OPTIMIZING THE GRAPH MINORS WEAK STRUCTURE THEOREM

ARCHONTIA C. GIANNOPOULOU[†][‡] AND DIMITRIOS M. THILIKOS[†][§]

Abstract. One of the major results of [*N. Robertson and P. D. Seymour. Graph minors. XIII. The disjoint paths problem. J. Combin. Theory Ser. B, 63(1):65–110, 1995*], also known as *the weak structure theorem*, reveals the local structure of graphs excluding some graph as a minor: each such graph G either has small treewidth or contains the subdivision of a planar graph (a wall) that can be arranged in a flat manner inside G, given that some small set of vertices is removed. We prove an optimized version of that theorem where (i) the relation between the treewidth of the graph and the height of the wall is linear (thus best possible) and (ii) the number of vertices to be removed is minimized.

Key words. Graph minors, Treewidth

1. Introduction.

The Graph Minors series of Robertson and Seymour appeared to be a rich source of structural results in graph theory with multiple applications in algorithms. One of the most celebrated outcomes of this project was the existence of an $O(n^3)$ step algorithm for solving problems such as the Disjoint Path and the Minor Containment. A basic ingredient of this algorithm is a theorem, proved in paper XIII of the series [12], revealing the local structure of graphs excluding some graph as a minor. This result, now called the *weak structure theorem*, asserts that there is a function $f : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ such that for every integer k, every h-vertex graph H, and every graph G, one of the following holds:

1. G contains H as a minor,
2. G has treewidth at most $f(k, h)$, or
3. G contains a set X of at most $\binom{h}{2}$ vertices (called *apices*) such that $G \setminus X$ contains as a subgraph the subdivision W of a wall of height k that is arranged inside G in a "flat" manner (*flatness condition*).

To make the above statement precise we need to clarify the flatness condition in the third statement above. We postpone this complicated task until Section 3 and instead, we roughly visualize W in a way that the part of $G \setminus X$ that is located inside the perimeter P of W can be seen as a set of graphs attached on a plane region where each of these graphs has bounded treewidth and its boundary with the other graphs is bounded by 3.

The algorithmic applications of the weak structure theorem reside in the fact that the graph inside P can be seen as a *bidimensional structure* where, for several combinatorial problems, a solution certificate can be revised so that it avoids the middle vertex of the subdivided wall W. This is known as the *irrelevant vertex technique* and can be seen as a reduction of an instance of a problem to an equivalent one where this "irrelevant vertex" has been deleted. The application of this technique has now gone much further than its original use in the Graph Minors series and has evolved to a standard tool in algorithmic graph minors theory (see [1, 2, 6, 7, 8, 9] for applications of this technique).

[†] Department of Mathematics, National and Kapodistrian University of Athens, Greece.

[‡] Supported by a grant of the Special Account for Research Grants of the National and Kapodistrian University of Athens (project code: 70/4/10311).

[§] Co-financed by the European Union (European Social Fund - ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) - Research Funding Program: "Thales. Investing in knowledge society through the European Social Fund".

In this paper we prove an optimized version of the weak structure theorem. Our improvement is twofold: first, the function f is now linear on k and second, the number of the apices is bounded by $h-5$. Both our optimizations are optimal as indicated by the graph J obtained by taking a $(k \times k)$-grid (for $k \geq 3$) and making all its vertices adjacent with a copy of K_{h-5}. Indeed, it is easy to verify that J excludes $H = K_h$ as a minor, its treewidth is $k + h - 5$ and becomes planar (here, this is equivalent to the flatness condition) after the removal of exactly $h-5$ vertices.

Our proof deviates significantly from the one in [12]. It builds on the (strong) *structure theorem* of the Graph Minors that was proven in paper XVI of the series [13]. This theorem reveals the global structure of a graph without a K_h as a minor and asserts that each such graph can be obtained by gluing together graphs that can "almost be embedded" in a surface where K_h cannot be embedded (see 2 for the exact statement). The proof exploits this structural result and combines it with the fact, proved in [5], that apex-free "almost embedded graphs" without a $(k \times k)$-grid have treewidth $O(k)$.

The organization of the paper is the following. In Section 2 we give the definitions of all the tools that we are going to use in our proof, including the Graph Minors structure theorem. The definition of the flatness condition is given in Section 3, together with the statement of our main result. Some lemmata concerning the invariance of the flatness property under certain local transformations are given in Section 4.1 and further definitions and results concerning apex vertices are given in Section 4.2. Finally, the proof of our main result is presented in Section 5.

2. Preliminaries. Let $n \in \mathbb{N}$. We denote by $[n]$ the set $\{1, 2, \ldots, n\}$. Moreover, for every $k \leq n$, if S is a set such that $|S| = n$ we say that a set $S' \subseteq S$ is a k-subset of S if $|S'| = k$.

A *graph* G is a pair (V, E) where V is a finite set, called the *vertex set* and denoted by $V(G)$, and E is a set of 2-subsets of V, called the *edge set* and denoted by $E(G)$. We denote by $\mathbf{n}(G)$ the number of vertices of G, that is, $\mathbf{n}(G) = |V(G)|$. If we allow E to be a subset of $\mathcal{P}(V)$, where by $\mathcal{P}(V)$ we denote the powerset of V, then we call the pair $H = (V(H), E(H))$ a *hypergraph.* The *incidence graph* of a hypergraph H is the bipartite graph $I(H)$ on the vertex set $V(H) \cup E(H)$ where $v \in V(H)$ is adjacent to $e \in E(H)$ if and only if $v \in e$, that is, v is incident to e in H. We say that a hypergraph H is *planar* if its incidence graph is planar. Unless otherwise stated, we consider finite undirected graphs without loops or multiple edges.

Let G be a graph. For a vertex v, we denote by $N_G(v)$ its *(open) neighborhood*, that is, the set of vertices which are adjacent to v. The *closed neighborhood* $N_G[v]$ of v is the set $N_G(v) \cup \{v\}$. For $U \subseteq V(G)$, we define $N_G[U] = \bigcup_{v \in U} N_G[v]$. We may omit the index if the graph under consideration is clear from the context. Given a vertex $v \in V(G)$, let $\deg_G(v)$ denote the degree of v in G, that is, $\deg_G(v) = |N_G(v)|$. We denote by $\Delta(G)$ the maximum degree over all vertices of G. If $U \subseteq V(G)$ (resp. $u \in V(G)$ or $E \subseteq E(G)$ or $e \in E(G)$) then $G - U$ (resp. $G - u$ or $G - E$ or $G - e$) is the graph obtained from G by removing of vertices of U (resp. of vertex u or edges of E or of the edge e). We say that a graph H is a subgraph of a graph G, denoted by $H \subseteq G$, if H can be obtained from G after deleting edges and vertices. Let $\mathcal{C}$ be a finite class of graphs and S be a set of vertices. We denote by $\bigcup\mathcal{C}$ the graph $\cup_{G \in \mathcal{C}} G$ (where the union is not disjoint) and by $\mathcal{C} \setminus S = \{G \setminus S \mid G \in \mathcal{C}\}$.

The complete graph on n vertices is denoted by K_n. Moreover, if S is a finite set, we denote by $K[S]$ the complete graph with vertex set S. Let G be a graph such that $K_3 \subseteq G$ and x, y, z be the vertices of K_3. The *ΔY-transformation of K_3 in G* is the following: We remove the edges $\{x, y\}, \{y, z\}, \{x, z\}$, add a new vertex w and then add the edges $\{x, w\}, \{y, w\}, \{z, w\}$.

Let G be a graph. We say that G is an apex graph if there exists a vertex $v \in V(G)$ such that $G \setminus v$ is planar. Moreover, we say that G is an $\alpha_{_G}$-apex graph if there exists an $S \subseteq V(G)$ such that $|S| \leq \alpha_{_G}$ and $G \setminus S$ is planar. We denote by $\mathbf{an}(G)$, the minimum $k \in \mathbb{N}$ such that G is a k-apex graph, that is, $\mathbf{an}(G) = \min\{k \in \mathbb{N} \mid \exists S \subseteq V(G) : (|S| \leq k \wedge G \setminus S \text{ is planar})\}$. Clearly, $\mathcal{G} = \{G \mid \mathbf{an}(G) = 1\}$ is the class of the apex graphs.

Let $S \subseteq V(G)$. We say that S is a separator of G if $G \setminus S$ contains strictly more connected components than G.

Observation 1. *Let T be a tree, $k \in \mathbb{N}$ and $w : V(T) \to \mathbb{N}$ such that there exists at least one vertex $v \in V(T)$ with $w(v) \geq k$. There exists a vertex $u \in V(T)$ with $w(u) \geq k$ such that at most one of the connected components of $T \setminus u$ contains a vertex u' with $w(u') > k$.*

Proof. Let $Y = \{v \in V(T) \mid w(v) \geq k\}$. Pick a vertex r of T and let v be a vertex of Y with maximum distance away from r. It is easy to verify that the lemma holds for v. $\square$

Surfaces. A *surface* Σ is a compact 2-manifold without boundary (we always consider connected surfaces). Whenever we refer to a Σ-*embedded graph* G we consider a 2-cell embedding of G in Σ. To simplify notations, we do not distinguish between a vertex of G and the point of Σ used in the drawing to represent the vertex or between an edge and the arc representing it. We also consider a graph G embedded in Σ as the union of the points corresponding to its vertices and edges. That way, a subgraph H of G can be seen as a graph H, where the points corresponding to H are a subset of the points corresponding to G. Recall that $\Delta \subseteq \Sigma$ is an open (resp. closed) disc if it is homeomorphic to $\{(x, y) : x^2 + y^2 < 1\}$ (resp. $\{(x, y) : x^2 + y^2 \leq 1\}$). The *Euler genus* of a non-orientable surface Σ is equal to the non-orientable genus $\tilde{g}(\Sigma)$ (or the crosscap number). The *Euler genus* of an orientable surface Σ is $2g(\Sigma)$, where $g(\Sigma)$ is the orientable genus of Σ. We refer to the book of Mohar and Thomassen [10] for more details on graphs embeddings. The *Euler genus* of a graph G (denoted by $\mathbf{eg}(G)$) is the minimum integer γ such that G can be embedded on a surface of the Euler genus γ.

Let C_1, C_2 be two disjoint cycles in a plane graph G. Let also Δ_i be the open disk of $\mathbb{S}_0 \setminus C_i$ that does not contain points of C_{3-i}, $i \in [2]$. The *annulus between* C_1 and C_2 is the set $\mathbb{S}_0 \setminus (\Delta_1 \cup \Delta_2)$ and we denote it by $A[C_1, C_2]$.

Contractions and minors. Given an edge $e = \{x, y\}$ of a graph G, the graph G/e is obtained from G by contracting the edge e, that is, the endpoints x and y are replaced by a new vertex v_{xy} which is adjacent to the old neighbors of x and y (except x and y). A graph H obtained by a sequence of edge-contractions is said to be a *contraction* of G. An alternative, and more useful for our purposes, definition of a contraction is the following.

Let G and H be graphs and let $\phi : V(G) \to V(H)$ be a surjective mapping such that

1. for every vertex $v \in V(H)$, its preimage $\phi^{-1}(v)$ induces a connected graph $G[\phi^{-1}(v)]$;
2. for every edge $\{v, u\} \in E(H)$, the graph $G[\phi^{-1}(v) \cup \phi^{-1}(u)]$ is connected;
3. for every $\{v, u\} \in E(G)$, either $\phi(v) = \phi(u)$, or $\{\phi(v), \phi(u)\} \in E(H)$.

We, then, say that H *is a contraction of* G *via* ϕ and denote it as $H \leq_c^{\phi} G$. Let us observe that H is a contraction of G if $H \leq_c^{\phi} G$ for some $\phi : V(G) \to V(H)$. In this case we simply write $H \leq_c G$. If $H \leq_c^{\phi} G$ and $v \in V(H)$, then we call the preimage $\phi^{-1}(v)$ the *model of* v in G.

Let $G = G_0 \cup G^+$, where G_0 is a graph embedded in a surface Σ of Euler genus γ and let G^+ be another graph that might share common vertices with G_0. Let also H be a graph and $v \in V(H)$. We say that G *contains* H *as a* v*-smooth contraction* if $H \leq_c^{\phi} G$ for some $\phi : V(G) \to V(H)$ and there exists a closed disk D in Σ such that the vertices of G^+ are not inside of D and all the vertices of G that are not inside the disk D are exactly the model of v, that is, $\phi^{-1}(v) = V(G) \setminus (V(G) \cap D)$.

A graph H is a *minor* of a graph G, denoted by $H \leq_m G$, if H is the contraction of some subgraph of G. If we restrict the contraction to edges whose one endpoint has degree exactly two, also called *disolving* that vertex, then we say that H is a *topological minor* of G and we denote it by $H \leq_{tm} G$. Moreover, we say that a graph G is a *subdivision* of a graph H, if H can be obtained from G by disolving vertices.

We say that a graph G is *H-minor-free* when it does not contain H as a minor. We also say that a graph class $\mathcal{G}$ is *H-minor-free* (or, excludes H as a minor) when all its members are H-minor-free.

We would like to state here the following folklore result on minors and topological minors. (See, for example, Proposition 1.7.2 in [3].)

PROPOSITION 2.1. *Let G and H be graphs such that $H \leq_m G$. If $\Delta(H) \leq 3$, then $H \leq_{tm} G$.*

Graph Minors structure theorem. The proof of our result is using the Excluded Minor Theorem from the Graph Minors. Before we state it, we need some definitions.

DEFINITION 2.2 (*h-nearly embeddable graphs*). *Let Σ be a surface and $h > 0$ be an integer. A graph G is h-nearly embeddable in Σ if there is a set of vertices $X \subseteq V(G)$ (called apices) of size at most h such that graph $G - X$ is the union of (possibly empty) subgraphs $G_0, \ldots, G_h$ with the following properties:*

i) There is a set of cycles $C_1, \ldots, C_h$ in Σ such that the cycles C_i are the borders of open pairwise disjoint discs Δ_i in Σ;

ii) G_0 has an embedding in Σ in such a way that $G_0 \cap \bigcup_{i=1,\ldots,h} \Delta_i = \emptyset$;

iii) graphs $G_1, \ldots, G_h$ (called vortices) are pairwise disjoint and for $1 \leq i \leq h$, $V(G_0) \cap V(G_i) \subset C_i$;

iv) for $1 \leq i \leq h$, let $U_i := \{u_1^i, \ldots, u_{m_i}^i\}$ be the vertices of $V(G_0) \cap V(G_i) \subset C_i$ appearing in an order obtained by clockwise traversing of C_i. We call vertices of U_i bases *of G_i. Then G_i has a path decomposition $\mathcal{B}_i = (B_j^i)_{1 \leq j \leq m_i}$, of width at most h such that for $1 \leq j \leq m_i$, we have $u_j^i \in B_j^i$.*

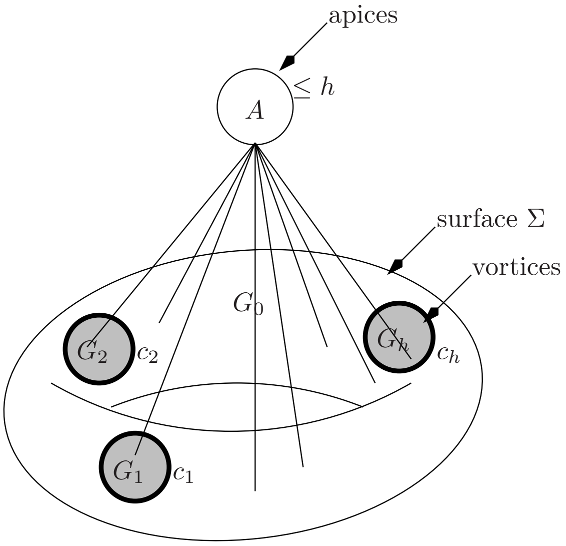

FIGURE 2.1. *An h-nearly embeddable graph*

A *tree decomposition* of a graph G is a pair $(\mathcal{X}, T)$ where T is a tree and $\mathcal{X} = \{X_i \mid i \in V(T)\}$ is a collection of subsets of $V(G)$ such that:

1. $\bigcup_{i \in V(T)} X_i = V(G)$;
2. for each edge $\{x, y\} \in E(G)$, $\{x, y\} \subseteq X_i$ for some $i \in V(T)$, and
3. for each $x \in V(G)$ the set $\{i \mid x \in X_i\}$ induces a connected subtree of T.

The *width* of a tree decomposition $(\{X_i \mid i \in V(T)\}, T)$ is $\max_{i \in V(T)} \{|X_i| - 1\}$. The *treewidth* of a graph G is the minimum width over all tree decompositions of G. Furthermore, we call the subsets X_i *bags* of the decomposition and for every $X_i, i \in V(T)$, we denote by $\overline{X}_i$ its closure, that is, $\overline{X}_i$ is the graph $G[X_i] \cup \left(\cup_{j \in N_T(i)} K[X_i \cap X_j]\right)$. (We would like to mention here that the graph $\overline{X}_i$ is also referred as *torso* at node i.)

Observation 2. *Let G be a graph and $(\mathcal{X}, T)$ be a tree decomposition of G. Then there exists an $X \in \mathcal{X}$ such that $\mathbf{tw}(\overline{X}) \geq \mathbf{tw}(G)$.*

We also need the following simple result.

Lemma 2.3. *If G is a graph and $X \subseteq V(G)$, then $\mathbf{tw}(G - X) \geq \mathbf{tw}(G) - |X|$.*

We say that a tree decomposition $(\mathcal{X}, T)$ of a graph G is *small* if for every $i, j \in V(T)$, $i \neq j$, $X_i \nsubseteq X_j$.

A simple proof of the following lemma can be found in [4].

Lemma 2.4. *The following two statements hold:*

1. *Let G be a graph and $(\mathcal{X}, T)$ be a small tree decomposition of G. Then $|V(T)| \leq |V(G)|$.*
2. *Every graph G has a small tree decomposition of width $\mathbf{tw}(G)$.*

The following proposition is known as the Graph Minors structure theorem [13]. (For an example[1] of an H-minor-free graph, see Figure 2.2)

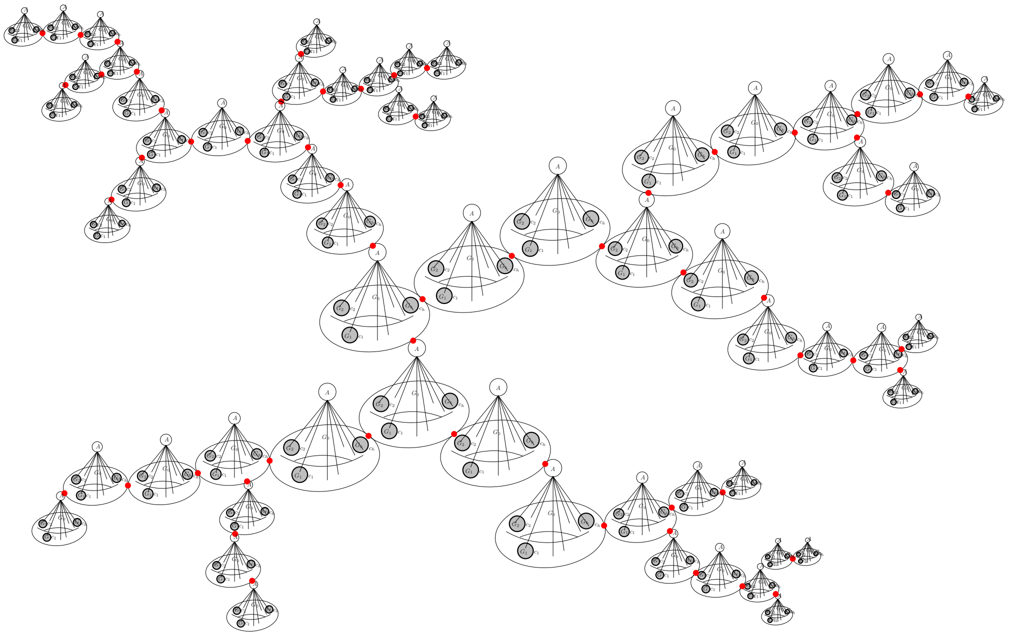

Figure 2.2. *An example of an H-minor-free graph*

---

[1] As the figures only aim to facilitate intuition the surface Σ is depicted as a torus. Note, however, that we work on general surfaces of bounded genus.

Proposition 2.5. *There exists a computable function* $f : \mathbb{N} \to \mathbb{N}$ *such that, for every non-planar graph* H *with* h *vertices and every graph* G *excluding* H *as a minor there exists a tree decomposition* $(\mathcal{G} = \{G_i \mid i \in V(T)\}, T)$ *where for every* $i \in V(T)$, $\overline{G}_i$ *is an* $f(h)$*-nearly embeddable graph in a surface* Σ *on which* H *cannot be embedded.*

3. Statement of the main result. Let k and r be positive integers where $k, r \geq 2$. The $(k \times r)$-*grid* is the Cartesian product of two paths of lengths $k-1$ and $r-1$ respectively. A vertex of a $(k \times r)$-grid is a *corner* if it has degree 2. Thus each $(k \times r)$-grid has 4 corners. A vertex of a $(k \times r)$-grid is called *internal* if it has degree 4, otherwise it is called *external*.

We define Γ_k as the following (unique, up to isomorphism) triangulation of the $(k \times k)$-grid. Let Γ be a plane embedding of the $(k \times k)$-grid such that all external vertices are on the boundary of the external face. We triangulate internal faces of the $(k \times k)$-grid such that, in the obtained graph, all the internal vertices have degree 6 and all non-corner external vertices have degree 4. The construction of Γ_k is complete when we make adjacent one corner of degree two to all vertices of the external face (we call this corner *loaded*). For an example, see Γ_6 in Figure 3.1. We also use notation Γ_k^* for the graph obtained from Γ_k if we remove all edges incident to its loaded vertex that do not exist in its underlying grid.

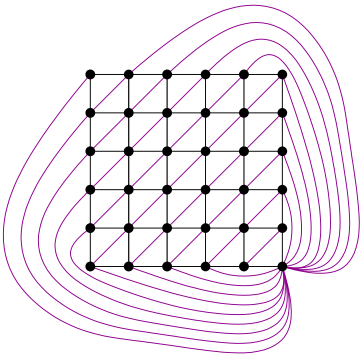

Figure 3.1. *The graph* Γ_6

We define the (k, l)-pyramid to be the graph obtained by taking the disjoint union of a $(k \times k)$-grid and a clique K_l and then adding all edges between the vertices of the clique and the vertices of the grid. We denote the (k, l)-pyramid by $\Pi_{k,l}$.

Walls. A *wall of height* k, $k \geq 1$, is the graph obtained from a $((k+1) \times (2 \cdot k+2))$-grid with vertices (x, y), $x \in \{1, \ldots, 2 \cdot k+4\}$, $y \in \{1, \ldots, k+1\}$, 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 1. We denote such a wall by W_k.

The *corners* of the wall W_k are the vertices $c_1 = (1, 1)$, $c_2 = (2 \cdot k+1, 1)$, $c_3 = (2 \cdot k+1+(k+1 \mod 2), k+1)$ and $c_4 = (1 + (k+1 \mod 2), k+1)$. (The square vertices in Figure 3.2) We let $C = \{c_1, c_2, c_3, c_4\}$ and we call the pairs $\{c_1, c_3\}$ and $\{c_2, c_4\}$ *anti-diametrical*.

A *subdivided wall* W *of height* k is a wall obtained from W_k after replacing some of its edges by paths without common internal vertices. We call the resulting graph W a *subdivision* of W_k and the vertices that appear in the wall after the replacement *subdivision* vertices.

The non-subdivision vertices of W are called *original* vertices. Consider an embedding of W in the plane. The cycle of W which bounds the infinite face is called the *perimeter* P of W.

(From now we will always denote the perimeter of a wall by P.)

The *layers* of a subdivided wall W of height k are recursively defined as follows. The first layer of W is its perimeter. For $i = 2, \cdots, \lfloor \frac{k}{2} \rfloor$, the i-th layer of W is the $(i-1)$-th layer of the subwall W' obtained from W after removing from W its perimeter and all occurring vertices of degree 1 recursively (see Figure 3.2).

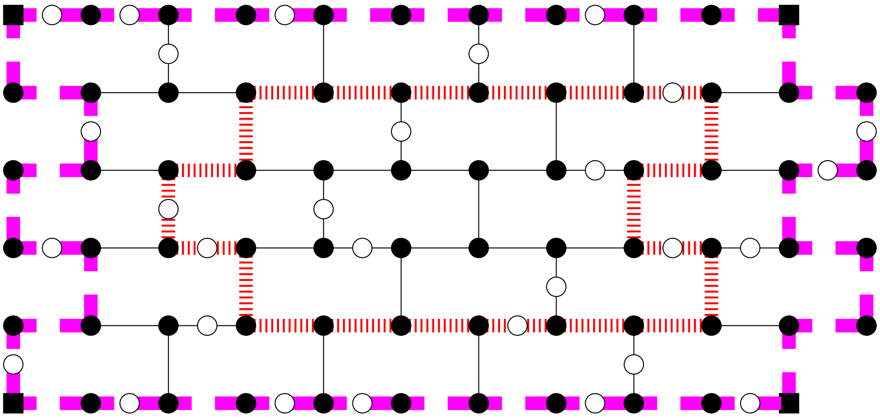

FIGURE 3.2. *The first (magenta–dashed) and second (red–dotted) layers of a wall of height 5*

If W is a subdivided wall of height k, we call a *brick* of W any facial cycle whose non-subdivided counterpart in W_h has length 6. We say that two bricks are *neighbors* if their intersection contains an edge.

Let W_k be a wall. We denote by $P_j^{(h)}$ the shortest path connecting vertices $(1, j)$ and $(2 \cdot k + 2, j)$, and by $P_i^{(v)}$ the shortest path connecting vertices $(i, 1)$ and $(i, k+1)$ with the assumption that for $i < 2 \cdot k + 2$, $P_i^{(v)}$ contains only vertices (x, y) with $x = i, i+1$. We call the paths $P_{k+1}^{(h)}$ and $P_1^{(h)}$ the *southern path of* W_k and *northern part of* W_k respectively. Similarly, we call the paths $P_1^{(v)}$ and $P_{2 \cdot k+1}^{(v)}$ the *western part of* W_k and the *eastern part of* W_k respectively. If W is a subdivision of W_k, we will use the same notation for the paths obtained by the subdivisions of the corresponding paths of W_k.

Compasses and rural divisions. Let W be a subdivided wall in G. Let K' be the connected component of $G \setminus P$ that contains $W \setminus P$. The *compass* K of W in G is the graph $G[V(K') \cup V(P)]$. Observe that W is a subgraph of K and K is connected. We say that a wall is *flat* in G if its compass K in G has no (c_1, c_3)-path and (c_2, c_4)-path that are vertex-disjoint.

OBSERVATION 3. *Let W be a flat wall. Then any subdivision of W is also flat.*

If J is a subgraph of K, we denote by $\partial_K J$ the set of all vertices v such that either $v \in C$ or v is incident with an edge of K that is not in J, that is,

$$\partial_K J = \{v \in V(J) \mid v \in C \text{ or } \exists e \in E(K) \setminus E(J) : v \in e\}.$$

A *rural division* $\mathcal{D}$ of the compass K is a collection $(D_1, D_2, \ldots, D_m)$ of subgraphs of K with the following properties:

1. $\{E(D_1), E(D_2), \ldots, E(D_m)\}$ is a partition of $E(K)$ into non-empty subsets,
2. for $i, j \in [m]$, if $i \neq j$ then $\partial_K D_i \neq \partial_K D_j$ and $V(D_i) \cap V(D_j) = \partial_K D_i \cap \partial_K D_j$,

3. for each $i \in [m]$ and all $x, y \in \partial_K D_i$ there exists a (x, y)-path in D_i with no internal vertex in $\partial_K D_i$,
4. for each $i \in [m]$, $|\partial_K D_i| \leq 3$, and
5. the hypergraph $H_K = (\bigcup_{i \in [m]} \partial_K D_i, \{\partial_K D_i \mid i \in [m]\})$ is planar, its incidence graph can be embedded in a closed disk Δ such that c_1, c_2, c_3 and c_4 appear in this order on the boundary of Δ and for each hyperedge e of H there exist $|e|$ mutually vertex-disjoint paths between e and C in K.

We call the elements of $\mathcal{D}$ *flaps.* A flap $D \in \mathcal{D}$ is *internal* if $V(D) \cap V(P) = \emptyset$.

We can now state the main result of this paper.

THEOREM 3.1. *There exists a computable function f such that for every two graphs H and G and every $k \in \mathbb{N}$, one of the following holds:*

1. *H is a minor of G,*
2. *$\mathbf{tw}(G) \leq f(h) \cdot k$, where $h = \mathbf{n}(H)$*
3. *$\exists A \subseteq V(G)$ with $|A| \leq \mathbf{an}(H) - 1$ such that $G \setminus A$ contains as a subgraph a flat subdivided wall W where*
    - *W has height at least k and*
    - *the compass of W has a rural division $\mathcal{D}$ such that each internal flap of $\mathcal{D}$ has treewidth at most $f(h) \cdot k$.*

We postpone the proof of Theorem 3.1 until Section 5 and we conclude this section with a brief description of the proof. A main ingredient is the Strong Structural Theorem, that is, Proposition 2.5, asserting that every H-minor free graph G can be seen as a tree decomposition such that, for every node G_i, the graph $\overline{G}_i$ is a $f(h)$-nearly embeddable graph. Given that the graph G has treewidth greater than $f(h) \cdot k$ where $f(h)$ is "big enough" (depending on the excluded graph H), there should exist a node G_i of the tree decomposition such that the treewidth of $\overline{G}_i$ is still big enough while the graphs induced by all other nodes have smaller treewidth. This, according to a result of [5], implies that the $f(h)$-nearly embeddable graph $\overline{G}_i$ contains as a subgraph the subdivision W of a "big enough" (but still depending linearly on k) wall that is flatly embedded in a surface, in the sense that its perimeter is the boundary of a disk whose interior, the compass of W, contains, among other parts of $\overline{G}_i$, the rest of W. Our next, and more technical, step is to extract from this "local" structure, concerning $\overline{G}_i$, a wall and a rural division of its compass in the general graph G. For this, we treat all other parts of the tree decomposition as flaps of bounded treewidth that contain at most three non-apex vertices of $\overline{G}_i$. At this point, it follows that the third assertion of Theorem 3.1 holds for the augmented graph G' that is obtained from G if we add all "virtual" edges that are edges of $\overline{G}_i$ but not of G_i. The proof concludes by showing that, even though the removal of these virtual edges may harm parts of the subdivided wall W and the corresponding rural division, these parts can be reconstructed by collections of paths inside the flaps that are now attached on the compass of $\overline{G}_i$.

4. Some auxiliary lemmata.

The main results in this section are Lemmata 4.4 and 4.9 that will be used for the proof of our main result in Section 5.

4.1. An invariance lemma for flatness.

LEMMA 4.1. *Let k be a positive integer and G be a graph that is h-nearly embedded in a surface of Euler genus γ without apices and contains $\Gamma_{2 \cdot k+8}$ as a v-smooth contraction, where v is the loaded corner of $\Gamma_{2 \cdot k+8}$. Then G contains as a* *subgraph a subdivided wall of height k whose compass can be embedded in a closed disk Δ such that the perimeter of W is identical to the boundary of Δ.*

Proof. Assume that $\Gamma_{2\cdot k+8}$ is a v-smooth contraction of G via ϕ, where v is the loaded corner of $\Gamma_{2\cdot k+8}$. Without loss of generality, let $V(\Gamma_{2\cdot k+8}) = \{1, \ldots, 2\cdot k+8\}^2$, where $v = (2\cdot k+8, 2\cdot k+8)$. Let R be the set of external vertices of $\Gamma_{2\cdot k+8}$ and let $G' = G \setminus \bigcup_{x\in R} \phi^{-1}(x)$. As G contains $\Gamma_{2\cdot k+8}$ as a v-smooth contraction and $v \in R$, it follows that G' is embedded inside an open disk Δ'. Moreover G' can be contracted to $\Gamma^*_{2\cdot k+6}$ via the restriction of ϕ to $V(G')$. From the definition of a wall, it follows that $\Gamma^*_{2\cdot k+6}$ contains W_{k+2} as a subgraph. As G' contains $\Gamma^*_{2\cdot k+6}$ as a minor, it follows that G' contains W_{k+2} as a minor. Furthermore, as W_{k+2} has maximum degree 3, from Proposition 2.1, it is also a topological minor of G'. Therefore G' contains as a subgraph (embedded in Δ') a subdivided wall of height $k + 2$. Among all such subdivided walls, let W_{ex} be the one whose compass has the minimum number of faces inside the annulus $\Phi = \Delta_{\text{ex}} \setminus \Delta \subseteq \Delta'$ where Δ_{ex} and Δ are defined as the closed disks defined so that the boundary of Δ_{ex} is the first layer of W_{ex} and the boundary of Δ is the second one.

Let W be the subdivided wall of G' whose perimeter is the boundary of Δ. By definition, all vertices of the compass K of W are inside Δ. It now remains to prove that the same holds also for the edges of K. Suppose for contradiction that $\{x, y\}$ is an edge outside Δ. Clearly, both x and y lie on the boundary of Δ and $\{x, y\}$ is inside the disk Δ^* defined by some brick of W_{ex} that is inside Φ. We distinguish two cases:

Case 1: $\{x, y\}$ are in the same brick, say A of W. Then, there is a path of this brick that can be replaced in W by $\{x, y\}$ and substitute W by a new subdivided wall corresponding to an annulus with less faces, a contradiction. (See, Figure 4.1.)

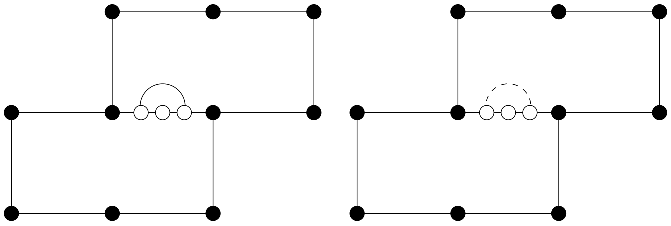

FIGURE 4.1. *Example of Case 1 in Lemma 4.1*

Case 2: $\{x, y\}$ are not in the same brick of W.

Then x and y should belong to neighboring bricks, say B and C respectively. Let A be the unique brick of W_{ex} that contains x and y and w be the unique common vertex in A, B and C. Observe that there is a path P_B of B connecting x and w and a path P_C of C connecting y and w. Then we substitute W by a new wall as follows: we replace w by x, P_C by $\{x, y\}$, and see P_B as a subpath of the common path between B and C. (See, Figure 4.2.) Again, the new wall corresponds to an annulus with less faces, a contradiction. ◻

LEMMA 4.2 ([5]). *There is a function $f : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ such that if G is a graph h-nearly embedded in a surface of Euler genus γ without apices, where $\mathbf{tw}(G) \geq f(\gamma, h)\cdot k$, then G contains as a v-smooth contraction the graph Γ_k with the loaded corner v.*

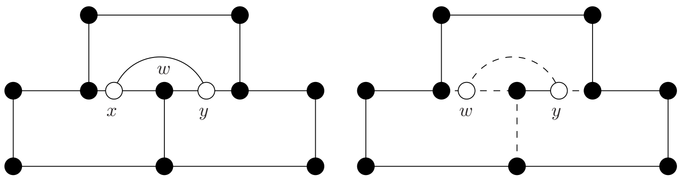

FIGURE 4.2. *Example of Case 2 in Lemma 4.1*

LEMMA 4.3. *Let h be a positive integer and G be a graph that contains a flat subdivided wall W of height h. If K_3 is a subgraph of the compass of W then after applying a ΔY-transformation in K_3 the resulting graph also contains a flat subdivided wall W' of height h as a subgraph. Moreover, W' is isomorphic to a subdivision of W.*

Proof. We examine the non-trivial case where $E(K_3) \cap E(W) \neq \emptyset$. Observe that, as W does not contain triangles, $|E(K_3) \cap E(W)| < 3$. In what follows we denote by x, y, z the vertices of K_3, w the vertex that appears after the transformation, and distinguish the following cases.

Case 1. K_3 and W have exactly one common edge, say $\{x, y\}$. As w is a new vertex, the path (x, w, y) that appears after the ΔY-transformation has no common internal vertices with W. In this case, we replace the edge $\{x, y\}$ in W by the edges $\{x, w\}, \{w, y\}$.

Case 2. K_3 and W have exactly two common edges, say $\{x, y\}$ and $\{x, z\}$.

We distinguish the following two subcases.

Subcase 2.1. x is an original vertex and x is not a corner of W. Let q be the third vertex in the neighborhood of x. Observe that the ΔY-transformation is equivalent to removing the edge $\{y, z\}$, which is not an edge of the wall, and subdividing the edge $\{x, q\}$. Then the lemma follows from Observation 3. (See, Figure 4.3)

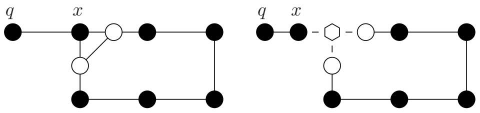

FIGURE 4.3. *Example of Case 2.1 in Lemma 4.3*

Subcase 2.2. x is not an original vertex or x is a corner. Then the ΔY-transformation is equivalent to removing the edge $\{y, z\}$, which is not an edge of W, and then substituting $\{y, x\}$ by $\{y, w\}$ and $\{x, z\}$ by $\{w, z\}$. (See, Figure 4.4)

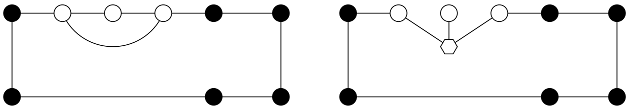

FIGURE 4.4. *Example of Case 2.2 in Lemma 4.3*

Observe that in all the above cases the resulting wall W' remains flat and is isomorphic to a subdivision of W and the lemma follows. ☐

By applying inductively Observation 3 and Lemma 4.3 we derive the following.

LEMMA 4.4. *Let h be a positive integer and G be a graph that contains a flat subdivided wall W of height h as a subgraph. If we apply a sequence of subdivisions or ΔY-transformations* *in G, then the resulting graph will contain a flat subdivided wall W' of height h as a subgraph. Moreover, W' is isomorphic to a subdivision of W.*

4.2. Pyramids and treewidth. Let us first state the next.

PROPOSITION 4.5 ([11]). *Let n be a positive integer. If H is a planar graph with $|V(H)| + 2|E(H)| \leq n$, then H is isomorphic to a minor of the $2n \times 2n$-grid.*

Combining Proposition 4.5 with Eűler's formula for planar graphs we obtain the following.

LEMMA 4.6. *If G is a planar graph then G is isomorphic to a minor of the $(14 \cdot \mathbf{n}(G) - 24) \times (14 \cdot \mathbf{n}(G) - 24)$-grid.*

From the above lemma we obtain the following.

LEMMA 4.7. *Let h be an integer. If G is an h-apex graph then G is isomorphic to a minor of $\Pi_{14\cdot(\mathbf{n}(G)-h)-24,h}$.*

LEMMA 4.8. *Let G be the graph obtained by a $((k + \lceil\sqrt{h}\rceil) \times (k + \lceil\sqrt{h}\rceil))$-grid if we make its vertices adjacent to a set A of h new vertices. Then G contains $\Pi_{k,h}$ as a minor.*

Proof. We denote by G' the grid used for constructing G and let G_1 and G_2 two disjoint subgraphs of G' where G_1 is isomorphic to a $(k \times k)$-grid and G_2 is isomorphic to a $(\alpha \times \alpha)$-grid where $\alpha = \lceil\sqrt{h}\rceil$. Remove from G all vertices not in $A \cup V(G_1) \cup V(G_2)$. Then remove all edges of G' incident to $V(G_2)$ and notice that in the remaining graph F, the vertices in $A \cup V(G_2)$ induce a graph isomorphic to K_{h,α^2} which, in turn, can be contracted to a clique on the vertices of A. Applying the same contractions in F one may obtain $\Pi_{k,h}$ as a minor of G. ☐

LEMMA 4.9. *Let G, H be graphs such that H is not a minor of G and there exists a set $A \subseteq V(G)$ such that $G \setminus A$ contains a wall W of height $g(h) \cdot (k+1) - 1$ as a subgraph, where $g(h) = 14 \cdot (h - \mathbf{an}(H)) + \lceil\sqrt{\mathbf{an}(H)}\rceil - 24$ and $h = \mathbf{n}(H)$. If $|A| \geq \mathbf{an}(H)$ then there exists an $A' \subseteq A$ such that $G \setminus A'$ contains a wall $W' \subseteq W$ of height k as a subgraph with the property that if K' is the compass of W' in $G \setminus A'$ then $V(K') \cap (A \setminus A') = \emptyset$. Moreover, $|A'| < |A|$.*

Proof. Let $A = \{\alpha_i \mid i \in [|A|]\}$ and $P_{g(h)} = \{W_{(m,l)} \mid (m,l) \in [g(h)]^2\}$ be a collection of $(g(h))^2$ disjoint subwalls $W_{(m,l)}, (m,l) \in [g(h)]^2$ of W with height k. For every $(m,l) \in [g(h)]^2$, we denote by $K_{(m,l)}$ the compass of $W_{(m,l)}$ in $G \setminus A$ and let $q_{(m,l)} = (q^1_{(m,l)}, q^2_{(m,l)}, \ldots, q^{|A|}_{(m,l)})$ be the binary vector where for every $j \in |A|$,

$$q^j_{(m,l)} = \begin{cases} 1 & \text{if } \exists v \in V(K_{(m,l)}) : \{v, \alpha_j\} \in E(G) \\ 0 & \text{if } \forall v \in V(K_{(m,l)}) : \{v, \alpha_j\} \notin E(G) \end{cases}$$

We claim that there exists an $(m', l') \in [g(h)]^2$ such that $q_{(m',l')} \neq (1, 1, \ldots, 1)$. Indeed, assume in contrary, that for every $(m,l) \in [g(h)]^2$, $q_{(m,l)} = (1, 1, \ldots, 1)$. We will arrive to a contradiction by showing that H is a minor of G. For this, consider the graph

$$G' = G[V(W) \cup \bigcup_{(m,l)\in[g(h)]^2} V(K_{(m,l)})] \subseteq G.$$

For every $(m,l) \in [g(h)]^2$, contract each $K_{(m,l)}$ to a single vertex and this implies the existence of a $(g(h) \times g(h))$-grid as a minor of G' and therefore of $G \setminus A$ as well. Moreover, for each vertex v of this grid it holds that each vertex in A is adjacent to some vertex of the model of v, therefore G contains the graph J obtained after we take a $(g(h) \times g(h))$-grid and connect all its vertices with $\mathbf{an}(H)$ new vertices. From Lemma 4.8, G contains $\Pi_{14\cdot(n(h)-\mathbf{an}(H))-24,\mathbf{an}(H)}$ as minor. Applying now Lemma 4.7, we obtain that G contains H as a minor, a contradiction. Therefore, there exist $(m', l') \in [g(h)]^2$ and $j_0 \in [|A|]$ such that $q^{j_0}_{(m',l')} = 0$. The lemma follows for $A' = A \setminus \{\alpha_{j_0}\}$ and $W' = W_{(m',l')}$. ☐

5. The main proof.

5.1. Notation. Below, we define the notation which will be useful in the proof of the main result.

Given a tree decomposition $\mathcal{T} = (\mathcal{X} = \{X_i \mid i \in V(T)\}, T)$ of a graph G a vertex $i_0 \in V(T)$ and a set of vertices $I \subseteq N_T(i_0)$, we define $\mathcal{T}_{i_0,I}$ as the collection of connected components of $T \setminus i_0$ that contain vertices of I. Given a subtree Y of T, we define $G_Y = G[\cup_{i \in V(Y)} X_i]$ and $\overline{G}_Y = \cup_{i \in V(Y)} \overline{X}_i$.

Observation 4. *Given a tree decomposition $\mathcal{T} = (\mathcal{X} = \{X_i \mid i \in V(T)\}, T)$ of a graph G, a vertex $i_0 \in V(T)$, and a set of vertices $I \subseteq N_T(i_0)$, it holds that for every $T_1, T_2 \in \mathcal{T}_{i_0,I}$, $\overline{G}_{T_1} \cap \overline{G}_{T_2}$ is a complete graph.*

Given a family of graphs $\mathcal{F}$, a graph G and a set of vertices $S \subseteq V(G)$, we define the class $\mathcal{F}^*_{S,G}$ as the collection of the connected components in the graphs of $\mathcal{F} \setminus S$ and the class $\mathcal{F}_{S,G}$ as the set of graphs in $\mathcal{F}^*_{S,G}$ that have some common vertex with $G \setminus S$. We say that two graphs $G_1, G_2 \in \mathcal{F}_{S,G}$ are G-equivalent if $V(G_1) \cap V(G \setminus S) = V(G_2) \cap V(G \setminus S)$ and let $\mathcal{F}^1_{S,G}, \ldots, \mathcal{F}^{\rho}_{S,G}$ be the equivalence classes defined that way. We denote by $\mathcal{P}_{\mathcal{F},S,G} = \{\bigcup \mathcal{F}^1_{S,G}, \ldots, \bigcup \mathcal{F}^{\rho}_{S,G}\}$, that is, for each equivalence class $\mathcal{F}^i_{S,G}$ we construct a graph in $\mathcal{P}_{\mathcal{F},S,G}$, by taking the union of the graphs in $\mathcal{F}^i_{S,G}$.

5.2. Proof of the Main Result. *Proof of Theorem 3.1.* Let G be a graph that excludes H as a minor. By Proposition 2.5, there is a computable function f_1 such that there exists a tree decomposition

$$\mathcal{T} = (\mathcal{X} = \{X_i \mid i \in V(T)\}, T)$$

of G, where for every $i \in V(T)$, the graphs $\overline{X}_i$ are $f_1(h)$-nearly-embeddable in a surface Σ of genus $f_1(h)$. Among all such tree-decompositions we choose $\mathcal{T} = (\mathcal{X}, T)$ such that

(i) $\mathcal{T}$ is small.
(ii) subject to (i) $\mathcal{T}$ has maximum number of nodes.
(iii) subject to (ii) the quantity $\sum_{\substack{i,j \in V(T) \\ i \neq j}} |X_i \cap X_j|$ is minimized.

Notice that, from Lemma 2.4, condition (i) guaranties the possibility of the choice of condition (ii). We use the notation $\overline{G}$ to denote the graph $\overline{G}_T$ and we call the edges of $E(\overline{G}) \setminus E(G)$ *virtual.*

Let $w : V(T) \to \mathbb{N}$ such that $w(i) = \mathbf{tw}(\overline{X}_i)$. Observation 1 and Observation 2 imply that there exists a vertex $i_0 \in V(T)$ such that $\mathbf{tw}(\overline{X}_{i_0}) \geq \mathbf{tw}(G)$ and at most one of the connected components of $T \setminus i_0$ contains a vertex j such that $w(j) > w(i_0)$. We denote by A_{i_0} the set of apices of the graph $\overline{X}_{i_0}$ and by F the graph $\overline{X}_{i_0} \setminus A_{i_0}$ (notice that $F \subseteq \overline{G}$ but F is not necessarily a subgraph of G as F may contain virtual edges).

From Lemma 2.3 and the choice of i_0 it holds that

$$\mathbf{tw}(F) = \mathbf{tw}(\overline{X}_{i_0} \setminus A_{i_0}) \geq \mathbf{tw}(\overline{X}_{i_0}) - |A_{i_0}| \geq \mathbf{tw}(G) - |A_{i_0}|. \tag{5.1}$$

Recall that $|A_{i_0}| \leq f_1(h)$. Let f_2 be the two-variable function of Lemma 4.2. We define the two-variable function f_3 and the one-variable functions f_4 and f_5 such that

$$\begin{aligned} f_5(h) &= 14 \cdot (h - \mathbf{an}(H)) + \lceil \sqrt{\mathbf{an}(H)} \rceil - 24 \\ f_4(h) &= f_5(h)^{|A_{i_0}| - \mathbf{an}(H) + 1} \\ f_3(h,k) &= f_2(f_1(h), f_1(h)) \cdot (4k \cdot f_4(h) + 12) + f_1(h) \end{aligned}$$

As F is $f_1(h)$-nearly embeddable in Σ and does not contain any apices, from (5.1) and Lemma 4.2, we obtain that if $\mathbf{tw}(G) \geq f_3(h,k)$ then F contains the graph $Q = \Gamma_{4k \cdot f_4(h)+12}$ as a v-smooth contraction, where v is the loaded corner of Q. (See Figure 5.1.)

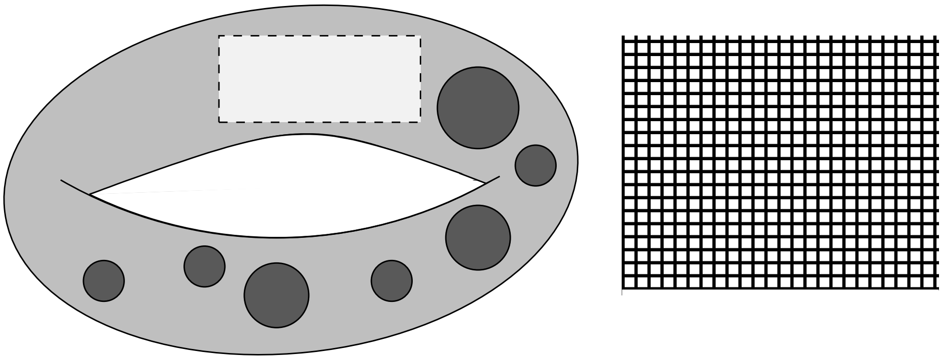

FIGURE 5.1. *On the left: The graph F nearly embedded on Σ (without apices), where the dashed parallelogram represents the disk of the v-smooth contraction. On the right: The content of the disk.*

By Lemma 4.1, it follows that F contains as a subgraph a flat subdivided wall W' of height $2k \cdot f_4(h) + 2$ whose compass K' in F can be embedded in a closed disk Δ such that the perimeter of W' is identical to the boundary of Δ. Furthermore, notice here, that W' is inside the same disk where the preimage of the vertices of the v-smooth contraction where embedded. (See Figure 5.2.)

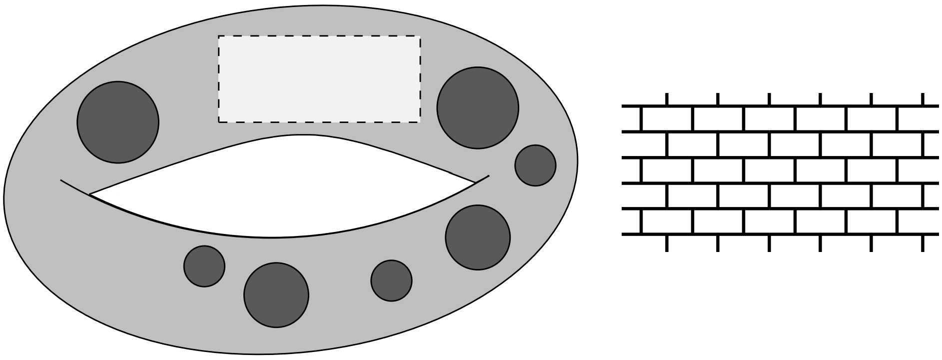

FIGURE 5.2. *On the left: The graph F nearly embedded on Σ (without apices), where the dashed parallelogram represents the perimeter of the wall. On the right: Part the content of the disk.*

Let

$$I' = \{ i \in N_T(i_0) \mid X_i \cap V(K') \neq \emptyset \}.$$

In other words, I' corresponds to all nodes of the tree decomposition $\mathcal{T}$, adjacent to i_0, that have vertices "inside" the compass of the subdivided wall W' in F (dashed parallelogram in Figure 5.2).

Claim 1. *For every* $i \in I', |V(K') \cap X_i| \leq 3$.

Proof of Claim 1. As K' is clearly planar, every clique in K' has size at most 4. Furthermore, K' does not contain K_4 as a subgraph.

Indeed, if so, one of the triangles of the clique would be a separator of G. We denote the vertices of this triangle by z_1, z_2 and, z_3. Let, then, z be the vertex of the clique that is different from z_1, z_2 and, z_3 and S be the vertices contained in the interior of the disk defined by the separating triangle, where $z \in S$.

We may assume that there exists a vertex $j \in I'$ that contains all the vertices of the clique, that is, $\{z, z_1, z_2, z_3\} \subseteq X_j$. Indeed, towards a contradiction assume that there is no such vertex. Then, we may modify $\mathcal{T}$ in order to construct a small tree decomposition $\mathcal{T}'$ of width $\mathbf{tw}(G)$ with more bags than $\mathcal{T}$ in the following way.

We add a new vertex j_0 in T, and a new bag X_{j_0} containing all vertices of the clique, the vertices belonging to the interior of the separating triangle of the clique and all apices of X_{i_0}, that is $X_{j_0} = \{z_1, z_2, z_3\} \cup S \cup A_{i_0}$. We then remove S from X_{i_0}. Finally we add the edge $\{i_0, j_0\}$, remove all the edges between the neighbors $i \in I'$ of i_0 in T, whose common vertices with X_{i_0} are either apices or vertices of S and the clique, and make them adjacent to j_0 instead. Formally, let $\mathcal{T}' = (\mathcal{Y}, T')$, with

$$\begin{aligned}
V(T') &= V(T) \cup \{j_0\}, \text{ where } j_0 \notin V(T),\\
E(T') &= (E(T) \setminus \{\{i, i_0\} \mid i \in N_T(i_0) \wedge (X_i \cap X_{i_0}) \subseteq S \cup \{z_1, z_2, z_3\} \cup A_{i_0}\}) \cup\\
&\qquad \{\{i, j_0\} \mid i \in N_T(i_0) \wedge (X_i \cap X_{i_0}) \subseteq S \cup \{z_1, z_2, z_3\} \cup A_{i_0}\} \cup \{i_0, j_0\} \text{ and}\\
\mathcal{Y} &= \{X_i \mid i \in V(T) \setminus \{i_0\}\} \cup \{Y_{i_0}, Y_{j_0}\},
\end{aligned}$$

where $Y_{i_0} = X_{i_0} \setminus S$ and $Y_{j_0} = \{z_1, z_2, z_3\} \cup S \cup A_{i_0}$. Notice that, as z_1, z_2 and, z_3 induce a separating triangle in K' there is no $i \in V(T)$ such that $X_i \cap \{z_1, z_2, z_3\} \neq \emptyset$ and $X_i \cap [V(K') \setminus (S \cup \{z_1, z_2, z_3\})] \neq \emptyset$. This implies that $\mathcal{T}'$ is indeed a tree decomposition. It is also easy to verify that $\mathcal{T}'$ is small. Indeed, notice that neither $Y_{i_0} \subseteq Y_{j_0}$ nor $Y_{j_0} \subseteq Y_{i_0}$. Moreover, for every $i \in V(T) \setminus \{i_0\}$, neither $X_i \subseteq Y_{j_0}$ nor $X_i \subseteq Y_{i_0}$, as this would imply $X_i \subseteq X_{i_0}$ which is a contradiction to the fact that $\mathcal{T}$ is small. Finally notice that if there exists an $i \in V(T) \setminus \{i_0\}$ such that $Y_{j_0} \subseteq X_i$, then $\{z, z_1, z_2, z_3\} \subseteq Y_{j_0} \subseteq X_i$, a contradiction to the hypothesis as then by the definition of a tree decomposition there exists a vertex $j \in I'$ such that X_j contains all the vertices of the clique. It is also easy to see that for every $i \in V(T) \setminus \{i_0\}$, $Y_{i_0} \nsubseteq X_i$. Thus, $\mathcal{T}'$ contradicts to the choice of $\mathcal{T}$ (condition (ii)). Therefore, there exists a vertex $j \in I'$, say j_0, such that $\{z, z_1, z_2, z_3\} \subseteq X_{j_0}$.

We will now prove that $\mathcal{T}$ does not satisfy condition (iii). Similarly, as above, we modify $\mathcal{T}$ into a tree decomposition $\mathcal{T}''$ in the following way. We remove S from X_{i_0} and add it to X_{j_0}. Then, we remove all the edges between the neighbors $i \in I' \setminus j_0$ of i_0 in T, whose common vertices with X_{i_0} are either apices or vertices of S and the clique, and make them adjacent to j_0 instead.

Formally, let $\mathcal{T}'' = (\mathcal{Z}, T'')$ be the tree decomposition of G with

$$\begin{aligned} &V(T'') = V(T), \\ &E(T'') = (E(T) \setminus \{\{i, i_0\} \mid i \in (N_T(i_0) \setminus \{j_0\}) \wedge (X_i \cap X_{i_0}) \subseteq S \cup \{z_1, z_2, z_3\} \cup A_{i_0}\}) \cup \\ &\qquad\qquad \{\{i, j_0\} \mid i \in N_T(i_0) \wedge (X_i \cap X_{i_0}) \subseteq S \cup \{z_1, z_2, z_3\} \cup A_{i_0}\} \text{ and} \\ &\qquad \mathcal{Z} = \{X_i \mid i \in V(T) \setminus \{i_0, j_0\}\} \cup \{Z_{i_0}, Z_{j_0}\}, \end{aligned}$$

where $Z_{i_0} = X_{i_0} \setminus S$ and $Z_{j_0} = X_{j_0} \cup S$. It is again easy to see that $\mathcal{T}''$ is a small tree decomposition of G. Notice also that $\mathcal{T}$ and $\mathcal{T}''$ contain the same amount of bags. Furthermore, it is easy to see that

$$\sum_{\substack{D, D' \in \mathcal{Z} \\ D \neq D'}} |D \cap D'| < \sum_{\substack{L, L' \in \mathcal{X} \\ L \neq L'}} |L \cap L'|,$$

a contradiction to the choice of $\mathcal{T}$ (condition (iii)). Therefore, K' does not contain any clique of size 4.

Recall, now, that for every $i \in \mathcal{I}'$, $F[V(K') \cap X_i] \subseteq K'$ is a clique. Thus, for every $i \in I', |V(K') \cap X_i| \leq 3$. $\square$

Recall now that $\mathcal{T}_{i_0,I'}$ is the collection of connected components of $T \setminus i_0$ that contain vertices of I' and recall, also, that there exists at most one tree in $\mathcal{T}_{i_0,I'}$, say T', that contains a vertex i_1 with $w(i_1) > w(i_0)$. Let $\mathcal{W}' = \{W'_1, W'_2, W'_3, W'_4\}$ be the collection of vertex-disjoint subwalls of W' of height $f_4(h) \cdot k$ not meeting the vertices of $P^{(h)}_{k \cdot f_4(h)+2}$ and $P^{(v)}_{2k \cdot f_4(h)+3}$ (see Figure 5.3).

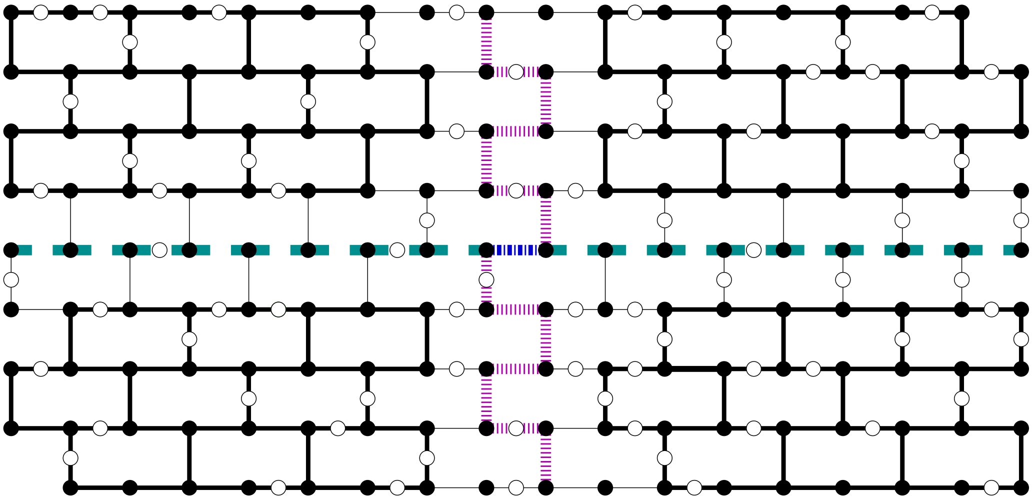

Figure 5.3. *The paths* $P^{(h)}_{k \cdot f_4(h)+2}$ *(cyan - dashed) and* $P^{(v)}_{2k \cdot f_4(h)+3}$ *(magenta - dotted) and the corresponding walls for* $k = 1$ *and* $f_4(h) = 3$.

From Claim 1, X_{i_1} has at most 3 vertices in common with K', therefore there exists a subwall $\tilde{W} \in \mathcal{W}'$ of height $f_4(h) \cdot k$, with compass $\tilde{K}$ in F such that $V(\tilde{K}) \cap V(G_{T'}) = \emptyset$. (For the "big picture", see Figure 5.4.)

Consequently, if we set

$$\tilde{I} = \{i \in N_T(i_0) \mid X_i \cap V(\tilde{K}) \neq \emptyset\}$$

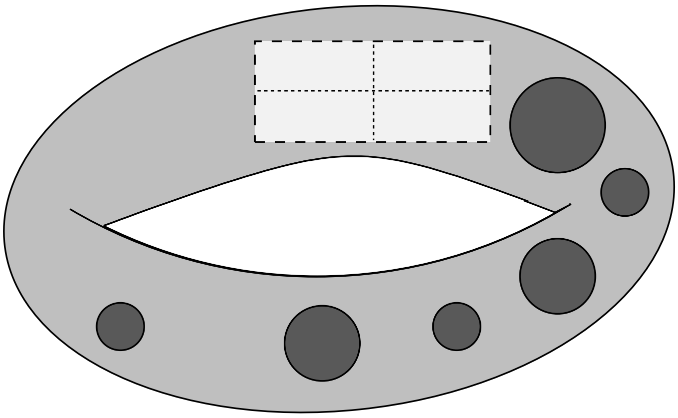

FIGURE 5.4. *The compass $\tilde{K}$ of the wall $\tilde{W}$ is inside one of the four parallelograms belonging to the interior of the dashed parallelogram.*

we have that $\tilde{I} \subseteq I' \setminus \{i_1\}$ and for every tree $\tilde{T} \in \mathcal{T}_{i_0,\tilde{I}} \subseteq \mathcal{T}_{i_0,I'} \setminus \{T'\}$ it holds that $\max\{w(i) \mid i \in V(\tilde{T})\} \leq f_3(h,k)$. Therefore, for every $\tilde{T} \in \mathcal{T}_{i_0,\tilde{I}}$, $\mathbf{tw}(\overline{G}_{\tilde{T}}) \leq f_3(h,k)$. As $G_{\tilde{T}}$ is a subgraph of $\overline{G}_{\tilde{T}}$, it follows that

$$\text{for every } \tilde{T} \in \mathcal{T}_{i_0,\tilde{I}}, \mathbf{tw}(G_{\tilde{T}}) \leq f_3(h,k). \tag{5.2}$$

From Claim 1, it follows that for every $\tilde{T} \in \mathcal{T}_{i_0,\tilde{I}}$, the vertices in $V(G_{\tilde{T}}) \cap V(\tilde{K})$ induce a clique in $\tilde{K}$ with at most 3 vertices where some of its edges may be virtual.

Let $\tilde{V} = V(F) \setminus V(\tilde{K})$ and $\mathcal{F}' = \{G_{\tilde{T}} \mid \tilde{T} \in \mathcal{T}_{i_0,\tilde{I}}\}$. Notice that $\tilde{K} = F \setminus \tilde{V}$. We denote by $\mathcal{F}$ the class $\mathcal{P}_{\mathcal{F}',\tilde{V},F}$. Our aim now is to use the graphs in $\mathcal{F}$ in order to define the compass of $\tilde{W}$ in the graph $\overline{G} \setminus A_{i_0}$ and use them to construct the rural division of that compass.

CLAIM 2. *For every connected component Y of $T \setminus \{i_0\}$ that contains a vertex i_Y of $\tilde{I}$, there is a vertex in $G_Y \setminus X_{i_0}$ connected with $V(\tilde{K}) \cap V(G_Y)$ with $|V(\tilde{K}) \cap V(G_Y)|$ vertex-disjoint paths whose internal vertices belong to $G_Y \setminus X_{i_0}$.*

Proof of Claim 2. First, observe that, G_Y has at least one connected component G'_Y that contains the vertices in $V(\tilde{K}) \cap V(G_Y)$ and such that $V(G'_Y) \setminus X_{i_0} \neq \emptyset$. Otherwise, observe that we may safely remove the vertices in $V(\tilde{K}) \cap V(G_Y)$ from the bags X_i, $i \in V(Y)$ and end up to a contradiction in the choice of $\mathcal{T}$ (condition (iii)). Moreover, condition (ii) implies that X_{i_0} is not a separator of G'_Y. Notice then that every vertex in $V(\tilde{K}) \cap V(G_Y)$ has a neighbor in $G'_Y \setminus X_{i_0}$, as if not, we would yield a contradiction in the choice of $\mathcal{T}$ (condition (iii)). As G'_Y is connected there exists a vertex $v \in V(G'_Y) \setminus X_{i_0}$ and vertex-disjoint paths from v to the vertices of $V(\tilde{K}) \cap V(G_Y)$. $\square$

We call the edges in $\tilde{E} = E(\tilde{K}) \setminus E(G)$ *useless.* We also call all vertices in $V(\bigcup \mathcal{F}) \setminus V(\tilde{K})$ *flying* vertices. The non-flying vertices of a graph R in $\mathcal{F}$ are the *base* of R. Notice that, by the definition of $\mathcal{F}$, each graph R in $\mathcal{F}$ is a subgraph of the union of some graphs of $\mathcal{F}'$. From Observation 4 and (5.2), It follows that

(a) all graphs in $\mathcal{F}$ have treewidth at most $f_3(h,k)$

Observation 4 and Claim 1 yields that

(b) the base vertices of each R induce a clique of size 1,2, or 3 in $\tilde{K}$.

Also, from Claim 2 and the fact that $\tilde{V} \cup A_{i_0} \subseteq X_{i_0}$, we have that

(c) each pair of vertices of some graph in $\mathcal{F}$ are connected in G by a path whose internal vertices are flying.

Note that each clique mentioned in (b) may contain useless edges. Moreover, from (c), all virtual edges of $\tilde{K}$ are edges of such a clique. Let $\tilde{G} = (V(G), \tilde{E} \cup E(G))$, that is, we add in G all useless edges.

It now follows that $\tilde{G} \setminus A_{i_0}$ contains the wall $\tilde{W}$ as a subgraph and the compass of $\tilde{W}$ in $\tilde{G} \setminus A_{i_0}$ is

$$\tilde{K}^{+} = \tilde{K} \cup \bigcup_{R \in \mathcal{F}} R$$

Notice that the wall $\tilde{W}$ remains flat in $\tilde{G}$. Indeed, suppose that Q_1 and Q_2 are two vertex-disjoint paths between the two anti-diametrical corners of $\tilde{W}$ such that the sum of their lengths is minimal. As not both of Q_1 and Q_2 may exist in $\tilde{K}$, some of them, say Q_1 contains some flying vertex. Let R be the graph in $\mathcal{F}$ containing that vertex. Then there are two vertices x and y of the base of R met by Q_1. From (b), $\{x, y\}$ is an edge of $\tilde{K}^{+}$ and we can substitute the portion of Q_1 that contains flying vertices by $\{x, y\}$, a contradiction to the minimality of the choice of Q_1 and Q_2.

Let $\tilde{E}^{+} = E(\tilde{K}^{+}) \setminus E(\bigcup \mathcal{F})$, that is, $\tilde{E}^{+}$ is the set of edges of $\tilde{K}$ not contained in any graph R of $\mathcal{F}$. It follows that all useless edges are contained in $\tilde{E}^{+}$, that is,

$$\tilde{E} \subseteq \tilde{E}^{+} \tag{5.3}$$

For every $e \in \tilde{E}^{+}$, we denote by $\tilde{G}_e$ the graph formed by the edge e (that is, the graph $\tilde{G}[e]$) and let $\mathcal{E} = \{\tilde{G}_e \mid e \in \tilde{E}^{+}\}$. We set $\tilde{\mathcal{D}}^{+} = \mathcal{F} \cup \mathcal{E}$. Notice that,

$$\text{For every graph } R \in \mathcal{F}, \partial_{\tilde{K}^{+}} R \text{ is the base of } R \tag{5.4}$$

$$\text{For every graph } \tilde{G}_e \in \mathcal{E}, \partial_{\tilde{K}^{+}} \tilde{G}_e = V(\tilde{G}_e) \tag{5.5}$$

Claim 3. *$\tilde{\mathcal{D}}^{+} = \mathcal{F} \cup \mathcal{E}$ is a rural division of $\tilde{K}^{+}$.*

Proof of Claim 3. Properties 1 and 2, follow from the construction of the graphs in $\mathcal{F}$ and $\mathcal{E}$. Moreover, Properties 3 and 4 follow from (c) and (b) respectively. For Property 5, recall that $\tilde{W}$ is a subwall of W' whose compass K' in F can be embedded in a closed disk Δ such that the perimeter of W' is identical to its boundary. This implies that $\tilde{K}$ can be embedded in a closed disk $\tilde{\Delta} \subseteq \Delta$ such that the corners c_1, c_2, c_3, and c_4 of $\tilde{W}$ appear in this order on its boundary. We now consider the following hypergraph:

$$\tilde{H}^{+} = (\bigcup \{\partial_{\tilde{K}^{+}} D \mid D \in \tilde{\mathcal{D}}^{+}\}, \{\partial_{\tilde{K}^{+}} D \mid D \in \tilde{\mathcal{D}}^{+}\}).$$

Notice that $V(\tilde{H}^{+}) = V(\tilde{K})$. We can now construct $I(\tilde{H}^{+})$ by applying, for each $D \in \tilde{\mathcal{D}}^{+}$, the following transformations on the planar graph $\tilde{K}$.

- If $|\partial_{\tilde{K}^{+}} D| = 1$, we add a new vertex and an edge that connects it with the unique vertex of $\partial_{\tilde{K}^{+}} D$.

- If $|\partial_{\tilde{K}^+}D| = 2$, we subdivide the edge of $\tilde{K}[\partial_{\tilde{K}^+}D]$ (recall that $\tilde{K}[\partial_{\tilde{K}^+}D]$ is isomorphic to K_2).
- If $|\partial_{\tilde{K}^+}D| = 3$, we apply a ΔY-transformation in $\tilde{K}[\partial_{\tilde{K}^+}D]$ (recall that $\tilde{K}[\partial_{\tilde{K}^+}D]$ is isomorphic to K_3).

From Observation 3 and Lemma 4.3 follows that the obtained graph remains embedded in $\tilde{\Delta}$ (thus, it is also planar). It now remains to show that for each $e \in E(\tilde{H}^+)$ there exist $|e|$ vertex-disjoint paths between $|e|$ and C in $\tilde{K}^+$. Notice that for each $e \in E(H^+)$ the vertices of e belong to $\tilde{K}$. Finally, there are $|e|$ paths between e and C, otherwise we would have a contradiction to the choice of the tree-decomposition. For this, notice that if there do not exist $|e|$ vertex-disjoint paths between e and C then there exists a separator of e and C of size strictly smaller that $|e|$. Then using similar arguments as in the proofs of Claim 1 and Claim 2 we end up contradicting the choice of $\mathcal{T}$. Therefore all conditions required for Claim 3 hold. □

Our aim now is to find in $G \setminus A_{i_0}$ a flat subdivided wall $\widehat{W}$ of height $f_4(h) \cdot k$. From (b),(c), and (5.4), all the useless edges of $\tilde{K}$ are induced by the sets $\partial_{\tilde{K}^+}R$, $R \in \mathcal{F}$ where $\tilde{K}[\partial_{\tilde{K}^+}R]$ is isomorphic to either K_2 or K_3. Our next step is to prove that, in both cases, we may find a flat subdivided wall in $G \setminus A_{i_0}$ of height $f_4(h) \cdot k$ that does not contain any useless edges.

Case 1. $\tilde{K}[\partial_{\tilde{K}^+}R]$ is isomorphic to K_2. Then, from (c), there exists a path in R whose endpoints are the vertices of $\partial_{\tilde{K}^+}R$ and such that its internal vertices are flying.

Case 2. $\tilde{K}[\partial_{\tilde{K}^+}R]$ is isomorphic to K_3. Claim 2, combined with the facts that $\tilde{V} \cup A_{i_0} \subseteq X_{i_0}$ and that $\forall_{R \in \mathcal{F}}\ \partial_{\tilde{K}^+}R \subseteq X_{i_0}$, imply that there exists a flying vertex v_R in R and vertex-disjoint paths between v_R and the vertices of $\partial_{\tilde{K}^+}R$ whose internal vertices are also flying.

The above case analysis implies that for each $R \in \mathcal{F}$ the edge $\{x, y\}$ or the triangle with vertices $\{x, y, z\}$, induced by $\partial_{\tilde{K}^+}R$ may be substituted, using subdivisions or ΔY-transformations by a flying path between x and y or by three flying paths from a flying vertex v_R to x, y, and z respectively. As all edges of these paths are flying, they cannot be useless and therefore they exist also in $G \setminus A_{i_0}$. We are now in position to apply Observation 3 and Lemma 4.4 and obtain that $\tilde{G} \setminus A_{i_0}$ contains a flat subdivided wall $\widehat{W}$ of height $f_4(h) \cdot k$ such that

(I) $E(\widehat{W}) \cap \tilde{E} = \emptyset$ (recall that $\tilde{E}$ is the set of the useless edges) and
(II) $\widehat{W}$ is isomorphic to a subdivision of $\tilde{W}$.

Therefore, from (I), $\widehat{W}$ is a flat subdivided wall of height $f_4(h) \cdot k$ in $G \setminus A_{i_0}$.

Let $\tilde{C}$ and $\widehat{C}$ be the corners of $\tilde{W}$ and $\widehat{W}$ respectively. We denote by σ be the bijection from $\tilde{C}$ to $\widehat{C}$ induced by the isomorphism in (II). We also enhance σ by defining $\phi = \sigma \cup \{(x, x) \mid x \in V(\tilde{W}) \setminus C(\tilde{W})\}$.

Let $\widehat{K}$ be the compass of $\widehat{W}$ in $G \setminus A_{i_0}$. We claim that

$$\widehat{\mathcal{D}} = \{D \cap \widehat{K} \mid D \in \widetilde{\mathcal{D}^+}\}$$

is a rural division of $\widehat{K}$. This is easy to verify in what concerns conditions (1–4). Condition (5) follows by the observation that the mapping ϕ, defined above, is an isomorphism between $H_{\tilde{K}^+}$ and $H_{\widehat{K}}$.

So far, we have found a flat subdivided wall $\widehat{W}$ in $G \setminus A_{i_0}$ and a rural division of its compass $\widehat{K}$. As each flap in $\widehat{\mathcal{D}}$ is a subgraph of a flap in $\widetilde{\mathcal{D}^+}$ we obtain that all flaps in $\widehat{\mathcal{D}}$ have treewidth at most $f_3(h, k)$. By applying Lemma 4.9 $|A_{i_0}| - \mathbf{an}(H) + 1$ times, it follows that there exists a set $A \subseteq A_{i_0}$, such that $|A| \leq \mathbf{an}(H) - 1$ and $G \setminus A$ contains a flat subdivided wall W of height

k such that $W \subseteq \widehat{W}$. Moreover, $V(K) \cap A_{i_0} = \emptyset$, where K is the compass of W in $G \setminus A$. As above,

$$\mathcal{D} = \{D \cap K \mid D \in \widehat{\mathcal{D}}\}$$

is a rural devision of K where all of its flaps have treewidth at most $f_3(h, k)$. The theorem follows as f_3 is a linear function of k. $\square$

The following corollary gives a more precise description of the structure of apex minor free graphs.

COROLLARY 5.1. *There exists a computable function f such that for every two graphs H and G, where H is an apex graph, and every $k \in \mathbb{N}$, one of the following holds:*

1. $\mathbf{tw}(G) \leq f(h) \cdot k$, *where* $h = |V(H)|$
2. *H is a minor of G,*
3. *G contains a flat subdivided wall W where*
    - *W has height k and*
    - *the compass of W has a rural division $\mathcal{D}$ such that each internal flap of $\mathcal{D}$ has treewidth at most $f(h) \cdot k$.*

REFERENCES

[1] ANUJ DAWAR, MARTIN GROHE, AND STEPHAN KREUTZER, *Locally excluding a minor*, in LICS'07, IEEE Computer Society, 2007, pp. 270–279.

[2] ANUJ DAWAR AND STEPHAN KREUTZER, *Domination problems in nowhere-dense classes*, in IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2009), 2009, pp. 157–168.

[3] REINHARD DIESTEL, *Graph Theory*, vol. 173 of Graduate Texts in Mathematics, Springer-Verlag, Berlin, third ed., 2005.

[4] J. FLUM AND M. GROHE, *Parameterized Complexity Theory*, Texts in Theoretical Computer Science. An EATCS Series, Springer-Verlag, Berlin, 2006.

[5] FEDOR V. FOMIN, PETR A. GOLOVACH, AND DIMITRIOS M. THILIKOS, *Contraction obstructions for treewidth*, J. Combin. Theory Ser. B, 101 (2011), pp. 302–314.

[6] PETR A. GOLOVACH, M. KAMINSKI, D. PAULUSMA, AND D. M. THILIKOS, *Induced packing of odd cycles in a planar graph*, in Proceedings of the 20th International Symposium on Algorithms and Computation (ISAAC 2009), vol. 5878 of Lecture Notes in Comput. Sci., Springer, Berlin, 2009, pp. 514–523.

[7] KEN-ICHI KAWARABAYASHI AND YUSUKE KOBAYASHI, *The induced disjoint path problem*, in Proceedings of the 13th Conference on Integer Programming and Combinatorial Optimization (IPCO 2008), vol. 5035 of Lecture Notes in Computer Science, Springer, Berlin, 2008, pp. 47–61.

[8] KEN-ICHI KAWARABAYASHI AND BRUCE REED, *Odd cycle packing*, in Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC 2010), New York, NY, USA, 2010, ACM, pp. 695–704.

[9] YUSUKE KOBAYASHI AND KEN-ICHI KAWARABAYASHI, *Algorithms for finding an induced cycle in planar graphs and bounded genus graphs*, in Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2009), ACM-SIAM, 2009, pp. 1146–1155.

[10] BOJAN MOHAR AND CARSTEN THOMASSEN, *Graphs on Surfaces*, Johns Hopkins University Press, Baltimore, MD, 2001.

[11] NEIL ROBERTSON, PAUL SEYMOUR, AND ROBIN THOMAS, *Quickly excluding a planar graph*, J. Combin. Theory Ser. B, 62 (1994), pp. 323–348.

[12] NEIL ROBERTSON AND P. D. SEYMOUR, *Graph minors. XIII. The disjoint paths problem*, J. Combin. Theory Ser. B, 63 (1995), pp. 65–110.

[13] ———, *Graph minors. XVI. Excluding a non-planar graph*, J. Combin. Theory Ser. B, 89 (2003), pp. 43–76.