



HAL
open science

A more accurate view of the Flat Wall Theorem

Ignasi Sau, Giannos Stamoulis, Dimitrios M. Thilikos

► **To cite this version:**

Ignasi Sau, Giannos Stamoulis, Dimitrios M. Thilikos. A more accurate view of the Flat Wall Theorem. Journal of Graph Theory, 2024, 107 (2), pp.263-297. <10.1002/jgt.23121>. <lirmm-04700544>

HAL Id: lirmm-04700544

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-04700544v1>

Submitted on 17 Sep 2024

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

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



HAL Authorization

A more accurate view of the Flat Wall Theorem¹

Ignasi Sau² Giannos Stamoulis³ Dimitrios M. Thilikos²

Abstract

We introduce a supporting combinatorial framework for the Flat Wall Theorem. In particular, we suggest two variants of the theorem and we introduce a new, more versatile, concept of wall homogeneity as well as the notion of regularity in flat walls. All proposed concepts and results aim at facilitating the use of the irrelevant vertex technique in future algorithmic applications.

Keywords: graph minors; treewidth; Flat Wall Theorem; parameterized algorithms; irrelevant vertex technique; homogeneous walls.

¹All authors were supported by the ANR projects DEMOGRAPH (ANR-16-CE40-0028), ANR project ELIT (ANR-20-CE48-0008), ESIGMA (ANR-17-CE23-0010), the French-German Collaboration ANR/DFG Project UTMA (ANR-20-CE92-0027). The third author was also supported by the *Ministère de l'Europe et des Affaires étrangères* (MEAE) and the *Ministère de l'Enseignement supérieur et de la Recherche* (MESR), via the Franco-Norwegian project PHC Aurora projet N° 51260WL (2024). Emails: ignasi.sau@lirmm.fr, giannos.stamoulis@lirmm.fr, sedthilk@thilikos.info.

²LIRMM, Université de Montpellier, CNRS, Montpellier, France.

³Institute of Informatics, University of Warsaw, Poland

Contents

1	Introduction	3
1.1	The Flat Wall Theorem and its variants	3
1.2	Our contribution	5
1.3	Organization of the paper	7
2	Definitions and preliminary results	7
2.1	Preliminaries	7
2.2	Renditions	9
2.3	Flatness pairs	12
3	Applications	16
3.1	Tilts of subwalls	16
3.2	Apex-walls with compasses of bounded treewidth	16
3.3	Homogeneous walls	21
3.4	Levelings and well-aligned flatness pairs	22
4	Proofs of Theorem 5 and Theorem 6	25
4.1	Stretchings	26
4.2	Classifying perimetric cells	26
4.3	The main lemma	27
4.4	Proofs of Theorem 5 and Theorem 6	32

1 Introduction

One of the cornerstone achievements of the Graph Minors series by Robertson and Seymour was the celebrated *Flat Wall Theorem*, proved in the 13th paper of the series [38]. It is a powerful graph structural result, revealing the local structure of H -minor-free graphs. The Flat Wall Theorem has important consequences and applications in structural graph theory and in graph algorithm design. It served as the combinatorial base for the design of an algorithm for the following two problems:

- **MINOR TESTING:** Given a graph G and a k -vertex graph H , decide whether G contains H as a minor.
- **DISJOINT PATHS:** Given a graph G with k pairs of terminals $(s_i, t_i), \dots, (s_k, t_k)$, decide whether G contains k vertex-disjoint paths joining s_i and t_i for every $i \in \{1, \dots, k\}$.

These algorithms run in time $f(k) \cdot n^3$ on n -vertex graphs, for some computable function $f : \mathbb{N} \rightarrow \mathbb{N}$ (see [28] for quadratic-time improvements). This, using the terminology of parameterized complexity, implies that both above problems, when parameterized by k , belong to the parameterized class FPT or, alternatively, *admit FPT-algorithms*. In order to obtain these algorithms, Robertson and Seymour introduced a powerful technique, called the *irrelevant vertex technique*, which has now become a standard technique in the design of parameterized algorithms (see e.g., Section 7 of the textbook [8]). Further algorithmic applications combining the Flat Wall Theorem and the irrelevant vertex technique appeared later in [2,9,12,17,26], while generalizations to directed graphs have recently appeared in [13,21].

1.1 The Flat Wall Theorem and its variants

The original statement of the Flat Wall Theorem, as appeared in [38], is the following.

Proposition 1. *There exist functions $f : \mathbb{N}^2 \rightarrow \mathbb{N}$ and $f' : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that if G is a graph and h and k are integers, then one of the following holds:*

1. G contains K_h as a minor¹.
2. G has treewidth at most $f(k, h)$.
3. G has a vertex set A with $|A| \leq f'(h)$, such that $G \setminus A$ contains a flat wall W of height k .

We postpone the formal definitions of “treewidth”, the related concept of “tree decomposition”, and “flat wall” to Section 2. One can get a quick idea of a wall by looking at Figure 1 and of flat wall by looking at Figure 3 and Figure 5. Intuitively, a flat wall W is contained in a larger graph, its *compass*, that is separated from the rest of the graph via a separator S that is a “suitably chosen” part of the perimeter of W . This compass is “flat” in the sense that it does not contain two disjoint paths whose endpoints are in S and are “crossing” with respect to the cyclic ordering induced in S by the perimeter of W . As proved by Kawarabayashi, Thomas, and Wollan [32], this flatness property can be certified by a concept called *rendition* (corresponding to the concept

¹I.e., some subgraph of G can be contracted to a complete graph on h vertices.

of *rural division* in [38]) that can be seen as a plane embedding inside a disk of a hypergraph with hyperedges of arity at most three (see Figure 2 for a visualization of a rendition). Then the compass is “embedded” inside the rendition so that it can be seen as the union of graphs called *flaps* bijectively mapped to the hyperedges of the rendition.

In its original version in [38], Proposition 1 was proved for $f'(h) = \binom{h}{2}$ with the additional assertion that $f(k, h)$ is a bound on the treewidth of the “internal flaps”, i.e., those that are not incident to the perimeter of W . Later, in [14], the same result was proved (without an algorithm) for $f'(h) = h - 5$ and $f(k, h) = \mathcal{O}_h(k)$.² The original result of Robertson and Seymour was accompanied with an $\mathcal{O}(n \cdot m)$ -time algorithm³ that outputs a certifying structure for each possible outcome. This algorithm was further improved to a linear one by Kawarabayashi, Kobayashi, and Reed in [28].

A recent wave of improvements of Proposition 1 appeared in the following form [7, 32].

Proposition 2. *There exist functions $f : \mathbb{N}^2 \rightarrow \mathbb{N}$ and $f' : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that if G is a graph and h and k are integers, and G contains a wall W of height $f(k, h)$ as a subgraph, then one of the following holds:*

1. G contains K_h as a minor.
2. G has a vertex set A with $|A| \leq f'(h)$, such that $G \setminus A$ contains a flat wall W' of height k .

Notice that Proposition 2 can indeed be seen as an extension of Proposition 1 because the exclusion of a wall of height k in a K_h -minor-free graph implies that its treewidth is bounded by $\mathcal{O}_h(k)$ [10, 27]. Moreover, according to [32, Theorem 1.9], Proposition 2 holds for $f'(h) = \mathcal{O}(h^{24})$ and $f(k, h) = \mathcal{O}(h^{24}(h^2 + k))$, and it enjoys the following additional features:

- (A) In the first case, the graph K_h is a minor of G in a way that is “grasped by the wall W ”⁴
- (B) In the second case, the flat wall W' is a *subwall* of W .
- (C) Proposition 2 comes with an algorithm that certifies one of the two outcomes in linear time, in particular, in $\mathcal{O}(h^{24} \cdot m + n)$ time.

Moreover, the same result with Features (A) and (B) is proved in [32, Theorem 1.7] with the optimal function $f'(h) = h - 5$ at the cost of a worse bound for $f(k, h)$. Also [32, Theorem 1.8] corresponds to Proposition 2 with the additional feature that the compass of the flat wall W' contains no wall of height $f(k, h) + 1$, again at the cost of a worse bound for $f(k, h)$.

Later, Chuzhoy [7] drastically improved the bounds of Proposition 2 with the extra Features (A) and (B) to $f'(h) = h - 5$ and $f(k, h) = \mathcal{O}(h \cdot (h + k))$. Moreover, Chuzhoy gives a polynomial-time algorithm for her improved variant, however she does not specify whether this algorithm can run in linear time, as the one in [32, Theorem 1.9].

²The notation ‘ $\mathcal{O}_h(\cdot)$ ’ means that the hidden constants depend only on h .

³In this paper we always denote by n and m the number of vertices and edges, respectively, of the graph under consideration.

⁴We avoid here the formal definition of “grasping by a wall” as we do not make use of it in this paper; see [32] for the details. However, we stress that it provides additional information that is used in further applications (see e.g., [33]).

1.2 Our contribution

In this paper we provide a series of enhanced algorithmic versions of the Flat Wall Theorem as well as a series of combinatorial tools related to the applicability of the irrelevant vertex technique. In our presentation we adopt the framework and the terminology of [32]. Our aim is to introduce a “more accurate” view of the Flat Wall Theorem that, we hope, will be useful for future algorithmic applications. Our contribution consists in the following.

(α) **Subwalls of flat walls are not always flat.** Our initial motivation comes from the fact⁵ that the claimed Feature (B) in Proposition 2, as stated in [32], needs some slight (but not neglectable) revision. This feature is based on [32, Lemma 6.1], asserting that if W is a flat wall and W' is a subwall of W that is disjoint from the perimeter of W , then W' is *also* a flat wall of G . As we observe in Subsection 2.3, there are some very marginal cases where a subwall of a flat wall is not flat anymore. This phenomenon is illustrated in the flat wall of Figure 3 (in Subsection 2.3).

(β) **A reparation framework.** Fortunately, the issue raised in (α) is just a minor formal mismatch that harms neither the spirit of the proofs of [32] nor the “essential” correctness of subsequent results that are based on [32]. The first contribution of our paper is to propose an extension of the framework of [32] that supports a formally correct statement of Feature (B) in Proposition 2. What we show (Theorem 5) is that if a wall W is a flat wall, whose flatness is certified by some rendition \mathfrak{R} , and W' is a subwall of W , then there is another, slightly different, subwall \tilde{W}' of W , which we call a W' -tilt, that is indeed flat⁶. By the term “slightly different” we mean that W' and its W' -tilt \tilde{W}' may differ *only* perimetrically. Moreover, the rendition certifying the flatness of \tilde{W}' maintains all the “internal” structure of the rendition \mathfrak{R} , relatively to W' . This implies that all the arguments based on Proposition 2 of [32] are essentially correct, and can become formally correct under the suggested framework. In our definitions and proofs we pay attention to all the necessary formalism so to facilitate dealing with future results that may use those of [32] (or [7]). We conclude with Proposition 7 that is a version of Proposition 2 translated into our framework.

(γ) **A Flat Wall Theorem with compasses of bounded treewidth.** Our next result, Theorem 8 in Subsection 3.2, is an improved version of Proposition 1 with the following additional features: (1) $f(k, h) = k \cdot 2^{\mathcal{O}(h^2 \log h)}$ and $f'(h) = \mathcal{O}(h^{24})$, (2) in the third case, the compass of the wall W comes with a tree decomposition of width at most $f(k, h)$, and (3) the result is accompanied by a $2^{\mathcal{O}_h(r^2)} \cdot n$ time algorithm. Notice that a non-algorithmic version of this result could be indirectly derived, with worse functions, combining [32, Theorem 1.8] and the main result of Kawarabayashi and Kobayashi in [27]. We present this result in this paper for the following reasons: first because it is new, second because it is in a form suitable for future applications where it is important that the compass has bounded treewidth, and third because its proof provides an indicative sample of the potential of the formalism of W' -tilts that we suggest in (β).

⁵This was first spotted in the conference article [41].

⁶In fact, Theorem 5 applies not only to subwalls W' of W , but also to every subwall W' of the compass of W that is not “contained” in a flap. See Subsection 2.3 for the details.

(δ) **An alternative concept of wall homogeneity.** As mentioned before, the Flat Wall Theorem has been the combinatorial base for the FPT-algorithms of [38] for MINOR TESTING and DISJOINT PATHS. One of the cornerstone ideas of [38] was to prove that the existence of a “big enough” flat wall W in the input graph G implies that a minor-model of H or a collection of k disjoint paths in G can be safely rerouted so to avoid the central vertices of this wall (see Figure 1 for a visualization of the central vertices of a wall). This permits us to declare parts of the wall “irrelevant” and find an equivalent instance of the problem with fewer vertices. In fact, avoiding the central vertices is not so straightforward when dealing with a flat wall W . This is because the rerouting has to be done inside the compass K of W where the paths should be rerouted through different, however “equivalent”, flaps of the compass. To deal with this, Robertson and Seymour defined in [38] the concept of wall *homogeneity*. Roughly speaking, when a wall is homogenous then the *variety of the ways* that paths may be routed through the flaps that are inside some “brick” of the wall is the same for all bricks. In [38] it was proved that every big enough flat wall contains a still big homogeneous subwall where the claimed rerouting is possible, with the help of later results of the Graph Minors series [39, 40].

The definition of wall homogeneity in [38] was based on the concept of the *vision of a flap* and was quite particular to the problems it was dealing with. To our knowledge, after [38], not much use of homogeneity, as defined in [38], was done for algorithmic purposes. Most of the results where the irrelevant vertex technique was applied concerned questions on surface-embeddable graphs where the wall is “already” disk-embedded and there is no need of homogeneity (see e.g. [16, 22–25, 29–31, 35, 36]). An indicative exception to this rule is the celebrated algorithm in [17, 18] for the problem of checking whether H is a topological minor of a graph G where some notion of homogeneity, tailor-made for this problem, was introduced (see [18, Theorem 5.8] and also [12]).

In this paper we introduce an alternative notion of wall homogeneity that is simpler and more versatile to use. This is done in Subsection 3.3 and is based on the framework introduced in (β). Our definition may help dealing with the wide variety of the problems as it permits any version of finite index flap equivalency (for instance, flap equivalency based on MSOL-expressibility). We accompany the definition with an FPT-algorithm that finds a homogeneous subwall. This, together with the main result of (γ), can permit us to find “big-enough” homogeneous walls with compasses of bounded treewidth. This, in turn, will permit the answer of MSOL-queries in parts of the compass and will allow more elaborated applications of the irrelevant vertex technique (such as those used for problems on surface-embeddable graphs in [11, 15]).

(ε) **Regular flatness pairs and plane representations.** We call a pair (W, \mathfrak{R}) *flatness pair* if W is a flat wall whose flatness is certified by the rendition \mathfrak{R} . Based on the framework of (β), in Subsection 3.4 we introduce a notion of *regularity* for flatness pairs, which roughly demands that the branching vertices of the wall are “internal” with respect to the flaps of the compass of W . Regular flatness pairs permit the representation of the compass of a flat wall by a graph embedded in a disk and a “well-arranged” wall inside it. This “plane” representation of flat walls will appear handy in other applications. For instance, it has been a useful tool for the proofs of the main combinatorial results of [5, 41] as it makes it possible to translate routing questions inside compasses to analogous questions on planar embeddings and deal with them in an easier way (using the new homogeneity concept of (δ)).

1.3 Organization of the paper

In [Section 2](#) we provide some definitions and preliminary results and we state the two main results of this paper ([Theorem 5](#) and [Theorem 6](#)), that assert the existence of an algorithm computing a tilt of a subwall of a flat wall and of an algorithm, that given a flatness pair outputs a regular flatness pair, respectively. We prove [Theorem 5](#) and [Theorem 6](#) in [Section 4](#). In [Section 3](#), we develop the tools to address the topics (β) , (γ) , (δ) , and (ε) listed above.

2 Definitions and preliminary results

2.1 Preliminaries

Sets and integers. We denote by \mathbb{N} the set of non-negative integers. Given two integers p, q , where $p \leq q$, we denote by $[p, q]$ the set $\{p, \dots, q\}$. For an integer $p \geq 1$, we set $[p] = [1, p]$ and $\mathbb{N}_{\geq p} = \mathbb{N} \setminus [0, p-1]$. For a set S , we denote by 2^S the set of all subsets of S and by $\binom{S}{2}$ the set of all subsets of S of size 2. If \mathcal{S} is a collection of objects where the operation \cup is defined, then we denote $\mathbf{U}\mathcal{S} = \cup_{X \in \mathcal{S}} X$.

Basic concepts on graphs. As a graph G we denote any pair (V, E) where V is a finite set and $E \subseteq \binom{V}{2}$, that is, all graphs of this paper are undirected, finite, and without loops or multiple edges. We also define $V(G) = V$ and $E(G) = E$. We say that a pair $(L, R) \in 2^{V(G)} \times 2^{V(G)}$ is a *separation* of G if $L \cup R = V(G)$ and there is no edge in G between $L \setminus R$ and $R \setminus L$. Given a vertex $v \in V(G)$, we denote by $N_G(v)$ the set of vertices of G that are adjacent to v in G . Also, given a set $S \subseteq V(G)$, we set $N_G(S) = \cup_{v \in S} N_G(v)$. A vertex $v \in V(G)$ is *isolated* if $N_G(v) = \emptyset$. For $S \subseteq V(G)$, we set $G[S] = (S, E \cap \binom{S}{2})$ and use $G \setminus S$ to denote $G[V(G) \setminus S]$. Given an edge $e = \{u, v\} \in E(G)$, we define the *subdivision* of e to be the operation of deleting e , adding a new vertex w , and making it adjacent to u and v . Given two graphs H, G , we say that H is a *subdivision* of G if H can be obtained from G by subdividing edges. The *contraction* of an edge $e = \{u, v\}$ of a simple graph G results in a simple graph G' obtained from $G \setminus \{u, v\}$ by adding a new vertex uv adjacent to all the vertices in the set $(N_G(u) \cup N_G(v)) \setminus \{u, v\}$. A graph G' is a *minor* of a graph G if G' can be obtained from a subgraph of G after a series of edge contractions.

Disk-embedded graphs. A *closed* (resp. *open*) *disk* is a set homeomorphic to the set $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq 1\}$ (resp. $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 < 1\}$). Let Δ be a closed disk. We use $\text{bd}(\Delta)$ to denote the boundary of Δ and $\text{int}(\Delta)$ to denote the open disk $\Delta \setminus \text{bd}(\Delta)$. When we embed a graph G in the plane or in a disk, we treat G as a set of points. This permits us to make set operations between graphs and sets of points. We say that a graph G is Δ -*embedded* if G is embedded in Δ without crossings such that the intersection of $\text{bd}(\Delta)$ and G (seen as a set of points of Δ) is a subset of $V(G)$.

A *circle* of Δ is any set homeomorphic to $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$. Given two distinct points $x, y \in D$, an (x, y) -*arc* of D is any subset of D that is homeomorphic to the closed interval $[0, 1]$, where the image of zero is x and the image of one is y .

Walls. Let $k, r \in \mathbb{N}$. The $(k \times r)$ -grid is the graph whose vertex set is $[k] \times [r]$ and two vertices (i, j) and (i', j') are adjacent if $|i - i'| + |j - j'| = 1$. An *elementary r -wall*, for some odd integer $r \geq 3$, is the graph obtained from a $(2r \times r)$ -grid with vertices $(x, y) \in [2r] \times [r]$, after the removal of the “vertical” edges $\{(x, y), (x, y + 1)\}$ for odd $x + y$, and then the removal of all vertices of degree one. This definition is slightly different than other definitions in the literature (i.e., we require r to be odd), but we adopt this one for technical reasons. Notice that, as $r \geq 3$, an elementary r -wall is a planar graph that has a unique (up to topological isomorphism) embedding in the plane \mathbb{R}^2 such that all its finite faces are incident to exactly six edges. The *perimeter* of an elementary r -wall is the cycle bounding its infinite face, while the cycles bounding its finite faces are called *bricks*. Also, the vertices in the perimeter of an elementary r -wall that have degree two are called *pegs*, while the vertices $(1, 1), (2, r), (2r - 1, 1)$, and $(2r, r)$ are called *corners* (notice that the corners are also pegs).




Figure 1: A 15-wall. The 3-branch vertices are depicted in cyan except from the corner and the central vertices that are depicted in red and orange respectively.

An r -wall is any graph W obtained from an elementary r -wall \bar{W} after subdividing edges (see Figure 1). A graph W is a *wall* if it is an r -wall for some odd $r \geq 3$ and we refer to r as the *height* of W . Given a graph G , a *wall of G* is a subgraph of G that is a wall. We insist that, for every r -wall, the number r is always odd.

We call the vertices of degree three of a wall W *3-branch vertices*. A cycle of W is a *brick* (resp. the *perimeter*) of W if its 3-branch vertices are the vertices of a brick (resp. the perimeter) of \bar{W} . We denote by $\mathcal{C}(W)$ the set of all cycles of W . We use $D(W)$ in order to denote the perimeter of the wall W . A brick of W is *internal* if it is disjoint from $D(W)$.

Subwalls. Given an elementary r -wall \bar{W} , some $i \in \{1, 3, \dots, 2r - 1\}$, and $i' = (i + 1)/2$, the i' -th *vertical path* of \bar{W} is the one whose vertices, in order of appearance, are $(i, 1), (i, 2), (i + 1, 2), (i + 1, 3), (i, 3), (i, 4), (i + 1, 4), (i + 1, 5), (i, 5), \dots, (i, r - 2), (i, r - 1), (i + 1, r - 1), (i + 1, r)$. Also, given

some $j \in [2, r - 1]$ the j -th horizontal path of \bar{W} is the one whose vertices, in order of appearance, are $(1, j), (2, j), \dots, (2r, j)$.

A *vertical* (resp. *horizontal*) path of W is one that is a subdivision of a vertical (resp. horizontal) path of \bar{W} . Notice that the perimeter of an r -wall W is uniquely defined regardless of the choice of the elementary r -wall \bar{W} . A *subwall* of W is any subgraph W' of W that is an r' -wall, with $r' \leq r$, and such the vertical (resp. horizontal) paths of W' are subpaths of the vertical (resp. horizontal) paths of W .

Tilts. The *interior* of a wall W is the graph obtained from W if we remove from it all edges of $D(W)$ and all vertices of $D(W)$ that have degree two in W . Given two walls W and \tilde{W} of a graph G , we say that \tilde{W} is a *tilt* of W if \tilde{W} and W have identical interiors.

2.2 Renditions

Paintings. Let Δ be a closed disk. Given a subset X of Δ , we denote its closure by \bar{X} and its boundary by $\text{bd}(X)$. A Δ -*painting* is a pair $\Gamma = (U, N)$ where

- N is a finite set of points of Δ ,
- $N \subseteq U \subseteq \Delta$, and
- $U \setminus N$ has finitely many arcwise-connected components, called *cells*, where, for every cell c ,
 - the closure \bar{c} of c is a closed disk and
 - $|\bar{c}| \leq 3$, where $\bar{c} := \text{bd}(c) \cap N$.

We use the notation $U(\Gamma) := U$, $N(\Gamma) := N$ and denote the set of cells of Γ by $C(\Gamma)$. For convenience, we may assume that each cell of Γ is an open disk of Δ .

Notice that, given a Δ -painting Γ , the pair $(N(\Gamma), \{\bar{c} \mid c \in C(\Gamma)\})$ is a hypergraph whose hyperedges have cardinality at most three and Γ can be seen as a plane embedding of this hypergraph in Δ .

Renditions. Let G be a graph, and let Ω be a cyclic permutation of a subset of $V(G)$ that we denote by $V(\Omega)$. By an Ω -*rendition* of G we mean a triple (Γ, σ, π) , where

- (a) Γ is a Δ -painting for some closed disk Δ ,
- (b) $\pi : N(\Gamma) \rightarrow V(G)$ is an injection, and
- (c) σ assigns to each cell $c \in C(\Gamma)$ a subgraph $\sigma(c)$ of G , such that
 - (1) $G = \bigcup_{c \in C(\Gamma)} \sigma(c)$,
 - (2) for distinct $c, c' \in C(\Gamma)$, $\sigma(c)$ and $\sigma(c')$ are edge-disjoint,
 - (3) for every cell $c \in C(\Gamma)$, $\pi(\bar{c}) \subseteq V(\sigma(c))$,
 - (4) for every cell $c \in C(\Gamma)$, $V(\sigma(c)) \cap \bigcup_{c' \in C(\Gamma) \setminus \{c\}} V(\sigma(c')) \subseteq \pi(\bar{c})$, and

- (5) $\pi(N(\Gamma) \cap \text{bd}(\Delta)) = V(\Omega)$, such that the points in $N(\Gamma) \cap \text{bd}(\Delta)$ appear in $\text{bd}(\Delta)$ in the same ordering as their images, via π , in Ω .

Given an Ω -rendition (Γ, σ, π) of a graph G , we call a cell c of Γ *trivial* if $\pi(\tilde{c}) = V(\sigma(c))$.

We say that an Ω -rendition (Γ, σ, π) of a graph G is *tight* if the following conditions are satisfied:

- (i) If there are two points x, y of $N(\Gamma)$ such that $e = \{\pi(x), \pi(y)\} \in E(G)$, then there is a cell $c \in C(\Gamma)$ such that $\sigma(c)$ is the two-vertex connected graph $(e, \{e\})$,
- (ii) for every $c \in C(\Gamma)$, every two vertices in $\pi(\tilde{c})$ belong to some path of $\sigma(c)$,
- (iii) for every $c \in C(\Gamma)$ and every connected component C of the graph $\sigma(c) \setminus \pi(\tilde{c})$, if $N_{\sigma(c)}(V(C)) \neq \emptyset$, then $N_{\sigma(c)}(V(C)) = \pi(\tilde{c})$,
- (iv) there are no two distinct non-trivial cells c_1 and c_2 such that $\pi(\tilde{c}_1) = \pi(\tilde{c}_2)$, and
- (v) for every $c \in C(\Gamma)$ there are $|\tilde{c}|$ vertex-disjoint paths in G from $\pi(\tilde{c})$ to the set $V(\Omega)$.




Figure 2: A graph G together with an Ω -rendition of G , where all tightness conditions are violated.

Lemma 3. *There is a linear-time algorithm that, given a graph G and an Ω -rendition (Γ, σ, π) of G , outputs a tight Ω -rendition of G .*

In order to prove [Lemma 3](#), and in particular condition (v), we need to define the notion of *triconnected components* (see [\[34, 43\]](#)).

Given a graph G , we say that G is *triconnected* if for each $\{u, v\} \in \binom{V}{2}$, there exist three pairwise internally disjoint (u, v) -paths of G , say P_1, P_2, P_3 , such that for each $\{i, j\} \in \binom{[3]}{2}$, $P_i \neq P_j$,

$V(P_i) \cap V(P_j) = \{u, v\}$. Let G be a graph and $S \subseteq V(G)$ and let V_1, \dots, V_q be the vertex sets of the connected components of $G \setminus S$. We define $\mathcal{C}(G, S) = \{G_1, \dots, G_q\}$ where, for $i \in [q]$, G_i is the graph obtained from $G[V_i \cup S]$ if we add all edges between vertices in S . We call the members of the set $\mathcal{C}(G, S)$ *augmented connected components*. Given a vertex $x \in V(G)$ we define $\mathcal{C}(G, x) = \mathcal{C}(G, \{x\})$.

Given a graph G , the set $\mathcal{Q}(G)$ of its *triconnected components* is recursively defined as follows:

- If G is triconnected or a clique of size at most three, then $\mathcal{Q}(G) = \{G\}$.
- If G contains a separator S where $|S| \leq 2$, then $\mathcal{Q}(G) = \bigcup_{H \in \mathcal{C}(G, S)} \mathcal{Q}(H)$.

Notice that all graphs in $\mathcal{Q}(G)$ are either cliques on at most three vertices or triconnected graphs (graphs without any separator of less than three vertices).

Proof of Lemma 3. We argue about how to transform (Γ, σ, π) to a tight Ω -rendition of G in $\mathcal{O}(n + m)$ time. See Figure 2 for an example of a graph G together with an Ω -rendition of G that *violates* each of the five tightness conditions (indicated in the figure).

For the first property, let $e = \{\pi(x), \pi(y)\} \in E(G)$ be an edge of G that belongs to some $\sigma(c)$ with $|V(\sigma(c))| > 2$. Then, we add a new cell c_{new} to the rendition, where $\pi(\tilde{c}_{\text{new}}) = \{\pi(x), \pi(y)\}$ and $\sigma(c_{\text{new}}) = (e, \{e\})$. Also, we remove the edge e from $\sigma(c)$.

For the second property, let c be a cell in $C(\Gamma)$ and let \mathcal{C} be the set of connected components of the graph $\sigma(c)$ and let $\mathcal{C}_\emptyset = \{C \in \mathcal{C} \mid V(C) \cap \pi(\tilde{c}) = \emptyset\}$. Observe that, because of condition (c.3) of the definition of rendition, for every $x \in \pi(\tilde{c})$ there is some $C \in \mathcal{C} \setminus \mathcal{C}_\emptyset$ such that $x \in V(C)$ and it can happen that two vertices in $\pi(\tilde{c})$ belong to the vertex set of the same connected component in $\mathcal{C} \setminus \mathcal{C}_\emptyset$. Therefore, if $|\mathcal{C} \setminus \mathcal{C}_\emptyset| = 1$, then (ii) holds. If $|\mathcal{C} \setminus \mathcal{C}_\emptyset| > 1$, then we obtain a new Δ -painting Γ' as follows: We remove c from $U(\Gamma)$ and we replace it with a new cell c_C for each $C \in \mathcal{C} \setminus \mathcal{C}_\emptyset$ (this can be done by taking a subset U' of $U(\Gamma)$ such that $(U' \setminus N(\Gamma)) \cap c$ has exactly $|\mathcal{C} \setminus \mathcal{C}_\emptyset|$ arcwise-connected components satisfying the conditions in the third item of the definition of a painting). Then, to obtain the new rendition, we update σ so that each new cell c_C is mapped to the graph C , except one arbitrarily chosen $C' \in \mathcal{C} \setminus \mathcal{C}_\emptyset$, for which $\sigma(c_{C'})$ is the union of C' with all connected components in \mathcal{C}_\emptyset .

For the third property, consider some $c \in C(\Gamma)$. Let \mathcal{C} be the set of connected components of the graph $\sigma(c) \setminus \pi(\tilde{c})$. We say that $C_1, C_2 \in \mathcal{C}$ are *equivalent* if $N_{\sigma(c)}(V(C_1)) = N_{\sigma(c)}(V(C_2))$. Notice that this equivalence relation has at most eight equivalence classes, each corresponding to a subset of $\pi(\tilde{c})$. For each subset X of $\pi(\tilde{c})$, we define the graph F_X as the subgraph of $\sigma(c)$ induced by the union of X and the (union of the) vertex sets of the graphs in the corresponding equivalence class. We then obtain a new Δ -painting Γ' as follows: We remove c from $U(\Gamma)$ and for every non-empty $X \in 2^{\pi(\tilde{c})}$ where F_X is non-null, we add a new cell c_X (this can be done by taking a subset U' of $U(\Gamma)$ such that $(U' \setminus N(\Gamma)) \cap c$ has exactly as many arcwise-connected components as the number of different sets X considered above, and these components satisfy the conditions in the third item of the definition of a painting). Then, to obtain the new rendition, for every non-empty $X \in 2^{\pi(\tilde{c})}$ where F_X is non-null, we update σ by mapping each c_X to the graph F_X . Observe that, for every cell c in the new rendition and for every connected component C of the graph $\sigma(c) \setminus \pi(\tilde{c})$, with $N_{\sigma(c)}(V(C)) \neq \emptyset$, it holds that $N_{\sigma(c)}(V(C)) = \pi(\tilde{c})$.

For property (iv), for every two distinct non-trivial cells c_1 and c_2 with $\pi(\tilde{c}_1) = \pi(\tilde{c}_2)$, we remove c_2 from the rendition and we update $\sigma(c_1) := \sigma(c_1) \cup \sigma(c_2)$.

The last property can be achieved as follows: we first construct an auxiliary planar graph G' by substituting in G each $\sigma(c)$ by a clique on $\pi(\tilde{c})$ (that is a vertex, an edge, or a triangle) and by adding a new vertex v_{new} adjacent to all the vertices in $V(\Omega)$. We first consider a minimal set $S \in V(G')$ of size at most two such that if C^* is the augmented connected component of $G' \setminus S$ that contains v_{new} , then C^* is a triconnected component of G' . The set S can be computed in time $\mathcal{O}(n + m)$ using the classic algorithm of Hopcroft and Tarjan [20] (see also [19]). We define a new Ω -rendition (Γ', σ', π') as follows: Let A be the set of all cells c of Γ such that $E(\sigma(c)) \cap E(C^*) = \emptyset$. We obtain Γ' from Γ by removing all cells $c \in A$ and replacing them by a single cell c' . We set σ' to be the function that maps every cell c such that $E(\sigma(c)) \subset E(C^*)$ to $\sigma(c)$ and c' to $\bigcup_{c \in A} \sigma(c)$. Also, we set π' to be the function that maps, for each cell c such that $E(\sigma(c)) \subset E(C^*)$, \tilde{c} to $\pi(\tilde{c})$ and \tilde{c}' to S . Observe that, by definition of S , if for some cell c of Γ there are no $|\tilde{c}|$ vertex-disjoint paths in G from $\pi(\tilde{c})$ to the set $V(\Omega)$, then c belongs to A . Also, since C^* is the triconnected component of G' that contains v_{new} , there are $|S|$ vertex-disjoint paths in G from S to $V(\Omega)$. Therefore, in (Γ', σ', π') , property (v) holds. \square

In the rest of this paper we use only conditions (i)–(iii) of the tightness definition. However, we adopt the above, more strict, version of tightness as it will be useful in further applications.

2.3 Flatness pairs

Let W be an r -wall, for some odd integer $r \geq 3$. We say that a pair $(P, C) \subseteq D(W) \times D(W)$ is a *choice of pegs and corners for W* if W is a subdivision of an elementary r -wall \bar{W} where P and C are the pegs and the corners of \bar{W} , respectively (clearly, $C \subseteq P$). To get more intuition, notice that a wall W can occur in several ways from the elementary wall \bar{W} , depending on the way the vertices in the perimeter of \bar{W} are subdivided. Each of them gives a different selection (P, C) of pegs and corners of W .

Let an odd integer $r \geq 3$ and W be an r -wall of some graph G . We say that W is a *flat r -wall* of G if there is a separation (X, Y) of G and a choice (P, C) of pegs and corners for W such that:

- $V(W) \subseteq Y$,
- $P \subseteq X \cap Y \subseteq V(D(W))$, and
- if Ω is the cyclic ordering of the vertices $X \cap Y$ as they appear in $D(W)$, then there exists an Ω -rendition (Γ, σ, π) of $G[Y]$.

Because of Lemma 3, we can assume (and we also demand) that the Ω -rendition (Γ, σ, π) of $G[Y]$ in the above definition is always tight. We mention here that Chuzhoy [7] uses a slightly different notion of flatness, where the separation (X, Y) consists of two edge-disjoint *subgraphs*, instead of two *vertex sets*, and where the graph Y may play the role of the compass.

Flatness pairs. Given the above, we say that the choice of the 7-tuple $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$ certifies that W is a flat wall of G . We call the pair (W, \mathfrak{R}) a *flatness pair* of G and define the *height* of the pair (W, \mathfrak{R}) to be the height of W . We use the term *cell of \mathfrak{R}* in order to refer to the cells of Γ .

We call the graph $G[Y]$ the \mathfrak{R} -compass of W in G , denoted by $\text{Compass}_{\mathfrak{R}}(W)$. We define the *flaps* of the wall W in \mathfrak{R} as $\text{Flaps}_{\mathfrak{R}}(W) := \{\sigma(c) \mid c \in C(\Gamma)\}$. Given a flap $F \in \text{Flaps}_{\mathfrak{R}}(W)$, we define its *base* as $\partial F := V(F) \cap \pi(N(\Gamma))$. A flap $F \in \text{Flaps}_{\mathfrak{R}}(W)$ is *trivial* if $|\partial F| = 2$ and F consists of one edge between the two vertices in ∂F . Keep in mind that trivial flaps are only flaps that consist of a single edge between the two vertices of their base. We call the edges of the trivial flaps *short edges of $\text{Compass}_{\mathfrak{R}}(W)$* . A cell c of \mathfrak{R} is *untidy* if $\pi(\tilde{c})$ contains a vertex x of W such that two of the edges of W that are incident to x are edges of $\sigma(c)$. Notice that if c is untidy then $|\tilde{c}| = 3$. A cell c is *tidy* if it is not untidy.




Figure 3: A flat 7-wall W in a graph G whose flatness is certified by some rendition \mathfrak{R} where the choice of pegs and corners in \mathfrak{R} corresponds to the squared vertices. We depict only the \mathfrak{R} -compass of W that consists of W and some “black paths” between the vertices of W . The 5-wall \tilde{W}' consisting of the fat edges (purple, green, blue) is a flat \mathfrak{R} -normal wall of $\text{Compass}_{\mathfrak{R}}(W)$. The flatness of \tilde{W}' is certified by the rendition $\tilde{\mathfrak{R}}' = (X', Y', P', C', \Gamma', \sigma', \pi')$, where X' contains all the vertices incident to at least one orange edge plus the non-depicted vertices in the grey area, Y' contains all vertices that are either in a “fat” black path or incident to at least two fat edges, the pegs are the diamond vertices, and the corners are the fat diamond vertices (that are also pegs). For the (tight) Ω' -rendition (Γ', σ', π') of $G[Y']$, see Figure 4.

In Figure 3 we depict a flat wall W in a graph G as well as the \mathfrak{R} -compass of W in G , for some rendition \mathfrak{R} certifying its flatness. Notice that there is a unique subwall W' of W that is disjoint from $D(W)$ and has height five. Interestingly, the subwall W' is *not* a flat wall of G , however there is a tilt \tilde{W}' of W' that is a flat wall of G . The wall \tilde{W}' is depicted in Figure 3 and the rendition certifying its flatness is depicted in Figure 4.




Figure 4: The painting of the rendition $\tilde{\mathfrak{R}}'$ certifying the flatness of the 5-wall \tilde{W}' of Figure 3. The $\tilde{\mathfrak{R}}'$ -compass of \tilde{W}' has two types of flaps: those whose base has three vertices (they are images of the blue cells) and those that are trivial (they are images of the purple cells).

Cell classification. Given a cycle C of $\text{Compass}_{\mathfrak{R}}(W)$, we say that C is \mathfrak{R} -normal if it is not a subgraph of a flap $F \in \text{Flaps}_{\mathfrak{R}}(W)$. Given an \mathfrak{R} -normal cycle C of $\text{Compass}_{\mathfrak{R}}(W)$, we call a cell c of \mathfrak{R} C -perimetric if $\sigma(c)$ contains some edge of C . Notice that if c is C -perimetric, then $\pi(\tilde{c})$ contains two points $p, q \in N(\Gamma)$ such that $\pi(p)$ and $\pi(q)$ are vertices of C where one, say P_c^{in} , of the two $(\pi(p), \pi(q))$ -subpaths of C is a subgraph of $\sigma(c)$ and the other, denoted by P_c^{out} , $(\pi(p), \pi(q))$ -subpath contains at most one internal vertex of $\sigma(c)$, which should be the (unique) vertex z in $\partial\sigma(c) \setminus \{\pi(p), \pi(q)\}$. We pick a (p, q) -arc A_c in $\hat{c} := c \cup \tilde{c}$ such that $\pi^{-1}(z) \in A_c$ if and only if P_c^{in} contains the vertex z as an internal vertex.

We consider the circle $K_C = \bigcup\{A_c \mid c \text{ is a } C\text{-perimetric cell of } \mathfrak{R}\}$ and we denote by Δ_C the closed disk bounded by K_C that is contained in Δ . A cell c of \mathfrak{R} is called C -internal if $c \subseteq \Delta_C$ and is called C -external if $\Delta_C \cap c = \emptyset$. Notice that the cells of \mathfrak{R} are partitioned into C -internal, C -perimetric, and C -external cells.

Let c be a tidy C -perimetric cell of \mathfrak{R} where $|\tilde{c}| = 3$. Notice that $c \setminus A_c$ has two arcwise-connected components and one of them is an open disk D_c that is a subset of Δ_C . If the closure \overline{D}_c of D_c contains only two points of \tilde{c} then we call the cell c C -marginal.

Influence. For every \mathfrak{R} -normal cycle C of $\text{Compass}_{\mathfrak{R}}(W)$ we define the set

$$\text{Influence}_{\mathfrak{R}}(C) = \{\sigma(c) \mid c \text{ is a cell of } \mathfrak{R} \text{ that is not } C\text{-external}\}.$$

A wall W' of $\text{Compass}_{\mathfrak{R}}(W)$ is \mathfrak{R} -normal if $D(W')$ is \mathfrak{R} -normal. Notice that every wall of W (and hence every subwall of W) is an \mathfrak{R} -normal wall of $\text{Compass}_{\mathfrak{R}}(W)$. We denote by $\mathcal{S}_{\mathfrak{R}}(W)$ the set of all \mathfrak{R} -normal walls of $\text{Compass}_{\mathfrak{R}}(W)$. Given a $W' \in \mathcal{S}_{\mathfrak{R}}(W)$ and a cell c of \mathfrak{R} we say that c is W' -perimetric/internal/external/marginal if c is $D(W')$ -perimetric/internal/external/marginal. We also use $K_{W'}$, $\Delta_{W'}$, $\text{Influence}_{\mathfrak{R}}(W')$ as shortcuts for $K_{D(W')}$, $\Delta_{D(W')}$, $\text{Influence}_{\mathfrak{R}}(D(W'))$.

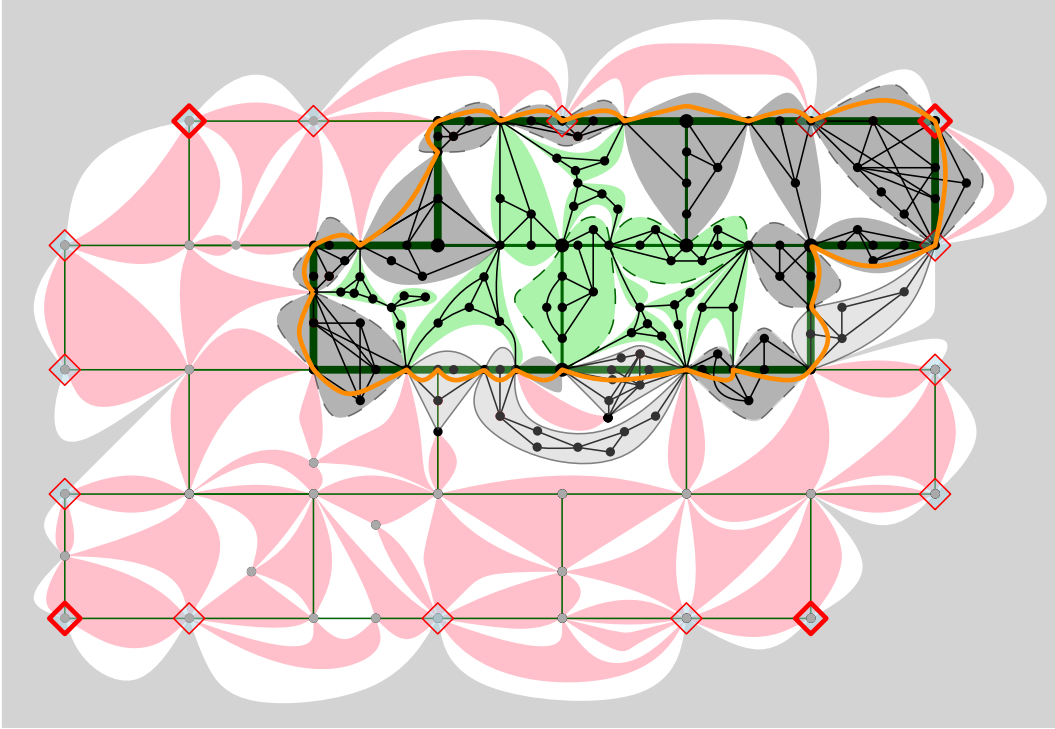


Figure 5: A flat wall W in a graph G , the painting of a rendition \mathfrak{R} certifying its flatness, a subwall W' of W , of height three, which is \mathfrak{R} -normal, and the \mathfrak{R} -flaps of W , that correspond to either W' -perimetric (depicted in grey) or W' -internal cells (depicted in green). The circle $K_{W'}$ is the fat orange cycle. The W' -marginal cells are depicted in light grey and the untidy cells are those with dashed boundary.

Regular pairs. Let (W, \mathfrak{R}) be a flatness pair of a graph G . We call a flatness pair (W, \mathfrak{R}) of a graph G *regular* if none of its cells is W -external, W -marginal, or untidy.

Tilts of flatness pairs. Let (W, \mathfrak{R}) and $(\tilde{W}', \tilde{\mathfrak{R}}')$ be two flatness pairs of a graph G and let $W' \in \mathcal{S}_{\mathfrak{R}}(W)$. We also assume that $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$ and $\tilde{\mathfrak{R}}' = (X', Y', P', C', \Gamma', \sigma', \pi')$. We say that $(\tilde{W}', \tilde{\mathfrak{R}}')$ is a W' -tilt of (W, \mathfrak{R}) if

- $\tilde{\mathfrak{R}}'$ does not have \tilde{W}' -external cells,
- \tilde{W}' is a tilt of W' ,
- the set of \tilde{W}' -internal cells of $\tilde{\mathfrak{R}}'$ is the same as the set of W' -internal cells of \mathfrak{R} and their images via σ' and σ are also the same,
- $\text{Compass}_{\tilde{\mathfrak{R}}'}(\tilde{W}')$ is a subgraph of $\mathbf{U}\text{Influence}_{\mathfrak{R}}(W')$, and
- if c is a cell in $C(\Gamma') \setminus C(\Gamma)$, then $|\tilde{c}| \leq 2$.

The next observation follows from the definitions of regular flatness pairs and tilts.

Observation 4. If (W, \mathfrak{R}) is a regular flatness pair, then for every $W' \in \mathcal{S}_{\mathfrak{R}}(W)$ every W' -tilt of (W, \mathfrak{R}) is also regular.

The main results of this paper are the following.

Theorem 5. *There exists an algorithm that given a graph G , a flatness pair (W, \mathfrak{R}) of G , and a wall $W' \in \mathcal{S}_{\mathfrak{R}}(W)$, outputs a W' -tilt of (W, \mathfrak{R}) in $\mathcal{O}(n + m)$ time.*

Theorem 6. *There is an algorithm that, given a graph G and a flatness pair (W, \mathfrak{R}) of G , outputs a regular flatness pair (W^*, \mathfrak{R}^*) of G , with the same height as (W, \mathfrak{R}) such that $\text{Compass}_{\mathfrak{R}^*}(W^*) \subseteq \text{Compass}_{\mathfrak{R}}(W)$. This algorithm runs in $\mathcal{O}(n + m)$ time.*

3 Applications

In this section we apply [Theorem 5](#) and [Theorem 6](#) in order to address the items (β) , (γ) , (δ) , and (ε) discussed in the introduction.

3.1 Tilts of subwalls

We present the following result from [\[32\]](#), stated in our new framework.

Proposition 7. *There are two functions $f_1 : \mathbb{N} \rightarrow \mathbb{N}$ and $f_2 : \mathbb{N} \rightarrow \mathbb{N}$ and an algorithm that receives as input a graph G , an odd integer $r \geq 3$, a $t \in \mathbb{N}_{\geq 1}$, and an $f_1(t) \cdot r$ -wall W in G , and outputs, in $\mathcal{O}(t^{24} \cdot m + n)$ time,*

- either that K_t is a minor of G or
- a set $A \subseteq V(G)$ where $|A| \leq f_2(t)$ and a flatness pair $(\tilde{W}', \tilde{\mathfrak{R}}')$ of $G \setminus A$ of height r , such that \tilde{W}' is a tilt of a subwall W' of W .

Moreover $f_1(t) = \mathcal{O}(t^{26})$ and $f_2(t) = \mathcal{O}(t^{24})$.

An alternative of the above where $f_1(t) = \mathcal{O}(t^2)$ and $f_2(t) = t - 5 = \mathcal{O}(t)$ has been proved by Chuzhoy in [\[7\]](#) with a running time that is polynomial in the input size. However, we prefer the version of Kawarabayashi, Thomas, and Wollan [\[32\]](#) as their algorithm is linear.

3.2 Apex-walls with compasses of bounded treewidth

We first define the notion of treewidth. A *tree decomposition* of a graph G is a pair (T, χ) where T is a tree and $\chi : V(T) \rightarrow 2^{V(G)}$ such that

1. $\bigcup_{t \in V(T)} \chi(t) = V(G)$,
2. for every edge e of G there is a $t \in V(T)$ such that $\chi(t)$ contains both endpoints of e , and
3. for every $v \in V(G)$, the subgraph of T induced by $\{t \in V(T) \mid v \in \chi(t)\}$ is connected.

The *width* of (T, χ) is defined as $w(T, \chi) := \max \{ |\chi(t)| - 1 \mid t \in V(T) \}$. The *treewidth* of G is defined as

$$\text{tw}(G) := \min \{ w(T, \chi) \mid (T, \chi) \text{ is a tree decomposition of } G \}.$$

This subsection is dedicated to the proof of the following result.

Theorem 8. *There is a function $f_3 : \mathbb{N} \rightarrow \mathbb{N}$ and an algorithm that receives as input a graph G , an odd integer $r \geq 3$, and a $t \in \mathbb{N}_{\geq 1}$, and outputs, in $2^{\mathcal{O}_t(r^2)} \cdot n$ time, one of the following:*

- a report that K_t is a minor of G ,
- a tree decomposition of G of width at most $f_3(t) \cdot r$, or
- a set $A \subseteq V(G)$, where $|A| \leq f_2(t)$, a regular flatness pair (W, \mathfrak{R}) of $G \setminus A$ of height r , and a tree decomposition of the \mathfrak{R} -compass of W of width at most $f_3(t) \cdot r$. (Here $f_2(t)$ is the function of Proposition 7 and $f_3(t) = 2^{\mathcal{O}(t^2 \log t)}$.)

Moreover, to obtain an explicit dependence on t , this algorithm can be modified to run in time $2^{2^{\mathcal{O}(t^2 \log t)} r \log r + \mathcal{O}(r^2)} \cdot n + 2^{2^{\mathcal{O}(t^2 \log t)} r^3 \log r}$.

We will need some additional results in order to prove Theorem 8. First we need the following result that is derived from [37]. For a detailed analysis of the results of [37], see [3].

Proposition 9. *There exists an algorithm with the following specifications:*

Input: A graph G and a non-negative integer k such that $|V(G)| \geq 12k^3$.

Output: A graph G^* such that $|V(G^*)| \leq (1 - \frac{1}{16k^2}) \cdot |V(G)|$ and:

- Either G^* is a subgraph of G such that $\text{tw}(G) = \text{tw}(G^*)$, or
- G^* is obtained from G after identifying the vertices of a matching in G .

Moreover, this algorithm runs in $2^{\mathcal{O}(k)} \cdot n$ time.

The following result of Kawarabayashi and Kobayashi [27], provides a linear relation between the treewidth and the height of a largest wall in a minor-free graph.

Proposition 10. *There is a function $f_4 : \mathbb{N} \rightarrow \mathbb{N}$ such that, for every $t, r \in \mathbb{N}$ and every graph G that does not contain K_t as a minor, if $\text{tw}(G) \geq f_4(t) \cdot r$, then G contains an r -wall. In particular, one may choose $f_4(t) = 2^{\mathcal{O}(t^2 \log t)}$.*

The following is the main result of [6]. We will use it to compute a tree decomposition of a graph of bounded treewidth.

Proposition 11. *There is an algorithm that, given a graph G and an integer k , outputs either a report that $\text{tw}(G) > k$, or a tree decomposition of G of width at most $5k + 4$. Moreover, this algorithm runs in $2^{\mathcal{O}(k)} \cdot n$ time.*

The following result is derived from [1]. We will use it in order to find a wall in a graph of bounded treewidth, given a tree decomposition of it.

Proposition 12. *There is an algorithm that, given a graph G , a graph H on h edges without isolated vertices, and a tree decomposition of G of width at most k , outputs, if it exists, a minor of G isomorphic to H . Moreover, this algorithm runs in $2^{\mathcal{O}(k \log k)} \cdot h^{\mathcal{O}(k)} \cdot 2^{\mathcal{O}(h)} \cdot m$ time.*

We start by proving the following “light version” of [Theorem 8](#).

Lemma 13. *There exists an algorithm as follows:*

Find-Wall(G, t, r)

Input: A graph G , an odd $r \in \mathbb{N}_{\geq 3}$, and a $t \in \mathbb{N}_{\geq 1}$.

Output: One of the following:

- a report that K_t is a minor of G ,
- a report that G has treewidth at most $f_4(t) \cdot r$, where f_4 is as in [Proposition 10](#), or
- an r -wall W of G .

Moreover, this algorithm runs in $2^{\mathcal{O}_t(r^2)} \cdot n$ time. To obtain an explicit dependence on t , this algorithm can be modified to run in time $2^{2^{\mathcal{O}(t^2 \log t)} r \log r + \mathcal{O}(r^2)} \cdot n + 2^{2^{\mathcal{O}(t^2 \log t)} r^3 \log r}$.

Proof. We set $c := f_4(t) \cdot r$. Notice that there is a constant c_t , depending on t , such that $c_t = \mathcal{O}(t\sqrt{\log t})$ and if $|E(G)| > c_t \cdot |V(G)|$, then G contains K_t as a minor [\[42\]](#). We therefore assume that $|E(G)| = \mathcal{O}(t\sqrt{\log t} \cdot n)$, otherwise we can immediately report that K_t is a minor of G and stop. We now describe a recursive algorithm as follows.

We first argue for the base case, namely when $|V(G)| < 12c^3$. To check whether K_t is a minor of G , we use the minor-containment algorithm of Robertson and Seymour [\[38\]](#), which runs in $\mathcal{O}_t(|V(G)|^3) = \mathcal{O}_t(r^3)$ time, and if this is the case, we report the same and stop. If not, then we check whether $\text{tw}(G) \leq c$, using the algorithm of Arnborg, Corneil, and Proskurowski [\[4\]](#), in time $\mathcal{O}(|V(G)|^{c+2}) = 2^{\mathcal{O}_t(r \log r)}$, and if this is the case, we report the same and stop. If not, we deal with the case where G does not contain K_t as a minor and $\text{tw}(G) > c$. By [Proposition 10](#) we know that G contains an r -wall. To find such a wall, we first consider an arbitrary ordering $(v_1, \dots, v_{|V(G)|})$ of the vertices of G . For each $i \in [|V(G)|]$, we set G_i to be the graph induced by the vertices v_1, \dots, v_i . We iteratively run the algorithm of [Proposition 11](#) on G_i and c for ascending values of i . This algorithm runs in $2^{\mathcal{O}(c)} \cdot |V(G)| = 2^{\mathcal{O}_t(r + \log r)}$ time. Let $j \in [|V(G)|]$ be the smallest integer such that the above algorithm outputs a report that $\text{tw}(G_j) > c$ and notice that there exists a tree decomposition (\mathcal{T}_j, χ_j) of G_j (obtained by the one of G_{j-1} by adding the vertex v_j in the appropriate bags) of width at most $5c + 5$. The fact that G_j does not contain K_t as a minor and $\text{tw}(G_j) > c$, implies that G_j contains an r -wall W , that is also a wall of G . To detect W , we run the algorithm of [Proposition 12](#) on G_j , W , and (\mathcal{T}_j, χ_j) . This algorithm runs in $2^{\mathcal{O}_t(r^2)} \cdot |V(G)| = 2^{\mathcal{O}_t(r^2 + \log r)}$ time. Therefore, in the case where $|V(G)| < 12c^3$, we obtain one of the three possible outputs in time $2^{\mathcal{O}_t(r^2)}$. Alternatively, to get an explicit dependence on t , instead of applying the minor-containment algorithm of Robertson and Seymour [\[38\]](#) in the beginning of the algorithm, we can do the following: first, apply the algorithm of [Proposition 11](#) on G and $12c^3$. Since $|V(G)| < 12c^3$ and therefore $\text{tw}(G) < 12c^3$, this algorithm outputs a tree decomposition (\mathcal{T}, χ) of G of width $62c^3 + 4$. Then, we apply the algorithm of [Proposition 12](#) on G , K_t , and (\mathcal{T}, χ) , to check whether K_t is a minor of G , in time $2^{2^{\mathcal{O}(t^2 \log t)} r^3 \log r}$.

If $|V(G)| \geq 12c^3$, then we call the algorithm of [Proposition 9](#) with input (G, c) , which outputs a graph G^* such that $|V(G^*)| \leq (1 - \frac{1}{16c^2}) \cdot |V(G)|$ and

- A. either G^* is a subgraph of G such that $\text{tw}(G) = \text{tw}(G^*)$, or
- B. G^* is obtained from G after identifying the vertices of a matching M of G .

In both cases, we recursively call the algorithm on G^* and we distinguish the following two cases.

Case A: G^* is a subgraph of G such that $\text{tw}(G) = \text{tw}(G^*)$. If the recursive call on G^* reports that K_t is a minor of G^* , then we report the same for G as well. If the recursive call on G^* reports that $\text{tw}(G^*) \leq c$, then we return that $\text{tw}(G) \leq c$. If it outputs an r -wall W of G^* , then we return W as a wall of G .

Case B: G^* is obtained from G after contracting the edges of a matching of G .

If the recursive call on G^* reports that $\text{tw}(G^*) \leq c$, then we do the following. We first notice that the fact that $\text{tw}(G^*) \leq c$ implies that $\text{tw}(G) \leq 2c$, since we can obtain a tree decomposition (\mathcal{T}, χ) of G from a tree decomposition (\mathcal{T}^*, χ^*) of G^* , by replacing, in every $t \in \mathcal{T}^*$, every occurrence of a vertex of G^* that is a result of an edge contraction by its endpoints in G . Thus, we can call the algorithm of [Proposition 12](#) on G , K_t , and (\mathcal{T}, χ) in order to check whether G contains K_t as a minor in $2^{\mathcal{O}_t(r \log r)} \cdot n$ steps and if this is the case, we report the same and stop (keep in mind that $c = \mathcal{O}_t(r)$). If not, then using the same algorithm we can also find in G , if it exists, an r -wall W of G in $2^{\mathcal{O}_t(r^2)} \cdot n$ time (this can be done by first finding an r -wall as a *minor* of G using [Proposition 12](#) and then, since every wall has maximum degree three, finding a subdivision of the obtained r -wall as a *subgraph* of G). If this is the case, i.e., we find an r -wall W of G , we report it and stop. In the remaining case, we can safely report, because of [Proposition 10](#), that $\text{tw}(G) \leq f_4(t) \cdot r = c$.

If the recursive call on G^* outputs an r -wall W^* of G^* , then by uncontracting the edges of M in W^* we can also return an r -wall of G . Finally, if the output is that K_t is a minor of G^* , then we return that the same holds for G .

It is easy to see that the running time of the above algorithm is

$$T(n, r, t) \leq T((1 - \frac{1}{12c^2}) \cdot n, r, t) + 2^{\mathcal{O}_t(r^2)} \cdot n,$$

where for $n < 12c^3$, $T(n, r, t) = 2^{2^{\mathcal{O}(t^2 \log t)} \cdot r \log r + \mathcal{O}(r^2)} + \mathcal{O}_t(r^3) = 2^{\mathcal{O}_t(r^2)}$. Recall that the alternative subroutine that we described above in the case where $n < 12c^3$ and we ask for an explicit dependence on t , runs in time $T(n, r, t) = 2^{2^{\mathcal{O}(t^2 \log t)} r^3 \log r}$. Therefore, we either have that $T(n, r, t) = 2^{\mathcal{O}_t(r^2)} \cdot n$ or $T(n, r, t) = 2^{2^{\mathcal{O}(t^2 \log t)} r \log r + \mathcal{O}(r^2)} \cdot n + 2^{2^{\mathcal{O}(t^2 \log t)} r^3 \log r}$, as claimed. \square

Given a flatness pair (W, \mathfrak{R}) of a graph G and a set $L \subseteq V(G)$, we say that (W, \mathfrak{R}) is L -avoiding if $L \cap V(\text{Compass}_{\mathfrak{R}}(W)) = \emptyset$. We now proceed to the proof of [Theorem 8](#).

Proof of [Theorem 8](#). Notice that there is a constant c_t , depending on t , such that $c_t = \mathcal{O}(t\sqrt{\log t})$ and if $|E(G)| > c_t \cdot |V(G)|$, then G contains K_t as a minor [\[42\]](#). We therefore assume that $|E(G)| = \mathcal{O}(t\sqrt{\log t} \cdot n)$, otherwise we can immediately report that K_t is a minor of G and stop. We first give an algorithm with the following specifications. This algorithm involves recursion assuming an

input with an additional set L that should be avoided by the desired flatness pair. For notational convenience, we define $z : \mathbb{N}^2 \rightarrow \mathbb{N}$ as $z(r, t) = 2 \cdot (\lceil \sqrt{f_2(t) + 2} \rceil + 1) \cdot f_4(t) \cdot (f_1(t) + 1) \cdot (r + 2)$.

Algorithm **Find_Low_TW_compass**(G, r, t, L).

Input: an odd $r \in \mathbb{N}_{\geq 3}$, a $t \in \mathbb{N}_{\geq 1}$, a graph G where $\text{tw}(G) > z(r, t)$, and a set $L \subseteq V(G)$ where $|L| \leq f_2(t) + 1$.

Output: either a report that K_t is a minor of G or a set $A \subseteq V(G)$, where $|A| \leq f_2(t)$, an L -avoiding flatness pair (W, \mathfrak{R}) of $G \setminus A$ of height r , and a tree decomposition of the \mathfrak{R} -compass of W of width at most $5 \cdot z(r, t) + 4$.

Step 1. We set ℓ as the smallest odd integer that is not smaller than $\sqrt{f_2(t) + 2}$. Also, let $\tilde{f}_1(t)$ be the smallest odd integer that is not smaller than $f_1(t)$. These augmentations are necessary in order to guarantee that the considered subwalls will be of odd height. We also set $r' = 2 \cdot (r + 2) + 1$. Run the algorithm of [Lemma 13](#) for G , $\ell \cdot \tilde{f}_1(t) \cdot r'$, and t . This takes time $2^{\mathcal{O}_t(r^2)} \cdot n$, or, for an explicit dependence on t , it can be modified to take time $2^{2^{\mathcal{O}(t^2 \log t)} r \log r + \mathcal{O}(r^2)} \cdot n + 2^{2^{\mathcal{O}(t^2 \log t)} r^3 \log r}$. If the output is a report that K_t is a minor of G , then return the same. Otherwise, because, $\text{tw}(G) > z(r, t) \geq \ell \cdot f_4(t) \cdot \tilde{f}_1(t) \cdot r'$, the algorithm returns an $\ell \cdot \tilde{f}_1(t) \cdot (2(r + 2) + 1)$ -wall W of G .

Step 2. Call the algorithm of [Proposition 7](#) on G , $\ell \cdot r'$, t , and W . This takes $\mathcal{O}(t^{25} \sqrt{\log t} \cdot n)$ time, since $|E(G)| = \mathcal{O}_t(n)$. If the output is a report that K_t is a minor of G , then return the same. Otherwise, we have a set $A \subseteq V(G)$, where $|A| \leq f_2(t)$, and a flatness pair $(\tilde{W}', \tilde{\mathfrak{R}}')$ of $G \setminus A$ of height $\ell \cdot r'$.

Step 3. Let W'' be a subwall of \tilde{W}' of height r' such that none of the vertices in L belongs to $\text{Influence}_{\tilde{\mathfrak{R}}'}(W'')$. The subwall W'' exists because $\ell^2 \geq f_2(t) + 2 \geq |L| + 1$ and \tilde{W}' has height $\ell \cdot r'$. We also consider four pairwise disjoint $(r + 2)$ -subwalls of W'' , namely W'_1, W'_2, W'_3 , and W'_4 , and observe that each W'_i is also a subwall of \tilde{W}' . For every $i \in [4]$, we call the algorithm of [Theorem 5](#) on $G \setminus A$, $(\tilde{W}', \tilde{\mathfrak{R}}')$, and W'_i which outputs, in $\mathcal{O}_t(n)$ time, a W'_i -tilt $(\tilde{W}'_i, \tilde{\mathfrak{R}}'_i)$ of $(\tilde{W}', \tilde{\mathfrak{R}}')$. Let K'_i be the compass of \tilde{W}'_i in $\tilde{\mathfrak{R}}'_i$. We finally fix i so that \tilde{W}'_i is a wall among W'_1, W'_2, W'_3 , and W'_4 where $|V(K'_i)|$ is minimized. Observe that $|V(K'_i)| \leq |V(G)|/4$ and that $(\tilde{W}'_i, \tilde{\mathfrak{R}}'_i)$ is L -avoiding. Indeed, since $(\tilde{W}'_i, \tilde{\mathfrak{R}}'_i)$ is a W'_i -tilt of $(\tilde{W}', \tilde{\mathfrak{R}}')$, $K'_i = \text{Compass}_{\tilde{\mathfrak{R}}'_i}(\tilde{W}'_i)$ is a subgraph of $\mathbf{UInfluence}_{\tilde{\mathfrak{R}}'}(W'_i)$ that, in turn, is a subgraph of $\mathbf{UInfluence}_{\tilde{\mathfrak{R}}'}(W'')$ and by definition of W'' , $\text{Influence}_{\tilde{\mathfrak{R}}'}(W'') \cap L = \emptyset$.

We update $W \leftarrow \tilde{W}'_i$, $\mathfrak{R} \leftarrow \tilde{\mathfrak{R}}'_i$ and we set $K = \text{Compass}_{\mathfrak{R}}(W)$. Recall that (W, \mathfrak{R}) is an L -avoiding flatness pair of $G \setminus A$ of height $r + 2$.

Step 4. We now consider the subwall W' of W obtained from $W \setminus D(W)$ after repeatedly removing vertices of degree one until no such vertices exist anymore. Notice that W' is an r -wall of $G \setminus A$. We call the algorithm of [Theorem 5](#) on $G \setminus A$, (W, \mathfrak{R}) , and W' which outputs, in $\mathcal{O}(t \sqrt{\log t} \cdot n)$ time, a W' -tilt $(\tilde{W}', \tilde{\mathfrak{R}}')$ of (W, \mathfrak{R}) . Let K' be the $\tilde{\mathfrak{R}}'$ -compass of \tilde{W}' . Clearly, $(\tilde{W}', \tilde{\mathfrak{R}}')$ is L -avoiding as well.

Step 5. Let G_D be the graph obtained from $G[V(K) \cup A]$ if we contract all the vertices of $D(W)$ to a single vertex v^* . Since $(\tilde{W}', \tilde{\mathfrak{R}}')$ is a W' -tilt of (W, \mathfrak{R}) , $K' = \text{Compass}_{\tilde{\mathfrak{R}}'}(\tilde{W}')$ is a subgraph of $\mathbf{UInfluence}_{\mathfrak{R}}(W')$, and therefore the perimeter of W and the graph K' do not have any vertex in common. This implies that K' is a subgraph of G_D .

Step 6. Call the algorithm of [Proposition 11](#) with input G_D and $z(r, t)$. This runs in $2^{2^{\mathcal{O}(t^2 \log t)} \cdot r} \cdot n$ time. If the output is a tree decomposition of G_D of width at most $5 \cdot z(r, t) + 4$, then, as K' is a

subgraph of G_D , we have that $(\tilde{W}', \tilde{\mathfrak{R}}')$ is an L -avoiding flatness pair of $G \setminus A$ of height r where the $\tilde{\mathfrak{R}}'$ -compass of \tilde{W}' has treewidth at most $5 \cdot z(r, t) + 4$. In this case, the algorithm outputs the pair $(\tilde{W}', \tilde{\mathfrak{R}}')$ and the corresponding tree decomposition of the $\tilde{\mathfrak{R}}'$ -compass K' of \tilde{W}' obtained from the one of G_D by removing the vertices in $V(G_D) \setminus V(K')$.

Step 7. Suppose now that $\text{tw}(G_D) > z(r, t)$. Notice that, by construction, if $G_D \setminus A$ has an $\{v^*\}$ -avoiding flatness pair (W^*, \mathfrak{R}^*) of height r , then (W^*, \mathfrak{R}^*) will also be an L -avoiding flatness pair of $G \setminus A$. Moreover, since G_D is a minor of G , if G_D contains K_t as a minor then also G does. Notice also that $|A \cup \{v'\}| \leq f_2(t) + 1$. Therefore, we can safely return **Find_Low_TW_compass** $(G_D, r, t, A \cup \{v'\})$. This completes the description of the algorithm and its correctness.

Notice that the running time of the above algorithm is

$$T(n, r, t) \leq T(n/4 + f_2(t), r, t) + 2^{\mathcal{O}_t(r^2)} \cdot n,$$

which implies that $T(n, r, t) = 2^{\mathcal{O}_t(r^2)} \cdot n$, and can be modified in order to obtain $T(n, r, t) = 2^{2^{\mathcal{O}(t^2 \log t)} r \log r + \mathcal{O}(r^2)} \cdot n + 2^{2^{\mathcal{O}(t^2 \log t)} r^3 \log r}$.

We define the function $f_3 : \mathbb{N} \rightarrow \mathbb{N}$ so that $f_3(t) = \min\{c \in \mathbb{N} \mid \forall r \geq 3, 5 \cdot z(r, t) + 4 \leq c \cdot r\}$. The algorithm claimed by the theorem calls first the algorithm of [Proposition 11](#) with input G and $z(r, t)$. This runs in $2^{2^{\mathcal{O}(t^2 \log t)} \cdot r} \cdot n$ time. If the output is a tree decomposition of G of width at most $5 \cdot z(r, t) + 4 \leq f_3(t) \cdot r$, then we report this and we are done. If the output is a report that $\text{tw}(G) > z(r, t)$, then we run Algorithm **Find_Low_TW_compass** (G, r, t, L) for $L = \emptyset$. This may provide either a report that K_t is a minor of G , or a set $A \subseteq V(G)$, where $|A| \leq f_2(t)$, a flatness pair (W, \mathfrak{R}) of $G \setminus A$ of height $f_2(t)$ that can be made regular by [Theorem 6](#), and a tree decomposition of the \mathfrak{R} -compass of W of width at most $5 \cdot z(r, t) + 4 \leq f_3(t) \cdot r$, and these are the possible outputs of the claimed algorithm. \square

3.3 Homogeneous walls

Palettes and homogeneity. Let $w \in \mathbb{N}$, let G be a graph, and let (W, \mathfrak{R}) be a flatness pair of G . A *flap-coloring* of (W, \mathfrak{R}) with w colors is any function $\zeta : \text{Flaps}_{\mathfrak{R}}(W) \rightarrow [w]$. For every \mathfrak{R} -normal cycle C of $\text{Compass}_{\mathfrak{R}}(W)$, we define ζ -palette(C) = $\{\zeta(F) \mid F \in \text{Influence}_{\mathfrak{R}}(C)\}$. We say that the flatness pair (W, \mathfrak{R}) of G is ζ -homogeneous if every *internal* brick of W (seen as a cycle of $\text{Compass}_{\mathfrak{R}}(W)$) has the *same* ζ -palette.

Finding a homogeneous flatness pair inside a flatness pair has a price, which is determined by the following lemma.

Lemma 14. *There is a function $f_5 : \mathbb{N}^2 \rightarrow \mathbb{N}$, whose images are odd integers, such that for every $w \in \mathbb{N}_{\geq 1}$ and every odd integer $r \geq 3$, if G is a graph, (W, \mathfrak{R}) is a flatness pair of G of height $f_5(r, w)$, and ζ is a flap-coloring of (W, \mathfrak{R}) with w colors, then W contains some subwall W' of height r such that every W' -tilt of (W, \mathfrak{R}) is ζ -homogeneous. Moreover, $f_5(r, w) = \mathcal{O}(r^w)$.*

Proof. Let $w \in \mathbb{N}$ and an odd integer $r \geq 3$. We define the function $f_5 : \mathbb{N}^2 \rightarrow \mathbb{N}$ so that, for every $x \in \mathbb{N}$, $f_5(x, 1) = x$ while, for $y \geq 2$, we set $f_5(x, y) = x \cdot (f_5(x, y - 1) - 1) + 1$. Notice that if x is odd, then $f_5(x, y)$ is also odd for every $y \in \mathbb{N}_{\geq 1}$.

Let G be a graph, (W, \mathfrak{R}) be a flatness pair of G of height $f_5(r, w)$, and ζ be a flap-coloring of (W, \mathfrak{R}) with w colors. We prove the lemma by induction on w . Clearly, if $w = 1$, then the lemma holds trivially as, in this case, for every brick B of W , $\zeta\text{-palette}(B) = \{1\}$, and therefore as W is a subwall of itself, every W -tilt of (W, \mathfrak{R}) is a flatness pair of G of height $f_5(r, 1) = r$ that is ζ -homogeneous.

Suppose now that $w \geq 2$ and that the lemma holds for smaller values of w . We set $q = f_5(r, w - 1)$. We define the subwall W' of W by taking the union of the i -th horizontal and the i -th vertical paths of W for all $i \in \{j \cdot (q - 1) + 1 \mid j \in [r]\}$. If for every brick B of W' it holds that $\zeta\text{-palette}(B) = [w]$, then consider a W' -tilt $(\tilde{W}', \tilde{\mathfrak{R}}')$ of (W, \mathfrak{R}) . The third property in the definition of a tilt of a flatness pair implies that for every internal brick \tilde{B} of \tilde{W}' there is an internal brick B of W' such that $\text{Influence}_{\mathfrak{R}}(B) = \text{Influence}_{\tilde{\mathfrak{R}}'}(\tilde{B})$. Therefore, for every internal brick \tilde{B} of \tilde{W}' , $\zeta\text{-palette}(\tilde{B}) = [w]$. Therefore, $(\tilde{W}', \tilde{\mathfrak{R}}')$ is a flatness pair of G of height r that is ζ -homogeneous. Otherwise, let \check{B} be some brick of W' such that $|\zeta\text{-palette}(\check{B})| < w$. Notice that \check{B} is the perimeter of a subwall \check{W} of W of height q . From the induction hypothesis applied to \check{W} , we have that \check{W} has a subwall W' (that is a subwall of W as well) such that every W' -tilt of (W, \mathfrak{R}) is a flatness pair of G of height r that is ζ -homogeneous. The lemma follows by observing that $f_5(r, w) = \mathcal{O}(r^w)$. \square

We now prove the main result of this subsection.

Lemma 15. *There is an algorithm that receives as input $w \in \mathbb{N}_{\geq 1}$, an odd integer $r \geq 3$, a graph G , a flatness pair (W, \mathfrak{R}) of G of height $f_5(r, w)$, and a flap-coloring ζ of (W, \mathfrak{R}) with w colors, and outputs a ζ -homogeneous flatness pair $(\check{W}, \check{\mathfrak{R}})$ of G of height r that is a W' -tilt of (W, \mathfrak{R}) for some subwall W' of W . This algorithm runs in time $2^{\mathcal{O}(wr \log r)} \cdot (n + m)$.*

Proof. Let \mathcal{W} be the collection of all r -subwalls of W . Clearly $|\mathcal{W}| = \binom{f_5(r, w)}{r} = 2^{\mathcal{O}(wr \log r)}$. For each $W' \in \mathcal{W}$, we call the algorithm of [Theorem 5](#) on G , (W, \mathfrak{R}) , and W' , which outputs, a W' -tilt $(\tilde{W}', \tilde{\mathfrak{R}}')$ of (W, \mathfrak{R}) . This algorithm runs in $\mathcal{O}(n + m)$ time. Then, for every $W' \in \mathcal{W}$, we check whether $(\tilde{W}', \tilde{\mathfrak{R}}')$ is ζ -homogeneous by computing the $\zeta\text{-palette}(\tilde{B})$ for every internal brick \tilde{B} of \tilde{W}' . This is done in linear time. [Lemma 14](#) guarantees that since the height of (W, \mathfrak{R}) is $f_5(r, w)$, W contains a subwall W' of height r such that every W' -tilt of (W, \mathfrak{R}) is ζ -homogeneous. Therefore, the above procedure will detect a flatness pair $(\check{W}', \check{\mathfrak{R}}')$ of G that is ζ -homogeneous and has height r , which we return. \square

3.4 Levelings and well-aligned flatness pairs

Let G be a graph and let (W, \mathfrak{R}) be a flatness pair of G . Let also $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$, where (Γ, σ, π) is an Ω -rendition of $G[Y]$ and $\Gamma = (U, N)$ is a Δ -painting. The *ground set* of W in \mathfrak{R} is $\text{ground}_{\mathfrak{R}}(W) := \pi(N(\Gamma))$ and we refer to the vertices of this set as the *ground vertices* of the \mathfrak{R} -compass of W in G . Notice that $\text{ground}_{\mathfrak{R}}(W)$ may contain vertices of $\text{Compass}_{\mathfrak{R}}(W)$ that are not necessarily vertices of W . For instance, in [Figure 3](#), all the ground vertices of the $\tilde{\mathfrak{R}}'$ -compass of \tilde{W}' are vertices of \tilde{W}' , while in [Figure 5](#), there are ground vertices of the \mathfrak{R} -compass of W that are not vertices of W .

We define the \mathfrak{R} -*leveling* of W in G , denoted by $W_{\mathfrak{R}}$, as the bipartite graph where one part is the ground set of W in \mathfrak{R} , the other part is a set $\text{vflaps}_{\mathfrak{R}}(W) = \{v_F \mid F \in \text{Flaps}_{\mathfrak{R}}(W)\}$ containing

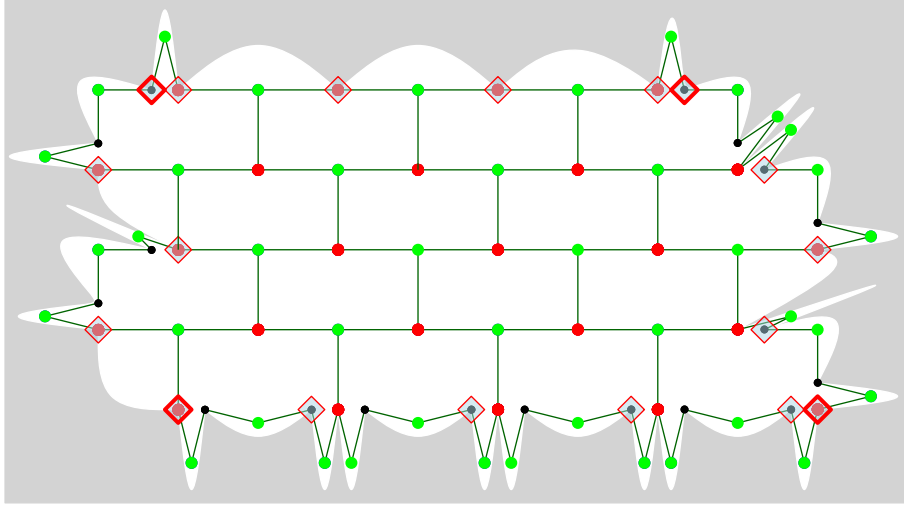


Figure 6: The $\tilde{\mathfrak{A}}'$ -leveling of the flat 5-wall \tilde{W}' of Figure 3.

one new vertex v_F for each flap F of W in \mathfrak{A} , and, given a pair $(x, F) \in \text{ground}_{\mathfrak{A}}(W) \times \text{Flaps}_{\mathfrak{A}}(W)$, the set $\{x, v_F\}$ is an edge of $W_{\mathfrak{A}}$ if and only if $x \in \partial F$. We call the vertices of $\text{ground}_{\mathfrak{A}}(W)$ (resp. $\text{vflaps}_{\mathfrak{A}}(W)$) *ground-vertices* (resp. *flap-vertices*) of $W_{\mathfrak{A}}$. Notice that the incidence graph of the plane hypergraph $(N(\Gamma), \{\tilde{c} \mid c \in C(\Gamma)\})$ is isomorphic to $W_{\mathfrak{A}}$ via an isomorphism that extends π and, moreover, bijectively corresponds cells to flap-vertices. This permits us to treat $W_{\mathfrak{A}}$ as a Δ -embedded graph where $\text{bd}(\Delta) \cap W_{\mathfrak{A}}$ is the set $X \cap Y$. As an example, see Figure 6 for the $\tilde{\mathfrak{A}}'$ -leveling of the flat 5-wall \tilde{W}' of Figure 3.

We denote by W^\bullet the graph obtained from W if we subdivide *once* every edge of W that is a short edge of $\text{Compass}_{\mathfrak{A}}(W)$. The graph W^\bullet is a “slightly richer variant” of W that is necessary for our definitions and proofs, namely to be able to associate every flap-vertex of an appropriate subgraph of $W_{\mathfrak{A}}$ (that we will denote by R_W) with a non-empty path of W^\bullet , as we proceed to formalize. We say that (W, \mathfrak{A}) is *well-aligned* if the following holds:

$W_{\mathfrak{A}}$ contains as a subgraph an r -wall R_W where $D(R_W) = D(W_{\mathfrak{A}})$ and W^\bullet is isomorphic to some subdivision of R_W via an isomorphism that maps each ground vertex to itself.

Suppose now that the flatness pair (W, \mathfrak{A}) is well-aligned. We call the wall R_W in the above condition a *representation* of W in $W_{\mathfrak{A}}$.

As an example, notice that the flatness pair $(\tilde{W}', \tilde{\mathfrak{A}}')$ of Figure 3 is well-aligned while the flatness pair (W, \mathfrak{A}) in Figure 5 is not since, for example, in the uppermost rightmost grey cell, the upper right ground vertex can not be mapped to itself in order to yield a subgraph R_W of $W_{\mathfrak{A}}$ as in the above property.

Lemma 16. *If a flatness pair (W, \mathfrak{A}) is regular, then it is also well-aligned. Moreover, there is an $\mathcal{O}(n)$ time algorithm that, given G and such a (W, \mathfrak{A}) , outputs a representation R_W of W in $W_{\mathfrak{A}}$.*

Proof. Let (W, \mathfrak{A}) be a regular flatness pair. All cells of \mathfrak{A} are tidy and there are no W -external or W -marginal cells. We also denote $\mathfrak{A} = (X, Y, P, C, \Gamma, \sigma, \pi)$. Recall that W^\bullet (whose edges are

depicted in orange in Figure 7) is the graph obtained from W if we subdivide once every short edge of $\text{Compass}_{\mathfrak{R}}(W)$ that is also an edge of W . Let ξ be the function mapping every “new” vertex of W^\bullet created by a subdivision of such a short edge of $\text{Compass}_{\mathfrak{R}}(W)$ (depicted by a cross in Figure 7) to the corresponding (trivial) flap-vertex of $W_{\mathfrak{R}}$ (that is depicted as one of the blue vertices of degree two).

Consider $R_W = (B \cup F_1 \cup F_2, E')$, where

$$\begin{aligned} B &= W \cap \text{ground}_{\mathfrak{R}}(W), \\ F_1 &= \{\xi(x) \mid x \text{ is a subdivision vertex of } W^\bullet\}, \text{ and} \\ F_2 &= \{v_F \in \text{vflaps}_{\mathfrak{R}}(W) \mid E(W \cap F) \neq \emptyset \text{ and } F \text{ is a non-trivial flap}\}. \end{aligned}$$

In Figure 7, the vertices in B are depicted in red in Figure 7 while the vertices in $F_1 \cup F_1$ are depicted in blue. Observe that $B \cup F_1 \cup F_2 \subseteq V(W_{\mathfrak{R}})$. We define E' as follows. For every $v_F \in F_1$ we include in E' both edges of $W_{\mathfrak{R}}$ that are incident to v_F . For every $v_F \in F_2$ such that $F \setminus \partial F$ contains a 3-branch vertex of W we include in E' the three edges of $W_{\mathfrak{R}}$ that are incident to v_F . Finally, for every $v_F \in F_2$ such that $F \setminus \partial F$ does not contain any 3-branch vertex of W we first consider the non-trivial path P_F in $W \cap F$ and we add in E' the edges of $W_{\mathfrak{R}}$ between the flap-vertex v_F and the endpoints of P_F . Notice that since $\sigma^{-1}(F)$ is tidy, P_F does not contain internal vertices in ∂F . Observe that R_W is indeed a wall of $W_{\mathfrak{R}}$, where $D(W_{\mathfrak{R}}) = D(R_W)$, that can be computed in $\mathcal{O}(n)$ time. We now define a mapping $\rho : V(R_W) \rightarrow V(W^\bullet)$ and a function τ mapping the edges in $E(R_W)$ (depicted as purple edges in Figure 7) to subpaths of W^\bullet as follows (an intuitive explanation follows the formal definition):

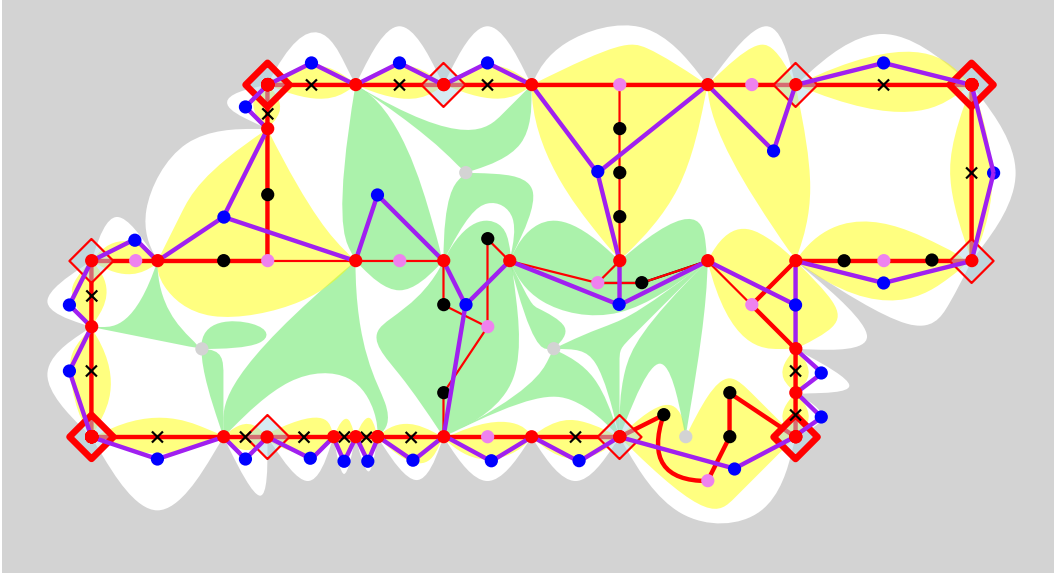


Figure 7: A well-aligned flatness pair (W, \mathfrak{R}) where W is a 3-wall, the wall W^\bullet (whose edges are depicted in red and the new subdivision vertices are depicted by small crosses), the leveling $W_{\mathfrak{R}}$ of W (whose edges are depicted in purple), and the subgraph R_W of $W_{\mathfrak{R}}$ (depicted by fat purple edges).

- If $x \in B$, then $\rho(x) = x$.
- If $v_F \in F_1$ and $\partial F = \{x, y\}$, then we set $\rho(v_F) = \xi^{-1}(v_F)$, $\tau(\{x, v_F\}) = \{x, \xi^{-1}(v_F)\}$, and $\tau(\{y, v_F\}) = \{y, \xi^{-1}(v_F)\}$.
- If $v_F \in F_2$ and v_F is a branch vertex of R_W , then assume first that $\partial F = \{x, y, z\}$. Because the cell $\sigma^{-1}(F)$ is tidy the graph $F \setminus \partial F$ contains a unique 3-branch vertex w of W (or equivalently of W^\bullet) and $F \cap W^\bullet$ consists of three internally disjoint paths $P_{w,x}$, $P_{w,y}$, and $P_{w,z}$ in F from w to x , y , and z , respectively. We set $\rho(v_F) = w$, $\tau(\{x, v_F\}) = P_{w,x}$, $\tau(\{y, v_F\}) = P_{w,y}$, and $\tau(\{z, v_F\}) = P_{w,z}$.
- If $v_F \in F_2$ and v_F is not a 3-branch vertex of R_W , then there exist two vertices x, y of R_W such that $N_{R_W}(v_F) = \{x, y\}$. Pick an internal vertex w of the (x, y) -path P_F and set $\rho(v_F) = w$ (recall that, as $\sigma^{-1}(F)$ is tidy, none of the internal vertices of the path P_F is a ground vertex). If $P_{w,x}$ is the (w, x) -subpath of P_F , and $P_{w,y}$ is the (w, y) -subpath of P_F , then set $\tau(\{x, v_F\}) = P_{w,x}$ and $\tau(\{y, v_F\}) = P_{w,y}$.

Intuitively, ρ maps each vertex of R_W to a vertex in W^\bullet in the following way: if $x \in B$, then, since x is already a vertex in W and therefore in its subdivision W^\bullet , x is mapped to itself. If $x \in F_1$, then $x = \xi(x)$, for some subdivision vertex of W^\bullet (recall that W^\bullet is obtained from W by subdividing once all short edges), and ρ maps x to x . Also, the function τ , which maps edges of R_W to paths in W^\bullet , maps the (two) edges adjacent to x to the (two) corresponding edges in W^\bullet that are adjacent to x . Finally, if $v_F \in F_2$, then v_F is a flap-vertex of $W_{\mathfrak{R}}$ such that the corresponding flap F is “traversed” by some path of the wall W and is non-trivial, i.e., contains non-boundary vertices. We assume that $\partial F = \{x, y, z\}$ and we consider two cases. First, if v_F is a 3-branch vertex of R_W , then tidiness of the cells implies that F contains a unique 3-branch vertex w of W , to which v_F is mapped via ρ . Then, the function τ maps the three edges of R_W between v_F and the boundary vertices x, y, z (that belong to B) to the corresponding (internally disjoint) paths of W from w to the boundary vertices x, y, z of F . Otherwise, i.e., if v_F is *not* a 3-branch vertex of R_W , then the intersection of W^\bullet and F is a single path connecting two boundary vertices of F , say x and y . We now choose an arbitrary internal vertex w of this path, we set $\rho(v_F) = w$ and we map, via τ , the (two) edges of R_W incident to v_F to the (two internally disjoint) paths connecting w to x and y in W^\bullet .

Consider a subdivision H of R_W where each edge $e \in E(R_W)$ is replaced with a path P_e of length $|\tau(e)|$. We extend ρ by mapping, via ρ , each vertex of P_e to the corresponding vertex in the corresponding path in W^\bullet . Then, the definition of the mappings ρ and τ above implies that this subdivision H of R_W is isomorphic to W^\bullet (see [Figure 7](#) for a visualization).

As all members of $B = W \cap \text{ground}_{\mathfrak{R}}(W)$ are, by definition, fixed points of ρ , then (W, \mathfrak{R}) is well-aligned. \square

4 Proofs of [Theorem 5](#) and [Theorem 6](#)

This section is devoted to the proofs of [Theorem 5](#) and [Theorem 6](#). We first present some definitions in [Subsection 4.1](#) and [Subsection 4.2](#), necessary for the proof of the main technical lemma of this paper, namely [Lemma 17](#), presented in [Subsection 4.3](#).

4.1 Stretchings

Let F be a graph and x and y be two distinct vertices belonging to the same connected component of F . We say that a sequence $\langle F_1, \dots, F_r \rangle$ of subgraphs of F is a *stretching of F along the pair (x, y)* if there is a shortest (x, y) -path P_F in F such that the sequence $\langle F_1, \dots, F_r \rangle$ consists of the (unique) minimum-sized collection of subpaths of P_F with the following properties:

- for every $i \in [r]$, F_i is a path where all internal vertices have degree two in F ,
- for every $i, j \in [r]$ such that $i \neq j$, F_i and F_j have no common edges,
- $F_1 \cup \dots \cup F_r = P_F$,
- for every $\{i, j\} \in \binom{[r]}{2}$, $F_i \cap F_j \neq \emptyset$ if and only if $|i - j| = 1$, and
- $x \in V(F_1)$ and $y \in V(F_r)$.

For an example of a stretching of a graph F along a pair (x, y) , see [Figure 8](#).




Figure 8: The stretching of a graph F along the pair (x, y) .

4.2 Classifying perimetric cells

Let G be a graph and let (W, \mathfrak{A}) be a flatness pair of G , where $\mathfrak{A} = (X, Y, P, C, \Gamma, \sigma, \pi)$. Let $W' \in \mathcal{S}_{\mathfrak{A}}(W)$. We now further refine the classification of the cells of \mathfrak{A} that we gave in [Subsection 2.3](#) with respect to W' . See [Figure 9](#) for an illustration of the ways a W' -perimetric cell c of Γ may intersect $\Delta_{W'}$. The simplest case when $|\tilde{c}| = 2$, depicted in the leftmost configuration of the figure. The remaining configurations correspond to the case where $\partial\sigma(c) = \{x, y, z\}$ where A_c is a $(\pi^{-1}(x), \pi^{-1}(y))$ -arc (see [Subsection 2.3](#) for the definition of the paths P_c^{in} and P_c^{out} , the arc A_c , and the vertex z). The second/fifth, third/sixth, and forth/seventh configurations correspond to the case where z is an internal vertex of P_c^{in} , P_c^{out} , or none of them, respectively. This permits a further classification of the W' -perimetric cells of Γ as follows. A cell c of Γ is *W' -inner-perimetric* (resp. *W' -outer-perimetric*) if $c \cap \Delta_{W'}$ is situated in c as indicated in the left (resp. right) part of [Figure 9](#).

We denote the set of cells of Γ that are W' -inner-perimetric, W' -outer-perimetric, W' -internal, and W' -strictly external by $C_{W'}^{\text{ip}}(\Gamma)$, $C_{W'}^{\text{op}}(\Gamma)$, $C_{W'}^{\text{in}}(\Gamma)$, and $C_{W'}^{\text{ex}}(\Gamma)$, respectively. See [Figure 10](#) for an example of this further classification (relatively to [Figure 5](#)). Notice that all W' -marginal cells of Γ are W' -outer-perimetric cells (corresponding to the last two cases of [Figure 9](#)).




Figure 9: Seven ways $\Delta_{W'}$ may traverse a cell. The arc A_c is depicted in orange.

4.3 The main lemma

Lemma 17. *There is an algorithm that, given a graph G , a flatness pair (W, \mathfrak{R}) , where $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$, and a wall $W' \in \mathcal{S}_{\mathfrak{R}}(W)$, outputs, in $\mathcal{O}(n + m)$ time, a flatness pair $(\tilde{W}', \tilde{\mathfrak{R}}')$ where $\tilde{\mathfrak{R}}' = (X', Y', P', C', \Gamma', \sigma', \pi')$ such that*

1. *all cells of $\tilde{\mathfrak{R}}'$ are \tilde{W}' -internal or \tilde{W}' -inner-perimetric,*
2. *\tilde{W}' is a tilt of W' ,*
3. *$\sigma'|_{C_{\tilde{W}'}^{\text{in}}(\Gamma')} = \sigma|_{C_{W'}^{\text{in}}(\Gamma)}$, i.e., the set of \tilde{W}' -internal cells of $\tilde{\mathfrak{R}}'$ is the same as the set of W' -internal cells of \mathfrak{R} and their images via σ' and σ are also the same, and*
4. *$\text{Compass}_{\tilde{\mathfrak{R}}'}(\tilde{W}')$ is a subgraph of $\mathbf{U}\text{Influence}_{\mathfrak{R}}(W')$.*

Moreover, if all W' -internal or W' -inner-perimetric cells of \mathfrak{R} are tidy, then the flatness pair $(\tilde{W}', \tilde{\mathfrak{R}}')$ is regular.

Proof. Since $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$ is a 7-tuple certifying that W is flat in G , we have that the triple (Γ, σ, π) is an Ω -rendition of $G[Y]$, where $\Gamma = (U, N)$ is a Δ -painting.

We define a series of ingredients that will permit us to define an alternative 7-tuple $\tilde{\mathfrak{R}}'$. We first deal with W' -inner-perimetric cells. For every W' -inner-perimetric cell $c \in C_{W'}^{\text{ip}}(\Gamma)$, we assume that $\pi(\tilde{c}) = \{x_c, y_c, z_c\}$ and that the two endpoints of the non-trivial path of $D(W') \cap \sigma(c)$ (by *non-trivial* we refer to the path that has distinct endpoints). Then, we consider an $(\pi^{-1}(x_c), \pi^{-1}(y_c))$ -arc Y_c of Δ that intersects the cell c only in the points $\pi^{-1}(x_c)$ and $\pi^{-1}(y_c)$ (see Figure 11, where Y_c is depicted in red). Also, we set $F_1^c = \sigma(c)$, $r_c = 1$, and $V_{\text{mid}}^c = \pi(\tilde{c}) \cap V(D(W'))$ (the vertices in V_{mid}^c are depicted in orange in Figure 11).

Next, we consider a W' -outer-perimetric cell $c \in C_{W'}^{\text{op}}(\Gamma)$. We assume that $\pi(\tilde{c}) = \{x, y, z\}$ and that x and y are the two endpoints of the non-trivial path of $D(W') \cap \sigma(c)$. We also define $V_{W'}^c$ as the set of all internal vertices of this path that are different from z . Let $\langle F_1^c, \dots, F_{r_c}^c \rangle$ be the stretching of $\sigma(c)$ along the pair (x, y) and let v_i , for $i \in [r_c - 1]$, be the common endpoint of F_i^c and F_{i+1}^c . Notice that by tightness property (i), $r_c \geq 2$. This permits us to set up a *special vertex* $v^c = v_1$. We also set

$$V_{\text{mid}}^c = \{x, v_1, \dots, v_{r_c-1}, y\}, \quad V_{\text{in}}^c = \bigcup \{V(F_i^c) \mid i \in [r_c]\} \setminus V_{\text{mid}}^c.$$

Let $p_0 = \pi^{-1}(x)$, $p_{r_c} = \pi^{-1}(y)$, and create a collection c_1, \dots, c_{r_c} of open disks in c and a set p_1, \dots, p_{r_c-1} of points in c such that

- $p_0 \in \text{bd}(c_1)$ and $p_{r_c} \in \text{bd}(c_{r_c})$, $p_0 \neq p_1$, and $p_{r_c} \neq p_{r_c-1}$,




Figure 10: A flat wall W in a graph G , the painting of a rendition \mathfrak{R} certifying its flatness, a subwall W' of W , of height three, which is \mathfrak{R} -normal, and the \mathfrak{R} -flaps of W , corresponding to the cells of \mathfrak{R} that are not W' -external. The edges and the non-boundary vertices of the flaps corresponding to the W' -external cells of \mathfrak{R} (depicted in pink) are not depicted (however their boundary vertices that are not in $D(W')$ are depicted in grey). There are nine W' -outer-perimetric cells of \mathfrak{R} (in blue) and seven W' -inner-perimetric cells (in yellow). Also, there are thirteen W' -internal cells of \mathfrak{R} (in green). Among the W' -inner-perimetric and W' -internal cells of \mathfrak{R} , those that are untidy are depicted with a dashed boundary. The orange cycle is the circle $K_{W'}$.

- for $i \in [r_c - 1]$, $\bar{c}_i \cap \bar{c}_{i+1} = \{p_i\}$, and
- for every $\{i, j\} \in \binom{[r_c]}{2}$, $\bar{c}_i \cap \bar{c}_j \neq \emptyset$ if and only if $|i - j| = 1$.

We define *the cell replacement* of c as the set $\mathbf{c}\text{-repl}(c) = \{c_1, \dots, c_{r_c}\}$, *the point replacement* of c as the set $\mathbf{p}\text{-repl}(c) = \{p_0, \dots, p_{r_c}\}$, and we set $C_{\text{new}}^c = \mathbf{U}\mathbf{c}\text{-repl}(c)$ and $N_{\text{new}}^c = \mathbf{U}\mathbf{p}\text{-repl}(c)$.

We also define the arc Y_c as an arc of c where $p_i \in Y_c$, $i \in [0, r_c]$, such that p_0, p_{r_c} are the extreme points of Y_c , and Y_c is traversing \bar{c} in a way that removing Y_c from c leaves $\pi^{-1}(z)$ and c_1, \dots, c_{r_c} in different arcwise-connected components of $c \setminus Y_c$ (as depicted by the red line in Figure 12). Observe that $\mathbf{U}\{Y_c \mid c \in C_{W'}^{\text{ip}}(\Gamma) \cup C_{W'}^{\text{op}}(\Gamma)\}$ is a “red” cycle of Δ . Let Δ' be the disk bounded by this cycle for which $\Delta' \subseteq \Delta$.




Figure 11: The four cases of the definition of the arc Y_c (depicted in red), for W' -inner-perimetric cells. The boundary of $\Delta_{W'}$ is depicted in orange and the boundary of Δ is depicted in purple.




Figure 12: The definition of the replacement sequence c_1, \dots, c_{r_c} and the arc Y_c for the three cases of W' -external cells of $C_{W'}^{\text{op}}(\Gamma)$.

We set

$$\begin{aligned}
 H &= \bigcup \{F_1^c \cup \dots \cup F_{r_c}^c \mid c \in C_{W'}^{\text{op}}(\Gamma)\}, & V_{W'} &= \bigcup \{V_{W'}^c \mid c \in C_{W'}^{\text{op}}(\Gamma)\}, \\
 V_{\text{mid}} &= \bigcup \{V_{\text{mid}}^c \mid c \in C_{W'}^{\text{ip}}(\Gamma) \cup C_{W'}^{\text{op}}(\Gamma)\}, & V_{\text{in}} &= \bigcup \{V_{\text{in}}^c \mid c \in C_{W'}^{\text{op}}(\Gamma)\}, \\
 N_{\text{new}} &= \bigcup \{N_{\text{new}}^c \mid c \in C_{W'}^{\text{op}}(\Gamma)\}, & U_{\text{new}} &= \bigcup \{C_{\text{new}}^c \cup N_{\text{new}}^c \mid c \in C_{W'}^{\text{op}}(\Gamma)\}.
 \end{aligned}$$

We now define the wall $\tilde{W}' = (W' \setminus V_{W'}) \cup H$, i.e., we extract from W' the internal vertices of the subpaths of W' that are intersected by images, via σ , of W' -outer-perimetric cells and we substitute them by the paths of their stretchings. Clearly this does not affect the interior of W' , and therefore \tilde{W}' is a tilt of W' , yielding Property 2 of the statement of the lemma. Next we define a separation (X', Y') of G so that

$$Y' = \bigcup \{V(\sigma(c)) \mid c \in C_{W'}^{\text{ip}}(\Gamma) \cup C_{W'}^{\text{in}}(\Gamma)\} \cup V_{\text{in}} \cup V_{\text{mid}}, \quad X' = (V(G) \setminus Y') \cup V_{\text{mid}}.$$

In other words, Y' consists of the images of the internal cells and the vertices of every path F_i^c , while X' consists of everything else, except from V_{mid} (that is, the set $X' \cap Y'$). Notice that

$$G[Y'] \text{ is a subgraph of } \bigcup \{\sigma(c) \mid c \in C_{W'}^{\text{in}}(\Gamma) \cup C_{W'}^{\text{ip}}(\Gamma) \cup C_{W'}^{\text{op}}(\Gamma)\} = \text{Influence}_{\mathfrak{A}}(W'). \quad (1)$$

We define the pair (P', C') as follows. Let c be a W' -outer-perimetric cell and $\sigma(c) \cap V(D(W'))$ contain a vertex w such that either w is a 3-branch vertex of W' or $w \in P$ (resp. $w \in C$). We distinguish two cases. If $w \in Y'$, then we include w in P' (resp. C'). If $w \notin Y'$, then we include the special vertex v^c in P' (resp. C').

We next define an Ω' -rendition (Γ', σ', π') of $G[Y']$ where $\Gamma' = (U', N')$ is a Δ' -painting. For this we set $\Gamma' = (U', N')$, where

$$U' = ((U \setminus \bigcup C_{\tilde{W}'}^{\text{op}}(\Gamma)) \cap \Delta') \cup U_{\text{new}} \quad \text{and} \quad N' = (N \cap \Delta') \cup N_{\text{new}}.$$

Let now K' be the set of the connected components of $U' \setminus N'$, which will form the cells of the new Ω' -rendition (Γ', σ', π') . We define the function σ' mapping the cells in C' to subgraphs of $G[Y']$ as follows. Notice that $c \in K' \cap C(\Gamma)$ if and only if $c \in C_{\tilde{W}'}^{\text{in}}(\Gamma) \cap C_{\tilde{W}'}^{\text{ip}}(\Gamma)$, and in this case we set $\sigma'(c) = \sigma(c)$. Suppose now that $c \in K' \setminus C(\Gamma)$. Then c should be one of the cells, say c_i , of $\mathbf{c}\text{-repl}(c^*) = \{c_1, \dots, c_{r_c}\}$ for some $c^* \in C_{\tilde{W}'}^{\text{op}}(\Gamma)$, and in this case we set $\sigma(c) = F_i^{c^*}$. It now remains to define $\pi' : N' \rightarrow Y'$. Similarly to the definition of σ' , we consider a $p' \in N'$ and if $p \in N \cap N'$ we set $\pi'(p) = \pi(p)$. Suppose now that $p \in N' \setminus N$. Then p should be one of the points, say p_i , of $\mathbf{p}\text{-repl}(c^*) = \{p_0, \dots, p_{r_c}\}$ for some $c^* \in C_{\tilde{W}'}^{\text{op}}(\Gamma)$ and such that $i \in [r_{c^*} - 1]$. In this case we define $\pi'(p)$ to be the unique common vertex of $F_i^{c^*}$ and $F_{i+1}^{c^*}$. It is now easy to verify that (Γ', σ', π') is a tight Ω' -rendition of $G[Y']$ and that the 7-tuple $\tilde{\mathfrak{R}}' := (X', Y', P', C', \Gamma', \sigma', \pi')$ certifies that \tilde{W}' is flat in G (see Figure 13). Moreover $K' = C(\Gamma')$.




Figure 13: The flatness pair $(\tilde{W}', \tilde{\mathfrak{R}}')$ created in the proof of Lemma 17. The wall \tilde{W}' is the tilt of W' where the updated part of \tilde{W}' correspond to the red paths in Figure 10 whose edges are drawn in the orange cells.

Recall now that all the cells in $C(\Gamma') \cap C(\Gamma)$ are either \tilde{W}' -inner-perimetric or \tilde{W}' -internal. Moreover, all the cells in $C(\Gamma') \setminus C(\Gamma)$ are cells as in the left part of Figure 9, therefore they are \tilde{W}' -inner-perimetric. This yields Property 1 in the statement of the lemma. Notice also that Property 3 follows directly from the definition of σ' , as it concerns the W' -internal cells of \mathfrak{R} , and these cells are the same as the \tilde{W}' -internal cells of $\tilde{\mathfrak{R}}'$. Finally, recall that $\text{Compass}_{\tilde{\mathfrak{R}}'}(\tilde{W}') = G[Y']$ and Property 4 follows because of (1).

On the other hand, notice that all \tilde{W}' -internal cells of $\tilde{\mathfrak{R}}'$ are also W -internal cells of \mathfrak{R} . Moreover, if a \tilde{W}' -inner-perimetric cell c of $\tilde{\mathfrak{R}}'$ is a cell of \mathfrak{R} , then c is either an W -inner-perimetric or

an W -internal cell of \mathfrak{R} . On the other hand, all \tilde{W}' -inner perimetric cells of $\tilde{\mathfrak{R}}'$ that are not cells of \mathfrak{R} are cells as in the left part of [Figure 9](#), therefore they are \tilde{W}' -inner-perimetric and tidy. We conclude that if all W' -internal or W' -inner-perimetric cells of \mathfrak{R} are tidy, then all cells of $\tilde{\mathfrak{R}}'$ are tidy as well. As $\tilde{\mathfrak{R}}'$ does not have any \tilde{W}' -outer-perimetric cells it also does not have \tilde{W}' -marginal cells. These two facts along with the fact that $\tilde{\mathfrak{R}}'$ does not have any \tilde{W}' -external cells imply that the flatness pair $(\tilde{W}', \tilde{\mathfrak{R}}')$ is regular.

The running time follows from the fact that the substitution of W' -outer-perimetric cells is based on the stretching operation on the corresponding flaps, and this requires the computation of shortest paths that, in total, takes $\mathcal{O}(n + m)$ time. \square

Lemma 18. *There is an algorithm that, given a graph G and a flatness pair (W, \mathfrak{R}) , outputs, in $\mathcal{O}(n + m)$ time, a flatness pair (W^*, \mathfrak{R}^*) of G with the same height as (W, \mathfrak{R}) , with $\mathfrak{R}^* = \mathfrak{R}$, and such that all the W^* -internal or W^* -inner-perimetric cells of \mathfrak{R}^* are tidy.*




Figure 14: An illustration of the proof of [Lemma 18](#), based on the flatness pair of [Figure 13](#). The new flatness pair is (W^*, \mathfrak{R}^*) where W^* is depicted in red and $\mathfrak{R}^* = \mathfrak{R}$.

Proof. Given a wall W and an $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$ as above, we denote by $C_W^{\text{utd}}(\Gamma)$ the set of all the W -internal or W -inner-perimetric cells of Γ that are untidy. Notice that for every $c \in C_W^{\text{utd}}(\Gamma)$, $|\pi(\tilde{c})| = 3$. In what follows, we explain how to update W , while leaving $(X, Y, P, C, \Gamma, \sigma, \pi)$ intact, in order to reduce $|C_W^{\text{utd}}(\Gamma)|$ by one. Repeating this procedure clearly yields the statement claimed in the lemma.

Let $c \in C_W^{\text{utd}}(\Gamma)$. We assume that $\pi(\tilde{c}) = \{x, y, z\}$ and that $z \in \pi(\tilde{c}) \cap V(W)$ is a vertex of W such that two of the edges of W incident to z are edges of $\sigma(c)$. This implies that $\bar{P} = W \cap \sigma(c)$ is an (x, y) -path containing z as an internal vertex. Moreover, none of the internal vertices of \bar{P} , except from z , is a 3-branch vertex of W . By tightness properties (i), (ii), and (iii), there is a vertex

$w \in \sigma(c) \setminus \pi(\tilde{c})$ and three internally vertex-disjoint paths P'_x , P'_y , and P'_z in $\sigma(c)$ such that P'_x is a (w, x) -path, P'_y is a (w, y) -path, and P'_z is a (w, z) -path. If z is a 3-branch vertex of W we update $W := (W \setminus V(\bar{P} \setminus \{x, y, z\})) \cup P'_x \cup P'_y \cup P'_z$ (see bottom yellow cell with dashed boundary in Figure 14 for an example), while, if not, we update $W := (W \setminus V(\bar{P} \setminus \{x, y\})) \cup P'_x \cup P'_y$ (see the leftmost green cell with dashed boundary in Figure 14 for an example) and observe that W is again a flat wall of G , certified by $(X, Y, P, C, \Gamma, \sigma, \pi)$. Moreover, in the first case, z is no longer a 3-branch vertex of W and is incident to only one edge of $\sigma(c) \cap W$, while, in the second case, z is no longer a vertex of W . This implies that c is tidy and $|C_W^{\text{utd}}(\Gamma)|$ is indeed reduced by one (see Figure 14 for an example). As for each cell c that we modify we need to identify the paths P'_x , P'_y , and P'_z in $\sigma(c)$, the construction of W' takes, in total, $\mathcal{O}(n + m)$ time. \square

4.4 Proofs of Theorem 5 and Theorem 6

We finally have all the ingredients to prove our two main results.

Proof of Theorem 5. Let (W, \mathfrak{R}) be a flatness pair of a graph G , where $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$ and $W' \in \mathcal{S}_{\mathfrak{R}}(W)$. We call the algorithm of Lemma 17 on G , (W, \mathfrak{R}) , and W' , which outputs, in $\mathcal{O}(n + m)$ time, a flatness pair $(\tilde{W}', \tilde{\mathfrak{R}}')$ where $\tilde{\mathfrak{R}}' = (X', Y', P', C', \Gamma', \sigma', \pi')$ such that all cells of $\tilde{\mathfrak{R}}'$ are \tilde{W}' -internal or \tilde{W}' -inner-perimetric (hence $\tilde{\mathfrak{R}}'$ does not have \tilde{W}' -external cells), \tilde{W}' is a tilt of W' , the set of \tilde{W}' -internal cells of $\tilde{\mathfrak{R}}'$ is the same as the set of W' -internal cells of \mathfrak{R} and their images via σ' and σ are also the same, and $\text{Compass}_{\tilde{\mathfrak{R}}'}(\tilde{W}')$ is a subgraph of $\mathbf{UInfluence}_{\mathfrak{R}}(W')$. We observe that $(\tilde{W}', \tilde{\mathfrak{R}}')$ is a W' -tilt of (W, \mathfrak{R}) and thus we return $(\tilde{W}', \tilde{\mathfrak{R}}')$. Notice that in the case where (W, \mathfrak{R}) is regular, all cells of \mathfrak{R} are tidy. Thus, by Lemma 17, $(\tilde{W}', \tilde{\mathfrak{R}}')$ is also regular. \square

Proof of Theorem 6. Given a flatness pair (W, \mathfrak{R}) of a graph G , we first apply Lemma 18 to (W, \mathfrak{R}) and obtain in time $\mathcal{O}(n + m)$ a flatness pair $(\hat{W}^*, \hat{\mathfrak{R}}^*)$ of G with the same height as (W, \mathfrak{R}) , with $\hat{\mathfrak{R}}^* = \mathfrak{R}$, and such that all \hat{W}^* -internal or \hat{W}^* -inner-perimetric cells of $\hat{\mathfrak{R}}^*$ are tidy.

We now apply Lemma 17 with input G , $(\hat{W}^*, \hat{\mathfrak{R}}^*)$, and \hat{W}^* and obtain, in $\mathcal{O}(n + m)$ time, a flatness pair (W^*, \mathfrak{R}^*) of G such that, if $\hat{\mathfrak{R}}^* = (\hat{X}, \hat{Y}, \hat{P}, \hat{C}, \hat{\Gamma}, \hat{\sigma}, \hat{\pi})$ and $\mathfrak{R}^* = (X, Y, P, C, \Gamma, \sigma, \pi)$, we have that all cells of \mathfrak{R}^* are W^* -internal or W^* -inner-perimetric (hence \mathfrak{R}^* does not have W^* -external cells), W^* is a tilt of \hat{W}^* , the set of W^* -internal cells of \hat{W}^* is the same as the set of \hat{W}^* -internal cells of $\hat{\mathfrak{R}}^*$ and their images via σ and $\hat{\sigma}$ are also the same, and $\text{Compass}_{\mathfrak{R}^*}(W^*)$ is a subgraph of $\mathbf{UInfluence}_{\hat{\mathfrak{R}}^*}(\hat{W}^*)$. Moreover, since all the \hat{W}^* -internal or \hat{W}^* -inner-perimetric cells of $\hat{\mathfrak{R}}^*$ are tidy, Lemma 17 implies that all $(W^*$ -internal or W^* -inner-perimetric) cells of \mathfrak{R}^* are tidy. Also, since none of the cells of \mathfrak{R}^* is W^* -outer-perimetric, none of the cells of \mathfrak{R}^* is W^* -marginal. These two facts together with the fact that none of the cells of \mathfrak{R}^* is W^* -external imply that (W^*, \mathfrak{R}^*) is a regular flatness pair of G with the same height as (W, \mathfrak{R}) , as required.

We now prove that $\text{Compass}_{\mathfrak{R}^*}(W^*) \subseteq \text{Compass}_{\mathfrak{R}}(W)$. First, keep in mind that $\text{Compass}_{\mathfrak{R}^*}(W^*) \subseteq \mathbf{UInfluence}_{\hat{\mathfrak{R}}^*}(\hat{W}^*)$. We observe that $\mathbf{UInfluence}_{\hat{\mathfrak{R}}^*}(\hat{W}^*) \subseteq \text{Compass}_{\hat{\mathfrak{R}}^*}(\hat{W}^*)$ and, since $\hat{\mathfrak{R}}^* = \mathfrak{R}$, $\text{Compass}_{\hat{\mathfrak{R}}^*}(\hat{W}^*) = \text{Compass}_{\mathfrak{R}}(W)$. Therefore, $\text{Compass}_{\mathfrak{R}^*}(W^*) \subseteq \text{Compass}_{\mathfrak{R}}(W)$.

Finally, the claimed running time follows from Lemma 17 and Lemma 18. \square

References

- [1] Isolde Adler, Frederic Dorn, Fedor V. Fomin, Ignasi Sau, and Dimitrios M. Thilikos. Faster parameterized algorithms for minor containment. *Theoretical Computer Science*, 412(50):7018–7028, 2011. doi:10.1016/j.tcs.2011.09.015. 17
- [2] Isolde Adler, Martin Grohe, and Stephan Kreutzer. Computing excluded minors. In *Proc. of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 641–650, 2008. URL: <http://portal.acm.org/citation.cfm?id=1347082.1347153>. 3
- [3] Ernst Althaus and Sarah Ziegler. Optimal Tree Decompositions Revisited: A Simpler Linear-Time FPT Algorithm. *CoRR*, abs/1912.09144, 2019. arXiv:1912.09144. 17
- [4] Stefan Arnborg, Derek G. Corneil, and Andrzej Proskurowski. Complexity of finding embeddings in a k -tree. *SIAM Journal on Algebraic Discrete Methods*, 8(2):277–284, 1987. doi:10.1137/0608024. 18
- [5] Julien Baste, Ignasi Sau, and Dimitrios M. Thilikos. A complexity dichotomy for hitting connected minors on bounded treewidth graphs: the chair and the banner draw the boundary. In *Proc. of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 951–970, 2020. doi:10.1137/1.9781611975994.57. 6
- [6] Hans L. Bodlaender, Pål Grønås Drange, Markus S. Dregi, Fedor V. Fomin, Daniel Lokshtanov, and Michal Pilipczuk. A $c^k n$ 5-approximation algorithm for treewidth. *SIAM Journal on Computing*, 45(2):317–378, 2016. doi:10.1137/130947374. 17
- [7] Julia Chuzhoy. Improved Bounds for the Flat Wall Theorem. In *Proc. of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 256–275, 2015. doi:10.1137/1.9781611973730.20. 4, 5, 12, 16
- [8] 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. 3
- [9] Anuj Dawar, Martin Grohe, and Stephan Kreutzer. Locally excluding a minor. In *Proc. of the 22nd IEEE Symposium on Logic in Computer Science (LICS)*, pages 270–279. IEEE Computer Society, 2007. doi:10.1109/LICS.2007.31. 3
- [10] Erik D. Demaine and Mohammad Taghi Hajiaghayi. Linearity of grid minors in treewidth with applications through bidimensionality. *Combinatorica*, 28(1):19–36, 2008. doi:10.1007/s00493-008-2140-4. 4
- [11] Fedor V. Fomin, Petr A. Golovach, and Dimitrios M. Thilikos. Modification to planarity is fixed parameter tractable. In *Proc. of the 36th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 126 of *LIPICs*, pages 28:1–28:17, 2019. doi:10.4230/LIPICs.STACS.2019.28. 6

- [12] Fedor V. Fomin, Daniel Lokshtanov, Fahad Panolan, Saket Saurabh, and Meirav Zehavi. Hitting topological minors is FPT. In *Proc. of the 52nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 1317–1326, 2020. doi:[10.1145/3357713.3384318](https://doi.org/10.1145/3357713.3384318). 3, 6
- [13] Archontia C. Giannopoulou, Ken-ichi Kawarabayashi, Stephan Kreutzer, and O-joung Kwon. The Directed Flat Wall Theorem. In *Proc. of the 31st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 239–258, 2020. doi:[10.1137/1.9781611975994.15](https://doi.org/10.1137/1.9781611975994.15). 3
- [14] Archontia C. Giannopoulou and Dimitrios M. Thilikos. Optimizing the graph minors weak structure theorem. *SIAM Journal on Discrete Mathematics*, 27(3):1209–1227, 2013. doi:[10.1137/110857027](https://doi.org/10.1137/110857027). 4
- [15] Petr A. Golovach, Marcin Kaminski, Spyridon Maniatis, and Dimitrios M. Thilikos. The parameterized complexity of graph cyclability. *SIAM Journal on Discrete Mathematics*, 31(1):511–541, 2017. doi:[10.1137/141000014](https://doi.org/10.1137/141000014). 6
- [16] Petr A. Golovach, Marcin Kaminski, Daniël Paulusma, and Dimitrios M. Thilikos. Induced packing of odd cycles in planar graphs. *Theoretical Computer Science*, 420:28–35, 2012. doi:[10.1016/j.tcs.2011.11.004](https://doi.org/10.1016/j.tcs.2011.11.004). 6
- [17] Martin Grohe, Ken-ichi Kawarabayashi, Dániel Marx, and Paul Wollan. Finding topological subgraphs is fixed-parameter tractable. In *Proc. of the 43rd ACM Symposium on Theory of Computing (STOC)*, pages 479–488. ACM, 2011. doi:[10.1145/1993636.1993700](https://doi.org/10.1145/1993636.1993700). 3, 6
- [18] Martin Grohe, Ken-ichi Kawarabayashi, Dániel Marx, and Paul Wollan. Finding topological subgraphs is fixed-parameter tractable. *CoRR*, abs/1011.1827, 2010. arXiv:[1011.1827](https://arxiv.org/abs/1011.1827). 6
- [19] Carsten Gutwenger and Petra Mutzel. A linear time implementation of spqr-trees. In Joe Marks, editor, *Graph Drawing*, pages 77–90, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg. doi:[10.1007/3-540-44541-2_8](https://doi.org/10.1007/3-540-44541-2_8). 12
- [20] John E. Hopcroft and Robert Endre Tarjan. Dividing a graph into triconnected components. *SIAM Journal on Computing*, 2(3):135–158, 1973. doi:[10.1137/0202012](https://doi.org/10.1137/0202012). 12
- [21] Tony Huynh, Felix Joos, and Paul Wollan. A Unified Erdős-Pósa Theorem for Constrained Cycles. *Combinatorica*, 39(1):91–133, 2019. doi:[10.1007/s00493-017-3683-z](https://doi.org/10.1007/s00493-017-3683-z). 3
- [22] Marcin Kaminski and Naomi Nishimura. Finding an induced path of given parity in planar graphs in polynomial time. In *Proc. of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 656–670, 2012. doi:[10.1137/1.9781611973099.55](https://doi.org/10.1137/1.9781611973099.55). 6
- [23] Marcin Kamiński and Dimitrios M. Thilikos. Contraction checking in graphs on surfaces. In *Proc. of the 29th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 14 of *LIPICs*, pages 182–193, 2012. doi:[10.4230/LIPICs.STACS.2012.182](https://doi.org/10.4230/LIPICs.STACS.2012.182). 6
- [24] Ken-ichi Kawarabayashi. Planarity allowing few error vertices in linear time. In *Proc. of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 639–648, 2009. doi:[10.1109/FOCS.2009.45](https://doi.org/10.1109/FOCS.2009.45). 6

- [25] Ken-ichi Kawarabayashi and Yusuke Kobayashi. Algorithms for finding an induced cycle in planar graphs. *Combinatorica*, 30(6):715–734, 2010. doi:[10.1007/s00493-010-2499-x](https://doi.org/10.1007/s00493-010-2499-x). 6
- [26] Ken-ichi Kawarabayashi and Yusuke Kobayashi. The edge disjoint paths problem in Eulerian graphs and 4-edge-connected graphs. In *Proc. of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 345–353, 2010. doi:[10.1137/1.9781611973075.29](https://doi.org/10.1137/1.9781611973075.29). 3
- [27] Ken-ichi Kawarabayashi and Yusuke Kobayashi. Linear min-max relation between the treewidth of an H -minor-free graph and its largest grid minor. *Journal of Combinatorial Theory, Series B*, 141:165–180, 2020. doi:[10.1016/j.jctb.2019.07.007](https://doi.org/10.1016/j.jctb.2019.07.007). 4, 5, 17
- [28] Ken-ichi Kawarabayashi, Yusuke Kobayashi, and Bruce Reed. The disjoint paths problem in quadratic time. *Journal of Combinatorial Theory, Series B*, 2011. doi:[10.1016/j.jctb.2011.07.004](https://doi.org/10.1016/j.jctb.2011.07.004). 3, 4
- [29] Ken-ichi Kawarabayashi, Stephan Kreutzer, and Bojan Mohar. Linkless and flat embeddings in 3-space. *Discrete & Computational Geometry*, 47(4):731–755, 2012. doi:[10.1007/s00454-012-9413-9](https://doi.org/10.1007/s00454-012-9413-9). 6
- [30] Ken-ichi Kawarabayashi, Zhentao Li, and Bruce A. Reed. Recognizing a Totally Odd K_4 -subdivision, Parity 2-disjoint Rooted Paths and a Parity Cycle Through Specified Elements. In *Proc. of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 318–328, 2010. doi:[10.1137/1.9781611973075.27](https://doi.org/10.1137/1.9781611973075.27). 6
- [31] Ken-ichi Kawarabayashi, Bojan Mohar, and Bruce A. Reed. A simpler linear time algorithm for embedding graphs into an arbitrary surface and the genus of graphs of bounded tree-width. In *Proc. of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 771–780, 2008. doi:[10.1109/FOCS.2008.53](https://doi.org/10.1109/FOCS.2008.53). 6
- [32] Ken-ichi Kawarabayashi, Robin Thomas, and Paul Wollan. A new proof of the flat wall theorem. *Journal of Combinatorial Theory, Series B*, 129:204–238, 2018. doi:[10.1016/j.jctb.2017.09.006](https://doi.org/10.1016/j.jctb.2017.09.006). 3, 4, 5, 16
- [33] Ken-ichi Kawarabayashi, Robin Thomas, and Paul Wollan. Quickly excluding a non-planar graph. *CoRR*, abs/2010.12397, 2021. arXiv:[2010.12397](https://arxiv.org/abs/2010.12397). 4
- [34] Saunders Mac Lane. A structural characterization of planar combinatorial graphs. *Duke Math. J.*, 3(3):460–472, 09 1937. URL: <http://dx.doi.org/10.1215/S0012-7094-37-00336-3>, doi:[10.1215/S0012-7094-37-00336-3](https://doi.org/10.1215/S0012-7094-37-00336-3). 10
- [35] Dániel Marx and Ildikó Schlotter. Obtaining a planar graph by vertex deletion. *Algorithmica*, 62(3-4):807–822, 2012. doi:[10.1007/s00453-010-9484-z](https://doi.org/10.1007/s00453-010-9484-z). 6
- [36] Bojan Mohar. A linear time algorithm for embedding graphs in an arbitrary surface. *SIAM Journal on Discrete Mathematics*, 12(1):6–26, 1999. doi:[10.1137/S089548019529248X](https://doi.org/10.1137/S089548019529248X). 6
- [37] Ljubomir Perkovic and Bruce A. Reed. An improved algorithm for finding tree decompositions of small width. *International Journal of Foundations of Computer Science*, 11(3):365–371, 2000. doi:[10.1142/S0129054100000247](https://doi.org/10.1142/S0129054100000247). 17

- [38] Neil Robertson and Paul D. Seymour. Graph Minors. XIII. The Disjoint Paths Problem. *Journal of Combinatorial Theory, Series B*, 63(1):65–110, 1995. doi:10.1006/jctb.1995.1006. 3, 4, 6, 18
- [39] Neil Robertson and Paul D. Seymour. Graph Minors. XXI. Graphs with unique linkages. *Journal of Combinatorial Theory, Series B*, 99(3):583–616, 2009. doi:10.1016/j.jctb.2008.08.003. 6
- [40] Neil Robertson and Paul D. Seymour. Graph Minors. XXII. Irrelevant vertices in linkage problems. *Journal of Combinatorial Theory, Series B*, 102(2):530–563, 2012. doi:10.1016/j.jctb.2007.12.007. 6
- [41] Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. An FPT-Algorithm for Recognizing k -Apices of Minor-Closed Graph Classes. In *Proc. of the 47th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 168 of *LIPICs*, pages 95:1–95:20, 2020. doi:10.4230/LIPICs.ICALP.2020.95. 5, 6
- [42] Andrew Thomason. The extremal function for complete minors. *Journal of Combinatorial Theory, Series B*, 81(2):318–338, 2001. doi:10.1006/jctb.2000.2013. 18, 19
- [43] William Thomas Tutte. *Connectivity in graphs*. University of Toronto Press, 1966. 10