



HAL
open science

Minor-Universal Graph for Graphs on Surfaces

Cyril Gavaille, Claire Hilaire

► **To cite this version:**

| Cyril Gavaille, Claire Hilaire. Minor-Universal Graph for Graphs on Surfaces. 2023. hal-04090935

HAL Id: hal-04090935

<https://hal.science/hal-04090935v1>

Preprint submitted on 10 May 2023

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.

Minor-Universal Graph for Graphs on Surfaces

Cyril Gavaille*

LaBRI

University of Bordeaux
France

Claire Hilaire†

LaBRI

University of Bordeaux
France

May 10, 2023

Abstract

We show that, for every n and every surface Σ , there is a graph \mathcal{U} embeddable on Σ with at most cn^2 vertices that contains as minor every graph embeddable on Σ with n vertices. The constant c depends polynomially on the Euler genus of Σ . This generalizes a well-known result for planar graphs due to Robertson, Seymour, and Thomas [*Quickly Excluding a Planar Graph*. J. Comb. Theory B, 1994] which states that the square grid on $4n^2$ vertices contains as minor every planar graph with n vertices.

Keywords: Minor, Universal Graphs, Bounded Genus Graphs

1 Introduction

Graph universality plays an important role in graph theory and in algorithmics. Rather than considering individually each graph of a given family, it is usually simpler to manipulate a reduced set of graphs, or even a single graph, whose properties are close enough from those of the graphs of the family.

For instance, the Strong Product Theorem [DJM⁺20], which states that every planar graph is a subgraph of the strong product of a graph of bounded treewidth and a path, has provided many new insights in planar graph theory, as it suffices to study such product to derive properties on all planar graphs. This includes a proof for bounding their queue number and other variant of chromatic number [DEJ⁺20]. Another example is the design of NCA-labeling schemes [AGKR04], whose goal is to associate with each vertex of each n -vertex tree a short label so that the label of the nearest-common ancestor between any two vertices u, v can be retrieved from the labels of u and of v . Interestingly, it is shown in [GKL⁺18] how a specific tree with n^c vertices containing every n -vertex trees as topological minor can be used for the design of succinct NCA-labeling schemes, namely with $c \log n$ -bit labels. Similar connections are known between adjacency-labeling schemes and induced-universal graphs [ADBTK17, DEG⁺21]. The last example is the refinement of celebrated

*gavoille@labri.fr. This work is partially funded by the French ANR projects ANR-17-CE40-0015 (DISTANCIA), ANR-18-CE47-0010 (QUDATA) and ANR-22-CE48-0001 (TEMPOGRAL).

†claire.hilaire@labri.fr

Excluded Grid Theorem of Robertson and Seymour [RST94] which states that the graphs excluding as minor a fixed planar graph have bounded treewidth. This result relies on the fact that every planar graph is a minor of some relatively small grid. So excluding such a grid as minor implies excluding as minor any small enough planar graph.

Our results. Motivated by the latter example, we focus in this paper on universality under minor¹ containment. More precisely, a graph \mathcal{U} is *minor-universal* for a given graph family \mathcal{F} if every graph of \mathcal{F} is a minor of \mathcal{U} . Keeping in mind that we want the graph \mathcal{U} to be representative of some properties of graphs in \mathcal{F} , most of the time we will impose that \mathcal{U} and the graphs in \mathcal{F} share a common property, say, for instance, planarity. Note that the Strong Product Theorem fails for this purpose as it provides for planar graphs a universal graph which is certainly of n^2 vertices but far from being planar in general. It can be shown², for instance, that the strong product of two paths of n vertices is already of genus $\Omega(n)$.

From [RST94, Theorem (1.3) & (1.4)], the $2n \times 2n$ -grid contains as minor every planar graph with n vertices. In other words, planar n -graphs have a minor-universal grid with at most $4n^2$ vertices. This quadratic bound is essentially optimal, as long as grids are concerned, since [BCE⁺19, Lemma 5] showed that any minor-universal grid for planar graphs must have $\Omega(n^2)$ vertices.

We generalize the result of [RST94] to graphs on surface as follows:

Theorem 1 *For every n and every surface Σ of Euler genus g , there is a graph $\mathcal{U}_{n,\Sigma}$ embedded on Σ with $O(g^2(n+g)^2)$ vertices that contains as minor every graph embeddable on Σ with n vertices.*

In other words, $\mathcal{U}_{n,\Sigma}$ is a minor-universal graph for n -vertex graphs embeddable on Σ . In fact, as we will see in Theorem 2, the graph $\mathcal{U}_{n,\Sigma}$ depends only on the Euler genus and orientability of Σ . Moreover, the embedding of the minor is preserved in $\mathcal{U}_{n,\Sigma}$.

Our proof is constructive. The minor-universal $\mathcal{U}_{n,\Sigma}$, as well as the witness of the minor in $\mathcal{U}_{n,\Sigma}$, can be constructed in polynomial time. Interestingly, along the way, it gives an alternative and self-contained proof of [RST94] for the minor-universal grid of planar graphs.

Related results. For trees, [Bod02] proved that the smallest minor-universal tree for n -vertex trees has between $\Omega(n \log n)$ and $O(n^{1.985})$ vertices. In fact, the result of [GKL⁺18] implies that this smallest minor-universal tree has between $\Omega(n^{1.724})$ and $O(n^{1.895})$ vertices³. From computational perspectives, the smallest minor-universal tree of two bounded degree trees can be computed in cubic time [NRT00], whereas it is NP-complete to decide whether or not a tree is minor of another one. Minor universality has also been investigated for infinite graphs [DK99, HMS⁺21]. However, as noticed by [DK99], the problem already becomes different in the planar case since the infinite grid is no longer minor-universal for infinite planar graphs.

¹Recall that H is a minor of G if H can be obtained from G by vertex removals, edge removals, and edge contractions.

²Using the Euler's Formula.

³For these results, rooted binary trees and topological minors were considered. However, this applies to our setting since, up to a factor of two, the rooted and unrooted versions are equivalent, and minor and topological minor containments are equivalent in binary trees.

2 Preliminaries

We consider compact connected surfaces without boundary. A *canonical system of loops* for a surface Σ of Euler genus g is a one-vertex graph embedded on Σ with g loop-edges such that by cutting Σ along these loops results in a surface homeomorphic to a polygon P with $2g$ sides. The way the pairs of sides of P are reattached to form Σ can be described by a *signature*, a word associating clockwise one symbol with each side of P . The polygon P and its signature form a *canonical polygonal schema*. Note that the loops are either two-sided (the surface Σ is orientable with oriented genus $g/2$) or one-sided (Σ is non-orientable with non-oriented genus g).

For $g > 0$, the signature is $a_1 a_2 \bar{a}_1 \bar{a}_2 \dots a_{g-1} a_g \bar{a}_{g-1} \bar{a}_g$ if Σ is orientable, and $a_1 a_1 \dots a_g a_g$ if Σ is non-orientable. For the sphere, $g = 0$, the signature is $a_0 \bar{a}_0$ by convention. Sides are reattached if they have the same (or opposite) symbol in the signature, given the convention that positive symbols correspond to a clockwise orientation of the sides, and negative symbols to an anticlockwise orientation.

Consider a graph G embedded⁴ on some surface Σ with a canonical system of loops L . Denote by $G \uplus L$ the graph embedded on Σ obtained from the union of G and L , and by adding a new vertex at each intersection between G and L .

A *polygonal embedding* for G is the planar embedding Π into a polygon obtained from $G \uplus L$ by cutting Σ along the edges of L . (We refer to Fig. 1 for an illustration.) The edges and vertices on L appear duplicated on the boundary of outerface of Π . More precisely, the boundary of the outerface of Π is a cycle that can be cut into a clockwise sequence of paths sharing their extremities. These paths, called *sides* of Π , correspond to the edges of the polygon P , and their extremities, called *corners* of Π , corresponds to the vertices of P . Note that corners have degree two in Π . The ordered sequence of corners is called the *border* of Π . The *signature* of Π , denoted by $\sigma(\Pi)$, describes how the sides of Π are reattached to form $G \uplus L$ on Σ . More precisely, the i th symbol of the signature is associated with the i th sides of Π , i.e., the path between the i th and $(i + 1)$ th corners. These sides are merged according to the orientation given by the symbols, and there are called *twin sides* of Π . Obviously, twin sides must contain the same number of vertices. Furthermore, the vertices and edges that are identified in this process are called respectively *twin vertices* and *twin edges*.

Given a polygonal embedding Π , the *sewing* of Π , denoted by $\text{sewing}(\Pi)$, is the graph obtained by reattaching the sides of Π according to its border and its signature. It is clear that, if Π comes from $G \uplus L$ embedded on Σ by cutting along the edges of L , then $\text{sewing}(\Pi)$ is isomorphic to $G \uplus L$. In particular, $\text{sewing}(\Pi)$ embeds on Σ and contains G as minor.

3 Main result

To prove Theorem 1, we consider a graph G embedded on the surface Σ of Euler genus g and having a canonical systems of loops L with vertex r and signature σ . Thus G has a polygonal embedding Π_G such that G is a minor of $\text{sewing}(\Pi_G)$, a graph that embeds on Σ . Note that $|V(\Pi_G)| \leq 2|V(G \uplus L)| + |\sigma|$ since every vertex of $G \uplus L$ appears at most on two sides of Π_G , and each corner is a copy of r . In fact, Π_G has at most n internal vertices due to G , and a

⁴We only consider simple graphs and cellular embeddings, i.e., where each face is homeomorphic to a disc.




Figure 1: The graph K_6 embedded on the torus (of Euler genus 2) with the system of loops $L = (\{r\}, \{\ell_1, \ell_2\})$, and a polygonal embedding Π of K_6 of signature $\sigma = a_1 a_2 \bar{a}_1 \bar{a}_2$. By construction, K_6 is a minor of the graph sewing(Π) = $K_6 \uplus L$.

certain number of vertices lying on the sides of Π_G due to r and to the intersections between G and L . We can bound this number thanks to the following result⁵:

Lemma 1 ([LPVV01, FHdM22]) *Given an embedding of G on a surface of Euler genus g , there is a polynomial time algorithm that computes a canonical system of g loops such that each loop intersects any edge of G in at most⁶ 30 points.*

It follows that $|V(\Pi_G)| = O(|V(G)| + g|E(G)|) = O(g(n + g))$ since each of the g loops of L crosses $O(1)$ times each of the $O(n + g)$ edges⁷ of G . Our strategy then is to transform step by step the initial polygonal embedding Π_G into some unique polygonal embedding \mathbb{U} that preserves signature and minor containment of its sewing. The number of vertices of this final embedding is $|V(\mathbb{U})| = O(|V(\Pi_G)|^2)$. Our minor-universal graph $\mathcal{U}_{n,\Sigma}$ as in Theorem 1 is nothing else than sewing(\mathbb{U}) that has less than $|V(\mathbb{U})| = O(g^2(n + g)^2)$ vertices as claimed.

To formalize signature and minor containment preserving, we introduce the following relation. A polygonal embedding Π is a p -minor of a polygonal embedding Π' if they have the same signature and if sewing(Π) is a minor of sewing(Π'). Moreover, we say that Π has size (m, n) if each side has at most m vertices, the corners excluded, and at most n internal vertices (those that are not lying on the boundary of the outerface of Π). In particular, every polygonal embedding Π of size (m, n) has at most $n + |\sigma(\Pi)|(m + 1)$ vertices⁸, and sewing(Π) has at most $n + |\sigma(\Pi)|(m + 1)/2$ vertices (this is $n + |\sigma|m/2 + 1$ vertices if $\sigma(\Pi)$ is canonical, see Section 4.3). A polygonal embedding of size of $(m, 0)$ is a Hamiltonian outerplane graph.

The final polygonal embedding is denoted hereafter by $\mathbb{U}_{\sigma,m}$, where σ is any signature and m an integral parameter. We will see later in Section 4.3 that $\mathbb{U}_{\sigma,m}$ has signature σ and size $(m, O(|\sigma|m^2))$. Roughly speaking, it is a square half-grid cut by its diagonal. The vertices on the boundary of the outerface are those of the diagonal, see Fig. 2 for an example.

⁵We refer to [CdV21][Theorem 8.1] for a reformulation of the original statement of [LPVV01].

⁶This is only 4 points in the case of orientable surface.

⁷From Euler's formula in simple graphs.

⁸ $|\sigma(\Pi)|$ denotes the number of symbols in $\sigma(\Pi)$.




Figure 2: The polygonal embedding $\mathbb{U}_{\sigma, m}$, with $m = 3$ and signature $\sigma = a_1 a_2 \bar{a}_1 \bar{a}_2$.

The technical theorem is the following. Note that it applies to any polygonal embedding, not only those of canonical signature.

Theorem 2 *Every polygonal embedding Π of size (m, n) is a p -minor of $\mathbb{U}_{\sigma(\Pi), m+2n}$.*

Before proving Theorem 2, let us show that it implies Theorem 1. Consider a graph G on n vertices embedded on a surface Σ of Euler genus g . By Lemma 1, G has a polygonal embedding Π_G whose signature σ depends only on Σ (as it corresponds to that of the canonical polygonal schema of Σ), and such that Π_G has size (m, n) with $m = O(n + g)$. Moreover, G is a minor⁹ of $\text{sewing}(\Pi_G)$. By Theorem 2, Π_G is a p -minor of $\mathbb{U}_{\sigma, m+2n}$, which by definition of p -minor implies that $\text{sewing}(\Pi_G)$ is a minor of $\text{sewing}(\mathbb{U}_{\sigma, m+2n})$. It follows that G is a minor of $\text{sewing}(\mathbb{U}_{\sigma, m+2n}) = \mathcal{U}_{n, \Sigma}$, that has the required properties: it is embedded on Σ and has $O(g^2(n + g)^2)$ vertices.

To prove Theorem 2, we proceed in two steps, which are summarized and formalized by the next two lemmas whose proofs are sketched hereafter and fully presented in Section 4:

Lemma 2 *Every polygonal embedding Π of size (m, n) is a p -minor of a polygonal embedding of size $(m + 2n, 0)$.*

Intuition is that all internal vertices of Π can all be somehow pushed to the sides of the outerface, without increasing much the number of vertices on each side. It follows we can consider that all the vertices belong to the outerface of the polygonal embedding making it outerplanar. This is done by first cutting along the edges of some spanning forest of the inner graph (rooted in a vertex of the outerface), creating an empty space and transforming the graph into an outerplanar graph. Along this process, each vertex that is not a leaf appears more than once (the number of new vertices is bounded by the number of edges of the forest, thus by n). To keep track of this vertex duplication, we add some edges that could be contracted later to get back to the original polygonal embedding. Other transformations are then needed to embed those new edges and obtain an outerplanar polygonal embedding of the right size and containing Π as a p -minor.

⁹In fact Π_G is a minor of $\text{sewing}(\Pi_G) \setminus \{r\}$ as the unique vertex r of system of loops is not part of G and can be removed.

Lemma 3 *Every polygonal embedding Π of size $(m, 0)$ is a p -minor of $\mathbb{U}_{\sigma(\Pi), m}$.*

Note that Theorem 2 is a straightforward combination of Lemma 2 and Lemma 3.

The proof of Lemma 3 is inspired from the embedding into a $n \times n$ -grid for hamiltonian planar graphs with n vertices [RST94, Theorem (1.3)]. From this embedding, one can easily show that outerplanar graphs with n vertices are minor of a half- $n \times n$ -grid. In our context, the diagonal of this half-grid is the place for the sides of our polygonal embedding, as depicted in Fig. 2. Interestingly, along the way, we give in Proposition 1 an alternative prove of a result of [RST94, Theorem (1.4)] that states that every planar graph with n vertices is minor of a Hamiltonian planar graph with $2n$ vertices.

4 Proofs of the lemmas

Along the proofs, if H is a minor of G , it will be sometimes more convenient to say that G is a *major* of H . Similarly, we say that polygonal embedding Π_G is a *p -major* of Π_H , if Π_H is a p -minor of Π_G .

Because plane graphs we will consider are actually polygonal embeddings, the boundary of the outerface is a cycle. Whenever we talk about outerface we refer to this cycle. We say that two plane graphs H, G such that H is minor of G have the *same outerface* if the minor H has been obtained from G without removing vertices or edges of the outerface of G , nor contracting edges with both endpoints on the outerface of G . It follows:

Property 1 *If a polygonal embedding Π_H is a minor of polygonal embedding Π_G with same outerface, border and signature, then Π_H is a p -minor of Π_G .*

Recall that a polygonal embedding is nothing else than a plane graph with a given cycle outerface, border and signature. Property 1 allows us to consider that we work on plane graph, as long as we keep the same outerface and border when constructing a major.

In Section 4.1, we consider plane graphs with a fixed cycle outerface, i.e., with an outerface whose boundary is a cycle. We give tools to construct a plane major that preserves this outerface. These tools will be used later in the construction of the p -major of Π in Section 4.2.

4.1 Considering plane graphs

Let G be a planar graph G with a given planar embedding and having a cycle outerface O . We denote by $n = |V(G)| - |V(O)|$ the number of internal vertices of G .

Because the boundary of the outerface of G is a connected graph (a cycle), we can assume w.l.o.g. that G is connected as well. If not, we can for instance triangulate all the faces except the outerface. By doing this we get a plane major of G without altering its number of vertices nor its cycle outerface.

Let \mathcal{F} be a forest of G , composed of k trees denoted by T_1, \dots, T_k , and constructed as follows:

- Start with a spanning tree T of G ;
- Remove one edge of any path in T that connects two vertices of O , if any;

- Remove all remaining isolated vertices that are in O ; and
- Root each T_i at the only one vertex of $T_i \cap O$, denoted by r_i .

Observe that \mathcal{F} spans all internal vertices of G , thus $|V(F)| = n + k$ and $|E(F)| = n$. Because we consider planar major of G , it is quite tempting to add edges to force \mathcal{F} to be a single tree. However, we cannot do that and suppose that the inner vertices induced a connected subgraph of G , as there could be an edge between two vertices of the outerface separating the internal vertices into two parts.

We now construct a major of G , denoted by G_1 , that consists in *blowing up* each tree T_i into a graph $T'_i = C_i \cup E_i$, composed of a cycle C_i plus a set E_i of extra edges. More precisely, blowing up T_i consists in traversing the tree according to a plane Euler tour from r_i (or in other words, a walk along the boundary of the outerface of T_i considered here as a single plane graph). The cycle C_i is constructed iteratively by adding a new vertex at each vertex of T_i visited along this tour, two consecutive vertices on this tour being connected by an edge. See Fig. 3 for an illustration. Since \mathcal{F} has no isolated vertices, T_i has at least one edge. If T_i has exactly one edge, then C_i consists of one single edge. The set E_i is composed of all the edges connecting any two vertices u, v of C_i if they corresponds to the same visited vertex of T_i and that, among them, appear consecutive during the visit. In other words, we add an edge between u, v in C_i if they corresponds to same visited vertex w of T_i and that none of the vertices between u and v in C_i corresponds to w .

This completes the description of blowing up T_i into T'_i .




Figure 3: Blowing up the tree T_i (with black edges) into $T'_i = C_i \cup E_i$ (with red and violet edges). The root r_i of T_i becomes an anchor a_i of C_i in T'_i .

The graph G_1 is then obtained from G where each T_i is replaced by T'_i which is outerplanar. This is possible because the walk along the boundary of the outerface of T_i and T'_i are isomorphic, so T'_i can be plugged into T_i by keeping the plane embedding of G .

Claim 1 *The graph G_1 has the following properties:*

- G_1 is a plane major of G with same outerface O ;
- C_1, \dots, C_k form a partition of internal vertices of G_1 .
- Each C_i contains exactly one vertex in O , called its anchor;
- G_1 has $2n - k$ internal vertices;

Proof. It is easy to check that T'_i is a major of T_i , as it suffices to contract all edges of E_i

to obtain a graph isomorphic to T_i . By doing this for every T'_i in G_1 we get exactly G (up to some isomorphism) since these are the only differences.

We have seen that the trees T_i 's of \mathcal{F} form a partition of the internal vertices of G . It follows that T'_i forms a partition of the internal vertices of G_1 , and C_i as well since $V(T'_i) = V(C_i)$.

Each T'_i intersects O in exactly one vertex, its anchor, the vertex corresponding to r_i in T_i . In particular, no edge of T'_i belongs to O . So contracting edges of E_i to get T_i cannot affect the outerface, and thus G_1 and G have the same outerface O .

To obtain C_i from T_i , each edge of T_i is traversed twice. It follows that $|V(C_i)| = 2|E(T_i)|$. We have seen that $|E(F)| = \sum_{i=1}^k |E(T_i)| = n$. Therefore, the number of internal vertices of G_1 is $\sum_{i=1}^k (|V(C_i)| - 1) = 2n - k$ as we need to remove the anchors. \square

We now construct from G_1 a major G_2 , that consists in applying a *splitting anchor* operation at every anchor of G_1 . More precisely, consider a cycle C_i with anchor a_i in G_1 . Let u, w_1, w_2, v be the neighbors of a_i , taken in this cyclically ordered around a_i , such that u, v belong to O and w_1, w_2 to T'_i . By construction u, a_i, v belongs to the boundary of a common face, as well as w_1, a_i, w_2 . The edge $a_i - w_1$ belongs to C_i , and $a_i - w_2$ belongs to either C_i or E_i , depending whether r_i has one or several children, and $w_1 = w_2$ if C_i consists in one edge. The splitting operation consists in replacing a_i by an edge $a'_i - a''_i$, called *edge-anchor* of a_i , and reconnecting to a'_i all the a_i 's neighbors going from u to w_1 , and reconnecting to a''_i all the a_i 's neighbors going from w_2 to v . We denote by P_i the path going from a'_i to a''_i in C_i , and by F_i the new set of edges E_i in G_2 . Note that C_i and P_i differ by a_i, a'_i, a''_i , and that E_i and F_i differ by a''_i in G_2 that plays the role of a_i in G_1 . See Fig. 4 for an illustration.




Figure 4: Splitting anchor a_i of C_i into the edge-anchor $a'_i - a''_i$.

Claim 2 *The graph G_2 is a Hamiltonian plane major of G_1 with $|V(O)| + 2n$ vertices.*

Proof. Clearly, G_2 is a plane major of G_1 , as it suffices to contract every edge-anchor of G_2 to obtain G_1 .

To construct a Hamiltonian cycle in G_2 , we start with the cycle outerface of G_2 . Then, each edge-anchor $a'_i - a''_i$ is replaced by P_i . Inherited from Claim 1 and from the fact that C_i 's partition all internal vertices of G_1 , all the internal vertices of G_2 are (disjointly) spanned by the P_i 's. Thus it forms a Hamiltonian cycle for G_2 .

The graph G_1 has $|V(O)|$ vertices on its outerface O and $2n - k$ internal vertices (by Claim 1). Splitting each of the k anchors adds k vertices in G_2 . Therefore,

$$|V(G_2)| = |V(G_1)| + k = (|V(O)| + (2n - k)) + k = |V(O)| + 2n. \quad \square$$

Note that the splitting anchor operation alters the outerface. So, for the proof of Lemma 2, that needs to consider polygonal embedding whose twin sides, in particular, have to be of same length, we will need to make extra transformation on the outerface as explained later in Section 4.2.

Before, we observe that the two previous transformations (by Claim 1 and Claim 2) provide an alternative proof of [RST94, Theorem (1.4)] that states that every n -vertex planar graph is a minor of a planar Hamiltonian graph with at most $2n$ vertices.

Recall that a *circuit* in a graph extends the notion of cycle to subgraph composed of one single vertex or edge. A circuit is *separating* if its deletion increases the number of connected component of the graph, and it is *non-separating* otherwise. It is not difficult to see that every graph has a non-separating circuit with at least one vertex.

Proposition 1 *Every planar graph with n vertices and with a non-separating circuit of k vertices is minor of a Hamiltonian planar graph with at most $2n - k$ vertices. In particular, every triangulation with $n \geq 4$ vertices is minor of a Hamiltonian planar graph with at most $2n - 4$ vertices.*

Proof. Let G be a plane graph with n vertices and with a non-separating circuit of k vertices, denoted by O . W.l.o.g. we can assume that O is a cycle, i.e., $k \geq 3$, since otherwise we can triangulate G and consider its triangle outerface as non-separating circuit. And, obviously the triangulation is a major of G and we will get a minor of $2n - 3$ vertices that is less than $2n - k$ if $k < 3$.

First, we embed G in the plane such that all the vertices of $G \setminus O$ are inside O . Note that O is not necessarily the boundary of the outerface of this embedding, since there can be some chords outside of O . Denote by C the set of these chords. Let G' be the plane graph obtained from $G \setminus C$ such that O is the boundary of its outerface which is a cycle.

Since G' is a plane graph with cycle outerface O , we can construct the major G_1 as in Claim 1, and the Hamiltonian plane major G_2 from G_1 (by Claim 2). Due to the edge-anchor, the outerface of G_2 is a subdivision of O . Therefore, we can add all the chords of C in G_2 that we can embed outside the outerface of G_2 , while preserving its planarity. Thus, $G_2 \cup C$ is the desired major of G .

The graph G' has $n' = n - k$ internal vertices by construction. So, by Claim 2, G_2 (and $G_2 \cup C$ as well) has $k + 2n' = 2n - k$ vertices as claimed.

If G is a triangulation with at least four vertices, then G has a non-separating circuit of length $k + 1$. Indeed, the outerface is a non-separating triangle that can be increased by one using any incident internal face. So the above property construction applies with $k = 4$. \square

Actually, the proof of [RST94, Theorem (1.4)] gives an upper bound of $2n - 4$ vertices (assuming $n \geq 4$), and relies on Whitney's Theorem¹⁰ [Whi31]. In contrast Proposition 1 gives an explicit and direct construction with the same upper bound. We believe that it gives an interesting construction of the major where the Hamiltonian cycle cuts the graph into two sides. Inside the cycle are the copies of edges that were initially in G , and outside of the

¹⁰Every planar triangulation without separating triangle has a Hamiltonian cycle.

cycle are the edges introduced by the operations of blowing up trees and splitting anchors: contracting those edges results in G .

4.2 Proof of Lemma 2

The goal of this section is to show that:

Lemma 2 *Every polygonal embedding Π of size (m, n) is a p-minor of a polygonal embedding of size $(m + 2n, 0)$.*

Let Π be a polygonal embedding of size (m, n) . Up to taking a p-major of Π , we assume that each corner of Π is on the boundary of a inner triangle: if not, we can connect by an edge the two neighbors of each corner. This does not alter the size of Π . This is to prevent against any further edge insertion that could increase the degree of corners which must be exactly two in any polygonal embedding. The boundary of the outerface of Π is a cycle, so the construction of G_2 as in Claim 1 applies. Let Π_1 be the polygonal embedding obtained from Claim 1, with Π in the role of G , Π_1 in the role of G_1 . The embedding Π_1 admits Π as a minor while preserving the cycle outerface (border and signature does not matter for Claim 1). Thus Π is a p-minor of Π_1 . Keeping the same notation as in Claim 1, Π_1 has exactly $2n - k$ internal vertices that are partitioned by the cycles C_1, \dots, C_k , each C_i intersecting the outerface in one vertex, its anchor a_i .

Unfortunately, as said previously the transformation of G_1 into G_2 as in Claim 2 does not preserve a polygonal embedding: twin edges in G_1 are not twin anymore in G_2 because of the new created edge-anchor. So it breaks the p-major sequence we are constructing. To preserve a polygonal embedding, we need to modify accordingly twin sides whenever we apply a splitting anchor so that each vertex of one side has a twin vertex on its twin side.

For this purpose we define a *twin splitting* operation that applies to each anchor of Π_1 . Roughly speaking, we apply a splitting anchor operation while subdividing one of the two twin edges that are incident to the twin vertex of the anchor. The edge to subdivide depends on a total ordering \prec defined on the vertices of the cycle outerface. It is based on their rank, starting with the very first corner of the border of Π_1 . So $u \prec v$ if vertex u is visited before v when traversing the cycle outerface from the first corner of the border.

Each step t of this twin process consists in applying one twin splitting operation which results into a new polygonal embedding Π_1^t obtained from Π_1^{t-1} , starting with $\Pi_1^0 = \Pi_1$. More precisely, consider $u - a - v$ be three consecutive vertices on a side of Π_1^{t-1} where a is an anchor and $u \prec a \prec v$. Let x, b, y be the twin vertices respectively of u, a, v . Because an anchor cannot be a corner, vertices u, v, x, y are well defined. The twin splitting operation consists in applying a splitting anchor operation on a followed by (cf. Fig. 5):

- a subdivision of the edge $b - y$, if $a \prec b$ or $y \prec b$; or
- a subdivision of the edge $x - b$, if $b \prec a$ and $b \prec y$.

In the former case, the edge $b - y$ is replaced by the path $b - s - y$ and $b - s$ becomes the twin edge of $a' - a''$, the edge-anchor of a . In the latter case, the edge $x - b$ is replaced by the path $x - s - b$ and $s - b$ becomes the twin of $a' - a''$. All the other twin relations are unchanged. See Fig. 5 for an illustration. It is important to observe that b may be an anchor that will be eventually split at some step $t' > t$.




Figure 5: The twin splitting operation for an anchor a : it combines a splitting anchor operation on a (in brown) and an edge subdivision (in blue) incident to its twin b . Twin relations are with dotted lines. In this representation, the focus on the twin sides creates the illusion that the inside part of the embedding (in gray) is a disconnected region, which is not true. However, in this representation choice, the inside part, that is locally planar, may appear twisted if fully represented in the case of a non-orientable surface.

After this twin process, denote by Π_2 the final embedding obtained from Π_1 , i.e., the embedding obtained from Π_1 by applying successively all twin splitting operations.

Claim 3 *The embedding Π_2 has the following properties:*

- Π_2 is a polygonal embedding that is p -major of Π_1 ;
- Π_2 has size $(m + k, 2n - k)$;
- P_1, \dots, P_k form a partition of the internal vertices of Π_2 ;
- During the twin process, no edge-anchor is subdivided and no two edge-anchors are twins.

Proof. Each twin splitting operation inserts edges and split anchors on the outerface of Π_1 . So, by contracting these new edges, we get back to Π_1 which shows that Π_2 is a major of Π_1 . Moreover, thanks to the twin splitting, eventually every vertex (and edge) of a side in Π_2 have a well defined twin. So, Π_2 is a polygonal embedding that is p -major of Π_1 .

At each twin splitting operation, the number of vertices of a side increases by at most one while the number of internal vertices does not change. Since Π_1 contains k anchors, after applying the twin process, each side of Π_2 contains at most $m + k$ vertices (corners excluded) and $2n - k$ internal vertices. So, the size of Π_2 is $(m + k, 2n - k)$.

The internal vertices in Π_1 and in Π_2 are the same. And after each splitting anchor operation the paths P_i 's span the same internal vertices than C_i 's. So by Claim 1, P_i 's also form a partition of the internal vertices of Π_2 .

It remains to show that during the twin splitting process no two edge-anchors can be twin or subdivided. So, consider any step t' that performs a twin splitting operation at some anchor b of $\Pi_1^{t'-1}$. We need to check that indeed the operation does not create twin edge-anchors or subdivide an edge-anchor. If the twin of b does not belong to an edge-anchor, then it is fine since the twin splitting of b will not result in a subdivision or a twin edge of an edge-anchor. So we only need to check the situation where the twin of b belongs to an edge-anchor. This occurs only if b and its twin a were both anchors in $\Pi_1^{t'-1}$ for some $t < t'$. W.l.o.g. assume that t is the step where a is split into $a' - a''$, and let s be the neighbor of b in Π_1^t resulting of the twin splitting of a in Π_1^{t-1} .

We remark that two edge-anchors cannot be incident, and that on a side with path $x-b-y$ that has been subdivided into $x-b-s-y$, then neither $x-b$ nor $b-s$ can be subdivided anymore (because the twin of b cannot be an anchor anymore). The path $x-b-s$ can only be replaced by an edge-anchor $x-b'-b''-s$ if b is an anchor.

We use the notations as above and as in Fig. 5. Let $u-a-v$ be the vertices on the twin side of b in Π_1^{t-1} , where a is an anchor and twin of b . W.l.o.g. assume $u \prec a \prec v$, and let x, y be the twins of u, v respectively. In Π_1^t , we have $u \prec a' \prec a'' \prec v$, and there is some vertex s adjacent to b coming from an edge subdivision incident to b (of either $b-y$ or $x-b$). Let us apply step t' on b which produces the edge-anchor $b'-b''$ in $\Pi_1^{t'}$.

To simplify notations we will reuse names u, v, x, y to denote the neighbors of a', a'', b, s in Π_1^t as well as in $\Pi_1^{t'-1}$.

Assume that $b \prec a$ and $b \prec y$, i.e., the case 2 holds for a . In that case, we have the path $x-s-b-y$ in $\Pi_1^{t'-1}$, and also $u \prec a' \prec a''$. By exchanging the roles of $s-b-y$ with $a'-a''-v$ the case 1 holds for b . This is because the condition $a \prec b$ rewrites in $b \prec a$ which is true (a'' plays the role of a). Therefore, the twin splitting of b subdivides the edge $a''-v$ which cannot be an edge-anchor by the previous remark. The edge-anchor $b'-b''$ is twin with a non edge-anchor as well. Thus, step t' does not twin or subdivide any edge-anchor in this case.

Assume that $a \prec b$ or $y \prec b$, i.e., the case 1 holds for a . In that case, we have the path $x-b-s-y$ in $\Pi_1^{t'-1}$, and also that $u \prec a' \prec a''$.

If $y \prec b$, then we have $y \prec s \prec b \prec x$. By exchanging the roles of $s-b-x$ with $u-a'-a''$ the case 1 holds for b . This is because the condition $u \prec a'$ rewrites in $y \prec b$ which is true.

If $b \prec y$, then we have $x \prec b \prec s$. By exchanging $x-b-s$ with $u-a'-a''$ the case 1 holds again for b . This is because the condition $a' \prec a''$ rewrites in $b \prec y$ which is true (s plays the role of y).

Thus, in both cases, the twin splitting of b subdivides the edge $u-a'$ which cannot be an edge-anchor by the previous remark. The edge-anchor $b'-b''$ is twin with a non edge-anchor as well. And therefore, in all the cases, step t' does not twin or subdivide any edge-anchor. \square

The last property of Claim 3 (no edge-anchors is subdivided) ensures that the k edge-anchors in Π_1 are still existing in Π_2 after applying the twin process on Π_1 .

In order to transform Π_2 into an outerplanar embedding, we will apply at each of its edge-anchor a *swapping* operation. So, each step t of this process consists in applying one such operation which results into a new polygonal embedding Π_2^t obtained from Π_2^{t-1} , starting with $\Pi_2^0 = \Pi_2$.

The swapping operation is defined as follows. (We reuse the same notations as in Fig. 4 describing the splitting anchor operation which has been used to make Π_2 .) Let $a'_i - a''_i$ be an edge-anchor of Π_2^{t-1} , P_i be the path connecting a'_i to a''_i , and F_i be the edges not in P_i and embedded inside the cycle $P_i \cup \{a'_i - a''_i\}$. Finally, let b', b'' be the twin vertices of respectively a'_i, a''_i .

To obtain Π_2^t , we first delete the edge $a'_i - a''_i$ and all edges of F_i . By this way, the side containing $a'_i - a''_i$ is extended by P_i . Then, on the outerface of Π_2^{t-1} , we connect b' to b'' by a copy P'_i of P_i and including a copy F'_i of all the edges of F_i such a way that the direction of P_i is preserved. It means that if u is traversed before v when going from a'_i to a''_i on P_i , then

u' is traversed before v' when going from b' in b'' in P'_i where u', v' are the copies in P'_i of u, v respectively. Furthermore, the twin of any vertex u in P_i is its copy of u' in P'_i . By this way, the side containing $b' - b''$ is extended by P'_i . See Fig. 6 for an illustration.




Figure 6: A swapping operation for the edge-anchor $a'_i - a''_i$ in Π_2^{t-1} (on the left) leading to a new embedding Π_2^t (on the right). Vertex u is traversed before v when going from a' to a'' on P_i , and so for their twin vertices u', v' when going from b' to b'' on P'_i . The vertices of P_i that were internal in Π_2^{t-1} belongs to the outerface of Π_2^t .

The last property of Claim 3 (no edge-anchors are twin) ensures that $b' - b''$ was not an edge-anchor in Π_2^{t-1} . So the swapping operation on the edge-anchor $a'_i - a''_i$ leaves unchanged all the other edge-anchors.

We denote by Π_3 the final embedding obtained from Π_2 by applying successively a swapping operation on all its edge-anchor. Let us denote by $A \cong B$ if A, B are isomorphic graphs or homomorphic embeddings.

Claim 4 *The embedding Π_3 is a polygonal embedding p -major of Π_2 with $\text{sewing}(\Pi_3) \cong \text{sewing}(\Pi_2)$ and size $(m + 2n, 0)$.*

Proof. Up to a permutation of the indices, we can assume that the swapping operation for $a'_i - a''_i$ is performed at step i from Π_2^{i-1} . It is easy to see that each operation does not create edge crossings, that the outerface remains a cycle, and that each vertex (edge) of each side has a well-defined twin vertex (edge) in its twin side. Therefore, Π_2^i is a polygonal embedding, and also Π_3 , that is Π_2^k , by transitivity.

To prove it has the expected size, let $n_i = |V(P_i) \setminus \{a'_i, a''_i\}|$ be the number of internal vertices of P_i in Π_2^{i-1} . Note that this number is the same as the number of vertices of P_i in $\Pi_2^0 = \Pi_2$ because P_i is not altered by any of the previous step $j < i$. Therefore, from Claim 3, $\sum_{j=1}^k n_j = 2n - k$.

The swapping operation for $a'_i - a''_i$ increases by n_i the number of vertices of its side and its twin side, and, at the same time, decreases by n_i the number of internal vertices of Π_2^{i-1} . By Claim 3, Π_2 has size $(m + k, 2n - k)$. So after step i , the size of Π_2^i is $(m + k + \sum_{j=1}^i n_j, n - 2k - \sum_{j=1}^i n_j)$. It follows that Π_3 , which is Π_2^k , has size $(m + 2n, 0)$ since $\sum_{j=1}^k n_j = 2n - k$.

It is clear that signature and border, i.e., the sequence of corners, are not altered by any swapping operation. So, it remains to show that $\text{sewing}(\Pi_3) \cong \text{sewing}(\Pi_2)$ (and by this way it will also show that Π_3 is a p-major of Π_2). By transitivity, it suffices to show that $\text{sewing}(\Pi_2^i) \cong \text{sewing}(\Pi_2^{i-1})$.

Let R_i be the plane graph composed of $P_i \cup \{a'_i - a''_i\} \cup F_i$ in Π_2^{i-1} such that the boundary of its outerface is the cycle $P_i \cup \{a'_i - a''_i\}$. Similarly, let R'_i be the plane graph composed of $P'_i \cup \{b' - b''\} \cup F'_i$ in Π_2^i with cycle outerface $P'_i \cup \{b' - b''\}$. Note that $V(R_i) = V(P_i)$ and $V(R'_i) = V(P'_i)$. For convenience, denote by $G^j = \text{sewing}(\Pi_2^j)$. We want to show that $G^{i-1} \cong G^i$.

Observe that in G^{i-1} and in G^i , whenever the twin side are merged, $a'_i = b'$, $a''_i = b''$ and thus the edge $a'_i - a''_i = b' - b''$ exists in both graphs (whereas $a'_i - a''_i$ exists only in Π_2^{i-1}). Also, the way the vertices of P_i are twin with the vertices P'_i (by preserving the order when going from a'_i to a''_i in P_i), ensures that $P'_i = P_i$ and $F'_i = F_i$ in G^i . It follows that $R_i \cong R'_i$, and since the swapping operation alters only R_i , $V(G^{i-1}) = V(G^i)$. And thus, if there is a difference between G^{i-1} and G^i , it must be an edge.

From previous equalities, if an edge $x - y$ lies in R_i then the edge exists in R'_i and thus in G^i . And conversely, if $x - y$ lies in R'_i , the edge exists in R_i and thus in G^{i-1} . If $x - y$ lies outside R_i (x, y may be both in $V(R_i)$ with an embedding not inside R_i), then $x - y$ lies also in Π_2^i since the swapping operation alters only R_i (and its inside). So $x - y \in E(G^{i-1})$ implies $x - y \in E(G^i)$ (even if $x - y$ belongs a some side of Π_2^{i-1}). Conversely, if $x - y$ lies outside R'_i in Π_2^i , then $x - y$ lies also in Π_2^{i-1} (and outside of R_i). It follows that edge sets of G^{i-1} and G^i are the same, proving that $G^{i-1} \cong G^i$. \square

By transitivity of the p-minor relation, every polygonal embedding Π of size (m, n) is a p-minor of Π_2 (by Claim 3) that is a p-minor of Π_3 of size $(m + 2n, 0)$ (by Claim 4), which completes the proof of Lemma 2.

4.3 Proof of Lemma 3

The goal of this section is to show that:

Lemma 3 *Every polygonal embedding Π of size $(m, 0)$ is a p-minor of $\mathbb{U}_{\sigma(\Pi), m}$.*

The proof is an adaptation of the proof of [RST94, Theorem (1.3)] to the case of outerplanar polygonal embeddings. The original proof uses grid major whereas we need to use $\mathbb{U}_{\sigma, m}$.

We first define the polygonal embedding $\mathbb{U}_{\sigma, m}$, which is defined for all signature σ and $m \in \mathbb{N}$. To construct $\mathbb{U}_{\sigma, m}$, we start from a square half-grid of dimension $|\sigma|m$, that is with $|\sigma|m$ rows and $|\sigma|m$ columns. The vertices are column-row pairs of integers (i, j) where $0 \leq j \leq i < |\sigma|m$, and (i, j) and (i', j') are adjacent if and only if $|i - i'| + |j - j'| = 1$. Then, the vertices of the diagonal are connected by a cycle, (i, i) being connected to $(i - 1, i - 1)$ modulo $|\sigma|m$. This cycle forms the cycle outerface of $\mathbb{U}_{\sigma, m}$. Finally, every m edges along this cycle is subdivided by one vertex, a corner. More precisely, for each $i \in \{0, \dots, |\sigma| - 1\}$, we add a vertex c_i between (im, im) and $(im - 1, im - 1)$ modulo $|\sigma|m$. The border of $\mathbb{U}_{\sigma, m}$ is the sequence $(c_0, \dots, c_{|\sigma|-1})$, and the i th side of $\mathbb{U}_{\sigma, m}$, associated with the i th symbol of σ , is the path between c_i and $c_{i+1 \bmod |\sigma|}$ on the cycle outerface. See Fig. 7 and Fig. 2 for illustrations.




Figure 7: The polygonal embedding $\mathbb{U}_{\sigma,m}$ for $m = 2$ and $\sigma = a_1 a_2 a_1 a_2$, a non-canonical signature for a non-orientable surface of genus 2. Its size is $(2, 28)$. The cycle outerface is $c_0 - (0, 0) - (1, 1) - c_1 - (2, 2) - (3, 3) - c_2 - (4, 4) - (5, 5) - c_3 - (6, 6) - (7, 7) - c_0$.

Clearly, the number of internal vertices of $\mathbb{U}_{\sigma,m}$ is $n = 1 + 2 + \dots + |\sigma|m - 1$. By merging the twin sides of $\mathbb{U}_{\sigma,m}$ to obtain $\text{sewing}(\mathbb{U}_{\sigma,m})$, we destroy at least half of the vertices of the cycle outerface of $\mathbb{U}_{\sigma,m}$. It follows that $\text{sewing}(\mathbb{U}_{\sigma,m})$ has no more than $n + |\sigma|(m + 1)/2$ vertices. However, depending on the signature, this number can be lower as the set of corners may collapse even more when merging the sides. E.g., if the signature is canonical, $\text{sewing}(\mathbb{U}_{\sigma,m})$ has only $n + |\sigma|m/2 + 1$ vertices since a canonical signature tells us that all the corners are twins, resulting into the single vertex of the canonical system of loops.

To summarize, we have:

Claim 5 *The embedding $\mathbb{U}_{\sigma,m}$ is a polygonal embedding of signature σ with $\frac{1}{2}|\sigma|m \cdot (|\sigma|m - 1)$ internal vertices, each of the $|\sigma|$ sides having m non-corner vertices. In particular, $\text{sewing}(\mathbb{U}_{\sigma,m})$ has no more than $\frac{1}{2}(|\sigma|^2 m^2 + |\sigma|)$ vertices.*

Let Π be a polygonal embedding of size $(m, 0)$. Up to taking a p-major of Π , we can suppose that each side of Π is composed of exactly m non-corner vertices. If needed, one can subdivide an edge and its twin of any too small side, resulting in a p-major of Π . So, w.l.o.g., the cycle outerface of Π is composed of $|\sigma(\Pi)|$ sides of m non-corner vertices each.

Since Π is outerplane, its cycle outerface contains all its vertices. The embeddings Π and $\mathbb{U}_{\sigma(\Pi),m}$ have same signature and also m non-corner vertices per side. So both embeddings have isomorphic border and a cycle outerface of $|\sigma(\Pi)|(m + 1)$ vertices.

We denote by $(u_0, \dots, u_{|\sigma(\Pi)|-1})$ the border of Π , and by $v_0, \dots, v_{|\sigma(\Pi)|m-1}$ its non-corner vertices ordered clockwise around its cycle outerface so that this cycle is

$$u_0 - v_0 - v_1 - \dots - v_{m-1} - u_1 - v_m - \dots - v_{im-1} - u_i - v_{im} - \dots$$

Up to taking a p-major of Π , we will assume that all inner faces of Π are triangulated such that corners have degree two. For each non-corner v_i , we define the two indices $a(i), b(i)$ such that $\{v_k : k \in [a(i), b(i)]\}$ is the minimal subset of vertices containing v_i and all its non-corner

neighbors. Note that $a(i) = i$ or $b(i) = i$ is possible. For technical reasons, we increase $b(0)$ to $b(0) = |\sigma(\Pi)|m$.

In order to show that Π is a minor of $\mathbb{U}_{\sigma(\Pi),m}$, we construct a *witness* of Π in $\mathbb{U}_{\sigma(\Pi),m}$ with the property that it preserves its outerface and its border. A witness for a minor Π in $\mathbb{U}_{\sigma(\Pi),m}$ is a collection $\{W(u)\}_{u \in V(\Pi)}$ of nonempty pairwise disjoint subsets of vertices of $\mathbb{U}_{\sigma(\Pi),m}$ such that each set $W(u)$ induces a connected subgraph of $\mathbb{U}_{\sigma(\Pi),m}$, and for each edge $u - v$ of Π , $W(u) \cup W(v)$ induces a connected component in $\mathbb{U}_{\sigma(\Pi),m}$.

The witness of Π in $\mathbb{U}_{\sigma(\Pi),m}$ is defined by:

- $W(u_i) = \{c_i\}$ for each corner u_i ; and
- $W(v_i) = \{(i, j)\} \cup C(v_i) \cup R(v_i)$ for each non-corner v_i ,
where $C(v_i) = \{(i, k) : k \in (a(i), i]\}$ and $R(v_i) = \{(k, i) : k \in [i, b(i))\}$.

Clearly, for corner vertices, $W(u_i) = \{c_i\}$ fulfills all desired properties: non-emptiness, disjointness and connectedness. For non-corner vertices, $W(v_i)$ induces a connected components of $\mathbb{U}_{\sigma(\Pi),m}$ since it is the union of a column subpath $C(v_i)$ and of a row subpath $R(v_i)$ of $\mathbb{U}_{\sigma(\Pi),m}$ intersecting in (i, i) (in the case where $C(v_i)$ and $R(v_i)$ are not empty).

Let us check that $W(v_i)$ and $W(v_j)$ are pairwise disjoint (witness for corner and non-corner are clearly disjoint). Assume $j < i$. A non-empty intersection is only possible between $C(v_i)$ and $R(v_j)$. And, this can only occur at (i, j) . Note that the case $j = 0$ provides disjoint witnesses, since $(i, 0) \notin C(v_i)$. If $(i, j) \in C(v_i)$, then $j \in (a(i), i]$, and thus v_i has a neighbor $v_{a(i)}$ with $a(i) < j$. If $(i, j) \in R(v_j)$, then $i \in [j, b(j))$, and thus v_j has a neighbor $v_{b(j)}$ with $i < b(j)$. It follows that $a(i) < j < i < a(j)$ and that $v_{a(i)} - v_i$ and $v_j - v_{b(j)}$ is a pair of crossing edges, which is impossible in the outerplanar embedding Π .

It remains to check that, for each $u - v \in E(\Pi)$, $W(u) \cup W(v)$ induces a connected graphs in $\mathbb{U}_{\sigma(\Pi),m}$. We can restrict our attention to non-corner neighbors, since $c_i \in W(u_i)$ have exactly two neighbors that are $(im - 1, im - 1) \in W(v_{im-1})$ and $(im, im) \in W(v_{im})$.

So consider an edge $v_i - v_j$ of Π with $j < i$. Let $W(v_i, v_j)$ be the set of vertices of the path connecting $(i, i) \in W(v_i)$ to $(j, j) \in W(v_j)$ inside the grid part of $\mathbb{U}_{\sigma(\Pi),m}$, and defined by

$$W(v_i, v_j) = \{(i, i), (i, i + 1), \dots, (i, j - 1), (i, j), (i - 1, j), \dots, (j, j)\} .$$

To show that $W(v_i) \cup W(v_j)$ induced a connected component in $\mathbb{U}_{\sigma(\Pi),m}$, it suffices to show that $W(v_i, v_j) \subseteq C(v_i) \cup R(v_j)$ since $W(v_i, v_j)$ induces a path in $\mathbb{U}_{\sigma(\Pi),m}$, and $C(v_i) \subset W(v_i)$ and $R(v_j) \subset W(v_j)$. If $(i', j') \notin W(v_i, v_j)$ for some $i' \in [j, i)$ or $j' \in [i, j)$, then the set $\{v_k : k \in [a(i), b(i)]\}$ cannot covers the non-corner neighbors of v_i or the set $\{v_k : k \in [a(j), b(j)]\}$ cannot covers the non-corner neighbors of v_j . It remains to check that $(i, j) \in C(v_i) \cup R(v_j)$, even if $(i, j) \notin C(v_i)$ or $(i, j) \notin R(v_j)$ is possible. If $j = 0$, then we are done because $b(0) = |\sigma(\Pi)|m$ and thus $(i, 0) \in R(0)$. We are also done if $|i - j| = 1$ since $(i, i) \in W(v_i)$ and $(j, j) \in W(v_j)$ are neighbors in $\mathbb{U}_{\sigma(\Pi),m}$. Since the inner faces of Π are triangulated, the edge $v_i - v_j$ (that lies inside Π since v_i and v_j are not consecutive in the cycle outerface) share a triangle $v_i - v_j - v_k$ such that $k \notin [i, j]$, the case $j = 0$ and $i = |\sigma(\Pi)|m - 1$ being excluded. If $k > i$, then $a(j) > i$ and $(i, j) \in C(v_j)$. If $k < j$, then $b(i) > j$ and $(i, j) \in R(v_i)$.

We have therefore proved that $W(v_i, v_j) \subseteq W(v_i) \cup W(v_j)$ and that $W(v_i) \cup W(v_j)$ induces a connected component in $\mathbb{U}_{\sigma(\Pi),m}$. It follows that Π is a minor of $\mathbb{U}_{\sigma(\Pi),m}$ with same outerface,

border and signature. By Property 1, Π is a p-minor of $\mathbb{U}_{\sigma(\Pi),m}$ that completes the proof of Lemma 3.

5 Conclusion

We have shown that Euler genus- g graphs have an Euler genus- g minor-universal graph of $O(g^2(n+g)^2)$ vertices. It is not clear whether the dependencies in g and in n are tight. Even for the planar case, it is not obvious that the n^2 term is required, noting that the $\Omega(n^2)$ known lower bound exists only if the minor-universal has to be a grid.

It would be interesting to extend the result to other graph families, in particular bounded treewidth graphs and minor-free graphs. Another direction could be to extend the result to topological minor-universal graphs.

Acknowledgement. We are grateful to Arnaud de Mesmay for preliminary discussions and suggestions for this problem.

References

- [ADBTk17] S. ALSTRUP, S. DAHLGAARD, AND M. BÆK TEJS KNUDSEN, *Optimal induced universal graphs and adjacency labeling for trees*, Journal of the ACM, 64 (2017), pp. Article No. 27, pp. 1–22. DOI: [10.1145/3088513](https://doi.org/10.1145/3088513).
- [AGKR04] S. ALSTRUP, C. GAVOILLE, H. KAPLAN, AND T. RAUHE, *Nearest common ancestors: A survey and a new algorithm for a distributed environment*, Theory of Computing Systems, 37 (2004), pp. 441–456. DOI: [10.1007/s00224-004-1155-5](https://doi.org/10.1007/s00224-004-1155-5).
- [BCE⁺19] T. C. BIEDL, E. W. CHAMBERS, D. EPPSTEIN, A. DE MESMAY, AND T. OPHELDERS, *Homotopy height, grid-major height and graph-drawing height*, in 27th International Symposium on Graph Drawing and Network Visualization (GD), vol. 11904 of Lecture Notes in Computer Science, Springer, September 2019, pp. 468–481. DOI: [10.1007/978-3-030-35802-0_36](https://doi.org/10.1007/978-3-030-35802-0_36).
- [Bod02] O. BODINI, *On the minimum size of a contraction-universal tree*, in 28th International Workshop on Graph-Theoretic Concepts in Computer Science (WG), vol. 2573 of Lecture Notes in Computer Science, Springer, June 2002, pp. 25–34. DOI: [10.1007/3-540-36379-3_3](https://doi.org/10.1007/3-540-36379-3_3).
- [CdV21] É. COLIN DE VERDIÈRE, *Algorithms for embedded graphs*, October 2021. monge.univ-mlv.fr/~colinde/cours/all-algo-embedded-graphs.pdf. Lectures Note.
- [DEG⁺21] V. DUJMOVIĆ, L. ESPERET, C. GAVOILLE, G. JORET, P. MICEK, AND P. MORIN, *Adjacency labelling for planar graphs (and beyond)*, Journal of the ACM, 68 (2021), pp. Article No. 42, pp. 1–33. DOI: [10.1145/3477542](https://doi.org/10.1145/3477542).
- [DEJ⁺20] V. DUJMOVIĆ, L. ESPERET, G. JORET, B. WALCZAK, AND D. R. WOOD, *Planar graphs have bounded nonrepetitive chromatic number*, Advances in Combinatorics, 5 (2020). DOI: [10.19086/aic.12100](https://doi.org/10.19086/aic.12100).

- [DJM⁺20] V. DUJMOVIĆ, G. JORET, P. MICEK, P. MORIN, T. UECKERDT, AND D. R. WOOD, *Planar graphs have bounded queue-number*, Journal of the ACM, 67 (2020), pp. Article No. 22, pp. 1–38. DOI: [10.1145/3385731](https://doi.org/10.1145/3385731).
- [DK99] R. DIESTEL AND D. KÜHN, *A universal planar graph under the minor relation*, Graph Theory, 32 (1999), pp. 191–206. DOI: [10.1002/\(SICI\)1097-0118\(199910\)32:2<191::AID-JGT8>3.0.CO;2-I](https://doi.org/10.1002/(SICI)1097-0118(199910)32:2<191::AID-JGT8>3.0.CO;2-I).
- [FHdM22] N. FULADI, A. HUBARD, AND A. DE MESMAY, *Short topological decompositions of non-orientable surfaces*, Tech. Rep. [2203.06659v1](https://arxiv.org/abs/2203.06659v1) [cs.CG], arXiv, March 2022.
- [GKL⁺18] P. GAWRYCHOWSKI, F. KUHN, J. ŁOPUSZAŃSKI, K. PANAGIOTOU, AND P. SU, *Labeling schemes for nearest common ancestors through minor-universal trees*, in 29th Symposium on Discrete Algorithms (SODA), ACM-SIAM, 2018, pp. 2604–2619. DOI: [10.1137/1.9781611975031.166](https://doi.org/10.1137/1.9781611975031.166).
- [HMS⁺21] T. HUYNH, B. MOHAR, R. ŠÁMAL, C. THOMASSEN, AND D. R. WOOD, *Universality in minor-closed graph classes*, Tech. Rep. [2109.00327v1](https://arxiv.org/abs/2109.00327v1) [math.CO], arXiv, September 2021.
- [LPVV01] F. LAZARUS, M. POCCHIOLA, G. VEGTER, AND A. VERROUST, *Computing a canonical polygonal schema of an orientable triangulated surface*, in 17th Annual ACM Symposium on Computational Geometry (SoCG), ACM-SIAM, June 2001, pp. 80–89. DOI: [10.1145/378583.378630](https://doi.org/10.1145/378583.378630).
- [NRT00] N. NISHIMURA, P. RAGDE, AND D. M. THILIKOS, *Finding smallest supertrees under minor containment*, International Journal of Foundations of Computer Science, 11 (2000), pp. 445–465. DOI: [10.1142/s0129054100000259](https://doi.org/10.1142/s0129054100000259).
- [RST94] N. ROBERTSON, P. D. SEYMOUR, AND R. THOMAS, *Quickly excluding a planar graph*, Journal of Combinatorial Theory, Series B, 62 (1994), pp. 323–348. DOI: [10.1006/jctb.1994.1073](https://doi.org/10.1006/jctb.1994.1073).
- [Whi31] H. WHITNEY, *A theorem on graphs*, Annals of Mathematics, 32 (1931), pp. 378–390. DOI: [10.2307/1968197](https://doi.org/10.2307/1968197).