \frac{0.8(1-\beta)}{8(\beta+\Delta)}.$$

To conclude the proof, we assume that $r = n^{\gamma/(\log \log n)^2}$ with $n := |V(G_\varphi)|^k$, and proceed to refute the ETH.

When run on H , $\mathcal{A}$ returns an independent set I_H of H of size at least $\alpha(H)/r$. By Lemma 9, we find in time $\mathcal{O}((m \log m)^{2k})$ an independent set I of G_φ of size at least

$$\left(\frac{\alpha(H)}{r} \right)^{1/k} = \frac{\alpha(G_\varphi)}{r^{1/k}}.$$

We report that φ is satisfiable if $|I| > \beta m + \delta m \log(3m)$, and that at most βm clauses of φ are satisfiable, otherwise.

Running time. Let us first verify that the overall process takes $2^{o(N/\log^c N)}$. Building H , optionally computing a 4-sequence of H , running $\mathcal{A}$ on H , finding I , and concluding whether φ is satisfiable takes time

$$(m \log m)^{\mathcal{O}(k)} + n^{\mathcal{O}(1)} + \mathcal{O}((m \log m)^{2k}) = m^{\mathcal{O}(m/\log^{c+2} m)} = 2^{\mathcal{O}(m/\log^{c+1} m)}$$

and, since $m \leq BN$ and $m \geq N/3$,

$$2^{\mathcal{O}(m/\log^{c+1} m)} = 2^{\mathcal{O}(BN/\log^{c+1}(N/3))} = 2^{o(N/\log^c N)}.$$

Correctness. We now check that we are correct on the satisfiability of φ , and thus conclude by Theorem 7. The crux of the argument is the following claim.

▷ **Claim 11.** $(m + \delta m \log(3m)) \cdot \frac{1}{r^{1/k}} > \beta m + \delta m \log(3m)$.

Proof. We first recall that

$$n = |V(G_\varphi)|^k = (3m + (e(\varphi) + 3m) \cdot 2\lceil \log(3m) \rceil)^k \leq (4(3B^2 + 6)m \log(3m))^k.$$

So,

$$\log n \leq k \log m + k \log (4(3B^2 + 6) \log(3m)) < 2k \log(3m), \quad (1)$$

where the last inequality holds by Item 1. And

$$\log \log n \geq \log k \geq \log m - \log \log^{c+2} m > \frac{1}{2} \log(3m),$$

where the last inequality holds by Item 2. Hence,

$$(\log \log n)^2 > \frac{1}{4} \log^2(3m). \quad (2)$$

On the one hand,

$$\log(r^{1/k}) = \log \left(\left(2^{\gamma \log n / (\log \log n)^2} \right)^{1/k} \right) = \frac{\gamma \log n}{k(\log \log n)^2} < \frac{\gamma \cdot 2 \log(3m)}{\frac{1}{4} \log^2(3m)} = \frac{0.8(1-\beta)}{(\beta + \Delta) \log(3m)},$$

where the inequality comes from the upper bound on the numerator by (1) and the lower bound on the denominator by (2).

On the other hand,

$$\begin{aligned} \log \left(\frac{m + \delta m \log(3m)}{\beta m + \delta m \log(3m)} \right) &= \log \left(1 + \frac{(1-\beta)m}{\beta m + \delta m \log(3m)} \right) \\ &= \log \left(1 + \frac{1-\beta}{\beta + \delta \log(3m)} \right) > \frac{0.8(1-\beta)}{\beta + \delta \log(3m)} > \frac{0.8(1-\beta)}{(\beta + \delta) \log(3m)} \geq \frac{0.8(1-\beta)}{(\beta + \Delta) \log(3m)}, \end{aligned}$$

where the first inequality holds because for every $x \in (0, 1/2)$, $\log(1+x) > 0.8x$, and by Item 3 $(1-\beta)/(\beta + \delta \log(3m)) \in (0, 1/2)$, and the second inequality holds because $\log(3m) > 1$.

Therefore,

$$\frac{m + \delta m \log(3m)}{\beta m + \delta m \log(3m)} > r^{1/k},$$

and we conclude. $\triangleleft$

By Claim 11, if $|I| \leq \beta m + \delta m \log(3m)$, then $\alpha(G_\varphi) < m + \delta m \log(3m)$. By Theorem 8, this indeed implies that φ is unsatisfiable.

If instead, $|I| > \beta m + \delta m \log(3m)$, there is a truth assignment that satisfies strictly more than βm clauses of φ . So, in this promise-problem setting, we correctly concluded that φ is satisfiable. $\blacktriangleleft$

4 Coloring Approximation Hardness in Graphs of Bounded Twin-Width

We show the same hardness of approximation for MIN COLORING as we showed in the previous section. The general strategy is the same but several elements need to be changed. The following consequence of [5, Thm. 14] and Theorem 7 is a convenient alternative to the mere Theorem 7 of Section 3.

► **Theorem 12.** *There are positive constants $\beta < 1, \Delta, c$ such that, for n -vertex graphs H of maximum degree at most Δ promised to satisfy one of the following two conditions*

- *H is 3-colorable, or*
- *no induced subgraph of H on at least βn vertices is 3-colorable*

determining which one holds requires $2^{\Omega(n/\log^c n)}$ time, unless the ETH fails.

Proof. We run the linear reduction from 3-SAT- B to 3-COLORING of [5, Thm. 14] on the hard instances of Theorem 7. To avoid a notational clash, we rename the constant β as $\beta' < 1$ in the latter theorem, and denote by N the number of variables of the 3-SAT- B instances, and by m the number of clauses in each instance.

The reduction of [5] (see the description preceding Theorem 14) produces graphs with at most κN vertices and maximum degree at most Δ for some constants $\kappa := \kappa(B), \Delta$. The

image of every satisfiable 3-SAT- B instance is 3-colorable, whereas in every (improper) 3-coloring of the image H of any 3-SAT- B instance φ such that $\text{val}(\varphi) \leq \beta' m$ at most a β'' fraction of the edges are properly colored (i.e., bichromatic) where $\beta'' := \frac{\beta' + 7}{8}$. We observe that $\beta'' < 1$, since $\beta' < 1$, and choose $\beta := 1 - \frac{1 - \beta''}{2\Delta} < 1$. Let $n := |V(H)|$.

Assume that there is a 3-colorable induced subgraph of H with at least βn vertices. Any extension to a complete 3-coloring of H improperly colors at most $(1 - \beta)\Delta n < (1 - \beta'')n$ edges, because H has maximum degree at most Δ . Here we observe that the reduction in [5] satisfies that $|V(H)| \leq |E(H)|$. Thus we reach the contradiction that more than a β'' fraction of the edges of H are properly colored.

This way, the two items of the current theorem correspond to the two outcomes of Theorem 7. We conclude by taking the same value for c as in Theorem 7 since the blow-up κ is constant. $\blacktriangleleft$

We now adapt Theorem 8. For that, we modify the subdivisions slightly without changing their twin-width upper bound of 4. One should also notice the hybrid nature of the last two items: in “positive” instances, we preserve a low *chromatic number*, whereas in “negative” instances, we upper-bound the *independence number* of the produced graphs.

► **Theorem 13.** *There is an $\mathcal{O}(\|G\|)$ -time algorithm that, given an n -vertex m -edge graph H , outputs a graph $G := G_H$ of twin-width at most 4 such that*

- $|V(G)| = 3n + (3n + 3m) \cdot 3\lceil \log(3n) \rceil$;
- if H is 3-colorable, then G is 3-colorable;
- for any number p , if no induced subgraph of H on at least p vertices is 3-colorable, then

$$\alpha(G) \leq p + (3n + 3m) \cdot \lceil \log(3n) \rceil.$$

Furthermore, a 4-sequence of G can be computed in $\mathcal{O}(\|G\|)$ time.

Proof. We first turn every vertex v of H into a triangle $(v, 1)(v, 2)(v, 3)$. For every $uv \in E(H)$, we add the three edges $(u, 1)(v, 1)$, $(u, 2)(v, 2)$, $(u, 3)(v, 3)$. So far, we get a graph J with $3n$ vertices and $3n + 3m$ edges.

To obtain G , we s -subdivide every edge e of J , for $s := 2\lceil \log(3n) \rceil$. We arbitrarily orient $e = uv$, say, from u to v , and we add a true twin to each odd-indexed vertex on the new path between u and v : $u(uv)_1(uv)_2(uv)_3 \dots (uv)_s v$. (Note that this “triangle-subdivision” was already used in a similar context in [8].) That concludes the construction of G .

From this description, one can build G in time proportional to its size. One can then observe that $|V(G)| = 3n + (3n + 3m) \cdot 3\lceil \log(3n) \rceil$. Indeed we add $\lceil \log(3n) \rceil$ true twins per subdivided edge of J . The $\mathcal{O}(\|G\|)$ -time algorithm to compute a 4-sequence of G follows the same argument as in the proof of Theorem 8 after contracting every pair of true twins. The latter has the same effect as removing the added true twins (without creating any red edges at this point).

Let us show that $\chi(H) \leq 3$ implies that $\chi(G) \leq 3$. Let $\mu: V(H) \rightarrow [3]$ be a proper 3-coloring of H . We define $\nu': V(J) \rightarrow [3]$ by $\nu'(u, i) := (\mu(u) + i) \bmod 3 + 1$. This is a proper coloring of J : each triangle $(u, 1)(u, 2)(u, 3)$ has all three colors, and a monochromatic edge $(u, i)(v, i)$ under ν' would imply that uv is monochromatic under μ . To extend $\nu': V(J) \rightarrow [3]$ to a proper coloring $\nu: V(G) \rightarrow [3]$ of G , give the color $\nu'(u)$ to every even-indexed vertex on the path: $u(uv)_1(uv)_2(uv)_3 \dots (uv)_s v$ (when uv was oriented from u to v). Finally, we assign the two remaining colors (the ones not already in their neighborhood) to each pair of true twins.

We move to the last item of the statement. First note that the set T_{uv} of $3\lceil \log(3n) \rceil$ new vertices introduced for each edge $uv \in E(J)$ can be vertex-partitioned into $\lceil \log(3n) \rceil$

triangles. Therefore, at most $\lceil \log(3n) \rceil$ of these vertices can be part of the same independent set. Similarly to the mere subdivision, there is a largest independent set of G that for every $uv \in E(J)$ contains at most one original vertex among u, v . Indeed, if both $u, v \in V(G)$ are in an independent set I of G , only $\lceil \log(3n) \rceil - 1$ vertices of T_{uv} can be part of I . This is because $T_{uv} \setminus N_G(\{u, v\})$ can be vertex-partitioned into $\lceil \log(3n) \rceil - 1$ triangles. One can thus remove, say, v from I , and have $\lceil \log(3n) \rceil$ vertices of T_{uv} in I . We conclude that

$$\alpha(G) \leq \alpha(J) + |E(J)| \lceil \log(3n) \rceil = \alpha(J) + (3n + 3m) \lceil \log(3n) \rceil.$$

Actually the equality holds.

We then claim that $\alpha(J) \leq \max\{|X| : X \subseteq V(H), H[X] \text{ is 3-colorable}\}$ (again, the equality holds). Let I be an independent set of J . Because of the triangles $(u, 1)(u, 2)(u, 3)$, I contains at most one vertex of the form (u, i) for every $u \in V(H)$. Furthermore, for every edge $uv \in E(H)$, it cannot be that both (u, i) and (v, i) are in I for some $i \in [3]$. Thus, the partial 3-coloring obtained by giving color i to vertex u for every $(u, i) \in I$ is well-defined, proper, and colors $|I|$ vertices of H . Therefore if no induced subgraph of H on at least p vertices is 3-colorable, then

$$\alpha(G) \leq p + (3n + 3m) \cdot \lceil \log(3n) \rceil. \quad \blacktriangleleft$$

Using these ingredients, we can now prove the main result of this section. The omitted proof is similar to the one of Theorem 1.

► **Theorem 2.** *Unless the ETH fails, there is a constant $\gamma > 0$ such that MIN COLORING does not admit a polynomial-time $n^{\gamma/(\log \log n)^2}$ -approximation algorithm on n -vertex graphs of twin-width at most 4, even when a 4-sequence is given as part of the input.*

Proof. Fix the constants $\beta < 1, \Delta, c$ satisfying Theorem 12. Every promised N -vertex instance H of Theorem 12 is mapped to the graph $G^{[k]}$ with $G := G_H$, as defined in Theorem 13, and $k := \lceil N / \log^{c+2} N \rceil$. We assume, without loss of generality, that $N \geq 2$ is large enough. In particular, the following inequalities hold:

1. $\lceil \log(3N) \rceil \leq 4 \log N$ (this is implied by $N \geq 2$);
2. $(1 - \beta) / (\beta + 3(\Delta + 1) \lceil \log(3N) \rceil) \in (0, 1/2)$;
3. $\log k \geq \log \sqrt{N} = 0.5 \log N$;
4. $12(\Delta + 1)N \geq \lceil \log(3N) \rceil$;
5. $N > \log^{c+2} N$, thus $k \geq 2$.

Building G takes $\mathcal{O}(\|G\|) = \mathcal{O}(\Delta N \log N) = \mathcal{O}(N \log N)$ time, and building $G^{[k]}$ takes time

$$\mathcal{O}(\|G^{[k]}\|) = 2^{\mathcal{O}\left(\frac{N \log N}{\log^{c+2} N}\right)} = 2^{\mathcal{O}\left(\frac{N}{\log^{c+1} N}\right)}.$$

By Theorem 13 and Lemma 10, $G^{[k]}$ indeed has twin-width at most 4, and a 4-sequence of $G^{[k]}$ can be computed in time polynomial in $\|G^{[k]}\|$. Thus the rest of the proof works even if the algorithm $\mathcal{A}$ has access to a 4-sequence of $G^{[k]}$.

Suppose there is an $n^{\mathcal{O}(1)}$ -time r -approximation algorithm $\mathcal{A}$ for MIN COLORING on n -vertex graphs of twin-width at most 4. We set $n := |V(G^{[k]})|$ and

$$\gamma := \frac{0.8(1 - \beta)}{768(\beta + 3(\Delta + 1))(\Delta + 1)}.$$

We now assume that $r = n^{\gamma/(\log \log n)^2}$ and refute the ETH. We run $\mathcal{A}$ on $G^{[k]}$ and obtain a proper coloring of this graph with at most $r \cdot \chi(G^{[k]})$ colors. Based on this coloring, we will

see that we can decide whether H is 3-colorable. This is, by Theorem 12, a contradiction to the ETH once we bound the running time of the reduction, which we now do.

Running time. Building $G^{[k]}$, optionally computing a 4-sequence $\mathcal{S}$ of $G^{[k]}$, and running $\mathcal{A}$ on $G^{[k]}, \mathcal{S}$ takes time

$$\mathcal{O}\left(2^{\mathcal{O}\left(\frac{N}{\log^c+1}\right)} + n^{\mathcal{O}(1)}\right) = 2^{\mathcal{O}\left(\frac{N}{\log^c+1}\right)} = 2^{o(N/\log^c N)},$$

as desired.

Case when H is 3-colorable. By Theorem 13, it holds that $\chi(G) \leq 3$. Therefore, $\chi(G^{[k]}) \leq 3^k$, by giving color $(\nu(v_1), \dots, \nu(v_k))$ to each vertex $(v_1, \dots, v_k) \in V(G^{[k]})$ for a fixed proper 3-coloring ν of G .

Case when H has no 3-colorable induced subgraph on at least βN vertices. We have

$$\chi(G^{[k]}) \geq \frac{|V(G^{[k]})|}{\alpha(G^{[k]})} = \frac{|V(G)|^k}{\alpha(G)^k} \geq \frac{|V(G)|^k}{\alpha(G)^k} = \left(\frac{|V(G)|}{\alpha(G)}\right)^k.$$

where the first inequality holds for any graph, and the second inequality uses Lemma 9.

We set

$$q := (3N + 3|E(H)|)\lceil \log(3N) \rceil \leq 3(\Delta + 1)N\lceil \log(3N) \rceil,$$

since H has maximum degree at most Δ . By Theorem 13,

$$\alpha(G) \leq \beta N + q \text{ and } |V(G)| = 3N + 3q.$$

Thus,

$$\chi(G^{[k]}) \geq \left(\frac{3N + 3q}{\beta N + q}\right)^k = 3^k \cdot \left(1 + \frac{(1 - \beta)N}{\beta N + q}\right)^k \geq 3^k \cdot \left(1 + \frac{1 - \beta}{\beta + 3(\Delta + 1)\lceil \log(3N) \rceil}\right)^k.$$

Approximation gap. The previous paragraphs have established that the approximation gap for $\chi(G^{[k]})$ is at least

$$\frac{3^k \cdot \left(1 + \frac{(1 - \beta)N}{\beta N + q}\right)^k}{3^k} = \left(1 + \frac{1 - \beta}{\beta + 3(\Delta + 1)\lceil \log(3N) \rceil}\right)^k.$$

Thus we just need to show that

$$r := n^{\gamma/(\log \log n)^2} < \left(1 + \frac{1 - \beta}{\beta + 3(\Delta + 1)\lceil \log(3N) \rceil}\right)^k$$

in order to conclude.

▷ **Claim 14.** It holds that $\log r < k \log \left(1 + \frac{1 - \beta}{\beta + 3(\Delta + 1)\lceil \log(3N) \rceil}\right)$.

Proof. We have $\log r = \frac{\gamma \log n}{(\log \log n)^2}$ and

$$\log n = k \log |V(G)| = k \log(3N + 3q) \leq k \log(12(\Delta + 1)N\lceil \log(3N) \rceil) \leq 2k \log(12(\Delta + 1)N),$$

where the last inequality uses Item 4.

Besides, $(\log \log n)^2 \geq \log^2 k$. Thus,

$$\begin{aligned} \log r &\leq \frac{\gamma \cdot 2k \log(12(\Delta+1)N)}{\log^2 k} \leq \frac{\gamma \cdot 24(\Delta+1)k \log N}{\log^2 k} < \gamma \cdot 48(\Delta+1) \cdot \frac{N}{\log^{c+1} N \cdot \log^2 k} \\ &\leq \gamma \cdot 48(\Delta+1) \cdot \frac{N}{\log^{c+1} N \cdot 0.25 \log^2 N} = \gamma \cdot 192(\Delta+1) \cdot \frac{N}{\log^{c+3} N}, \end{aligned}$$

where the penultimate inequality holds because of Item 5, the last inequality uses Item 3, whereas the other inequalities always hold.

On the other hand,

$$\begin{aligned} k \log \left(1 + \frac{1-\beta}{\beta + 3(\Delta+1)\lceil \log(3N) \rceil} \right) &\geq \frac{k \cdot 0.8(1-\beta)}{\beta + 3(\Delta+1)\lceil \log(3N) \rceil} \geq \frac{k \cdot 0.8(1-\beta)}{4(\beta + 3(\Delta+1)) \log N} \\ &\geq \frac{0.8(1-\beta)}{4(\beta + 3(\Delta+1))} \cdot \frac{N}{\log^{c+3} N}, \end{aligned}$$

where the first inequality holds because of Item 2, and the second uses Item 1.

And we conclude since $\gamma = \frac{0.8(1-\beta)}{4(\beta+3(\Delta+1))} \cdot \frac{1}{192(\Delta+1)}$. $\triangleleft$

This concludes the proof since, under the promise, H is 3-colorable if and only if $\mathcal{A}$ uses at most $r \cdot 3^k$ colors. $\blacktriangleleft$

5 k -Independent Set in graphs of bounded low-radius merge-width

In this section, we reduce GRID TILING to k -INDEPENDENT SET, to show the $W[1]$ -hardness of the latter problem in graphs of bounded low-radius merge-width. The reduction is inspired by the one of Dreier, Mählmann, and Siebertz. [16, Thm. 2] and we use their notations. The GRID TILING problem was introduced in [24]. Given some integers $k, n \in \mathbb{N}$ and a collection $\mathcal{S}$ of *tile* subsets $S_\pi \subseteq [n]^2$ for each *grid cell* $\pi \in [k]^2$, the goal is to decide whether there exists a selection of tiles $s(\pi) \in S_\pi$ for all $\pi \in [k]^2$ such that:

- The tiles $s(\pi)$ and $s(\pi_{\blacktriangleright})$ have the same first component for all cells $\pi = (i, j) \in [k-1] \times [k]$ and $\pi_{\blacktriangleright} := (i+1, j)$; and
- The tiles $s(\pi)$ and $s(\pi_{\blacktriangledown})$ have the same second component for all cells $\pi = (i, j) \in [k] \times [k-1]$ and $\pi_{\blacktriangledown} := (i, j+1)$.

It is known that GRID TILING is $W[1]$ -hard parameterized by the grid size k and unless the ETH fails it has no $f(k)n^{o(k)}$ -time algorithm for any computable function f [13, Thm. 14.28].

► **Theorem 3.** k -INDEPENDENT SET is $W[1]$ -hard in graphs of bounded $mw_{o(k)}$.


Proof. Let $r \geq 3$ be an odd integer. For a GRID TILING instance $(\mathcal{S}, k, n)$, we will construct a graph $G_{\mathcal{S}}^r$ with $mw_{r-1}(G_{\mathcal{S}}^r) \leq 4r-3$. For each grid cell $\pi \in [k]^2$, we construct a cell gadget C_π^r , illustrated in Figure 1. The cell gadget C_π^r consists of a horizontal path and a vertical path of r cliques:

- $B_\pi^{h_1}, B_\pi^{h_2}, \dots, B_\pi^{h_r}$ where $B_\pi^{h_i}$ has vertex set $\{b(\tau, h_i, \pi) \mid \tau \in S_\pi\}$; and
- $B_\pi^{v_1}, B_\pi^{v_2}, \dots, B_\pi^{v_r}$ where $B_\pi^{v_i}$ has vertex set $\{b(\tau, v_i, \pi) \mid \tau \in S_\pi\}$.

They cross in the middle, i.e., $B_\pi^{h_m} = B_\pi^{v_m}$ for $m = \frac{r+1}{2}$.

We then add a co-matching between any two consecutive blocks. Formally, for every $\tau, \tau' \in S_\pi$ with $\tau \neq \tau'$, and every $i \in [r-1]$, we add an edge

- between $b(\tau, h_i, \pi)$ and $b(\tau', h_{i+1}, \pi)$; and



■ **Figure 1** Construction of a gadget cell C_π^5 for a cell S_π , and its edge-relation with the right and down gadget cells. Any set of vertices in a black box forms a clique. Two black boxes linked with a black thick edge form a biclique minus the edges in blue, which are non-edges.

- between $b(\tau, v_i, \pi)$ and $b(\tau', v_{i+1}, \pi)$.

Let us name the endpoints of the paths after the four directions *up*, *down*, *left*, and *right*:

$$L_\pi = B_\pi^{h_1} \quad R_\pi = B_\pi^{h_r} \quad U_\pi = B_\pi^{v_1} \quad D_\pi = B_\pi^{v_r}$$

We call these blocks the *border blocks*, and any other block in C_π^r is called an *interior block*. We construct G_S^r by taking the disjoint union of all the C_π^r for $\pi \in [k]^2$, and adding edges between border blocks of adjacent cells. Specifically, we add an edge between

- $b(\tau, h_r, \pi) \in R_\pi$ and $b(\tau', h_1, \pi_\blacktriangleright) \in L_{\pi_\blacktriangleright}$ if and only if τ and τ' have a different first component, for all $\pi \in [k-1] \times [k]$, and
- $b(\tau, v_r, \pi) \in D_\pi$ and $b(\tau', v_1, \pi_\nabla) \in U_{\pi_\nabla}$ if and only if τ and τ' have a different second component, for all $\pi \in [k] \times [k-1]$.

This concludes the construction of G_S^r . Note that $|V(G_S^r)| = \mathcal{O}(rn^2k^2)$.

▷ **Claim 15.** $(\mathcal{S}, k, n)$ is a positive GRID TILING instance if and only if G_S^r is a positive $(2r-1)k^2$ -INDEPENDENT SET instance.

Proof. Assume there is a solution for $(\mathcal{S}, k, n)$ using tiles $s(\pi) \in S_\pi$ for every $\pi \in [k]^2$. We claim that

$$I = \bigcup_{\pi \in [k]^2} I_\pi \quad \text{with} \quad I_\pi := \{b(s(\pi), h_i, \pi), b(s(\pi), v_i, \pi) \mid i \in [r]\}$$

is an independent set of size $(2r-1)k^2$. Clearly, I has the desired size. For each cell π , I_π is an independent set since it comprises only copies of the same tile $s(\pi)$. By construction, two vertices in distinct sets $x \in I_\pi$ and $y \in I_{\pi'}$ can only be adjacent if, up to symmetry, either

- $\pi' = \pi_\blacktriangleright$, and $x \in R_\pi$ and $y \in L_{\pi_\blacktriangleright}$ have a different first component; or
- $\pi' = \pi_\nabla$, and $x \in D_\pi$ and $y \in U_{\pi_\nabla}$ have a different second component.

Both possibilities are ruled out due to the $s(\pi)$ forming a solution to GRID TILING. Hence, I is indeed an independent set.

Conversely, assume there exists an independent set I of size $(2r - 1)k^2$. Since we constructed the k^2 graphs C_π^r from $2r - 1$ cliques each, I must contain exactly $2r - 1$ vertices from each C_π^r . Since those $2r - 1$ vertices are pairwise non-adjacent, they must be of the form

$$I_\pi := \{b(s(\pi), h_i, \pi), b(s(\pi), v_i, \pi) \mid i \in [r]\}$$

for some common $s(\pi) \in S_\pi$. We now argue that the tiles $s(\pi)$ form a GRID TILING solution for $(\mathcal{S}, k, n)$. For every $\pi \in [k - 1] \times [k]$, we have that $b(s(\pi), h_r, \pi) \in R_\pi$ and $b(s(\pi), h_1, \pi) \in L_\pi$ are part of the same independent set I and hence non-adjacent. By construction, we must have that $s(\pi)$ and $s(\pi_\blacktriangleright)$ agree on the first component. Analogously, for every $\pi \in [k] \times [k - 1]$, the tiles $s(\pi)$ and $s(\pi_\nabla)$ agree on the second component. This concludes the proof of the correctness of the reduction. $\triangleleft$

We now upper-bound the radius- $(r - 1)$ merge-width of the produced graphs.

$\triangleright$ **Claim 16.** $\text{mw}_{r-1}(G_S^r) \leq 4r - 3$.

Proof. We will construct a merge sequence for $G := G_S^r$ of radius- $(r - 1)$ width $4r - 3$. First, we define the following sets:

- Let N_{in} be the set of all the non-edges between consecutive copies of B_π inside a cell gadget C_π^r for all $\pi \in [k]^2$. These are the blue edges within the same cell gadget in Figure 1.
- Let N_{out} be the set of all the non-edges between adjacent border blocks of adjacent cell gadgets in G . These are the blue edges between two cell gadgets in Figure 1.
- Define the *layers* of each border block:

$$L_\pi^i := \{b(\tau, h_1, \pi) \mid \tau \in S_\pi, \tau = (i, _)\}, \quad R_\pi^i := \{b(\tau, h_r, \pi) \mid \tau \in S_\pi, \tau = (i, _)\},$$

$$U_\pi^i := \{b(\tau, v_1, \pi) \mid \tau \in S_\pi, \tau = (_, i)\}, \quad D_\pi^i := \{b(\tau, v_r, \pi) \mid \tau \in S_\pi, \tau = (_, i)\}.$$

We can define the following merge sequence. We use the letter F for the set of resolved pairs, since R is taken by the right border blocks.

1. Let $\mathcal{P}_1$ be the partition of $V(G)$ into singletons and $F_1 := N_{in}$.
2. The parts of $\mathcal{P}_2$ are each interior block of each cell gadget, and for each $i \in [n]$ and $X \in \{L, R, U, D\}$, each layer X_π^i . We keep $F_2 := F_1$.
3. The parts of $\mathcal{P}_3$ are each block of each cell gadget. We set $F_3 := F_2 \cup N_{out}$.
4. We finally let $\mathcal{P}_4 := \{V(G)\}$ and $F_4 := F_3 \cup E(G)$.

By construction $\mathcal{P}_i$ is coarser than $\mathcal{P}_{i-1}$, $F_i \subseteq F_{i+1}$, and $\mathcal{P}_i$ is homogeneous modulo F_i . We now argue that the radius- $(r - 1)$ width of the sequence

$$(\mathcal{P}_1, F_1), (\mathcal{P}_2, F_2), (\mathcal{P}_3, F_3), (\mathcal{P}_4, F_4)$$

is at most $4r - 3$.

The radius- $(r - 1)$ (in fact the radius- ∞) width of $(\mathcal{P}_1, F_2 = N_{in})$ is at most $2r - 1$ since every connected component of $(V(G), F_2)$ has exactly $2r - 1$ vertices.

For $(\mathcal{P}_2, F_3)$, observe that for any vertices x and y in the same border block B_π of C_π^r but in different layers, $\text{dist}_{F_3}(x, y) \geq 2r$. Indeed, if there is a path of length at most $2r$ between them, then it must be of the form $x, \dots, x', z, y', \dots, y$ where x' and y' are the respective copies of x and y in another border block B'_π of C_π^r , and z is a vertex in the border block adjacent to B'_π such that $x'z, y'z \in N_{out}$. Since $\text{dist}_{F_3}(x, x') = \text{dist}_{F_3}(x, y') = r - 1$, it follows

that $\text{dist}_{F_3}(x, y) \geq 2r$. Therefore, for any $v \in V(G)$, the set $\text{Ball}_{F_3}^{r-1}(v)$ intersects at most one layer from each border block. Since each layer of a border block and each interior block is in one part of $\mathcal{P}_2$, it follows that for any $v \in V(G)$, we have $|\text{Ball}_{F_3}^{r-1}(v)/\mathcal{P}_2| \leq 4r - 3$.

Finally, for $(\mathcal{P}_3, F_4)$, the edges of G are only between consecutive blocks inside a cell gadget, or between adjacent border blocks. Since any block of G is in one part of $\mathcal{P}_3$, for any $v \in V(G)$, we have $|\text{Ball}_{F_4}^{r-1}(v)/\mathcal{P}_3| \leq 4r - 3$. This concludes the proof. $\triangleleft$

These claims readily give that k -IS is $W[1]$ -hard on graphs of bounded radius- r merge-width, where r is a fixed constant. However, we can extract more. Fix two computable functions $f(x) = o(x)$ and $g(x) \geq 4x$. Then we say that G has (f, g, k) -merge-width if $\text{mw}_{f(k)}(G) \leq g(k)$.

$\triangleright$ **Claim 17.** k -INDEPENDENT SET on graphs with (f, g, k) -merge-width is $W[1]$ -hard.

Proof. We design an FPT reduction from GRID TILING with parameter q to k -INDEPENDENT SET with $k := (2r - 1)q^2$, where $r := r(q)$ is the smallest odd integer at least equal to 3 such that

$$f((2r - 1)q^2) \leq r - 1.$$

This integer exists since $f(x) = o(x)$, and is computable since f is computable.

Let $(\mathcal{S}, q, n)$ be an instance of GRID TILING. Computing $G := G_{\mathcal{S}}^r$ can be done in time $\mathcal{O}(\|G\|) = \mathcal{O}(r^2 q^4 \cdot n^4)$. This is an FPT reduction in parameter q since r depends only on q . Besides, G has (f, g, k) -merge-width, since

$$\text{mw}_{f(k)}(G) = \text{mw}_{f((2r-1)q^2)}(G) \leq \text{mw}_{r-1}(G) \leq 4r - 3 \leq 4(2r - 1)q^2 \leq g((2r - 1)q^2) = g(k).$$

Where the first inequality comes from monotonicity of merge-width with respect to the radius, the second is due to Claim 16, and the last, by the assumption on g . We conclude by Claim 15. $\triangleleft$

In particular, for every $r = o(k)$, k -IS is $W[1]$ -hard on graphs of bounded mw_r . $\blacktriangleleft$


6 k -Dominating Set in graphs of bounded low-radius merge-width

In this section, we adapt the reduction from the previous section to show the $W[1]$ -hardness of k -DOMINATING SET in graphs of low-radius merge-width. The main difference is how the selection of copies of the same tile in a cell gadget is enforced. First, instead of cliques, the blocks will be independent sets. Then, we will achieve the vertex selection by inserting between every pair of consecutive blocks B and B' , an additional block A such that there is a co-matching between B and A , and a matching between A and B' . The only way to dominate this new block A with two vertices will be to pick copies of the same tile in B and B' . A similar construction can be done to enforce that selected tiles in adjacent cells agree on their corresponding component (see Figure 2).

$\blacktriangleright$ **Theorem 4.** k -DOMINATING SET is $W[1]$ -hard in graphs of bounded $\text{mw}_{o(k)}$.

Proof. Let $r > 3$ be an odd integer, for a GRID TILING instance $(\mathcal{S}, k, n)$ we will construct a graph $G_{\mathcal{S}}^r$ with $\text{mw}_{2r-2}(G_{\mathcal{S}}^r) = \mathcal{O}(r)$. For each grid cell $\pi \in [k]^2$, we construct a cell gadget C_{π}^r . The cell gadget C_{π}^r consists of a horizontal path and a vertical path of $2r - 1$ edgeless graphs:

$$\blacksquare B_{\pi}^{h_1}, B_{\pi}^{h_2}, \dots, B_{\pi}^{h_{2r-1}} \text{ where } B_{\pi}^{h_i} \text{ have vertex set } \{b(\tau, h_i, \pi) \mid \tau \in S_{\pi}\}; \text{ and}$$



■ **Figure 2** Left: Co-matching and matching between consecutive blocks in a gadget cell with i even. Right: construction between two border blocks of adjacent cells. Any set of vertices in a black box forms an edgeless graph. Vertices incident to a black box are complete to the set of vertices contained in that box.

- $B_\pi^{v_1}, B_\pi^{v_2}, \dots, B_\pi^{v_{2r-1}}$ where $B_\pi^{v_i}$ have vertex set $\{b(\tau, v_i, \pi) \mid \tau \in S_\pi\}$.

They are crossing in the middle, i.e., $B_\pi^{h_r} = B_\pi^{v_r}$.

For every even block we add a co-matching towards the previous block and a matching towards the next block. That is, for $\tau, \tau' \in S_\pi$ with $\tau \neq \tau'$ and even $i \in [2r-1]$, add to C_π^r edges between

- $b(\tau, h_i, \pi)$ and $b(\tau', h_{i-1}, \pi)$; $b(\tau, h_i, \pi)$ and $b(\tau, h_{i+1}, \pi)$;
- $b(\tau, v_i, \pi)$ and $b(\tau', v_{i-1}, \pi)$; $b(\tau, v_i, \pi)$ and $b(\tau, v_{i+1}, \pi)$.

Additionally, for each odd $i \in [2r-1]$, add

- vertices $a_1(h_i, \pi)$ and $a_2(h_i, \pi)$ complete to $B_\pi^{h_i}$; and
- vertices $a_1(v_i, \pi)$ and $a_2(v_i, \pi)$ complete to $B_\pi^{v_i}$.

Note that since r is odd, a pair of such vertices is created for $B_\pi^{h_r} = B_\pi^{v_r}$. This is done to guarantee that a vertex is picked in each odd block. The endpoints of the paths are called *border blocks*, and any other block in C_π^r is called an *interior block*.

We construct $G_\mathcal{S}^r$ by taking the disjoint union of all the C_π^r for $\pi \in [k]^2$. Additionally, to ensure that the same first or second component is selected between adjacent gadget cells, we add the following vertices and edges.

- For each $\pi \in [k-1] \times [k]$, add a set of vertices $\{r(i, \pi) \mid i \in [n]\}$ and edges between $b((i, _), h_{2r-1}, \pi) \in B_\pi^{h_{2r-1}}$ and $r(j, \pi)$ for $i, j \in [n]$ with $i \neq j$; and $b((i, _), h_1, \pi_\blacktriangleright) \in B_\pi^{h_1}$ and $r(i, \pi)$ for $i \in [n]$.
- For each $\pi \in [k] \times [k-1]$, add a set of vertices $\{d(i, \pi) \mid i \in [n]\}$ and edges between $b((_, i), v_{2r-1}, \pi) \in B_\pi^{v_{2r-1}}$ and $d(j, \pi)$ for $i \neq j \in [n]$; and $b((_, i), v_1, \pi_\blacktriangledown) \in B_\pi^{v_1}$ and $d(i, \pi)$ for $i \in [n]$.

This concludes the construction of $G_\mathcal{S}^r$ and note that $|V(G_\mathcal{S}^r)| = \mathcal{O}(rn^2k^2)$.

The following claims can be shown similarly to their counterparts in Theorem 3. For the equivalence of GRID TILING and k -DOMINATING SET instances, it follows from the fact that the only dominating sets of size $2r-1$ of a gadget cell C_π^r are of the form

$$\text{Dom}_\pi := \{b(\tau, h_{2i-1}, \pi), b(\tau, v_{2i-1}, \pi) \mid i \in [r]\}$$

for some $\tau \in S_\pi$. Moreover, the vertices added between adjacent gadget cells force any dominating set to select the corresponding component in a dominating set of size $(2r-1)k^2$.

▷ **Claim 18.** $(\mathcal{S}, k, n)$ is a positive GRID TILING instance if and only if $G_\mathcal{S}^r$ is a positive $(2r-1)k^2$ -DS instance.

Then, for the bound on the radius- $(2r - 2)$ merge-width, the merge sequence is done similarly: first resolve all matchings and co-matchings inside gadget cells; then resolve the co-matchings and the matchings with the vertices between adjacent gadget cells; then add every remaining edge of G_S^r .

▷ **Claim 19.** $\text{mw}_{2r-2}(G_S^r) = \mathcal{O}(r)$.

Using the same arguments as in Claim 17, it follows that for any functions $f(k) = o(k)$ and $g(k)$, solving k -DOMINATING SET on graphs with $\text{mw}_{f(k)} \leq g(k)$ is $W[1]$ -hard. ◀

References

- 1 Anna Adamaszek, Sarel Har-Peled, and Andreas Wiese. Approximation schemes for independent set and sparse subsets of polygons. *J. ACM*, 66(4):29:1–29:40, 2019. doi:10.1145/3326122.
- 2 Brenda S. Baker. Approximation algorithms for np-complete problems on planar graphs. *J. ACM*, 41(1):153–180, 1994. doi:10.1145/174644.174650.
- 3 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.
- 4 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.
- 5 Andrej Bogdanov, Kenji Obata, and Luca Trevisan. A lower bound for testing 3-colorability in bounded-degree graphs. In *43rd Symposium on Foundations of Computer Science, FOCS 2002, Vancouver, BC, Canada, November 16-19, 2002, Proceedings*, pages 93–102. IEEE Computer Society, 2002. doi:10.1109/SFCS.2002.1181886.
- 6 Marthe Bonamy and Colin Geniet. χ -boundedness and neighbourhood complexity of bounded merge-width graphs. *CoRR*, abs/2504.08266, 2025. URL: <https://doi.org/10.48550/arXiv.2504.08266>, arXiv:2504.08266, doi:10.48550/ARXIV.2504.08266.
- 7 Édouard Bonnet. *Twin-Width and Contraction Sequences*. 2024. URL: <https://tel.archives-ouvertes.fr/tel-04990502>.
- 8 Édouard Bonnet. Coloring hardness on low twin-width graphs. *Information Processing Letters*, 194:106646, 2026. URL: <https://www.sciencedirect.com/science/article/pii/S002001902600027X>, doi:https://doi.org/10.1016/j.ipl.2026.106646.
- 9 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.
- 10 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.
- 11 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.
- 12 Timothy M. Chan and Sarel Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. *Discret. Comput. Geom.*, 48(2):373–392, 2012. URL: <https://doi.org/10.1007/s00454-012-9417-5>, doi:10.1007/S00454-012-9417-5.
- 13 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3.

- 14 Hugues Déprés. *Twin-width: lower bounds and approximation algorithms. (Twin-width: bornes inférieures et algorithmes d'approximation)*. PhD thesis, École normale supérieure de Lyon, France, 2024. URL: <https://tel.archives-ouvertes.fr/tel-04648829>.
- 15 Irit Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3):12, 2007. doi:10.1145/1236457.1236459.
- 16 Jan Dreier, Nikolas Mählmann, and Sebastian Siebertz. The parameterized complexity of independent set and more when excluding a half-graph, co-matching, or matching. *CoRR*, abs/2602.07606, 2026. URL: <https://doi.org/10.48550/arXiv.2602.07606>, arXiv:2602.07606, doi:10.48550/ARXIV.2602.07606.
- 17 Jan Dreier and Szymon Toruńczyk. Merge-width and first-order model checking. In Michal Koucký and Nikhil Bansal, editors, *Proceedings of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025*, pages 1944–1955. ACM, 2025. doi:10.1145/3717823.3718259.
- 18 Uriel Feige. Approximating maximum clique by removing subgraphs. *SIAM J. Discret. Math.*, 18(2):219–225, 2004. doi:10.1137/S089548010240415X.
- 19 Jacob Fox and János Pach. Computing the independence number of intersection graphs. In Dana Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 1161–1165. SIAM, 2011. doi:10.1137/1.9781611973082.87.
- 20 Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- 21 Martin Grohe. Local tree-width, excluded minors, and approximation algorithms. *Comb.*, 23(4):613–632, 2003. URL: <https://doi.org/10.1007/s00493-003-0037-9>, doi:10.1007/S00493-003-0037-9.
- 22 Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. In *37th Annual Symposium on Foundations of Computer Science, FOCS '96, Burlington, Vermont, USA, 14-16 October, 1996*, pages 627–636, 1996. doi:10.1109/SFCS.1996.548522.
- 23 Pasin Manurangsi. Almost-polynomial ratio eth-hardness of approximating densest k-subgraph. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 954–961. ACM, 2017. doi:10.1145/3055399.3055412.
- 24 Dániel Marx. On the optimality of planar and geometric approximation schemes. In *48th Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 338–348. IEEE Computer Society, 2007. doi:10.1109/FOCS.2007.26.
- 25 Svatopluk Poljak. A note on stable sets and colorings of graphs. *Commentationes Mathematicae Universitatis Carolinae*, 15(2):307–309, 1974.
- 26 David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3(1):103–128, 2007.