THEORY OF COMPUTING, Volume 13 (5), 2017, pp. 1–47
www.theoryofcomputing.org

SPECIAL ISSUE: APPROX-RANDOM 2013

A Pseudo-Approximation for the Genus of Hamiltonian Graphs

Yury Makarychev* Amir Nayyeri† Anastasios Sidiropoulos‡

Received October 1, 2013; Revised May 15, 2017; Published August 31, 2017

Abstract: The genus of a graph is a basic parameter in topological graph theory that has been the subject of extensive study. Perhaps surprisingly, despite its importance, the problem of approximating the genus of a graph is very poorly understood. Thomassen (1989) showed that computing the exact genus is NP-complete, and the best known upper bound for general graphs is an $O(n)$-approximation that follows by Euler's characteristic.

We give a polynomial-time pseudo-approximation algorithm for the orientable genus of Hamiltonian graphs. More specifically, on input a graph G of orientable genus g and a Hamiltonian path in G, our algorithm computes a drawing on a surface of either orientable or non-orientable genus $O(g^7)$.

ACM Classification: F.2.2, G.2.2

AMS Classification: 68W25

Key words and phrases: approximation algorithms, graph genus, Hamiltonian graphs

1 Introduction

A *drawing* of a graph G on a surface $\mathcal{S}$ is an injective mapping φ that sends every vertex $v \in V(G)$ into a point $\varphi(v) \in \mathcal{S}$, and every edge $\{u,v\}$ into a simple curve between $\varphi(u)$ and $\varphi(v)$, so that the images of

A preliminary version of this paper appeared in the Proceedings of the 15th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2013) [9].

*Supported in part by the NSF Career Award CCF-1150062.

†Supported by the NSF under Grants No. CCF 1566624 and CCF 1617951.

‡Supported by the NSF under award CAREER 1453472 and grant CCF 1423230.

 DOI: 10.4086/toc.2017.v013a005

different edges are allowed to intersect only on their endpoints. A surface is called *orientable* if it can be embedded into $\mathbb{R}^3$, and *non-orientable* otherwise. The *genus* of a graph G is the minimum $g \geq 0$, such that G can be drawn on a surface of genus g. Similarly, the *orientable* (resp. *non-orientable*) genus of a graph is the minimum genus of an orientable (resp. non-orientable) surface on which G can be drawn.

Drawing graphs on various surfaces is central to graph theory (e. g., [4, 13]), and of importance to topology and to mathematics in general (e. g., [17]), and have been the subject of intense study. Graphs of small genus are also of importance to computer science and engineering, since they can be used to model a wide variety of natural objects. For further background, we refer the reader to Gross and Tucker [4] for topological graph theory and to Hatcher [5] for algebraic topology.

Computing the genus of a graph exactly. It was shown by Thomassen that computing the orientable genus [15], and the non-orientable genus of a graph [16], are both NP-complete problems. Deciding whether a graph has genus 0, i. e., planarity testing, can be done in linear time by the seminal result of Hopcroft and Tarjan [6]. Filotti et al. [3] were the first to obtain an algorithm for computing a drawing of an n-vertex graph of genus g, on a minimum-genus surface, in time $n^{O(g)}$. Robertson and Seymour [14], as part of their Graph Minor project, gave an $O(f(g) \cdot n^3)$ time algorithm for determining the genus of a graph. This was improved by a breakthrough result of Mohar [11] who gave a linear-time algorithm for computing a minimum-genus drawing, for any fixed g. A relatively simpler linear-time algorithm has subsequently been obtained by Kawarabayashi, Mohar and Reed [7]. The running time of the above algorithms is exponential in g.

Approximating the genus of a graph. Perhaps surprisingly, the problem of approximating the genus of a graph is very poorly understood. In general, the genus of a graph can be as large as $\Omega(n^2)$ (e. g., for the complete graph K_n). By Euler's characteristic, there is a trivial $O(1)$-approximation for sufficiently dense graphs (i. e., of average degree at least $6+\varepsilon$, for some fixed $\varepsilon > 0$). For graphs of bounded degree, Chen, Kanchi, and Kanevsky [2] described a simple $O(\sqrt{n})$-approximation, which follows by the fact that graphs of small genus have small balanced vertex-separators. Following the present paper, Chekuri and Sidiropoulos [1] obtained a polynomial-time algorithm which, given a graph G of bounded degree and of genus g, outputs a drawing on a surface of genus $O(g^{12}\log^{19/2} n)$. Combined with the result of Chen et al., this implies a $n^{1/2-\alpha}$ approximation for bounded-degree graphs, for some constant $\alpha > 0$. Subsequently, Kawarabayashi and Sidiropoulos [8] obtained a polynomial-time algorithm that computes a drawing on a surface of genus $O(g^{256}\log^{189} n)$ for general graphs (that is, without the condition that the maximum degree is bounded). Combined with Euler's characteristic, this also implies a $O(n^{1-\beta})$-approximation for general graphs, for some constant $\beta > 0$. We also remark that better bounds are known for 1-apex graphs [12, 8] (that is, graphs that can be made planar by removing a single vertex).

Our results. We present a pseudo-approximation algorithm for the orientable genus of Hamiltonian graphs.[1] More specifically, we obtain a polynomial-time algorithm which, given a graph G and a Hamiltonian path P in G, computes a drawing of G on a surface of either orientable or non-orientable genus $O(g^7)$, where g is the orientable genus of G. The dependence of the running time is polynomial in

[1] We use the term *pseudo-approximation* to denote the fact that even though we approximate the orientable genus of the input graph, the output surface is allowed to be non-orientable. In this paper, we call a graph *Hamiltonian* if it has a Hamiltonian *path.*

both g and n. Combined with the simple $O(n/g)$-approximation described above, our result immediately gives a $O(n^{6/7})$-pseudo-approximation for the orientable genus of graphs with a known Hamiltonian path. These bounds are significantly stronger than the approximation guarantee achieved for general graphs [8]. We note that the algorithm for bounded-degree graphs given in [1] is not applicable in our setting, since we are dealing with graphs of unbounded degree. The following summarizes our main result.

Theorem 1.1. *There exists a polynomial-time algorithm, which, given a graph G of orientable genus g and a Hamiltonian path P in G, outputs a drawing of G on a surface of either orientable or non-orientable genus $O(g^7)$. The running time is polynomial in n and g.*

1.1 Overview

In this section, we present a very informal and somewhat imprecise overview of our algorithm. We are given a Hamiltonian graph G of genus g and a Hamiltonian path P in G. Our high-level approach is to *cover* the graph G by $O(g)$ subgraphs of constant genus, then independently draw each of them on a surface of small genus and finally combine all drawings. More precisely, our algorithm consists of the following steps:

Step 1: Cover G by $O(g)$ toroidal subgraphs $G_1,\dots,G_k$.

Step 2: For each pair of graphs G_i, G_j, compute a drawing of $G_i \cup G_j$.

Step 3: Combine the drawing of all pairs $G_i \cup G_j$, to obtain a drawing of G.

Step 1: Greedy band covering. The goal in this step is to cover G by $O(g)$ toroidal graphs. Fix an optimal drawing of G on a surface of genus g. Let us say that an edge $e \in E(G) \setminus E(P)$ is *global* if after contracting P, e becomes a noncontractible loop. Otherwise, we say that e is *local*. We can show that G can be covered by a collection of toroidal graphs $G_1^*,\dots G_k^*$, $k = O(g)$, such that every local edge appears in exactly one G_i^*, and every global edge appears in exactly two graphs G_i^*, G_j^*. Roughly speaking, we walk along the path, and we find maximal edge-disjoint subpaths $Q_1,\dots,Q_k$ in P, such that all edges on either side of each Q_i are homotopic (after contracting P). See Figure 1 for an example; for clarity, local edges are omitted from the figure.

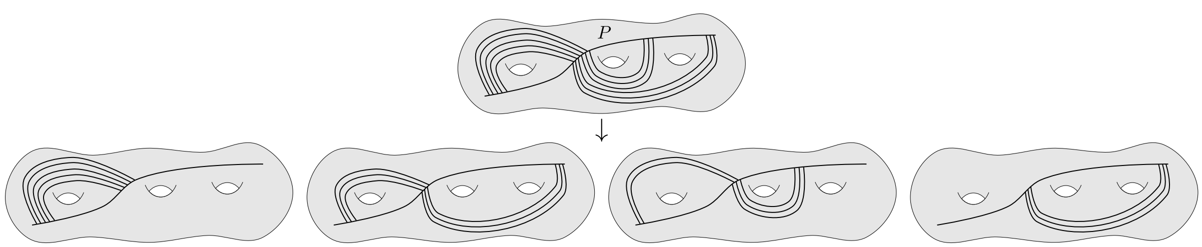

Figure 1: A greedy band covering.

Although, we do not have access to the optimal embedding of G, we still can compute such a covering $G_1,\dots,G_\ell$, for some $\ell = O(g)$. We refer to the graphs $G_1,\dots,G_\ell$ as *elementary bands*. Of course, this means that the elementary bands $G_1, \ldots, G_\ell$ we compute might differ from $G_1^*, \ldots, G_k^*$. This introduces certain complications as we explain in Section 2.

Step 2: Drawing a pair of bands. Even though we have $O(g)$ graphs G_i and each of them has genus at most 1, we cannot naively combine their drawings. Roughly speaking, the problem is that graphs G_i share global edges: the way an edge e is drawn in the drawing of G_i may be inconsistent with the way it is drawn in the drawing of G_j. This can happen because in Step 1, when we greedily compute the covering by elementary bands, we can only keep track of the local edges in the current band. As a result, when we try to combine the drawings of two bands, the local edges can introduce conflicts between the two drawings. Figure 2 depicts an example of such a conflict. In this example, the graph G is covered by two

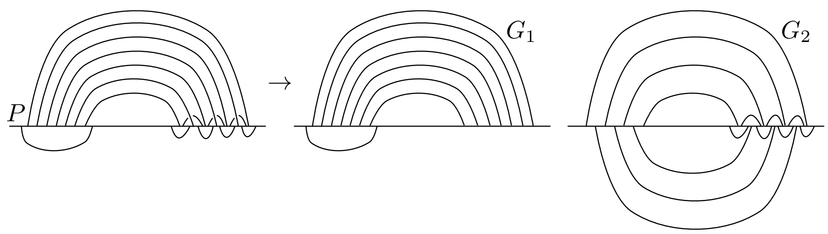

Figure 2: Edge conflicts in the greedy cover.

toroidal (in fact, planar) graphs G_1, G_2. However, in the drawing of G_1, all global edges are drawn on one side of P, while in the drawing of G_2 the global edges alternate between the two sides of P. We overcome this obstacle by showing that, roughly speaking, for every pair of bands G_i, G_j, there are drawings that are *nearly consistent*. This is a technical part of the paper, and we refer the reader to Section 5. We just note here that, at the high level, we decompose each band into $g^{O(1)}$ subgraphs, such that each subgraph has a certain structure allowing us to find a drawing on a surface of genus $g^{O(1)}$ in polynomial time.

Step 3: Combining the drawings of all pairs. In this last step we combine the drawings of all pairs of bands into a drawing of G. The key idea is that for every pair $G_i \cup G_j$, we can modify its drawing, such that the global edges that are shared between either G_i, or G_j, and some other elementary band G_ℓ, are contained in a small number of homotopy classes (after contracting P). Intuitively, this means that we can combine the drawings of all pairs by introducing a small number of "interface" handles between them. A precise definition appears in Section 7.

1.2 Why a pseudo-approximation?

Our current algorithm is a pseudo-approximation, which means that given a graph G of orientable genus g, it outputs a drawing on a surface of either orientable, or non-orientable genus $g^{O(1)}$. However, we believe that our approach can be generalized to obtain a (true) approximation algorithm. That is, given a graph of genus g, compute a drawing on a surface of genus $g^{O(1)}$, with the same orientability type as G. Doing so, requires an extension of the definition of an *elementary band*, to account for graphs that can be drawn on the projective plane (i. e., elementary bands of non-orientable genus at most 1). This small modification causes the number of different types of elementary bands to grow by a constant factor. Unfortunately, as a consequence, the (already lengthy) case analysis in the proof becomes dauntingly long. The rest of our proof remains essentially unchanged. We are not aware of a way to simplify this case analysis, so in the interest of clarity we have decided to omit it from the present paper.

1.3 Organization

In Section 2, we show how to find the elementary band covering. The proof of the main technical step required for Section 2 is given in the subsequent two sections. Section 3 introduces the main notions used in the proof, and Section 4 presents the main case analysis required for the proof. In Section 5, we explain how to combine two drawings, by decomposition into smaller subgraphs; the drawing of these subgraphs is described in Section 6. In Section 7, we put the all pieces of our algorithm together, and show how to find a drawing of a Hamiltonian graph.

2 Band coverings of Hamiltonian graphs

In this section, we describe a greedy algorithm that partitions the input graph to a set of $O(g)$ *simple* toroidal graphs, which we call *bands*. One of the tools that we have developed for this section is based on the notion of ribbons and petals. Intuitively, ribbons and petals describe minimal topological subspaces that contain the edges of a band.

Definition 2.1 (Bands in Hamiltonian graphs)**.** Let G be a graph, and let P be a Hamiltonian path in G. Let $B \subseteq E(G) \setminus E(P)$. Then, B is called a *band*. Let Q be a subpath of P, such that every edge $e \in B$ has at least one endpoint in $V(Q)$. We also say that B has *spine* P, and *primary segment* Q. An edge in B is called *global* if exactly one of its endpoints is in $V(Q)$; it is called *local* if both of its endpoints are in $V(Q)$. Note that every set $B \subseteq E(G) \setminus E(P)$ is a band with spine P, and primary segment P.

Definition 2.2 (Elementary bands)**.** Let G be a graph, and let P be a Hamiltonian path in G. Let $B \subseteq E(G) \setminus E(P)$ be a band with spine P, and primary segment $Q \subseteq P$. We define the following types of bands, which we refer to as *elementary*.

Type-1. We say that B is of *type-1* (see Figure 3(a)) if the following conditions are satisfied. There exist subpaths $P_1, P_2, P_3, P_4 \subset P$, with P_1, P_2, P_3, P_4, Q being pairwise edge-disjoint, and such that the set M of global edges in B can be decomposed into $M = M_1 \cup M_2 \cup M_3 \cup M_4$, such that for every $i \in \{1, \dots, 4\}$, every edge in M_i has one endpoint in $V(Q)$, and one in $V(P_i)$. Moreover, there exists a planar drawing φ of the graph $H = B \cup Q \cup P_1 \cup \dots \cup P_4$, satisfying the following. For every $i \in \{1, \dots, 4\}$, $\varphi(P_i)$ lies in the outer face of φ, and all the curves $\varphi(e)$, $e \in M_i$, are attached to the same side of $\varphi(P_i)$. We say that the paths $P_1, \dots, P_4$ are the *outlets* of B.

Type-2. We say that B is of *type-2* (see Figure 3(b)) if the following conditions are satisfied. There exist subpaths $P_1, P_2, P_3 \subset P$, with P_1, P_2, P_3, Q being pairwise edge-disjoint, and such that the set M of global edges in B can be decomposed into $M = M_1 \cup M_2 \cup M_3$, such that for every $i \in \{1, 2, 3\}$, every edge in M_i has one endpoint in $V(Q)$, and one in $V(P_i)$. Moreover, there exists a planar drawing φ of the graph $H = B \cup Q \cup P_1 \cup P_2 \cup P_3$, satisfying the following. If we denote by φ' the

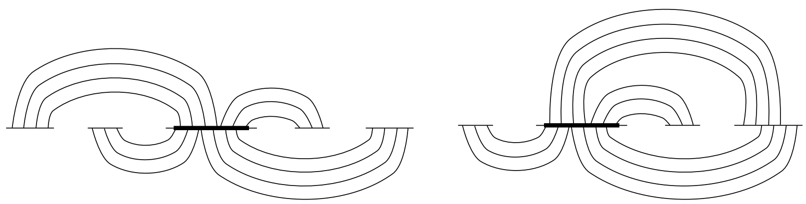

(a) Type-1. (b) Type-2.

Figure 3: The different types of elementary bands. The primary segment is in bold.

drawing induced by φ on $H \setminus V(P_3)$, then for every $i \in \{1,2\}$, $\varphi'(P_i)$ lies in the outer face. Also, there exists $v \in V(P_3)$, such that $\varphi(v)$ also lies in the outer face. Moreover, for every $i \in \{1,2\}$, all the curves $\varphi(e)$, $e \in M_i$, are attached to the same side of $\varphi(P_i)$. Note that the curves $\varphi(e)$, $e \in M_3$ are allowed to be attached to both sides of $\varphi(P_3)$. We say that the paths P_1, and P_2 are the *outlets* of B, and that the path P_3 is the *double outlet* of B.

We say that φ is a *canonical drawing* of B (or a canonical drawing of H, when B is clear from the context). For an elementary band of type-2, we can pick φ so that the curves $\varphi(Q)$, $\varphi(P_1)$, $\varphi(P_2)$, $\varphi(P_3)$ become segments of a horizontal line ℓ, with $\varphi(Q)$ appearing to the left of $\varphi(P_3)$. We call ℓ the *canonical line* of φ. Let $\prec$ be the total ordering of $V(Q) \cup V(P_3)$ induced by a left-to-right traversal of ℓ. We say that $\prec$ is the *canonical ordering* of φ. We extend $\prec$ to B as follows. Let $\{x,y\},\{x',y'\} \in B$, with $x,x' \in V(Q)$. Then, we define $\{x,y\} \prec \{x',y'\}$ if either $x \prec x'$, or $(x = x') \wedge (y \prec y')$. Since G does not contain any parallel edges, it follows that $\prec$ is a total ordering of B.

We remark that a band can be of both type-1 and type-2 (e. g., the trivial band). Figure 3 depicts examples of type-1 and type-2 elementary bands.

Definition 2.3 (Band coverings for Hamiltonian graphs)**.** Let G be a graph, and let P be a Hamiltonian path in G. A *band covering with spine P for G* is a collection $\mathcal{B} = \{(B_i, Q_i)\}_{i=1}^{t}$ satisfying the following conditions:

1. For every $i \in \{1,\ldots,t\}$, B_i is an elementary band with spine P, and primary segment Q_i.
2. $\bigcup_{i \in \{1,\ldots,t\}} B_i = E(G) \setminus E(P)$.
3. For every $i \neq j \in \{1,\ldots,t\}$, we have $V(Q_i) \cap V(Q_j) = \emptyset$.

We remark that every edge in $E(G) \setminus E(P)$ is contained in at least one, and at most two bands in the band covering $\mathcal{B}$. If it is contained in exactly one band, then we say that it is *local*, and otherwise we say that it is *global*.

First we show that a band covering that is composed of $O(g)$ bands exists. The intuition behind this proof comes from looking at the homotopy classes of the edges in $E(G) \setminus P$ in an optimal embedding. Moreover, we show that we can compute a band covering of size $O(g)$. Note that this band covering that we compute is not necessarily optimal. However, its size is within a constant factor of the optimal band covering size. The main tool that our algorithm uses is a so-called *ribbon-petal covering*. See Section 3 for the description of ribbon petal covering and the proof of the following lemmas.

Lemma 2.4 (Existence of a small band covering)**.** *Let G be a graph of orientable genus g, and P a Hamiltonian path in G. Then, there exists a band covering $\mathcal{B} = \{(B_i, Q_i)\}_{i=1}^{t}$, with spine P for G, with $t = O(g)$.*

The proof of the above Lemma is rather lengthy, and is presented in the next two sections. Assuming Lemma 2.4, we can now prove the main result of this section.

Lemma 2.5 (Computing a small band covering)**.** *Let G be a graph of orientable genus g, and P a Hamiltonian path in G. Then, given G and P, we can compute in polynomial time a band covering for G with spine P, of size $O(g)$.*

Proof. Fix an ordering $\prec$ of $V(P)$, induced by a traversal of P. We inductively define a partition of P into vertex disjoint subpaths $Q_1, \dots, Q_t$ as follows. Let $Q_0 = \emptyset$. For $i \geq 1$, given $Q_0, \dots, Q_{i-1}$, we inductively define Q_i to be the maximal prefix (with respect to $\prec$) of $P \setminus \bigcup_{j=0}^{i-1} Q_i$, such that if we set

$$B_i = \{\{u, v\} \in E(G) \setminus E(P) : \{u, v\} \cap V(Q_i) \neq \emptyset\},$$

then B_i is an elementary band with spine P, and primary segment Q_i.

Finally, we set t to be the smallest integer such that $V(P) \setminus \bigcup_{j=1}^{t} V(Q_j) = \emptyset$. It is straightforward to check that $\mathcal{B}$ is a band covering for G with spine P. It is also clear that $\mathcal{B}$ can be computed in polynomial time.

It remains to bound $|\mathcal{B}| = O(g)$. By Lemma 2.4 there exists a band covering $\mathcal{B}' = \{(B'_i, Q'_i)\}_{i=1}^{t'}$ for G, with $t' = O(g)$. Since for every $i \in \{1, \dots, t\}$ we pick a maximal prefix Q_i such that B_i is an elementary with spine P, and primary segment Q_i, it follows that $\bigcup_{j=1}^{i} V(Q_i) \supseteq \bigcup_{j=1}^{i} V(Q'_i)$. Therefore, $t \leq t' = O(g)$, as required. □

3 Ribbons and petals

In this section we define a decomposition of an embedded graph into a small collection of topologically simpler subgraphs.

Definition 3.1 (Ribbon)**.** Let γ be a simple open curve in a surface $\mathcal{F}$. A set $A \subset \mathcal{F}$ is called a *γ-ribbon* if it satisfies one of the following conditions:

1. The set A is the image of a simple curve with endpoints $x, y \in \gamma$, and such that $A \cap \gamma = \{x, y\}$.

2. Intuitively, A is a deformed triangle in $\mathcal{F}$ that intersects γ only on a vertex and its opposite edge. Formally, let T be 2-simplex, let a be a vertex of T, and let ℓ be the edge of T opposite to a. Then, there exists a continuous mapping $f : T \to \mathcal{F}$ such that f is a homeomorphism on $T \setminus \{a\}$, and on ℓ. Moreover,
$$f(a \cup \ell) \subseteq \gamma, \qquad f(T \setminus (a \cup \ell)) \cap \gamma = \emptyset, \quad \text{and} \quad f(T) = A.$$

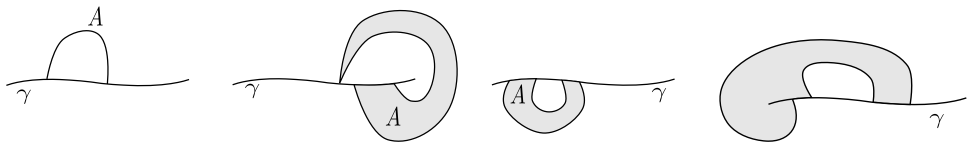

(a) Examples of sets A that are γ-ribbons. (b) Not a γ-ribbon.

Figure 4: Ribbons.

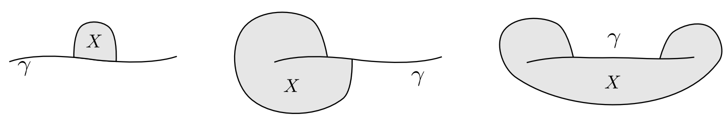

Figure 5: From left to right: A single γ-petal, a double γ-petal, and a triple γ-petal X.

3. Intuitively, A is a deformed rectangle in $\mathcal{F}$ that intersects γ only on two opposite edges. Formally, there exists an $f : [0,1]^2 \to \mathcal{F}$ be a continuous mapping such that f is a homeomorphism on $(0,1) \times [0,1]$, on $\{0\} \times [0,1]$, and on $\{1\} \times [0,1]$. Moreover,

$$f((0,1) \times [0,1]) \cap \gamma = \emptyset\,, \qquad f(\{0,1\} \times [0,1]) \subseteq \gamma, \quad \text{and} \quad f([0,1]^2) = A\,.$$

We say that the points $f((0,0))$, $f((0,1))$, $f((1,0))$, $f((1,1))$ are *endpoints* of A. Figure 4(a) depicts examples of ribbons.

Definition 3.2 (Petal)**.** Let γ be a simple open curve in a surface $\mathcal{F}$. A set $X \subset \mathcal{F}$ is called a *γ-petal* if there exists a continuous mapping $f : [0,1]^2 \to \mathcal{F}$, so that f is a homeomorphism on $[0,1] \times (0,1]$, with

$$f((0,1] \times [0,1]) \cap \gamma = \emptyset\,, \qquad f(\{0\} \times [0,1]) \subseteq \gamma, \quad \text{and} \quad f([0,1]^2) = X\,.$$

We distinguish between the following three types of petals:

Single. The γ-petal X is called *single* if f is a homeomorphism on $[0,1]^2$.

Double. The γ-petal X is called *double* if it is not single, and if there exists $x \in [0,1]$, so that f is a homeomorphism on $\{0\} \times [0,x]$, and on $\{0\} \times [x,1]$.

Triple. The γ-petal X is called *triple* if it is not single, and it is not double, and if there exist $x < y \in [0,1]$, so that f is a homeomorphism on $\{0\} \times [0,x]$, on $\{0\} \times [x,y]$, and on $\{0\} \times [y,1]$.

We say that the points $f((0,0))$, and $f((0,1))$ are *endpoints* of X. Figure 5 depicts examples of the three different types of petals.

Definition 3.3 (Ribbon-petal covering)**.** Let G be a graph of genus g, and let P be a Hamiltonian path in G. Let φ be a drawing of G on a surface $\mathcal{S}$ of genus g. Let $\mathcal{X}$ be a collection of subsets of $\mathcal{S}$. Then, we say that $\mathcal{X}$ is a *ribbon-petal covering* for (G,P,φ), if the following conditions are satisfied:

1. Every $X \in \mathcal{X}$ is either a $\varphi(P)$-ribbon, or a $\varphi(P)$-petal.
2. For every $X, X' \in \mathcal{X}$, with $X \neq X'$, we have $X \cap X' \subseteq \varphi(P)$.
3. For every $e \in E(G) \setminus E(P)$, there exists $X \in \mathcal{X}$, such that $\varphi(e) \subseteq X$.

Lemma 3.4 (Malnič and Mohar [10]). *Let $\mathcal{M}$ be an orientable surface of genus g, and let $x \in \mathcal{M}$. Let $\mathcal{X}$ be a collection of noncontractible simple loops such that for every $C, C' \in \mathcal{X}$, we have $C \cap C' = x$. Suppose that all curves in $\mathcal{X}$ are pairwise nonhomotopic with respect to the basepoint x (i. e., in $\pi_1(\mathcal{M}, x)$). Then, $|\mathcal{X}| \leq 6g - 3$.*

Lemma 3.5. *Let G be a graph of genus g, and let P be a Hamiltonian path in G. Let φ be a drawing of G on a surface $\mathcal{S}$ of genus g. For every $e \in E(G) \setminus E(P)$, let P_e be the subpath of P between the endpoints of e, and define the cycle $C_e = P_e \cup \{e\}$; let also γ_e be the loop obtained from $\varphi(C_e)$ by contracting $\varphi(P)$ to the basepoint. Let $E^* = \{e \in E(G) \setminus E(P) : \gamma_e$ is non-contractible in $\mathcal{S}\}$. Let $X \subseteq E^*$, such that for every $e \neq e' \in X$, the cycles γ_e and $\gamma_{e'}$ are non-homotopic. Then, $|X| \leq 6g - 3$.*

Proof. Let $\mathcal{S}'$ be the surface obtained from $\mathcal{S}$ by contracting $\varphi(P)$ to a single point. Let G' be the graph obtained from G by contracting P to a single vertex p, and thus G' is drawn on a surface of genus g. The assertion now follows immediately by Lemma 3.4. □

Lemma 3.6 (Existence of small ribbon-petal coverings). *Let G be a graph of genus g, and let P be a Hamiltonian path in G. Let φ be a drawing of G on a surface $\mathcal{S}$ of genus g. Then, there exists a ribbon-petal covering $\mathcal{X}$ for (G, P, φ), with $|\mathcal{X}| \leq 24g - 12$.*

Proof. For every $e \in E(G) \setminus E(P)$, let P_e be the subpath of P between the endpoints of e, and define the cycle $C_e = P_e \cup \{e\}$; let also γ_e be the loop obtained from $\varphi(C_e)$ by contracting $\varphi(P)$ to the basepoint. Let

$$E^* = \{e \in E(G) \setminus E(P) : \gamma_e \text{ is non-contractible in } \mathcal{S}\}.$$

Consider the partition $E^* = Y_1 \cup \ldots \cup Y_k$, such that for every $i \in \{1, \ldots, k\}$, for every $e, e' \in Y_i$, the loops γ_e and $\gamma_{e'}$ are homotopic, and for every $i \neq j \in \{1, \ldots, k\}$, for every $e \in Y_i$, $e' \in Y_j$, the loops γ_e and $\gamma_{e'}$ are non-homotopic. By Lemma 3.5, we have $k \leq 6g - 3$.

Let $i \in \{1, \ldots, k\}$. Since for all $e \in Y_i$, all the loops γ_e, $e \in Y_i$ are homotopic, it follows that there exists an ordering $e_1, \ldots, e_{k_i}$ of the edges in Y_i, and a continuous mapping $f_i : [0,1]^2 \to \mathcal{S}$, with $f_i([0,1] \times \{0,1\}) \subseteq \varphi(P)$, and moreover for every $e_j \in Y_i$, there exists $x_j \in [0,1]$, so that $f_i(x_j \times [0,1]) = \varphi(e_j)$. Since for every $i' \neq i \in \{1, \ldots, k\}$ and for every $e \in Y_i$, $e' \in Y_{i'}$, the loops γ_e and $\gamma_{e'}$ are non-homotopic, it follows that we can pick the maps $f_1, \ldots, f_k$ so that for every $i \neq i' \in \{1, \ldots, k\}$ the sets $f_i([0,1]^2)$ and $f_{i'}([0,1]^2)$ have disjoint interiors. It therefore suffices to show how every set $A_i = f_i([0,1]^2)$ can be decomposed into at most three $\varphi(P)$-ribbons with disjoint interiors. To that end, we consider the following cases: Let a, b be the two endpoints of $\varphi(P)$.

1. If $f_i([0,1] \times \{0\})$, and $f_i([0,1] \times \{1\})$ are both single points, then the set $f_i([0,1]^2)$ is clearly a $\varphi(P)$-ribbon.

2. If $f_i([0,1]\times\{0\})$ is a single point, and $f_i([0,1]\times\{1\})$ is not a single point, then we can pick f_i so that there exist at most two values $x<y\in[0,1]$ such that for every $z\in[0,1]\setminus\{x,y\}$, we have $f((z,1))\notin\{a,b\}$. It follows that the sets $f_i([0,x]\times[0,1])$, $f_i([x,y]\times[0,1])$, $f_i([y,1]\times[0,1])$ are the desired $\varphi(P)$-ribbons.

3. If $f_i([0,1]\times\{0\})$ is not a single point, and $f_i([0,1]\times\{1\})$ is a single point, we can decompose $f_i([0,1]^2)$ into at most three $\varphi(P)$-ribbons as in the previous case.

4. If $f_i([0,1]\times\{0\})$ is not a single point, and $f_i([0,1]\times\{1\})$ is not a single point, then we can pick f_i so that there exist at most two values $x<y\in[0,1]$ such that for every $z\in[0,1]\setminus\{x,y\}$, we have $f((z,1))\notin\{a,b\}$, and $f((z,0))\notin\{a,b\}$. It follows that the sets $f_i([0,x]\times[0,1])$, $f_i([x,y]\times[0,1])$, $f_i([y,1]\times[0,1])$ are the desired $\varphi(P)$-ribbons.

We perform the above decomposition to each set $f_i([0,1]^2)$, $i\in\{1,\dots,k\}$.

For every $e\in E(G)\setminus(E(P)\cup E^*)$, the cycle $\varphi(C_e)$ is contractible in $\mathcal{S}$. Therefore, it bounds a disk $D_e\subset\mathcal{S}$, and this disk is a $\varphi(P)$-petal. Let $e,e'\in E^*$. Since $\varphi(e)\cap\varphi(e')\subset\varphi(P)$, we have that either $D_e\subset D_{e'}$, or $D_{e'}\subset D_e$, or $D_e\cap D_{e'}\subset\varphi(P)$. It follows that we can cover all the disks $\{D_e\}_{e\in E(G)\setminus E^*}$ by using at most one $\varphi(P)$-petal between every two consecutive ribbons on every side of $\varphi(P)$, and possibly one more $\varphi(P)$-petal for every one of the two endpoints of $\varphi(P)$. Since every ribbon intersects $\varphi(P)$ in at most two segments, in total we obtain a collection of at most k petals satisfying the assertion. □

4 From ribbons and petals to bands

In this section we show that a ribbon-petal covering can be used to cover the graph with a small number of elementary bands.

Lemma 4.1. *Let G be a graph, and let P be a Hamiltonian path in G. Let $v\in V(P)$, and let*

$$B=\{\{x,y\}\in E(G)\setminus E(P): x=v, \text{ or } y=v\}.$$

Then, B is an elementary band of type-1, with spine P, and primary segment v.

Proof. Let R_1, R_2 be the two connected components of $P-v$, where any of R_1, and R_2, can be empty. Then, it is immediate to check that B is an elementary band of type 1, with primary segment v, by setting in Definition 2.2(Type-1) $P_1=P\setminus R_2$, and $P_2=R_2$. □

Lemma 4.2 (From ribbons and petals to bands)**.** *Let G be a graph, and let P be a Hamiltonian path in G. Let φ be a drawing of G on an orientable surface $\mathcal{S}$. Let $\mathcal{X}$ be a ribbon-petal covering for (G,P,φ). Let A be the set of all points that are endpoints of all the ribbons, and all the petals in $\mathcal{X}$. Recall that all these points lie in $\varphi(P)$. Let p,q be the endpoints of the curve $\varphi(P)$. Let $p_1,\dots,p_t$ be the ordering of the points in $A\cup\{p,q\}$ induced by a traversal of $\varphi(P)$. Let $i\in\{1,\dots,t-1\}$, and let α be the segment of $\varphi(P)$ between p_i, and p_{i+1}. Let*

$$B=\{e\in E(G)\setminus E(P):\varphi(e)\cap\alpha\neq\emptyset\}.$$

Let also Q be the maximal subpath of P such that $\varphi(Q)\subseteq\alpha$. Then, there exists $k\leq 3$, so that the following conditions are satisfied:

1. *There exist pairwise vertex-disjoint subpaths $Q_1,\dots,Q_k \subseteq Q$, such that $V(Q) = \bigcup_{i=1}^{k} V(Q_i)$.*

2. *There exist pairwise disjoint subsets $B_1,\dots,B_k \subseteq B$, such that for any $j \in \{1,\dots,k\}$, B_j is an elementary band with spine P, and with primary segment Q_j.*

Proof. For any pair of points $x,y \in \varphi(P)$, let $\lambda[x,y]$ be the segment of $\varphi(P)$ between x, and y. Let also $\lambda[x,y) = \lambda[x,y] \setminus \{y\}$, $\lambda(x,y] = \lambda[x,y] \setminus \{x\}$, and $\lambda(x,y) = \lambda[x,y] \setminus \{x,y\}$. Let $P[x,y]$ denote the maximal subpath $P' \subseteq P$, such that $\varphi(P')$ is contained in $\lambda[x,y]$. Let $B[x,y]$ be the subset of edges in B with at least one endpoint in $V(P([x,y]))$. Define similarly $P[x,y)$, $P(x,y]$, $P(x,y)$, $B[x,y)$, $B(x,y]$, and $B(x,y)$.

Let $\prec$ be the total ordering of the points in the curve $\varphi(P)$ induced by a traversal of $\varphi(P)$, with $p_1 \prec \cdots \prec p_t$.

We consider the following cases. (For the sake of clarity, in the accompanying figures, we only draw the global edges of the resulting bands.)

Case 1: There exists a double $\varphi(P)$-petal $X \in \mathcal{X}$, such that both sides of α are in X. Assume w.l.o.g. that p_i is an endpoint of $\varphi(P)$, and p_{i+1} is an endpoint of X. Let also r be the other endpoint of X.

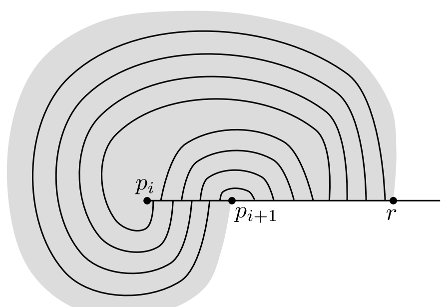

Then, B is an elementary band of type-1, with primary segment Q, and outlet $P[p_{i+1},r]$.

Case 2: There exists a triple $\varphi(P)$-petal $X \in \mathcal{X}$, such that both sides of α are in X. Assume w.l.o.g. that p_i is an endpoint of $\varphi(P)$, and p_{i+1} is an endpoint of X. Let also r be the other endpoint of $\varphi(P)$.

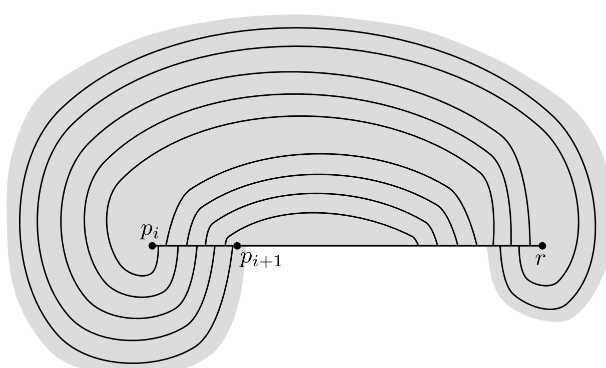

Then, B is an elementary band of type-2, with primary segment Q, and double outlet $P[p_{i+1},r]$.

Case 3: There exists a $\varphi(P)$-petal $X \in \mathcal{X}$, and a single $\varphi(P)$-petal $Y \in \mathcal{X}$, such that one side of α is in X, and the other is in Y. The cases where X is a single, or a double $\varphi(P)$-petal are special sub-cases of the case where X is a triple $\varphi(P)$-petal. It therefore suffices to consider the latter case. Let p_j, p_l, with $j \leq i < i+1 \leq l$ be the endpoints of X.

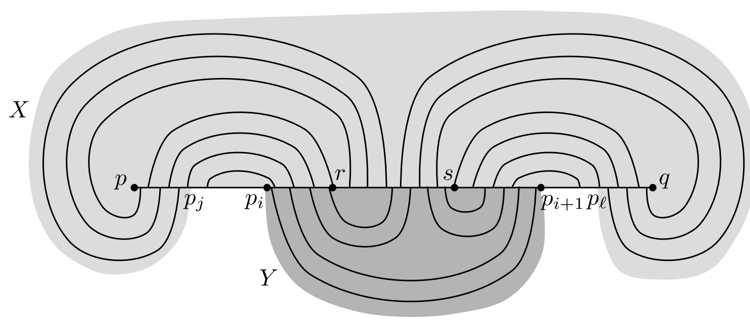

Let $r \in V(P[p_i, p_{i+1}])$ be minimal with respect to $\prec$ such that all edges in $B(r, p_{i+1}]$ that have an endpoint in $P[p_1, p_j]$, are incident to $P[p, p_j]$ on the same side as Y is incident to P. Similarly, let $s \in V(P[p_i, p_{i+1}])$ be maximal with respect to $\prec$ such that all edges in $B(p_i, s]$ that have an endpoint in $P[p_l, q]$, are incident to $P[p_l, q]$ on the same side as Y is incident to P. Then, $B[p_i, r)$ is an elementary band of type-1, with primary segment $Q[p_i, r)$, and outlets $P[p, p_i]$, and $P[r, p_{i+1}]$. Also, $B(r, s)$ is an elementary band of type-1, with primary segment $Q(r, s)$, and outlets $P[p, p_i]$, $P[p_i, r]$, $P[s, p_{i+1}]$, and $P[p_{i+1}, q]$. Also, $B(s, p_{i+1})$ is an elementary band of type-1, with primary segment $Q(s, p_{i+1})$, and outlets $P[p_i, s]$, and $P[p_{i+1}, q]$.

Case 4: There exists a triple $\varphi(P)$-petal $X \in \mathcal{X}$, and a $\varphi(P)$-ribbon $Y \in \mathcal{X}$, such that one side of α is in X, and the other is in Y. This case is identical to Case 3.

Case 5: There exists either a single, or a double $\varphi(P)$-petal $X \in \mathcal{X}$, and a $\varphi(P)$-ribbon $Y \in \mathcal{X}$, such that one side of α is in X, and the other is in Y. The case where X is a single $\varphi(P)$-petal is a special case of the case where X is a double $\varphi(P)$-petal. We can therefore assume w.l.o.g. that X is a double $\varphi(P)$-petal. Moreover, the case were both X and Y are double $\varphi(P)$-petals is a special case of this one. If Y is attached only to one side of $\varphi(P)$, then this case is identical to Case 4 above. Therefore, it remains to consider the case where Y is attached to both sides of $\varphi(P)$. Let $a, a', b, b' \in \mathcal{S}$ be the endpoints of Y, and $r, r' \in \mathcal{S}$ be the endpoints of X. We may assume w.l.o.g. that either

$$r \preceq a \preceq a' \preceq r' \preceq b \preceq b', \qquad \text{or} \qquad r \preceq a \preceq r' \preceq a' \preceq b \preceq b'.$$

Let s be the minimal point in $\{a', r'\}$ with respect to $\prec$, and let s' be the maximal one.

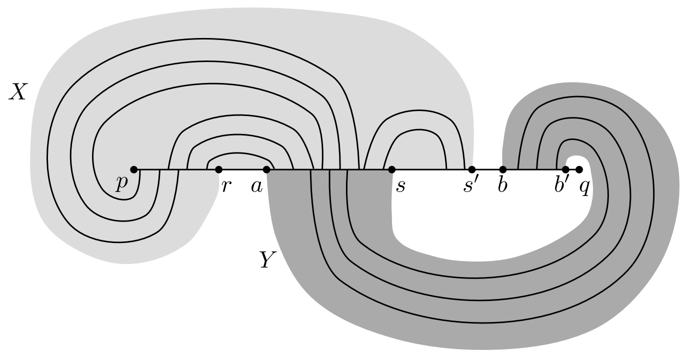

Then, B is an elementary band of type-2, with primary segment $P[a,s]$, outlets $P[s,s']$, and $P[b,b']$, and with double outlet $P[p,a]$.

Case 6: There exists a $\varphi(P)$-ribbon $X \in \mathcal{X}$, such that both sides of α are in X. Let $a,a',b,b' \in \mathcal{S}$ be the endpoints of X. We may assume w.l.o.g. that either

$$a \preceq b \preceq a' \preceq b', \qquad \text{or} \qquad b \preceq a \preceq a' \preceq b'.$$

Let r be the minimal in $\{a,b\}$, and r' be the maximal in $\{a,b\}$. Let also s be the minimal in $\{a',b'\}$, and s' be the maximal in $\{a',b'\}$.

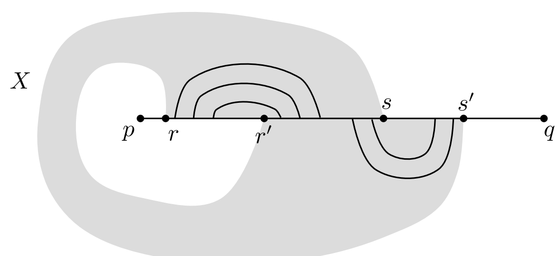

Then, B is an elementary band of type-1, with primary segment $P[r',s]$, and outlets $P[r,r']$, and $P[s,s']$.

Case 7: There exists a $\varphi(P)$-ribbon $X \in \mathcal{X}$, and a $\varphi(P)$-ribbon $Y \in \mathcal{X}$, such that one side of α is in X, and the other is in Y. The $\varphi(P)$-ribbon X is attached to a single side of $\varphi(P)$, and the $\varphi(P)$-ribbon Y is also attached to a single side of $\varphi(P)$. Let $a_1,\dots,a_4 \in \mathcal{S}$ be the endpoints of X, and let $b_1,\dots,b_4 \in \mathcal{S}$ be the endpoints of Y. Assume w.l.o.g. that $a_1 \preceq a_2 \preceq a_3 \preceq a_4$, and $b_1 \preceq b_2 \preceq b_3 \preceq b_4$.

We consider the following subcases:

Case 7-1: Suppose that $a_2 \prec b_1$, and $a_4 \prec b_3$. Let r be the minimal with respect to $\prec$ in $\{b_1,a_3\}$, and let r' be maximal in $\{b_1,a_3\}$. Similarly, let s be the minimal in $\{b_2,a_4\}$, and let s' be the maximal in $\{b_2,a_4\}$.

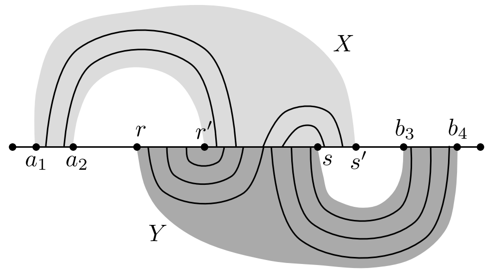

Then, B is an elementary band of type-1, with primary segment $P[r',s]$, and with outlets $P[a_1,a_2]$, $P[r,r')$, $P(s,s']$, and $P[b_3,b_4]$.

Case 7-2: Suppose that $a_2 \prec b_3$, and $b_2 \prec a_3$. Let r be the minimal with respect to $\prec$ in $\{a_1,b_1\}$, and let r' be maximal in $\{a_1,b_1\}$. Similarly, let s be the minimal in $\{a_2,b_2\}$, and let s' be the maximal in $\{a_2,b_2\}$. Let z be the minimal in $\{a_3,a_4,b_3,b_4\}$, and let z' be the maximal in $\{a_3,a_4,b_3,b_4\}$. Assume w.l.o.g. that $p_i = r'$, and $p_{i+1} = s$.

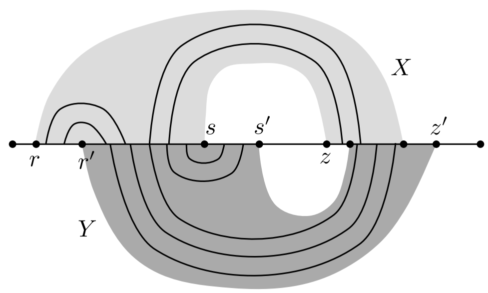

Then, B is an elementary band of type-2, with primary segment $P[r',s]$, with outlets $P[r,r')$, $P(s,s']$, and with double outlet $P[z,z']$.

Case 7-3: Suppose that $b_1 \preceq a_2 \preceq a_3 \preceq b_2 \preceq b_3 \preceq a_4$. Let r be the minimal with respect to $\prec$ in $\{a_1,b_1\}$, and let r' be maximal in $\{a_1,b_1\}$. Similarly, let s be the minimal in $\{a_4,b_4\}$, and let s' be the maximal in $\{a_4,b_4\}$. If $p_i = r'$, and $p_{i+1} = a_2$, then this configuration can be handled in the same way as case of Case 7-2. So it suffices to consider the case $p_i = a_3$, and $p_{i+1} = b_2$. Let z be the maximal with respect to $\prec$ vertex in $P[a_3,b_2]$, such that z has an incident edge $e = \{z,w\}$, with $w \in V(P[r,a_2])$.

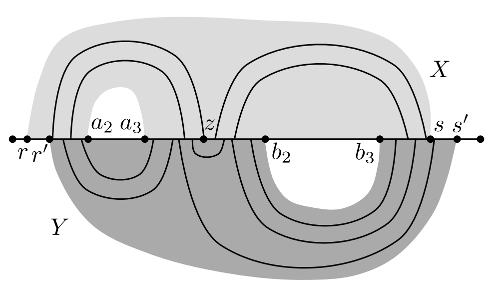

We can decompose B into three elementary bands as follows. First, $B[a_3,z)$ is an elementary band of type-2, with primary segment $P[a_3,z)$, outlet $P[z,s']$, and double outlet $P[r,a_3)$. Next,

$B[z,z]$ is an elementary band of type-1 (by Lemma 4.1), with primary segment z. Finally, $B(z,b_2]$ is an elementary band of type-2, with primary segment $P(z,b_2]$, outlet $P[r,a_3]$, and double outlet $P(b_2,s']$.

Case 8: There exists a $\varphi(P)$-ribbon $X \in \mathcal{X}$, and a $\varphi(P)$-ribbon $Y \in \mathcal{X}$, such that one side of α is in X, and the other is in Y. The $\varphi(P)$-ribbon X is attached to a single side of $\varphi(P)$, and the $\varphi(P)$-ribbon Y is attached to both sides of $\varphi(P)$. Let $a_1,\ldots,a_4 \in \mathcal{S}$ be the endpoints of X, and let $b_1,\ldots,b_4 \in \mathcal{S}$ be the endpoints of Y. Assume w.l.o.g. that $a_1 \preceq a_2 \preceq a_3 \preceq a_4$, and $b_1 \preceq b_2 \preceq b_3 \preceq b_4$.

We consider the following sub-cases:

Case 8-1: Suppose that $X \cap Y$ is a contiguous segment of $\varphi(P)$, and that the two segments where Y intersects $\varphi(P)$, are disjoint. Assume w.l.o.g. that $X \cap Y = \lambda[a_3,a_4] \cap \lambda[b_1,b_2]$. Let r be the minimal with respect to $\prec$ in $\{a_3,b_1\}$, and let r' be the maximal in $\{a_3,b_1\}$. Similarly, let s be the minimal in $\{a_4,b_2\}$, and let s' be the maximal in $\{a_4,b_2\}$.

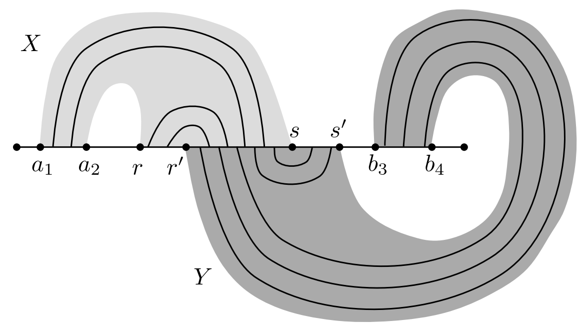

Then, B is an elementary band of type-1, with primary segment $P[r',s]$, and with outlets $P[a_1,a_2]$, $P[r,r')$, $P(s,s']$, and $P[b_3,b_4]$.

Case 8-2: Suppose that $X \cap Y$ consists of two disjoint contiguous segments of $\varphi(P)$, and that the two segments where Y intersects $\varphi(P)$, are disjoint. Let r be the minimal with respect to $\prec$ in $\{a_1,b_1\}$, and let r' be the maximal in $\{a_1,b_1\}$. Similarly, let s be the minimal in $\{a_4,b_2\}$, and let s' be the maximal in $\{a_4,b_2\}$. Assume w.l.o.g. that $X \cap Y = \lambda[r',a_2] \cup \lambda[a_3,s]$. We may further assume w.l.o.g. that $p_i = a_3$, and $p_{i+1} = s$, since all remaining cases can be handled in the exact same way.

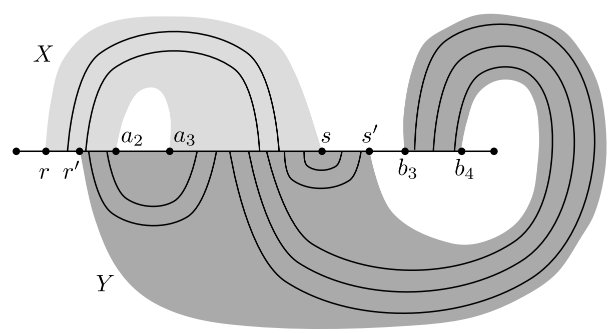

Then, B is an elementary band of type-2, with primary segment $P[a_3,s]$, and with outlets $P(s,s']$, $P[b_3,b_4]$, and with double outlet $P[r,a_3)$.

Case 8-3: Suppose that $X \cap Y$ is a contiguous segment of $\varphi(P)$, and that the two segments where Y intersects $\varphi(P)$, are not disjoint. We can assume w.l.o.g. that $a_1 \preceq a_2 \preceq a_3 \preceq b_1 \preceq a_4 \preceq b_2 \preceq b_3 \preceq b_4$, since all other cases can be handled in the exact same way.

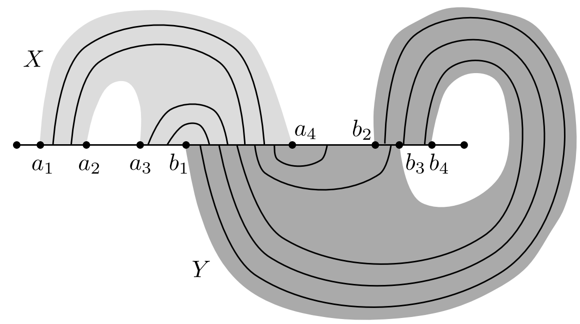

Then, B is an elementary band of type-2, with primary segment $P[b_1,a_4]$, and with outlets $P[a_1,a_2]$, $P[a_3,b_1)$, and double outlet $P(a_4,b_4]$. We remark that the outlets $P[a_1,a_2]$ and $P[a_3,b_1]$ can also be merged into a single outlet.

Case 8-4: Suppose that $X \cap Y$ consists of two disjoint contiguous segments of $\varphi(P)$, and that the two segments where Y intersects $\varphi(P)$, are not disjoint. We can assume w.l.o.g. that $a_1 \preceq b_1 \preceq a_2 \preceq b_2 \preceq b_3 \preceq a_3 \preceq a_4 \preceq b_4$, and that $p_i = b_1$, $p_{i+1} = a_2$, since all other cases can be handled in the exact same way.

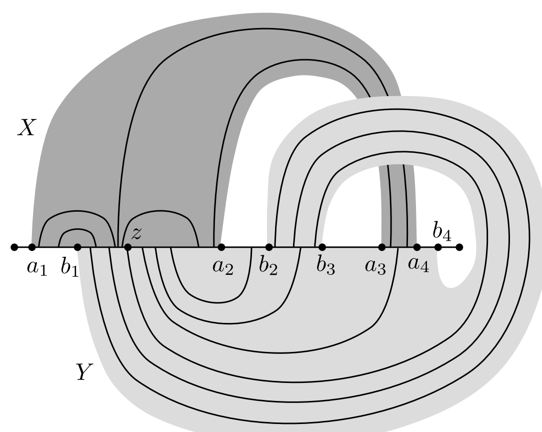

Let z be the maximal vertex with respect to $\prec$ such that there exists an edge e with one endpoint in z, and such that $\varphi(e)$ is incident to both sides of $\varphi(P)$. Then, B can be covered by three elementary bands as follows. First, $B[b_1,z)$ is an elementary band of type-1, with primary segment $P[b_1,z)$, and outlets $P[a_1,b_1)$, $P[z,a_2]$, $P[b_2,b_3]$, and $P[a_3,a_4]$. Next, $B[z,z]$ is an elementary band of type-1, with primary segment z. Finally, $B(z,a_2]$ is an elementary band of type-2, with primary segment $P(z,a_2]$, outlet $P[a_1,z]$, and with double outlet $P(a_2,b_4]$.

Case 9: There exists a $\varphi(P)$-ribbon $X \in \mathcal{X}$, and a $\varphi(P)$-ribbon $Y \in \mathcal{X}$, such that one side of α is in X, and the other is in Y. The $\varphi(P)$-ribbon X is attached to both sides of $\varphi(P)$, and the $\varphi(P)$-ribbon Y is also attached to both sides of $\varphi(P)$. Let $a_1, \ldots, a_4 \in \mathcal{S}$ be the endpoints of X, and let $b_1, \ldots, b_4 \in \mathcal{S}$ be the endpoints of Y. Assume w.l.o.g. that $a_1 \preceq a_2 \preceq a_3 \preceq a_4$, and $b_1 \preceq b_2 \preceq b_3 \preceq b_4$.

We consider the following subcases:

Case 9-1: Suppose that $X \cap \varphi(P)$ consists of two disjoint segments. Assume that $b_2 \prec a_3$, $a_2 \prec b_3$, $b_1 \preceq a_2$, and $a_1 \preceq b_2$. Let r be the minimal with respect to $\prec$ in $\{a_1, b_1\}$, and let r' be the maximal in $\{a_1, b_1\}$. Let s be the minimal in $\{a_2, b_2\}$, and let s' be the maximal in $\{a_2, b_2\}$. Suppose that $p_i = r'$, and $p_{i+1} = s$.

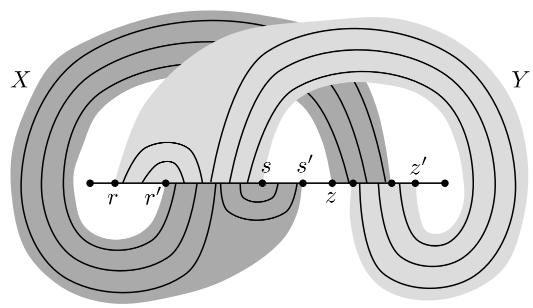

Let z be the minimal in $\{a_3, b_3, a_4, b_4\}$, and let z' be the maximal in $\{a_3, b_3, a_4, b_4\}$. Then, B is an elementary band of type-2, with primary segment $P[r', s]$, with outlets $P[r, r')$, and $P[s, s')$, and with double outlet $P[z, z']$.

Case 9-2: Suppose that $X \cap \varphi(P)$ consists of a single disjoint segment. Assume that $\lambda[a_1, a_3]$, and $\lambda[a_2, a_4]$ are sides of X, and that $a_3 \preceq b_3 \preceq a_4$. Let r be the minimal with respect to $\prec$ in $\{a_1, b_1\}$, and let r' be the maximal in $\{a_1, b_1\}$. Let s be the minimal in $\{a_4, b_4\}$, and let s' be the maximal in $\{a_4, b_4\}$. Suppose that $p_i = r'$, and $p_{i+1} = s$.

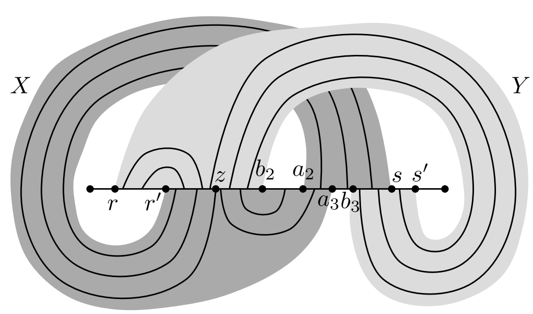

Let z be the maximal point in $P[r', b_2]$, such that there exists an edge $e \in B$ having z as an endpoint, and such that $\varphi(e)$ is incident to both sides of $\varphi(P)$. Then, B can be covered by three elementary bands as follows. First, $B[r', z)$ is an elementary band of type-1, with primary segment $P[r', z)$, and with outlets $P[r, r')$, $P[z, b_2]$, and $P[a_2, a_4]$. Next, $B[z, z]$ is an elementary band of type-1, with primary segment v. Finally, $B(z, b_2]$ is an elementary band of type-1, with primary segment $P(z, b_2]$, and with outlets $P[r, z]$, $P(b_2, a_3]$, and $P[b_3, b_4]$.

The above cases exhaust all possibilities for the following reason: Each side of α is covered by at most one element in $\mathcal{X}$, and each such element is either a $\varphi(P)$-ribbon or a single, double, or triple $\varphi(P)$-petal. In Cases 1 and 2, both sides of α are covered by the same $\varphi(P)$-petal; such a $\varphi(P)$-petal must be either double (Case 1) or triple (Case 2), since single $\varphi(P)$-petals are attached only to one side of $\varphi(P)$. In Case 3, a distinct $\varphi(P)$-petal covers each side of α, and one of them is single; the other petal can be either single, double, or triple, and the most general case is the one where it is triple. The case where both sides of α are covered by double $\varphi(P)$-petals is a special subcase of Case 5. This exhausts all possibilities for the cases where both sides of α are covered by φ-petals. In Case 4, one side of α is covered by some $\varphi(P)$-petal and the other side is covered by some $\varphi(P)$-ribbon; If the petal is triple, then the ribbon must be attached to a single side of $\varphi(P)$, and thus this case can be treated as a special subcase of Case 3. This exhausts all possibilities for the cases where one side of α is covered by a $\varphi(P)$-petal, and the other by a $\varphi(P)$-ribbon. It remains to consider the case where both sides of α are covered by φ-ribbons. In Case 6, both sides of α are covered by the same $\varphi(P)$-ribbon. In the remaining Cases 7–9, each side of α is covered by a distinct $\varphi(P)$-ribbon; each one of these two ribbons is attached either to a single or to both sides of $\varphi(P)$, which leads to three possible cases. In Case 7, each one of the two $\varphi(P)$-ribbons are attached to a single side of $\varphi(P)$. In Case 8, one $\varphi(P)$-ribbon is attached to a single side of $\varphi(P)$, and the other one is attached to both sides of $\varphi(P)$. Finally, in Case 9, both $\varphi(P)$-ribbons are attached to both sides of $\varphi(P)$. Thus every possible configuration is isomorphic to one of the cases considered above. This concludes the proof. □

We are now ready to prove the existence of small band coverings.

Proof of Lemma 2.4. Let φ be a drawing of G on an orientable surface $\mathcal{S}$ of genus g. By Lemma 3.6 there exists a ribbon-petal covering $\mathcal{X}$ for (G,P,φ), with $|\mathcal{X}| \leq 24g-12$. Let A be the set of all endpoints of all ribbons, and all petals in $\mathcal{X}$. Since any $X \in \mathcal{X}$ can have at most 4 endpoints, we have $|A| \leq 96g-48$. Let $p,q \in \mathcal{S}$ be the endpoints of the curve $\varphi(P)$, and let $A' = A \cup p,q$. Let $p_1,\ldots,p_k$ be an ordering of A' induced by a traversal of the curve $\varphi(P)$, where $k \leq |A|+2 \leq 96g-46$. For any $i \in \{1,\ldots,k-1\}$, let α_i be the segment of $\varphi(P)$ between p_i, and p_{i+1}. Let also Z_i be the maximal subpath of P with $\varphi(Z_i) \subseteq \alpha_i$. Let

$$F_i = \{\{x,y\} \in E(G) \setminus E(P) : \{x,y\} \cap V(Z_i) \neq \emptyset\}\,.$$

By Lemma 4.2 there exists a collection $\{(B_{i,j},Q_{i,j})\}_{j=1}^{k_i}$, with $k_i \leq 3$, such that

$$V(Z_i) = \bigcup_{j=1}^{k_i} V(Q_{i,j})\,, \qquad \text{and} \qquad F_i = \bigcup_{j=1}^{k_i} B_{i,j}\,.$$

It follows by Definition 2.3 that the final collection

$$\{(B_{i,j},Q_{i,j})\}_{i\in\{1,\ldots,k-1\},j\in\{1,\ldots,k_i\}}$$

is a band covering with spine P for G, of size at most $3(k-1) \leq 288g-141$. □

5 Compatible pairs of planar drawings

In the previous sections, we show how to decompose the input graph to a collection of toroidal graphs. We can use Mohar's algorithm to find an embedding for all those toroidal graphs in linear time. Nevertheless, the computed embeddings are not necessarily consistent. To fix this issue, we consider all pairs of bands and corresponding toroidal embeddings and change them to make their intersections consistent. We obtain a drawing of each pair of bands on a surface of genus $O(g^3)$. In the next section, we show how to resolve the remaining slight inconsistencies to get the final embedding.

Consider two bands B_1 and B_2 with disjoint primary segments Q_1 and Q_2, respectively. Let $B = B_1 \cap B_2$ (the set of edges going from Q_1 to Q_2). Denote the set of local edges incident on vertices in Q_1 by L_1, and the set of edges incident on vertices in Q_2 by L_2 (see Figure 6.).

The main result of this section is Theorem 5.1.

Theorem 5.1. *There is a polynomial-time algorithm that finds a drawing of $P \cup B \cup L_1 \cup L_2$ on a surface of genus $O(g^3)$.*

We will use the following notion of a restriction of a drawing.

Definition 5.2 (Combinatorial restriction)**.** Let G_1 be a graph, and let G_2 be a subgraph of G_1. Let φ_1 be a drawing of G_1 on some surface $\mathcal{S}_1$. The *combinatorial restriction* of φ_1 on G_2 is defined to be the combinatorial drawing φ_2 of G_2 induced by setting for every $v \in V(G_2)$, the ordering of the edges incident to v in G_2 to be as in φ_1. By gluing a disk along every facial walk in φ_2 we obtain a surface $\mathcal{S}_2$. When this does not cause confusion, we will naturally identify f_2 with the induced drawing of G_2 on $\mathcal{S}_2$. Note that the genus of $\mathcal{S}_2$ can be smaller than the genus of $\mathcal{S}_1$.

We first prove the following decomposition theorem and then show that it implies Theorem 5.1.

Theorem 5.3. *There is a polynomial-time algorithm that does the following. It divides segments Q_1 and Q_2 into consecutive segments $Q_1^1, \dots, Q_1^s$ and $Q_2^1, \dots, Q_2^s$, respectively, where $s = O(g)$. These segments do not share any vertices except possibly for endpoints. Also it partitions all edges in B into p disjoint sets $T_1, \dots, T_s$ such that all edges in T_i go between Q_1^i and Q_2^i. For each $i \in \{1, \dots, s\}$, the algorithm finds a drawing ψ_i of $Q_1 \cup Q_2 \cup T_i \cup L_1 \cup L_2$ on a surface of genus $O(g^2)$. Additionally, the combinatorial restriction of ψ_i to $Q_1 \cup Q_2 \cup L_1 \cup L_2$ is planar and canonical (as in Definition 2.2).*

We consider canonical drawings φ_1 and φ_2 of B_1 and B_2, respectively. Each of the drawings draws paths Q_1 and Q_2 on a horizontal line ℓ and hence defines a total ordering of vertices in Q_1 and Q_2. Let $\prec_1$ be the ordering defined by φ_1 and $\prec_2$ be the ordering defined by φ_2. By changing the orientation of the line ℓ in one of the drawings, if necessary, we may assume that $\prec_1$ and $\prec_2$ agree on $V(Q_1)$. However, orderings $\prec_1$ and $\prec_2$ may define the same or opposite order on $V(Q_2)$. In the former case, we say that the band intersection is *regular*; in the latter case, we say that the band intersection is *irregular*. In the Appendix, we show that in the irregular case the intersection of bands B_1 and B_2 can be drawn on a surface of genus 8. Thus Theorems 5.1 and 5.3 follow (with $s = 1$).

Lemma 5.4. *In the irregular case, the intersection of bands B_1 and B_2 can be drawn on a surface of genus 8.*

Before we give the proof of the above lemma, we introduce some auxiliary results. Define two conflict graphs $\mathcal{C}_1$ and $\mathcal{C}_2$ on the set $B \cup L_1 \cup L_2$. Two nodes in $\mathcal{C}_1$ are connected with an edge if they are in conflict with respect to $\prec_1$; two nodes in $\mathcal{C}_2$ are connected with an edge if they are in conflict with respect to $\prec_2$.

Lemma 5.5. *Let $B^* \subset B$ be a subset of global edges such that no two edges in B^* share an endpoint (i. e., B^* is a matching). Then $|B^*| \leq 4$.*

Proof. Similarly to Lemma 5.8, we have that $\mathcal{C}_1[B \cup L_1]$ and $\mathcal{C}_2[B \cup L_2]$ are bipartite graphs. So their subgraphs $\mathcal{C}_1[B^*]$ and $\mathcal{C}_2[B^*]$ are also bipartite. However, note that two edges $e_1, e_2 \in B^*$ are in conflict with respect to $\prec_1$ if and only if they are not in conflict with respect to $\prec_2$. Thus e_1 and e_2 are connected with an edge in $\mathcal{C}_1[B^*]$ if and only if they are not connected with an edge in $\mathcal{C}_2[B^*]$. Now consider the bipartition (X,Y) of $\mathcal{C}_1[B^*]$. Since no two vertices in X are connected with an edge in $\mathcal{C}_1[B^*]$, X forms a clique in $\mathcal{C}_2[B^*]$. Since $\mathcal{C}_2[B^*]$ is a bipartite graph, $|X| \leq 2$. Similarly, $|Y| \leq 2$. We conclude $|B^*| = |X| + |Y| \leq 4$. □

Lemma 5.6. *There exist sets of vertices $X \subset V(Q_1)$ and $Y \subset V(Q_2)$ with $|X| \leq 4$, $|Y| \leq 4$ such every edge in B is incident to at least one vertex in $X \cup Y$.*

Proof. Let B^* be a maximal matching in B. By Lemma 5.5, $|B^*| \leq 4$. Let X be the set of endpoints of edges in B^* that lie in Q_1 and let Y be the set of endpoints of edges in B^* that lie in Q_2. We have $|X| \leq 4$ and $|Y| \leq 4$. Since B^* is a maximal matching, every edge in B is incident to a vertex in $X \cup Y$. □

Now we are ready to prove Lemma 5.4.

Proof of Lemma 5.4. We find sets X and Y as guaranteed by Lemma 5.6. Let B' be the subset of edges in B that are incident to a vertex in X; let $B'' = B \setminus B'$. By Lemma 5.6, every edge in B'' is incident to a vertex in Y. We consider the drawing φ_1 restricted to $Q_1 \cup L_1 \cup Y \cup B''$ and the drawing φ_2 restricted to $Q_2 \cup L_2 \cup X \cup B'$. We assume w.l.o.g. that these two drawings do not overlap. For every point $x \in X$ we make two punctures in the plane, one in a small neighborhood of $\varphi_1(x)$, the other in a small neighborhood of $\varphi_2(x)$, and attach a handle $\mathcal{H}_x$ between them. We perform the same operation for every $y \in Y$. Now we define the desired drawing ψ of $Q_1 \cup Q_2 \cup B \cup L_1 \cup L_2$. The drawing ψ coincides with φ_1 on $Q_1 \cup L_1$ and coincides with φ_2 on $Q_2 \cup L_2$. We explain now how we define the drawing ψ of an edge $e \in B''$. Denote $e = \{x,y\}$ where $y \in Y$. Let $\gamma_e^1 = \varphi_1(e)$; the curve γ_e^1 connects $\varphi_1(x) = \psi(x)$ and $\varphi_1(y)$. Let γ_e^2 be a curve drawn on the handle $\mathcal{H}_y$ that connects $\varphi_1(y)$ and $\varphi_2(y) = \psi(y)$. Then we let $\psi(e)$ to be the concatenation of curves γ_e^1 and γ_e^2. Similarly, we define $\psi(e)$ for edges $e \in B'$. We obtain a drawing ψ of $Q_1 \cup Q_2 \cup B \cup L_1 \cup L_2$ on a surface of genus at most 8. Edges that share a common endpoint may cross or overlap in this drawing. We eliminate all such crossings by slightly perturbing edge drawings and uncrossing all crossings. □

In the rest of this section, we will analyze the regular case. Since in the regular case orderings $\prec_1$ and $\prec_2$ agree on $V(Q_1) \cup V(Q_2)$, we will just denote this ordering by $\prec$.

Definition 5.7. Let us say that two edges $\{x,y\}, \{x',y'\} \in B \cup L_1 \cup L_2$ (with $x \prec y$ and $x' \prec y'$) are *in conflict* if $x \prec x' \prec y \prec y'$ or $x' \prec x \prec y' \prec y$. We consider an auxiliary *conflict* graph $\mathcal{C}$ on the set of edges

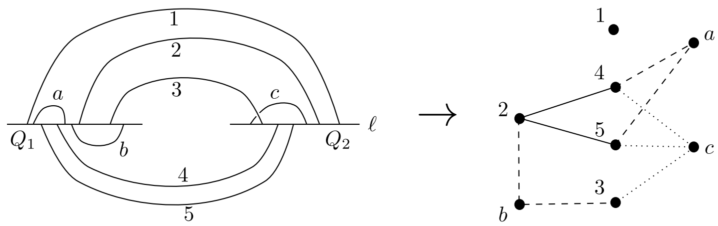

Figure 6: The figure on the left shows two bands with primary segments Q_1 and Q_2, global edges $B = \{1,2,3,4,5\}$, local edges $L_1 = \{a,b\}$ and $L_2 = \{c\}$. The right figure shows the corresponding conflict graph $\mathcal{C}$. Note that graphs $\mathcal{C}[B \cup L_1]$ and $\mathcal{C}[B \cup L_2]$ are bipartite as claimed by Lemma 5.8. However, the graph $\mathcal{C}$ is not bipartite since it contains an odd cycle $2 \to 5 \to c \to 3 \to b \to 2$. The graph $\mathcal{C}[B]$ has three connected components $S_1 = \{1\}$, $S_2 = \{2,4,5\}$ and $S_3 = \{3\}$ with $S_1 \lhd S_2 \lhd S_3$.

$B \cup L_1 \cup L_2$ in which e and e' are connected with an auxiliary edge if e and e' are in conflict. We denote the set of connected components of $\mathcal{C}[B]$ (the subgraph of $\mathcal{C}$ induced by B) by $\mathcal{S}$. (See Figure 6.)

The motivation for Definition 5.7 is that if two edges e and e' are not conflict, we can draw them in the plane on one side of ℓ or on opposite sides of ℓ. However, if two edges are in conflict we can only draw them on opposite sides of ℓ. Accordingly, if the conflict graph is bipartite, then we can simultaneously draw all the edges (the nodes of the conflict graph) in the plane. The following lemma shows that for every $S \in \mathcal{S}$, $\mathcal{C}[S \cup L_1 \cup L_2]$ is bipartite.

Lemma 5.8. *Graphs $\mathcal{C}[B \cup L_1]$ and $\mathcal{C}[B \cup L_2]$ are bipartite. Moreover, for every $S \in \mathcal{S}$, $\mathcal{C}[S \cup L_1 \cup L_2]$ is bipartite.*

Proof. Consider a canonical drawing of the band B_1. Let U be the set of edges in $B \cup L_1$ that lie above the line ℓ and D be the set of edges in $B \cup L_1$ that lie below the line ℓ. Note that if two edges e and e' lie on one side of the line ℓ—either $e, e' \in U$ or $e, e' \in D$—then e and e' are *not in conflict* since otherwise they would intersect. That is, e and e' are not connected with an edge in $\mathcal{C}$. Thus $U \cup D$ is a valid bipartition of $\mathcal{C}[B \cup L_1]$. We showed that $\mathcal{C}[B \cup L_1]$ is a bipartite graph. Similarly, $\mathcal{C}[B \cup L_2]$ is a bipartite graph.

Since S is connected, there is only one bipartition of S, which coincides with bipartitions of S in $\mathcal{C}[S \cup L_1]$ and in $\mathcal{C}[S \cup L_2]$. Thus bipartitions of $\mathcal{C}[S \cup L_1]$ and $\mathcal{C}[S \cup L_2]$ can be combined into a bipartition of $\mathcal{C}[S \cup L_1 \cup L_2]$ (here we use that there are no edges between L_1 and L_2). $\square$

Definition 5.9. We define a partial order $\lhd$ on edges as follows: $e \lhd e'$ for two edges $e = \{x,y\}$ and $e' = \{x',y'\}$ if $x \preceq x'$ and $y \succeq y'$ (here, we assume w.l.o.g. that $x \prec y$ and $x' \prec y'$), where one of the two inequalities is strict. We write $S \lhd S'$ for two sets of edges S and S' if $e \lhd e'$ for every $e \in S$ and $e' \in S'$.

Note that two edges $e, e' \in B$ are comparable with respect to $\lhd$ if and only if e and e' are not in conflict.

Claim 5.10. *The set $\mathcal{S}$ is totally ordered by $\lhd$.*

Proof. Consider two distinct connected components $S, S' \in \mathcal{S}$. Every two edges $e \in S$ and $e' \in S'$ are not in conflict, and therefore either $e \lhd e'$ or $e' \lhd e$. If $e \rhd e'$ for all $e \in S$ and $e' \in S'$ then $S \rhd S'$ and we are done.

So suppose that $e \lhd e'$ for some $e \in S$ and $e' \in S'$. Consider $e'' \in S'$ adjacent to e' in $\mathcal{C}$. We show that $e \lhd e''$. Indeed, assume to the contrary that $e \lhd e'$ and $e'' \lhd e$. Then from the transitivity of partial order $\lhd$, we get that $e'' \lhd e'$, which contradicts to the fact that e' and e'' are in conflict. Since S' is connected in $\mathcal{C}$, we conclude that for all $e'' \in S'$, $e \lhd e''$. Applying the same argument to S, we get that $S \lhd S'$. □

We order elements of $\mathcal{S}$ (connected components of $\mathcal{C}[B]$) with respect to to the order $\lhd$: $S_1 \lhd \cdots \lhd S_r$. We note that the sets $S_1,\dots,S_r$ satisfy most properties we require in Theorem 5.3. Since sets $S_1,\dots,S_r$ are ordered with respect to $\lhd$, endpoints of edges in sets S_i divide Q_1 and Q_2 into r consecutive segments. By Lemma 5.8, each graph $\mathcal{C}[S \cup L_1 \cup L_2]$ is bipartite and therefore the graph $Q_1 \cup Q_2 \cup S \cup L_1 \cup L_2$ is planar. The only obstacle is that the number r of sets S_i can be arbitrarily large (it is not bounded by a function of g). To resolve this issue, we will join together some consecutive sets S_i and obtain the desired sets T_j. We do that using a simple greedy algorithm. We define numbers $t_0,t_1,t_2,\dots$ by induction. First, we let $t_0 = 0$. Let t_{q+1} be the largest t such that one of the following two conditions holds:

1. The graph $\mathcal{C}[S_{t_q+1} \cup \cdots \cup S_t \cup L_1 \cup L_2]$ is bipartite.

2. All edges in $S_{t_q+1},\dots,S_t$ are in conflict with some local edge $\hat{e} \in L_1 \cup L_2$.

We stop this procedure when we process all sets S_i. We obtain numbers $t_0,\dots,t_s$ (for some $s > 0$). Let $T_q = S_{t_{q-1}+1} \cup \cdots \cup S_{t_q}$ for every $q \in \{1,\dots,s\}$.

Since each set T_j is the union of consecutive sets S_i, we have that $T_1 \lhd T_2 \cdots \lhd \cdots \lhd T_s$. For every p consider all vertices in Q_1 incident to edges in T_p. Let Q_1^p be the segment between the leftmost and rightmost among such vertices; define Q_2^p similarly. Then $Q_1^1 \preceq Q_1^2 \preceq \cdots \preceq Q_1^s$, and $Q_2^1 \succeq Q_2^2 \succeq \cdots \succeq Q_1^s$. Note that Q_1^i and Q_1^{i+1} share at most one vertex and Q_1^i and Q_1^j are disjoint if $|i-j| > 1$.

Recall that we need to find a drawing of each graph $T_q \cup Q_1 \cup Q_2 \cup L_1 \cup L_2$ on a surface of genus $O(g^2)$ and prove that $s = O(g)$. Note that if $\mathcal{C}[S_{t_{q-1}+1} \cup \cdots \cup S_{t_q} \cup L_1 \cup L_2]$ is bipartite (the first stopping condition holds) then $T_q \cup Q_1 \cup Q_2 \cup L_1 \cup L_2$ is planar. So besides proving that $s = O(g)$, we only need to analyze the case when all edges in $S_{t_{q-1}+1},\dots,S_{t_q}$ are in conflict with some local edge $\hat{e} \in L_1 \cup L_2$. We need the following definition that captures this case.

Definition 5.11 (Comb)**.** Let B_1 and B_2 be two bands. Let $B' \subset B_1 \cap B_2$. Suppose that all edges in B' are in conflict with a local edge $\hat{e} \in L_1$. Let φ_1 and φ_2 be canonical drawings of $H_1 = Q_1 \cup Q_2 \cup B' \cup L_1$ and $H_2 = Q_1 \cup Q_2 \cup B' \cup L_2$ respectively (as in Definition 2.2). Note that all edges in B' are drawn on one side of ℓ in φ_1 since all edges in B' are in conflict with $\hat{e}$; we assume w.l.o.g. that all edges in B' are drawn above ℓ. Then, we say that $((B_1,Q_1,\varphi_1),(B_2,Q_2,\varphi_2),B')$ is a *comb* with *spine* P, or simply a comb, when P is clear from the context.

Lemma 5.12. *For every set T_p, we have*

- *there is a canonical planar drawing φ of $Q_1 \cup Q_2 \cup L_1 \cup L_2 \cup T_p$, or*

- *there are drawings φ_1 and φ_2 such that $((B_1,Q_1,\varphi_1),(B_2,Q_2,\varphi_2),T_p)$ is a comb, or*

- *there are drawings φ_1 and φ_2 such that $((B_2,Q_2,\varphi_2),(B_1,Q_1,\varphi_1),T_p)$ is a comb.*

Proof. Consider a set T_p. We know that either $\mathcal{C}[T_p \cup L_1 \cup L_2]$ is bipartite or all edges in T_p are in conflict with some local edge $\hat{e} \in L_1 \cup L_2$.

First, consider the former case. Let (U,D) be the bipartition of $\mathcal{C}[T_p \cup L_1 \cup L_2]$. There are no conflicts between edges in U, and no conflicts between edges in D.

We arrange all vertices of $Q_1 \cup Q_2$ on horizontal line ℓ according to the order $\prec$. We draw each edge in U in the top half plane with respect to ℓ and each edge in D in the bottom half plane so that no two edges intersect. Specifically, we consider a coordinate frame in which ℓ is defined by $\{(x,y) : y = 0\}$. We draw every edge $e \in U$ connecting points with coordinates $(x_1,0)$ and $(x_2,0)$ as a polygonal chain

$$(x_1,0) \leftrightarrow ((x_1+x_2)/2, |x_2-x_1|/2) \leftrightarrow (x_2,0)\,;$$

we draw every edge $e \in D$ connecting points $(x_1,0)$ and $(x_2,0)$ as

$$(x_1,0) \leftrightarrow ((x_1+x_2)/2, -|x_2-x_1|/2) \leftrightarrow (x_2,0)\,.$$

It is easy to see that no two edges intersect. We obtain the desired drawing φ.

Now consider the case when all edges in T_p are in conflict with some local edge $\hat{e} \in L_1 \cup L_2$. Without loss of generality, $\hat{e} \in L_1$. Consider a canonical planar drawing φ_1 of $Q_1 \cup L_1 \cup T_p$. Since every edge $e \in T_p$ is in conflict with $\hat{e}$, edge e lies on the other side of ℓ than $\hat{e}$. Thus all edges in T_p lie on the same side of ℓ. Let φ_2 be a canonical planar drawing of $Q_2 \cup L_2 \cup T_p$. We get that $((B_1,Q_1,\varphi_1),(B_2,Q_2,\varphi_2),T_p)$ is a comb. $\square$

We present an algorithm that finds a drawing of the graph $H_1 \cup H_2$, where H_1 and H_2 are as in Definition 5.11, on a surface of genus $O(g^2)$ in Section 6.

It remains to show that $s = O(g)$. First, we give a high-level overview of the proof. Roughly speaking, we observe that $T_i \cup T_{i+1} \cup Q_1 \cup Q_2 \cup L_1 \cup L_2$ is not planar because if it was planar our greedy algorithm would include T_{i+1} in T_i. Then we consider $s/2$ graphs $T_i \cup T_{i+1} \cup Q_1 \cup Q_2 \cup L_1 \cup L_2$ for $i \in \{1,3,5,\dots\}$. Each of them is not planar and there are $s/2$ of them. Note that the union of k *disjoint* non-planar graphs has genus at least k. So we want to argue that $s/2 \le g$ as otherwise the union of graphs $T_i \cup T_{i+1} \cup Q_1 \cup Q_2 \cup L_1 \cup L_2$ would have genus at least $s/2 > g$, which would contradict the fact that the genus of G is g. However, this approach does not work as stated because graphs $T_i \cup T_{i+1} \cup Q_1 \cup Q_2 \cup L_1 \cup L_2$ share local edges. To overcome this obstacle, for each set T_i we define sets $\Lambda_1^i \subset L_1$ and $\Lambda_2^i \subset L_2$ such that $T_i \cup T_{i+1} \cup Q_1 \cup Q_2 \cup \Lambda_1^i \cup \Lambda_2^i$ is not planar and sets Λ_1^i and Λ_1^j do not intersect if $|i-j| > 2$. This resolves the problem in the argument we outlined above (we choose indices i with some sufficiently large constant gap). We get that $s = O(g)$. Now we present the formal proof.

Claim 5.13. *Suppose that $\hat{e} \in L_1 \cup L_2$ is in conflict with edges $e \in S_a$ and $e' \in S_b$ and $b - a > 1$. Then $\hat{e}$ is in conflict with edges in all sets S_c for $c \in \{a+1,\dots,b-1\}$. Moreover, each set S_c contains only one edge.*

Proof. Assume without loss of generality that $\hat{e} = (\hat{x},\hat{y}) \in L_1$ where $\hat{x} \prec \hat{y}$. Consider a canonical drawing of B_1. Let $e = (x,y)$ and $e' = (x',y')$ so that $x,x' \in Q_1$. From $S_a \lhd S_b$, we get that $x \prec x'$. Since $\hat{e}$ is in conflict with e and e' (and $y,y' \succ \hat{y}$), we have $\hat{x} \prec x \preceq x' \prec \hat{y}$. Now if $e'' = (x'',y'') \in S_c$ then $\hat{x} \prec x \preceq x'' \preceq x' \prec \hat{y}$, and therefore e'' is in conflict with $\hat{e}$.

Note that if S_c contains at least two edges then it also contains two edges e_1 and e_2 that are adjacent in $\mathcal{C}$ (since $\mathcal{C}[S_c]$ is connected). Both of them are adjacent to $\hat{e}$ in S_C. So $\hat{e}$, e_1, and e_2 form a triangle in $\mathcal{C}$. This contradicts to the fact that $\mathcal{C}[B \cup L_1]$ is bipartite. $\square$

We now show that the genus of G is at least $\Omega(s)$. Define subsets $\Lambda_1^i \subset L_1$ and $\Lambda_2^i \subset L_2$ such that $T_i \cup T_{i+1} \cup Q_1^i \cup Q_2^i \cup \Lambda_1^i \cup \Lambda_2^i$ are disjoint and non-planar for $\Omega(s)$ values of i. The union of these graphs has genus at most g; on the other hand, it is at least $\Omega(s)$. Thus we conclude that $s = O(g)$.

For a graph G and a Hamiltonian cycle C, let $\prec_C$ be an ordering of vertices induced by a traversal of C (starting at an arbitrary vertex). We say that two edges $\{x_1, y_1\}$ and $\{x_2, y_2\}$ (with $x_1 \prec_C y_1$ and $x_2 \prec_C y_2$) of $E(G) \setminus E(C)$ are *in conflict* if either $x_1 \prec_C x_2 \prec_C y_1 \prec_C y_2$ or $x_2 \prec_C x_1 \prec_C y_2 \prec_C y_1$. Equivalently, this happens when the vertices x_1, x_2, y_1, y_2 are distinct and x_2 and y_2 are in distinct components of $C \setminus \{x_1, y_1\}$.

Lemma 5.14. *Let G be a graph, and let C be a Hamiltonian cycle in G. The graph G is planar if and only if there exists a partition $E(G) \setminus E(C) = E_1 \cup E_2$ such that no two edges in E_1 are in conflict and no two edges in E_2 are in conflict (with respect to C). In other words, G is planar if and only if and only if the corresponding conflict graph is bipartite.*

Proof. The assertion follows immediately by the Jordan Curve Theorem. $\square$

Before we define sets Λ_1^i and Λ_2^i, we define auxiliary sets L_1^i and L_2^i. Let

$$L_1^i = \{e \in L_1 : e \text{ is in conflict with some } e' \in T_i\},$$
$$L_2^i = \{e \in L_2 : e \text{ is in conflict with some } e' \in T_i\}.$$

Claim 5.15. *Sets L_1^i and L_1^j are disjoint if $|i - j| > 2$.*

Proof. Assume without loss of generality that $j > i$. Suppose to the contrary that there exist $\hat{e} \in L_1^i \cap L_1^j$ then, by Claim 5.13, $\hat{e}$ is in conflict with with all edges in $T_{i+1} \cup \cdots \cup T_{j-1}$. Thus by the construction of sets T_p, all edges in $T_{i+1} \cup \cdots \cup T_{j-1}$ must lie in one set T_p, which contradicts to the fact that $j - i > 2$. $\square$

Claim 5.16. *Consider two edges $e = (x, y)$ and $e' = (x', y')$ in $L_1^p \cup L_1^{p+1}$ that are connected with a path π_1 in $\mathcal{C}[L_1]$. Then they are connected with a path $\pi_2 : e = e_1 \to \cdots \to e_l = e'$ such that the endpoints of every e_r lie between the rightmost point of Q_1^{p-10} and leftmost point of Q_1^{p+10}. Since the graph $\mathcal{C}[L_1]$ is bipartite, the lengths of paths π_1 and π_2 have the same parity.*

Proof. Consider a shortest path $e = e_1 \to \cdots \to e_l = e'$ in $\mathcal{C}[L_1]$ between e and e'. Denote it by π_2. Note that e_a and e_b are not in conflict if $|a - b| > 1$ (otherwise, e_a and e_b would be adjacent in $\mathcal{C}[L_1]$, and we would be able to shortcut π_2).

We claim that π_2 satisfies the conditions of the lemma. Indeed, assume to the contrary that the leftmost among endpoints of edges $e_1, \dots, e_l$ lies to the left of the rightmost point of Q_1^{p-7} (the other case is similar). Denote this endpoint by u; let us say that u is a endpoint of e_t. Let v be the other end of e_t. Then v lies to the left of the rightmost point of Q_1^{p-5}. Let $e_{t-1} = (u', v')$ and $e_{t+1} = (u'', v'')$ so that $u' \prec v'$ and $u'' \prec v''$. Consider two cases.

Since e_{t-1} and e_{t+1} are in conflict with e_t (and $u \prec u'$, $u \prec v'$), we have $u' \prec v \prec v'$ and $u'' \prec v \prec v''$. Without loss of generality assume that $u' \succ v'$ and $v' \prec v''$ (note that e_{t-1} and e_{t+1} are not in a conflict).

All edges $e_1,\ldots, e_{t-1}$ are not in conflict with e_{t+1}. Thus if one of their endpoints lies between u'' and v'' then the other point also lies between u'' and v''. Now if both endpoints of $e_{q+1} = (x_{q+1}, y_{q+1})$ (where $q+1 < t+1$) lie between u'' and v'' then one of the endpoints of e_q lies between (x_{q+1}, y_{q+1}) (since e_q and e_{q+1} conflict with each other), and thus it lies between u'' and v''. Consequently, both endpoints of e_q lie between u'' and v''. We conclude that all endpoints of $e_1,\ldots, e_{t-1}$ lie between u'' and v''. Thus x lies between u' and v'. That contradicts to the fact that u' lies in Q_1^p. □

Let Λ_1^i be the set edges in L_1 whose endpoints lie between the rightmost point of Q_1^{i-11} and leftmost point of Q_1^{i+11}. By definition, $\mathcal{C}[T_i \cup T_{i+1} \cup L_1 \cup L_2]$ is not bipartite. From Claim 5.16, we get the following corollary.

Corollary 5.17. *The graph $\mathcal{C}[T_p \cup T_{p+1} \cup \Lambda_1^p \cup \Lambda_2^p]$ is not bipartite.*

Lemma 5.18. *We have $s = O(\mathsf{genus}(G))$.*

Proof. For every $q \in \{1,\ldots,\lfloor s/40 \rfloor\}$ let f_q and f'_q be two edges in T_{40q-38} and T_{40q-1}, respectively. Also, let

$$H_q = Q_i^{40q-38} \cup \cdots \cup Q_i^{40q-1} \cup Q_j^{40q-38} \cup \cdots \cup Q_j^{40q-1} \cup T_{40q-20} \cup T_{40q-21} \cup \Lambda_1^{40q-20} \cup \Lambda_2^{40q-20} \cup \{f_q, f'_q\}.$$

Note that graphs H_q share no edges since for all sets Q_1^p, T_p, Λ_1^{40q-20} and Λ_2^{40q-20} are disjoint. Moreover, all vertices of H_q lie on the union of paths

$$Q_i^{40q-38} \cup \cdots \cup Q_i^{40q-1} \cup Q_j^{40q-38} \cup \cdots \cup Q_j^{40q-1}.$$

Therefore, all graphs H_q are disjoint. Since $\bigcup_q H_q$ is a subgraph of G, the genus of $\bigcup_q H_q$ is at most $\mathsf{genus}(G)$. Thus the sum of genera of graphs H_q is at most $\mathsf{genus}(G)$.

Note that edges f_q and f'_q and segments of paths Q_1 and Q_2 that lie between the endpoints of f_q and f'_q form a cycle. Endpoints of all local and global edges of H_q lie on this cycle. By Corollary 5.17, the graph

$$T_{40q-20} \cup T_{40q-21} \cup \Lambda_1^{40q-20} \cup \Lambda_2^{40q-20}$$

is not bipartite, hence by Lemma 5.14, H_q is not planar.

We conclude that $s = O(\mathsf{genus}(G))$. □

This completes the proof of Theorem 5.3, an algorithm that finds a drawing ψ_i of $Q_1 \cup Q_2 \cup T_i \cup L_1 \cup L_2$ on a surface of genus $O(g^2)$ for each i. Now we are ready to prove Theorem 5.1 by combining drawings ψ_i to obtain one drawing ψ.

Proof of Theorem 5.1. We assume that the drawing of ℓ and $Q_1 \cup Q_2$ are the same in all drawings ψ_i (we can do that without loss of generality since all vertices in $Q_1 \cup Q_2$ are ordered with respect to $\prec$ in all drawings ψ_i).

First we take care of global edges in T_i. Consider the drawing ψ_i. Recall that if T_i is not a comb then the drawing ψ_i is planar. Otherwise, it is a drawing on a plane with attached handles. In the latter case, all handles are attached to the plane above or below either Q_1^i or Q_2^i. We make four punctures in the plane: one above Q_1^i (sufficiently far away from ℓ), one below Q_1^i, one above Q_2^i and one below Q_2^i. We attach a handle H^U between two punctures above ℓ, and another handle H^D between two punctures below ℓ. Now we redraw all global edges that go above ℓ on the handle H_U, and all edges that go below ℓ on the handle H_D. We cut the part of the plane that lies above and below Q_1^i and Q_2^i. We denote this part by $\mathcal{T}_i$. The boundary of $\mathcal{T}_i$ consists of 4 vertical lines that pass through the left end of Q_1^i, the right end of Q_1^i, the left end of Q_2^i, and the right end of Q_2^i. We combine these parts $\mathcal{T}_i$ together and get a surface $\mathcal{T}$ of genus at most $O(g^3)$; we do not identify boundaries of $\mathcal{T}_i$ except for points of ℓ that belong to boundaries of several sets $\mathcal{T}_i$ (at this point, the surface might be disconnected). We also add line ℓ to $\mathcal{T}$. Now we partially define the drawing ψ on $\mathcal{T}$. We draw Q_1 and Q_2 in ψ in the same way as they are drawn in ψ_i. We draw each global edge $e \in T_i$ in the same way it is drawn in ψ_i (on $\mathcal{T}_i$). We perform this step for all i and obtain a drawing of all global edges.

Now we take care of local edges. Consider a local edge e whose drawing ψ_i partially lies in $\mathcal{T}_i$. We draw the segment of e that lies in $\mathcal{T}_i$ on $\mathcal{T}$ in the same way it is drawn on $\mathcal{T}_i$. It remains to draw missing segments of local edges and connect all segments together. We describe how we do that for edges in L_1; we process edges in L_2 in exactly the same way.

Consider two consecutive sets T_i and T_{i+1}. Let u be the the rightmost vertex of Q_1^i and v be the leftmost vertex of Q_{i+1}^i. Let e_i be a global edge in T_i incident on u, and e_{i+1} be a global edge in T_{i+1} incident on v. Let h_u be the line perpendicular to u in $\mathcal{T}_i$ and h_v be the line perpendicular to v in $\mathcal{T}_{i+1}$. There are two possibilities: either $u = v$ or $u \prec v$.

First, we consider the case $u = v$. Let A be the set of edges $e = (x, y)$ with $x \prec u$ and $u \prec y$. All edges in A are in conflict with both e and e'. Thus all edges in A are drawn on one side of ℓ in ψ_i and ψ_{i+1} (in particular, no two edges in A are in conflict). Consider all crossing points of edges in A and line h_u ordered descendingly by their distance from ℓ. Since the drawing of local edges in ψ_i is combinatorially planar, crossing points are ordered in the same way as corresponding edges ordered by $\prec$: if $e_1 \prec e_2$ then the crossing point of e_1 is further away from ℓ than the crossing point of e_2. Similarly, the crossing points of edges in A and line h_v are ordered in the same way as edges in A. Thus edges cross lines h_u and h_v in the same order. We attach a handle between $\mathcal{T}_i$ and $\mathcal{T}_{i+1}$ and then for every edge $e \in A$ draw a curve that connects the segment of e in $\mathcal{T}_i$ and the segment of e in $\mathcal{T}_{i+1}$.

Now consider the case when $u \prec v$. Let A be the set of edges $e = (x, y)$ with $x \prec u$ and $v \prec y$. We treat edges in A in exactly the same way as before; we attach one handle and connect segments of edges in A drawn on $\mathcal{T}_i$ and on $\mathcal{T}_{i+1}$. It remains to draw edges in the set

$$D = \{e = (x, y) : x \prec u \prec y \prec v \text{ or } u \prec x \prec v \prec y \text{ or } u \prec x \prec y \prec v\}.$$

Note that $A \prec D$ thus all crossing points of edges in D with h_u and h_v lie closer to ℓ than crossing points of A with h_u and h_v, respectively.

Denote the conflict graph $\mathcal{C}[D \cup \{e_i, e_{i+1}\}]$ by H.

Lemma 5.19. *The graph H is bipartite.*

Proof. Consider connected components of $\mathcal{C}[D]$. We show that there is at most one connected component C that is connected with both e_i and e_{i+1} in H. Assume to the contrary that there are two such connected components C_1 and C_2. Repeating the proof of Claim 5.10, we get that if two connected components C_1 and C_2 of $\mathcal{C}[D]$ are connected with e_i then either $C_1 \triangleleft C_2$ or $C_2 \triangleleft C_1$. Without loss of generality, $C_1 \triangleleft C_2$.

Then for every edge $e=(x,y)\in C_1$ (with $x\prec y$) we have $x\preceq C_2$ and $C_2\preceq y$. Thus $x\prec u$ and $v\prec y$, which contradicts to the fact that $e\notin A$.

Let C be a connected component of $\mathcal{C}[D]$. By Lemma 5.8, graphs $C\cup\{e_i\}$ and $C\cup\{e_{i+1}\}$ are bipartite. Since there is no edge between e_i and e_{i+1} in H, the graph $C\cup\{e_i\}\cup\{e_{i+1}\}$ is also bipartite. We color the graph H with two colors as follows. If there is a connected component of $\mathcal{C}[D]$ which is connected to both e_i and e_{i+1}, we first color it with 2 colors. Otherwise, we arbitrarily color nodes e_i and e_{i+1} of H. Every other connected component of $\mathcal{C}[D]$ is connected to at most one of nodes e_i and e_{i+1}. We color it with 2 colors so that its coloring agrees with the coloring of e_i and e_{i+1}. We obtain a valid 2-coloring of H. □

Since H is bipartite there exist a canonical drawing γ of edges of D, in which all edges that are in conflict with e_i lie on one side of ℓ and all edges that are in conflict with e_{i+1} lie on one side of ℓ. We attach a slab $\mathcal{T}_{i,i+1}$ above and below the segment between u and v of Q_1 to $\mathcal{T}$. Then we draw segments of all edges in D on $\mathcal{T}_{i,i+1}$ in the same way they are drawn in γ. Now the leftmost vertex of $\mathcal{T}_{i,i+1}$ is u and the rightmost vertex of $\mathcal{T}_{i,i+1}$ is v. So we can use the argument we used above to connect drawings of segments of $e\in D$ drawn on $\mathcal{T}_i$, $\mathcal{T}_{i,i+1}$, and $\mathcal{T}_{i+1}$. □

6 Drawing a comb

In this section, we will show how to find a drawing of all the edges participating in a comb. Throughout this section G denotes a Hamiltonian graph with a Hamiltonian path P. $B_1,B_2\subseteq E(G)$ are two bands of type-2, with spine P and $B=B_1\cap B_2$. For $i\in\{1,2\}$, Q_i is the primary segment of B_i, L_i is the set of its local edges. We will assume that $((B_1,Q_1,\gamma_1),(B_2,Q_2,\gamma_2),B)$ is a comb.

6.1 Twists and minimally-twisting drawing

Definition 6.1 (Minimally-twisting drawing)**.** Let $\varphi_1=\gamma_1$. We inductively define a canonical drawing φ_2 of H_2 (note that there could be several such drawings, and we shall specify one of them). We begin by drawing Q_1 and Q_2 on a line ℓ as in Definition 2.2. We then define the drawing of B. Since $((B_1,Q_1,\gamma_1),(B_2,Q_2,\gamma_2),B)$ is a comb, it follows that $\prec$ is a total ordering on B_1, and therefore also on B. Let $e_1,\dots,e_t$ be that ordering of B. We draw e_1 above ℓ. Given the drawing of $e_1,\dots,e_{i-1}$, we define the drawing of e_i as follows. If there exists a planar drawing of H_2 that extends the current planar drawing, and such that e_i appears on the same side of ℓ as e_{i-1}, then we draw e_i on the same side of ℓ as e_{i-1}. Otherwise, we draw e_i on the opposite side of ℓ. Since $B\cup L_2$ is an elementary band of type-2, it follows that one of the two drawings always exists. Finally, after defining the drawing on all edges in B, we extend it to a canonical drawing of H_2. We say that the resulting planar drawing φ_2 of H_2 is a *minimally-twisting* drawing for the comb $((B_1,Q_1,\gamma_1),(B_2,Q_2,\gamma_2),B)$. For every $i\in\{1,\dots,t-1\}$, such that e_i and e_{i+1} are drawn on opposite sides of ℓ in φ_2, we say that (e_i,e_{i+1}) is a *twist*.

Definition 6.2 (Short and long twists)**.** Let G, P, B_1, B_2, B, γ_1, γ_2, L_1, L_2, H_1, H_2, φ_1, and φ_2 be as in Definition 6.1. As in Section 5, let $\mathcal{C}[B\cup L_1]$ be the conflict graph for $B\cup L_1$. Let (e,e') be a twist in φ_2. Note that since $\mathcal{C}[B\cup L_1]$ is a bipartite graph every path P between e and e' in $\mathcal{C}[B\cup L_1]$ is of even length.

We say that a simple path P in $\mathcal{C}[B\cup L_1]$ is a *left blocking* path for (e,e') if all internal vertices of P are edges in L_1 and no two non-consecutive nodes of P are connected by an edge in $\mathcal{C}[B\cup L_1]$. We say that a left blocking path P is long if it has length at least 4 (i. e., $|V(P)| \geq 5$). We say that a left blocking path P is short if it has length 2 (i. e., $|V(P)| = 3$). Accordingly, we say that (e,e') is a *long twist* if there exists a long blocking path for (e,e'). We say that (e,e') is a *short twist* if there is no such path.

Lemma 6.3. *Consider a short twist (e,e'). Let $X_1 \subset L_1$ be the set of all local edges $f \in L_1$ such that the path $e \to f \to e'$ is a short blocking path for (e,e'). Then the edges of X_1 are totally ordered by $\lhd$.*

Proof. Let $f_1, f_2 \in X_1$. Note that f_1 and f_2 are in conflict with e. Since the conflict graph $\mathcal{C}[B \cup L_1]$ is bipartite, f_1 and f_2 are not in conflict with each other. Let $e = \{a,b\}$, $f_1 = \{x_1,y_1\}$, and $f_2 = \{x_2,y_2\}$. Assume w.l.o.g. that $x_1 \prec y_1$, and $x_2 \prec y_2$. Since each f_i is a left blocking edge of (e,e'), it follows that $x_i \preceq a \preceq y_i$ (for $i \in \{1,2\}$). Since f_1 and f_2 are not in conflict, we have that either $x_1 \preceq x_2 \preceq a \preceq y_2 \preceq y_1$ or $x_2 \preceq x_1 \preceq a \preceq y_1 \preceq y_2$. Therefore, either $f_1 \lhd f_2$ or $f_2 \lhd f_1$. □

Definition 6.4 (Inner-most left blocking edge of a short twist)**.** Let (e,e') be a short twist. Let $X_1 \subset L_1$ be the set of all local edges $f \in L_1$ such that the path $e \to f \to e'$ is a short blocking path for (e,e'). Let $f^* \in X_1$ be maximal with respect to $\lhd$. Then we say that f^* is the *inner-most left blocking edge* of the twist (e,e'). By Lemma 6.3, f^* is uniquely defined.

Lemma 6.5 (Right blocking paths)**.** *Let (e,e') be a twist in φ_2. Then, there exists a path P in $\mathcal{C}[B \cup L_2]$ between e and e', such that all internal vertices in P are edges in L_2, and such that $|V(P)|$ is even. We say that P is the* right blocking path *of (e,e').*

Proof. If there exists no right blocking path of (e,e'), then when computing φ_2, the edges e and e' end up on the same side of Q_1, and Q_2. Thus, since (e,e') is a twist in φ_2, it follows that there exists a right blocking path of (e,e'). □

We denote the set of all long twists in φ_2 by T_L. For every long twist (e,e'), we choose one left long blocking path. We denote it by $P^L_{(e,e')}$. Let f be the node next to e on $P^L_{(e,e')}$ and f' be the node next to e' on $P^L_{(e,e')}$. We denote the segment of Q_1 between the left end of f and the right end of f' by $Q_{1,(e,e')}$.

For every twist (e,e'), we find the shortest among all right blocking paths. We denote it by $P^R_{(e,e')}$. Let f be the node next to e on $P^R_{(e,e')}$ and f' be the node next to e' on $P^R_{(e,e')}$. We denote the segment of Q_2 between the left end of f' and the right end of f by $Q_{2,(e,e')}$.

Claim 6.6. *Consider a long twist (e,e'). Let $e = \{a,b\}$ and $e' = \{a',b'\}$ with $a,a' \in Q_1$. Let $f_0 = 1 \to f_1 \to \cdots \to f_{k+1} = e'$ be the left blocking path $P^L_{(e,e')}$. Let $f_i = \{x_i,y_i\}$ for $i \in \{1,\ldots,k\}$ where $x_i \prec y_i$. Then $x_{i+1} \prec y_i$ and $y_i \preceq x_{i+2}$. Thus*

$$x_1 \prec a \prec x_2 \prec y_1 \preceq x_3 \prec y_2 \preceq x_4 \prec \cdots \prec x_k \prec y_{k-1} \prec a' \prec y_k\,.$$

In particular, endpoints of all edges f_i lie on $Q_{1,(e,e')}$.

Proof. Consider two consecutive edges f_i and f_{i+1}. First, we prove that $x_{i+1} \prec y_i$ and $y_i \prec x_{i+2}$. Since f_i and f_{i+1} are in conflict either $x_i \prec x_{i+1} \prec y_i$ or $x_i \prec y_{i+1} \prec y_i$. In the former case, $x_{i+1} \prec y_i$. In the latter case, $x_{i+1} \prec y_{i+1} \prec y_i$. We are done.

Let us now prove that $y_i \preceq x_{i+2}$. Note that since no two non-consecutive nodes of $P^L_{(e,e')}$ are connected with an edge in the conflict graph, edges $f_1,\ldots,f_{k-1}$ are not in conflict with e'. Therefore, $y_i \preceq a'$ for

$i \le k-1$. Suppose to the contrary that $x_{i+2} \prec y_i$ for some i. Consider all pairs (i, j) such that $x_j \prec y_i$ and $j \ge i+2$. By our assumption this set is not empty. Consider pairs with the smallest value of i, and among them choose the pair with the largest value of j. Denote this pair by (i^*, j^*). Note that $j^* \ge i^*+2$, and therefore, f_{i^*} and f_{j^*} are not in conflict. Thus not only $x_{j^*} \prec y_{i^*}$ but also $y_{j^*} \prec y_{i^*}$. Since $y_{j^*} \prec y_{i^*} \prec a'$ and $y_k \succ a'$, we know that $j^* < k$. Using that f_{j^*} and f_{j^*+1} are in conflict, we get that $x_{j^*+1} \prec y_{j^*} \prec y_{i^*}$. That contradicts to our choice of (i^*, j^*). □

Claim 6.7. *Consider a twist (e,e'). Let $e = \{a,b\}$ and $e' = \{a',b'\}$ with $a,a' \in Q_1$. Let $f_0 = 1 \to f_1 \to \cdots \to f_{k+1} = e'$ be the right blocking path $P^R_{(e,e')}$. Let $f_i = \{x_i, y_i\}$ for $i \in \{1,\dots,k\}$ where $x_i \succ y_i$. Then $x_{i+1} \succ y_i$ and $y_i \succeq x_{i+2}$. Thus*

$$x_1 \succ a \succ x_2 \succ y_1 \succeq x_3 \succ y_2 \succeq x_4 \succ \cdots \succ x_k \succ y_{k-1} \succ b \succ y_k.$$

In particular, endpoints of all edges f_i lie on $Q_{2,(e,e')}$.

Proof. Since $P^R_{(e,e')}$ is the shortest among all right blocking paths, no two consecutive nodes of $P^R_{(e,e')}$ are connected with an edge in the conflict graph. The proof of the claim repeats the proof of Claim 6.6. □

Claim 6.8. *Consider two twists (e_1,e'_1) and (e_2,e'_2). Denote the endpoints of e_i by a_i and b_i, the endpoints of e'_i by a'_i and b'_i, so that b_i, b'_i belong to Q_2 (for $i \in \{1,2\}$). Suppose that $e_1 \lhd e'_1 \lhd e_2 \lhd e'_2$. Then, $V(Q_{2,(e_1,e'_1)}) \succeq b'_2$ and $V(Q_{2,(e_2,e'_2)}) \preceq b_1$.*

Proof. Let $f = \{x,y\}$ be the node of $P^R_{(e_1,e'_1)}$ adjacent to e'_1 (with $x \prec y$). Note that e_2 and e'_2 are connected with $P^R_{(e_2,e'_2)}$, a path of odd length in $\mathcal{C}[B \cup L_2]$. Since $\mathcal{C}[B \cup L_2]$ is a bipartite graph, every path between e_2 and e'_2 must have odd length. Thus f cannot be in conflict with both e_2 and e'_2. Moreover, since $b'_2 \preceq b_2 \preceq b'_1 \prec y$, f cannot be in conflict with e'_2 but not with e_2. Thus f is not in conflict with e'_2. Thus $x \succeq b'_2$ and by Lemma 6.7,

$$V(Q_{2,(e_1,e'_1)}) \succeq b'_2\,.$$

Similarly,

$$V(Q_{2,(e_2,e'_2)}) \preceq b_1\,.$$

□

Claim 6.9. *Consider five twists $\{(e_i,e'_i)\}_{i=1,\dots,5}$. Denote the endpoints of e_i by a_i and b_i, the endpoints of e'_i by a'_i and b'_i, so that b_i, b'_i belong to Q_2 (for $i \in \{1,\dots,5\}$). Suppose that $e_i \lhd e_{i+1}$. Then*

$$V(Q_{2,(e_1,e'_1)}) \succ V(Q_{2,(e_5,e'_5)})$$

and therefore we have the following properties.

1. *The paths $P^R_{(e_1,e'_1)}$ and $P^R_{(e_5,e'_5)}$ are disjoint.*

2. *The paths $Q_{2,(e_1,e'_1)}$ and $Q_{2,(e_5,e'_5)}$ are disjoint.*

Proof. From Claim 6.8, we get

$$V(Q_{2,(e_1,e'_1)}) \succeq b'_2 \qquad \text{and} \qquad b_3 \succeq V(Q_{2,(e_5,e'_5)})\,.$$

Note that $b'_2 \neq b_2$ as otherwise there would be no right blocking path between e_2 and e'_2. Therefore, $b'_2 \succ b_2$. We have,

$$V(Q_{2,(e_1,e'_1)}) \succeq b'_2 \succ b_2 \succeq b_3 \succ V(Q_{2,(e_5,e'_5)})\,.$$

Parts (1) and (2) immediately follow. □

Lemma 6.10 (A non-planar arrangement of five twists)**.** *Let $\{(e_i,e'_i)\}_{i=1}^5$ be a collection of disjoint twists in φ_2. For every $i \in \{1,\ldots,5\}$, let $P^L_{(e_i,e'_i)}$ and $P^R_{(e_i,e'_i)}$ be a left and a right blocking path of (e_i,e'_i), respectively. Suppose further that for any $i \neq j \in \{1,\ldots,5\}$, we have*

$$V(P^L_{(e_i,e'_i)}) \cap V(P^L_{(e_j,e'_j)}) = \emptyset\,, \qquad \textit{and} \qquad V(P^R_{(e_i,e'_i)}) \cap V(P^R_{(e_j,e'_j)}) = \emptyset\,.$$

Then, the graph

$$Q_1 \cup Q_2 \cup \bigcup_{i=1}^{5} \left(\{e_i,e'_i\} \cup P^L_{(e_i,e'_i)} \cup P^R_{(e_i,e'_i)}\right)$$

is non-planar.

Proof. For every $i \in \{1,\ldots,5\}$, let $e_i = \{a_i,b_i\}$, $e'_i = \{a'_i,b'_i\}$. Observe that for any $\{z,w\} \in E(P^L_{(e_3,e'_3)})$, with $z \prec w$, we have

$$a_1 \preceq a'_1 \preceq z \prec w \preceq a_5 \preceq a'_5.$$

Similarly, for any $\{z,w\} \in E(P^R_{(e_3,e'_3)})$, with $z \prec w$, we have

$$b_1 \preceq b'_1 \preceq z \prec w \preceq b_5 \preceq b'_5.$$

Thus, $J = Q_1[a_1,a'_5] \cup Q_2[b_1,b'_5] \cup e_1 \cup e'_5 \cup e_3 \cup e'_3 \cup P^L_{(e_3,e'_3)} \cup P^R_{(e_3,e'_3)}$ is a graph with Hamiltonian cycle $C = Q_1[a_1,a'_5] \cup Q_2[b_1,b'_5] \cup e_1 \cup e'_5$. Since $P^L_{(e_3,e'_3)}$ a left blocking path of (e_3,e'_3) it follows that in any partition of the edges in $E(J) \setminus E(C)$ as in Lemma 5.14, the edges e_3, and e'_3 have to be in the same set. On the other hand, since $P^R_{(e_3,e'_3)}$ is a right blocking path of (e_3,e'_3), it follows that e_3 and e'_3 have to be in different sets of the partition. Since this is impossible, we conclude that J is non-planar. Since J is a subgraph of G, it follows that G is also non-planar. □

Lemma 6.11. *Let T be a set of twists in φ_2, satisfying the following conditions.*

1. *For every $(e_1,e'_1) \neq (e_2,e'_2) \in T$, we have $\{e_1,e'_1\} \cap \{e_2,e'_2\} = \emptyset$.*
2. *There exists a collection $\{P^L_{(e,e')}\}_{(e,e')\in T}$ of left blocking paths, where $P^L_{(e,e')}$ is a left blocking path for (e,e') such that segments $Q_{1,(e,e')}$ are disjoint.*

Then, $|T| = O(\mathsf{genus}(G))$.

Proof. By Claim 6.9, we can remove at most three quarters of all twists in T so that for every two distinct remaining twists,

$$V\left(Q_{2,(e_i,e_i')}\right) \cap V\left(Q_{2,(e_j,e_j')}\right) = \emptyset .$$

Therefore, we will assume below that for $i \neq j$,

$$V\left(Q_{2,(e_i,e_i')}\right) \cap V\left(Q_{2,(e_j,e_j')}\right) = \emptyset$$

(and consequently that paths $P^R_{(e_i,e_i')}$ and $P^R_{(e_j,e_j')}$ do not intersect).

Suppose that $|T| \geq 6$, since otherwise the assertion is trivial (recall that we assume that G is non-planar). For any $i \in \{0,\dots,\lfloor |T|/6 \rfloor - 1\}$, let

$$G_i = Q_{1,i} \cup Q_{2,i} \cup \bigcup_{j=1}^{5} \left(e_{5i+j} \cup e'_{5i+j} \cup P^L_{(e_{5i+j},e'_{5i+j})} \cup P^R_{(e_{5i+j},e'_{5i+j})} \right),$$

where $Q_{1,i}$ (resp. $Q_{2,i}$) is the minimal subpath of Q_1 (resp. Q_2) containing all the vertices in

$$\bigcup_{j=1}^{5} \left(e_{5i+j} \cup e'_{5i+j} \cup P^L_{(e_{5i+j},e'_{5i+j})} \cup P^R_{(e_{5i+j},e'_{5i+j})} \right).$$

Observe that for any $i \neq i' \in \{0,\dots,\lfloor |T|/6 \rfloor - 1\}$, we have $V(G_i) \cap V(G_{i'}) = \emptyset$. Therefore, by Lemma 6.11, for any $i \neq i' \in \{0,\dots,\lfloor |T|/6 \rfloor - 1\}$, we have $V(Q_{1,i}) \cap V(Q_{1,i'}) = \emptyset$, and $V(Q_{2,i}) \cap V(Q_{2,i'}) = \emptyset$.

By Lemma 6.10 it follows that each graph G_i is non-planar. Since all the graphs G_i are pair-wise vertex-disjoint subgraphs of G, it follows that $\mathsf{genus}(G) \geq \lfloor |T|/6 \rfloor$, as required. $\square$

We define a partial order on the set of long twists. Let us say that

$$(e_1,e_1') \ll (e_2,e_2') \qquad \text{if} \qquad Q_{1,(e_1,e_1')} \subset Q_{2,(e_2,e_2')} .$$

Then $(T_L,\ll)$ is a partially ordered set. We will prove now that the width of $(T_L,\ll)$ is $O(\mathsf{genus}(G))$ and the depth of $(T_L,\ll)$ is $O(\mathsf{genus}(G))$. Then by Dilworth's theorem, we will conclude that $|T_L| = O(\mathsf{genus}(G)^2)$.

Lemma 6.12. *Let (e_1,e_1') and (e_2,e_2') be incomparable twists in $(T_L,\ll)$. Let $e_1 = \{a_1,b_1\}$, $e_1' = \{a_1',b_1'\}$, $e_2 = \{a_2,b_2\}$, and $e_2 = \{a_2',b_2'\}$ so that $a_1,a_1',a_2,a_2' \in Q_1$. Suppose that $a_1 \prec a_1' \preceq a_2 \prec a_2'$. Then $V(Q_{1,(e_1,e_1')}) \preceq a_2'$ and $a_1 \preceq V(Q_{1,(e_2,e_2')})$.*

Proof. Let f_1 be the node on $P^L(e_1,e_1')$ adjacent to e_1'. Similarly, let f_2 be the node on $P^L(e_2,e_2')$ adjacent to e_2'. Let $f_1 = \{x_1,y_1\}$ with $x_1 \prec y_1$ and $f_2 = \{x_2,y_2\}$ with $x_2 \prec y_2$. Note that y_i is the right end of $Q_{1,(e_i,e_i')}$ for $i \in \{1,2\}$. We want to prove that $y_1 \preceq a_2'$. Assume to the contrary that $a_2' \prec y_1$. Note that then edges f_1 and f_2 are in conflict with e_2'. Therefore, either $f_1 \lhd f_2$ or $f_2 \lhd f_1$. Since e_2 and f_2 are not consecutive nodes on $P^L_{(e_2,e_2')}$, e_2 is not in conflict with f_2. Therefore, it is impossible that $f_2 \lhd f_1$. Thus $f_1 \lhd f_2$. Then $y_2 \preceq y_1$. Now let f_3 be the node on $P^L(e_2,e_2')$ adjacent to e_2. Similarly, since f_1, f_3 are in conflict with e_2,

f_1 is in conflict with e'_2 but f_3 is not in conflict with e'_2, we conclude that $f_1 \lhd f_3$. Therefore, $Q_{1,(e_2,e'_2)}$ lies between x_1 and y_1. In particular, $(e_2,e'_2) \ll (e_1,e'_1)$, which contradicts to our assumption that (e_1,e'_1) and (e_2,e'_2) are not comparable. We conclude that $y_1 \preceq a'_2$. Therefore, $y_1 \preceq V(Q_{1,(e_2,e'_2)})$.

Similarly, we prove that $a_1 \preceq V(Q_{1,(e_2,e'_2)})$. □

Lemma 6.13. *The width of* $(T_L, \ll)$ *is* $O(\mathsf{genus}(G))$.

Proof. Let $(e_1,e'_1),\dots,(e_r,e'_r)$ be a maximal antichain in $(T_L,\ll)$. Assume that $e_1 \lhd e'_1 \lhd e_2 \lhd e'_2 \lhd \dots \lhd e_r \lhd e'_r$. By Lemma 6.12, segments

$$Q_{1,(e_1,e'_1)}, \quad Q_{1,(e_3,e'_3)}, \quad Q_{1,(e_5,e'_5)}, \quad \dots$$

are disjoint. Thus we can apply Lemma 6.11 to twists $(e_1,e'_1),(e_3,e'_3),(e_5,e'_5),\dots$. We get that $r = O(\mathsf{genus}(G))$. □

Lemma 6.14. *Suppose that* $(e_1,e'_1) \gg (e_2,e'_2) \gg (e_3,e'_3)$ *and* $e_1 \lhd e_2$. *Then* $e_1 \lhd e_3$.

Proof. Let $e_i = \{a_i,b_i\}$ and $e'_i = \{a'_i,b'_i\}$ with $a_i,a'_i \in Q_1$. Consider the long left blocking path $P^L_{(e_1,e'_1)}$ for (e_1,e'_1). Let $f = \{x,y\}$ be the node on the path adjacent to e'_1. Suppose w.l.o.g. $x \prec y$. Since $(e_1,e'_1) \gg (e_2,e'_2)$ and $e_1 \lhd e_2$, we have that $Q_{1,(e_2,e'_2)}$ lies between x and y on P. Therefore, $Q_{1,(e_2,e'_2)}$ lies to the right of a_1. Since $Q_{1,(e_3,e'_3)}$ is a subset of $Q_{1,(e_2,e'_2)}$, $Q_{1,(e_3,e'_3)}$ also lies to the right of a_1. Hence a_3 lies to the right of a_1. We conclude that $e_1 \lhd e_3$. □

Lemma 6.15. *The depth of* $(T_L, \ll)$ *is* $O(\mathsf{genus}(G))$.

Proof. Let $(e_1,e'_1) \gg \dots \gg (e_r,e'_r)$ be a maximal chain in $(T_L,\ll)$. Let us say (e_i,e'_i) (for $i<r$) is of the first type if $e_i \lhd e_{i+1}$ and of the second type if $e_i \rhd e_{i+1}$. We will assume that (e_r,e'_r) is both of the first and the second type.

Note that by Lemma 6.14 if e_i is of the first type, then $e_i \lhd e_j$ for every $j > i$. Either there are at least $r/2$ twists of the first type or there are at least $r/2$ twists of the second type. Assume w.l.o.g. that there are at least $r/2$ twists of the first type. Denote them by $(\tilde{e}_1,\tilde{e}'_1) \gg \dots \gg (\tilde{e}_k,\tilde{e}'_k)$ (where $k \geq r/2$).

Now we are going to construct $\lfloor k/10 \rfloor$ disjoint non-planar subgraphs of G. By doing so, we will prove that $r = O(\mathsf{genus}(G))$. (This is similar to what we did in Lemma 6.11.)

We show how to construct the first non-planar graph G_1. Consider the left blocking paths $P^L_{(e_5,e'_5)}$ and $P^L_{(e_6,e'_6)}$. Let f_5 and the node adjacent to e'_5 on $P^L_{(e_5,e'_5)}$, and f_6 be the node adjacent to e'_6 on $P^L_{(e_6,e'_6)}$. Let h be the segment of Q_1 that connects the right endpoints of f_3 and f_5.

The graph G_1 is formed by

- edges e'_1, e_5, e'_5, e_{10},
- the segment of Q_1 between e'_1 and e_{10}, the left blocking path $P^L_{(e_5,e'_5)}$, edges f_5, f_6 and path h,
- the segment of Q_2 between e'_1 and e'_{10}, the right blocking path $P^R_{(e_5,e'_5)}$.

Now we replace the path formed by f_5, h and f_6 with a single edge $\tilde{f}_5$ and obtain a graph G'_1. The graph G'_1 has a Hamiltonian cycle formed by edges e'_1 and e_{10} and segments of paths Q_1 and Q_2 between these two edges. Note that $P^L_{(e_5,e'_5)} - f_5 + \tilde{f}_5$ is a left blocking path for (e_5,e'_5) in G'_1; also $P^L_{(e_5,e'_5)}$ is a right blocking path for (e_5,e'_5) in G'_1. Thus by Lemma 5.14 the graph G'_1 is not planar. So G_1 is not planar.

We construct graphs G_i for $2 \le i \le k/5$ similarly. All graphs G_i are disjoint. Therefore, $r = O(\mathsf{genus}(G))$. □

Lemma 6.16 (Bounding the number of long twists)**.** *We have* $|T_L| = O(\mathsf{genus}(G)^2)$.

Proof. The width of $(T_L,\ll)$ is $O(\mathsf{genus}(G))$ and the depth of $(T_L,\ll)$ is $O(\mathsf{genus}(G))$. Thus by Dilworth's theorem, $|T_L| = O(\mathsf{genus}(G)^2)$. □

Lemma 6.17 (Bounding the number of inner-most left blocking edges)**.** *Let G, P, B_1, B_2, B, γ_1, γ_2, L_1, L_2, H_1, H_2, φ_1, and φ_2 be as in Definition 6.1. Let T_S be the set of all short twists in φ_2. For every twist $(e,e') \in T_S$, let $f_{(e,e')}$ be the inner-most left blocking edge of (e,e'). Let $W = \{f_{(e,e')} : (e,e') \in T_S\}$. Then,* $|W| = O(\mathsf{genus}(G)^2)$.

Proof. For every $f \in W$, pick $(e_f,e'_f) \in T_S$, such that f is the inner-most left blocking edge of (e_f,e'_f).

We begin by constructing a subset $W' \subseteq W$, as follows. Initially, we set $W' = \emptyset$. We use an auxiliary set X, which we initialize to $X = W$. While $X \neq \emptyset$, we pick an edge $f^* \in X$, which is maximal with respect to $\lhd$. Suppose that $f^* = \{x^*,y^*\}$. We add f^* to W', and we remove it from X. For every remaining $g \neq f^* \in X$, with $e_g = \{a,b\}$, $e'_g = \{a',b'\}$, where $a,a' \in V(Q_1)$, we proceed as follows. If either

$$a \preceq x^* \preceq a' \preceq y^* , \tag{6.1}$$

or

$$x^* \preceq a \preceq y^* \preceq a' , \tag{6.2}$$

then we remove g from X. We repeat the above process, until $X = \emptyset$. Observe that every time we consider some $f^* \in X$, we remove at most three edges from X and we add one edge to W' (specifically, we remove edge f^*, possibly one edge for which condition (6.1) holds, and possibly one edge for which condition (6.2) holds; we add f^* to W'). Therefore, for the resulting W' we have

$$|W'| \ge \lceil |W|/3 \rceil .$$

Moreover the set W' has the following property.

(P1) For any distinct $f,g \in W'$, with $f = \{x,y\}$, $e_g = \{a,b\}$, $e'_g = \{a',b'\}$, where $a,a' \in V(Q_1)$, we have that at least one of the following three conditions is satisfied.

$$a \preceq a' \preceq x \prec y$$
$$x \preceq a \preceq a' \preceq y$$
$$x \prec y \preceq a \preceq a'$$

Let ψ be a drawing of G on a surface $\mathcal{S}$ of genus $\gamma = \mathsf{genus}(G)$. Let G' be the graph obtained from G by contracting Q_1 to a single vertex q_1. Let ψ' be the drawing of G' on a surface $\mathcal{S}' = \mathcal{S}/\psi(Q_1)$ of genus γ, obtained by contracting the simple curve $\psi(Q_1)$ to a point x. For every $f \in W'$, let $C_f = \psi'(f)$. Observe that C_f is a cycle in $\mathcal{S}'$, which passes through x. Consider the partition

$$W' = \bigcup_{i=1}^{k} W_i',$$

satisfying the following conditions.

1. For any $i \in \{1,\ldots,k\}$, for any $f,g \in W_i'$, the cycles C_f, and C_g are homotopic.

2. For any $i \neq j \in \{1,\ldots,k\}$, for any $f \in W_i'$, $g \in W_j'$, the cycles C_f and C_g are non-homotopic.

For any $i \neq j \in \{1,\ldots,k\}$, and for any $f \in W_i'$, $g \in W_j'$, we have $C_f \cap C_g = x$. By Lemma 3.4, the number of integers $i \in \{1,\ldots,k\}$ such that the cycles C_f, $f \in W_i'$, are noncontractible is at most $6\gamma - 3$. Moreover there exists at most one $i_0 \in \{1,\ldots,k\}$ such that the cycles C_f, $f \in W_{i_0}'$, are contractible. It follows that

$$k \leq 6\gamma - 2.$$

Thus, there exists $i^* \in \{1,\ldots,k\}$ such that

$$|W_{i^*}'| \geq \frac{|W'|}{6\gamma - 2} \geq \frac{|W|}{18\gamma - 6}.$$

Observe that since no two edges in W_{i^*} are in conflict, it follows that the directed graph on W_{i^*} induced by $\lhd$ is a forest. More concretely, let us consider a collection of rooted trees $\mathcal{F} = \{F_i\}_i$, such that $\{V(F_i)\}_i$ is a partition of W_{i^*}, and such that for any $f,g \in W_{i^*}$, we have that f is a child of g in some tree F_i, if and only if $g \lhd f$, and for any $h \in W_{i^*}$ with $h \neq g$, if $h \lhd f$, then $h \lhd g$.

Let $X \subseteq W_{i^*}'$ be the set of all edges $f \in W_{i^*}'$, such that either f is a leaf, or it has exactly one child in the tree F_i that it belongs to. Note that

$$|X| > \frac{|W_{i^*}'|}{2} \geq \frac{|W|}{36\gamma - 12}.$$

We examine all edges $f \in X$, such that f is an internal vertex in some tree $F_i \in \mathcal{F}$. Let g be the unique child of f. Let $f = \{x,y\}$, $g = \{x',y'\}$. We have $x \prec x' \prec y' \prec y$. Consider the cycle K formed by $Q_1[x,x']$, g, $Q_1[y',y]$ and f. Since C_f, and C_g are homotopic, it follows that K is contractible, and therefore bounds a disk $\mathcal{D}$ in $\mathcal{S}$. We remove from G all edges r, such that the interior of the curve $\psi(r)$ is contained in the interior of $\mathcal{D}$ (note that we do not remove any edges from X at this step since f has only one child in the tree). Let $e_f = \{a,b\}$, $e_f' = \{a',b'\}$, with $a,a' \in V(Q_1)$. By property (P1), and by the fact that f is the inner-most left blocking edge of (e_f, e_f'), it follows that either

$$x \preceq a \preceq a' \preceq x' \qquad \text{or} \qquad y' \preceq a \preceq a' \preceq y.$$

If $x \preceq a \preceq a' \preceq x'$, then we replace f with a new edge $f' = \{x,x'\}$. If the edge $\{x,x'\}$ already appears in G we add another copy of $\{x,x'\}$ (we allow multiple edges). We modify the drawing ψ so that the new edge

f' is drawn inside the disk $\mathcal{D}$. This can be done since we have deleted all edges whose image intersects the interior of $\mathcal{D}$. Notice that after doing this, the new inner-most left-blocking edge of (e_f, e'_f) is f'.

We repeat the above process for all edges in X that are internal vertices in some tree $F_i \in \mathcal{F}$. At the end, we obtain a new graph G^*, and a drawing ψ^* of G^* on the same surface $\mathcal{S}$ of genus γ. Let

$$X^* = \{(e_f, e'_f) : f \in X\}.$$

The set X^* is a collection of short twists such that the corresponding left-most blocking edges in G^* appear consecutively on Q_1 (with respect to $\prec$). Formally, for any $(e,e') \in X^*$, let $f'_{(e,e')}$ denote the inner-most left blocking edge of (e,e') in G^*. Recall that, in general, $f'_{(e,e')}$ can be different than $f_{(e,e')}$. Let $(e,e'),(r,r') \in X^*$, and suppose that $f_{(e,e')} = \{x,y\}$, $f_{(r,r')} = \{z,w\}$, with $x \prec y$, $z \prec w$. Then, either $x \preceq y \preceq z \preceq w$, or $z \preceq w \preceq x \preceq y$. Thus, $\prec$ induces (in the natural way) a total ordering of X^*. Let $Y \subseteq X^*$ be the subset of X^* containing all odd-indexed elements in this total ordering. It follows that for any $(e,e'),(r,r') \in Y$, we have that the edges $f'_{(e,e')}$, and $f'_{(r,r')}$ have disjoint endpoints. We apply Lemma 6.11 and get

$$|Y| = O(\gamma).$$

Since $|Y| \geq |X^*|/2$, we conclude that $|W| = O(|X| \cdot \gamma) = O(\gamma^2) = O(\mathsf{genus}(G)^2)$, as required. □

6.2 Untwisting

In this subsection we will show how to draw a comb. We will first transform the drawing φ_1 to a drawing ψ_1. In drawing ψ_1 all edges will be drawn in the same way as in φ_2. Then we will combine these two drawings and obtain a drawing of the comb.

Rerouting left blocking edges. Let φ_1 be a canonical drawing of $Q_i \cup L_i \cup B$ (as in the definition of the comb). We assume that the line ℓ is horizontal. We also assume that the drawing of every edge crosses every vertical line at most once.

Let W be the set of edges f such that f is the inner-most left blocking edge for some short twist. We perform the following transformation for every $f \in W$. We process edges from W one by one. Every time we choose an edge that is minimal with respect to $\lhd$ (if there are several minimal edges, we choose any of them).

Let $f = \{x,y\}$ with $x \preceq y$. Consider the set of all short twists whose inner-most left blocking edge is f. Denote it by $X = \{(e_1, e'_1), \ldots, (e_r, e'_r)\}$, where $e_1 \lhd e'_1 \lhd \cdots \lhd e_r \lhd e'_r$.

Let p'_x be a point on ℓ between $\varphi_1(x)$ and the drawing of the next vertex to the right of x on Q_1; similarly let p'_y be a point on ℓ between $\varphi_1(x)$ and the drawing of the next vertex to the left of y on Q_1. Consider two vertical lines $h_{p'_x}$ and $h_{p'_y}$ that cross ℓ at points p'_x and p'_y, respectively. Let $h^-_{p'_x}$ and $h^-_{p'_y}$ be subrays of $h_{p'_x}$ and $h_{p'_y}$, respectively, that consist of vertices that lie below ℓ. Let Y be the set of all local edges g such that $g \lhd f$ or $g = f$. Note that the drawing of every edge in Y crosses $h^-_{p'_x}$ and $h^-_{p'_y}$. Let f' be the minimal edge in Y with respect to $\lhd$. Each of the edges f and f' crosses the ray h_x at one point. Let h'_x be the segment of $h^-_{p'_x}$ between these crossing points. Similarly, we define h'_y. We puncture the plane along segments h'_x and h'_y and attach a handle that connects these two punctures. Then we redraw the segments of all edges in Y between h'_x and h'_y on this handle.

Then we process the next edge in W (if later we process another edge with a endpoint in x, we choose another point p''_x to the right of p'_x; similarly, if we process another edge with a endpoint in y, we choose another point p''_y to the left of p'_y).

We first reroute all left blocking edges using the procedure described above and obtain a drawing φ'_1. Note that all global edges are drawn in one plane in φ'_1. We describe now how to "untwist" all twists in this drawing and obtain a drawing compatible with φ_2.

x-untwisting. Let ξ be a canonical drawing of H_1 on a plane with attached handles, as above. Let $x \in V(Q_1)$. The *x-untwisting* of ξ is a pair $(\mathcal{S}, \xi')$, where $\mathcal{S}$ is a surface, and ξ is a drawing of H_1 on $\mathcal{S}$, defined as follows. Let h be the vertical line passing through $\xi(x)$. Suppose that the x-coordinate of $\xi(x)$ is χ. Let h^+, h^- be the open subrays of h that are chosen as follows. The ray h^+ intersects all edges in B that intersect h above ℓ, and it does not intersect any other edges. Similarly, h^- intersects all edges in B that intersect h below ℓ, and it does not intersect any other edges (see Figure 7(a)).

If the line h does not cross any local edges then rays h^+ and h^- start at point x (but since the rays are open, point x does not belong to them). In this case, we say that the untwisting is trivial. Otherwise, we say that the untwisting is non-trivial.

Let $\mathcal{S}'$ be the surface obtained by cutting $\mathbb{R}^2$ along h^+, and along h^-. Formally, let $\mathcal{S}'$ be the topological closure of $\mathbb{R}^2 \setminus (h^+ \cup h^-)$ (see Figure 7(b)). Intuitively, $\mathcal{S}'$ is obtained by cutting the plane along h, and by connecting the two resulting surfaces on the point where h meets ℓ. Cutting $\mathbb{R}^2$ along the ray h^+ gives rise to a surface with a boundary consisting of two copies of h^+, which we denote by $\mathsf{left}(h^+)$, and $\mathsf{right}(h^+)$ respectively. Similarly, cutting along h^- gives rise to two boundary rays $\mathsf{left}(h^-)$, and $\mathsf{right}(h^-)$ respectively. Moreover, every point $p \in h^+$, naturally corresponds to two points $\mathsf{left}(p) \in \mathsf{left}(h^+)$, $\mathsf{right}(p) \in \mathsf{right}(h^+)$. Similarly, every point $p \in h^-$, corresponds to two points $\mathsf{left}(p) \in \mathsf{left}(h^-)$, $\mathsf{right}(p) \in \mathsf{right}(h^-)$.

Let X be the set of edges $e \in B$, such that the interior of the curve $\xi(e)$ intersects h. Note that X is a prefix of B with respect to $\prec$. Consider the partition $X = X^+ \cup X^-$, where

$$X^+ = \{e \in X : \xi(e) \cap h^+ \neq \emptyset\}, \qquad \text{and} \qquad X^- = \{e \in X : \xi(e) \cap h^- \neq \emptyset\}.$$

For every $e \in X$, let $p(e)$ be the point where $\xi(e)$ intersects h. Notice that after cutting the plane along two rays, the drawing ξ does not correspond to a valid graph drawing on $\mathcal{S}'$. This is because the edges with images that cross the two deleted rays, are now drawn only partially. Let $X^+ = \{b_1^+, \dots, b_r^+\}$, where the edges are ordered with respect to the y-coordinates of the corresponding intersection points $p(b_i^+)$, $i \in \{1, \dots, r\}$. Similarly, let $X^- = \{b_1^-, \dots, b_s^-\}$, where again the points are ordered with respect to the y-coordinate of the intersection points $p(b_i^-)$, $i \in \{1, \dots, s\}$. For every $e \in X$, we flip around ℓ the part of the image of $\xi(e)$ that lies to the right of h (see Figure 7(c)). Notice that since X is a prefix of B (with respect to $\prec$), this can be done without introducing any crossings. For a point $p \in h$, let $-p$ denote the point in h obtained by negating the y-coordinate of p. Let A_1 be the segment in $\mathsf{left}(h^+)$ between $\mathsf{left}(p(b_1^+))$, and $\mathsf{left}(p(b_r^+))$, and let A_2 be the segment in $\mathsf{right}(h^-)$ between $\mathsf{right}(-p(b_1^+))$, and $\mathsf{right}(-p(b_r^+))$. Similarly, let A_3 be the segment in $\mathsf{left}(h^-)$ between $\mathsf{left}(p(b_1^-))$, and $\mathsf{left}(p(b_s^-))$, and let A_4 be the segment in $\mathsf{right}(h^+)$ between $\mathsf{right}(-p(b_1^-))$, and $\mathsf{right}(-p(b_s^-))$. We add a Möbius band M_1 connecting A_1 with A_2, and a Möbius band M_2 connecting A_3 with A_4. We can now connect the

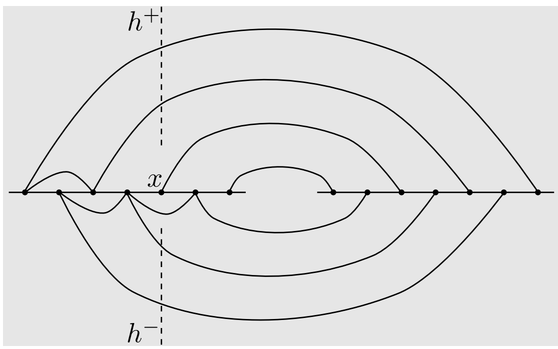

(a) A canonical drawing of H_1, and the rays h^+, and h^-.

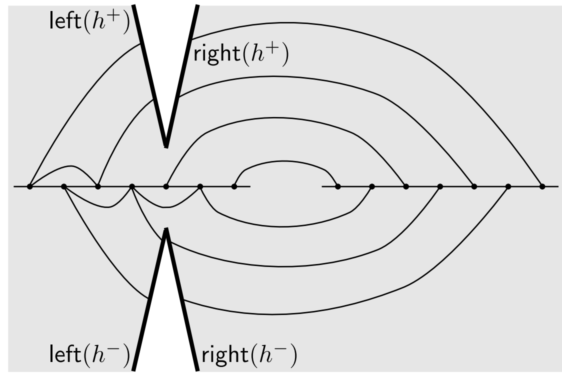

(b) Cutting $\mathbb{R}^2$ along the rays h^+, and h^-.

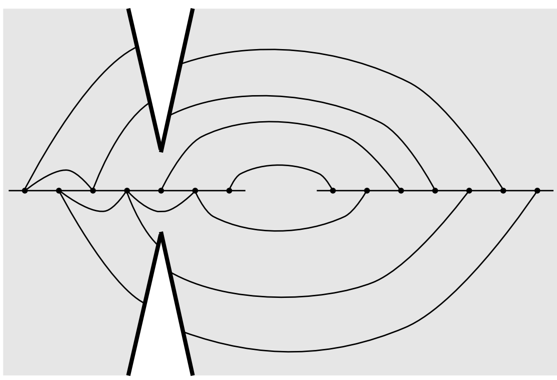

(c) Flipping around ℓ the right parts of the edges that cross the rays.

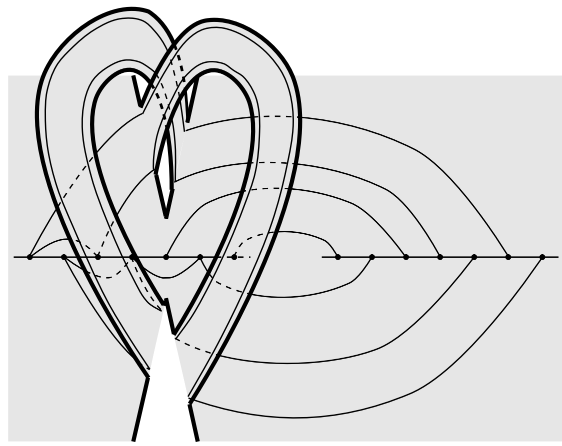

(d) Adding the Möbius bands M_1, and M_2, and completing the drawing of the edges in X along them.

Figure 7: An example of a x-untwisting.

partial drawings of the edges in X^+, by drawing them in M_1. Similarly, we can draw the edges in X^- in M_2 (see Figure 7(d)). Let $\mathcal{S}$ be the resulting surface, and ξ' the resulting drawing of H_1 on $\mathcal{S}$.

$\{x_1,\ldots,x_t\}$-multi-untwisting. Let ξ be a canonical drawing of H_1 on a plane with attached handles (as above), and let $x_1,\ldots,x_t \in V(Q_1)$ be distinct vertices. Then, the $\{x_1,\ldots,x_t\}$*-multi-untwisting* of ξ is a pair $(\mathcal{S},\xi')$, where $\mathcal{S}$ is a surface, and ξ' is a drawing of H_1 on $\mathcal{S}$, defined as follows. Suppose that after reordering of the indices, we have $x_1 \prec x_2 \prec \cdots \prec x_t$. Then, intuitively, the drawing ξ' is obtained starting from ξ, and inductively taking an x_i-untwisting of the current drawing, for $i=1,\ldots,t$. Note that after taking the x_1-untwisting of ξ, the graph is not necessarily drawn on a plane with attached handles, so the precise definition is a bit more subtle. For any $i \in \{1,\ldots,t\}$, let ℓ_t be the vertical line that intersects ℓ at $\xi(x_i)$. We inductively define a sequence $\{(\mathcal{S}_i,\xi_i')\}_{i=0}^{t}$, where $\mathcal{S}_i$ is a surface, and x_i' is a drawing of H_1

on $\mathcal{S}_i$, with $\mathcal{S}_0$ be the surface on which the drawing ξ is drawn, and $\xi'_0 = \xi$. For $i > 0$, $(\mathcal{S}_i, \xi'_i)$ is defined as follows. We cut $\mathcal{S}_{i-1}$ along ℓ_i^+, and ℓ_i^-, as in the untwisting paragraph, we flip along ℓ the parts of the images of the global edges crossing ℓ, to the right of ℓ_i, we add the two Möbius bands M_1, and M_2, and we draw the crossing edges along these two bands. The details are exactly the same as in the untwisting paragraph, and they are therefore omitted. After performing the above operations for all $i = 1, \ldots, t$, we obtain $\mathcal{S}_t$, and ξ_t. We set $\mathcal{S} = \mathcal{S}_t$, and $\xi' = \xi_t$.

Lemma 6.18. *Let φ'_1 be a drawing of H_1 on a surface $\mathcal{S}$, which is a plane with several attached handles (as above). Let $\mathcal{X} \subseteq V(Q_1)$, and let $(\mathcal{S}', \xi')$ be the $\mathcal{X}$-multi-untwisting of ξ. Let r be the number of non-trivial untwistings in $\mathcal{X}$. Then, the surface $\mathcal{S}'$ has genus* $\mathsf{genus}(\mathcal{S}) + O(r)$.

Proof. Note that when we do a trivial untwisting, we do not change the genus of a surface. Each time we do a non-trivial untwisting, we create 2 punctures, and then add 2 Möbius bands between boundary segments. Since adding one such Möbius band increases the genus by at most one, the assertion follows. □

Lemma 6.19. *Let (e, e') be a short twist. Let $e = \{x, y\}$ and $e' = \{x', y'\}$ with $x, x' \in V(Q_1)$ and $x \preceq x'$. Then there is a vertex w on Q_1 that lies between x and x', such that $\varphi'_1(w)$-untwisting of φ'_1 is trivial.*

Proof. Denote the inner-most left blocking edge of (e, e') by f. Let $f = \{a, b\}$ with $a \prec b$. Consider the set of edges $Y = \{g : f \lhd g\}$.

For every point $p \in \ell$, let h_p be the vertical line passing through p. Note that for every point $p \in \ell$ between $\varphi'_1(x)$ and $\varphi'_1(y)$, if the line h_p crosses a local edge g in φ'_1 then $f \lhd g$ since we routed f and all edges g with $g \lhd f$ on handles in φ'_1.

Consider the conflict graph $\mathcal{C}[Y \cup \{e, e'\}]$. Note that e and e' are not connected in $\mathcal{C}[Y \cup \{e, e'\}]$ by a path of length 2 since f is the inner-most left blocking edge for (e, e'); e and e' are not connected in $\mathcal{C}[Y \cup \{e, e'\}]$ by a path of length more than two since (e, e') is a short twist. Therefore, e and e' are disconnected in $\mathcal{C}[Y \cup \{e, e'\}]$. Denote the connected component of $\mathcal{C}[Y \cup \{e, e'\}]$ that e belongs to by C. Let w be the rightmost of endpoints of edges in C. We have, $x \preceq w \preceq y$. Note that if the drawing $\varphi'_1(g)$ of an edge $g \in Y$ crosses $h_{\varphi'_1(w)}$, then g must belong to C and at the same time the right endpoint of g must lie to the right of w, which is impossible. We conclude that $h_{\varphi'_1(w)}$ does not cross any edges in Y, and, therefore, any local edges in φ'_1. Thus $\varphi'_1(w)$-untwisting of φ'_1 is trivial. □

Lemma 6.20 (Drawing a comb)**.** *Let G, P, B_1, B_2, B, L_1, L_2, H_1, H_2, φ_1, and φ_2 be as in Definition 6.1. Then, there exists a polynomial-time algorithm that computes a drawing of $H_1 \cup H_2$ on a surface of genus $O(\mathsf{genus}(G)^2)$. The combinatorial restriction of this drawing to $Q_1 \cup Q_2 \cup L_1 \cup L_2$ is planar and canonical.*

Proof. We first reroute left blocking edges as described above. By Lemma 6.17, we get a drawing φ'_1 of H_1 on a surface of genus $O(\mathsf{genus}(G)^2)$.

Now we construct a set $\mathcal{X}$ so that $\mathcal{X}$-multi-untwisting untwists all twists. We do the following for every twist (e, e'). Let $e = \{x, y\}$ and $e' = \{x', y'\}$ with $x, x' \in V(Q_1)$ and $x \preceq x'$. If the twist (e, e') is a long twist, we choose an arbitrary point p on ℓ between $\varphi'_1(x)$ and $\varphi'_1(x')$ and add it to $\mathcal{X}$. If the twist $\{e, e'\}$ is a short twist, we find a vertex w such that $\varphi'_1(w)$-untwisting is trivial, using Lemma 6.19. Then we add $\varphi'_1(w)$ to $\mathcal{X}$.

Now we perform $\mathcal{X}$-multi-untwisting. The number of non-trivial untwistings equals the number of long twists, and by Lemma 6.16, it is $O(\mathsf{genus}(G)^2)$. Thus by Lemma 6.18, we obtain a drawing ψ_1 of the comb on a surface of genus at most $O(\mathsf{genus}(G)^2)$.

The drawing ψ_1 is consistent with φ_2 in the following sense: every global edge goes above Q_2 in ψ_1 if and only if it goes above Q_2 in φ_2. Thus we can combine drawings ψ_1 and φ_2. We obtain a drawing of the comb on a surface of genus $O(\mathsf{genus}(G)^2)$. □

7 Drawing a Hamiltonian graph

In previous sections, we showed how to obtain $O(g^2)$ embeddings that are almost consistent. In this section, we show how to resolve all remaining inconsistencies to obtain the final embedding.

Definition 7.1 (Internal and marginal edges)**.** Let G be a graph, and let $S \subseteq V(G)$. An edge $\{u,v\} \in E(G)$ is called *S-marginal* if and only if exactly one of u and v is in S; it is called *S-internal* if and only if both u and v are in S.

Definition 7.2 (Extended graph)**.** Let G be a graph, and let $S \subseteq V(G)$. Let $M \subseteq E(G)$ be the set of all S-marginal edges. We denote by $G[S]$ the subgraph of G induced by S. Let J be the graph obtained starting from $G[S]$, and adding for every $e = \{u,v\} \in M$, with $v \in S$, a new vertex u_e, and the edge $\{u_e,v\}$. We refer to e' as the *copy* of e in J. Formally, we define J to be the graph with

$$V(J) = S \cup \bigcup_{e\in M} \{u_e\}, \qquad \text{and} \qquad E(J) = E(G[S]) \cup \bigcup_{e=\{u,v\}\in M, v\in S} \{u_e,v\}.$$

With this notation, the edge $\{u_e,v\}$ is the *copy* of e in H. We refer to J as the *S-extended* graph (with respect to G).

Definition 7.3 (Extended drawing)**.** Let G be a graph, let H be a subgraph of G, let $S \subseteq V(H)$, and let φ be some drawing of H. Let φ' be a drawing obtained from φ by embedding each $V(H)$-marginal edge in a face that contains one of its endpoints (recall that each marginal edge is attached to a leaf). Let φ'' be the drawing obtained by φ' by removing the vertices in $V(H)\setminus S$. We say φ'' is a S-extended drawing induced from φ.

Definition 7.4 (Agreement)**.** Let G be a graph, and let $S_1, S_2 \subseteq V(G)$. For any $i \in \{1,2\}$, let ψ_i be an S_i-extended drawing. We say that ψ_1 and ψ_2 *agree* on a vertex $v \in S_1 \cap S_2$ if and only if the cyclic orderings assigned to edges incident to v in ψ_1 and ψ_2 become identical after identifying each edge in the extended graph with its copy in G. We say that ψ_1 and ψ_2 agree if they agree on all vertices in $S_1 \cap S_2$.

Definition 7.5 (Interval of marginal edges)**.** Let G be a graph, and let P be a Hamiltonian path in G. Let $S \subseteq V(G)$, let M be the set of S-marginal edges, and let ψ be an S-extended drawing. Let $M' \subseteq M$, and let P' be a subpath of P. We say that (M',P') is an *(ψ,P)-interval of S-marginal edges* if and only if the following conditions are satisfied:

1. All edges in M' are incident to P'.
2. All edges in M' appear on the same side of P in ψ.

3. All edges in M' are drawn inside a single face of ψ.

4. Let $e \in M \setminus M'$. Then, either e is not incident to P', or it is incident to one of the endpoints of P', or it appears on the opposite side of P than the edges in M' in the drawing ψ.

5. Let F be the face in ψ that contains all edges in M', and let $\prec$ be the total ordering of the edges incident to P', induced by a traversal of F. Then, all edges in M' form a contiguous interval of $\prec$.

The path P' is called the *subpath* of the interval (M', P'). Two intervals are called *disjoint* if their subpaths are vertex disjoint.

Definition 7.6 (Essential Drawing of a sub-band)**.** Let G be a graph of genus g, and let P be a Hamiltonian path in G. For $1 \le i \le 2$, let (B_i, Q_i) be a band, L_i be its local edges. Let $B = B_1 \cap B_2$, $H = Q_1 \cup Q_2 \cup B \cup L_1 \cup L_2$ and φ be an embedding of H. For $1 \le i \le 2$ we make the following definitions.

1. $E_a[i]$ and $E_b[i]$ to be the subset of edges of B that are incident to Q_i from above and below, respectively.

2. $x_a[i]$ and $y_a[i]$ to be the first and last vertex incident to $E_a[i]$, with respect to $\prec$ on Q_i; similarly, $x_b[i]$ and $y_b[i]$ are the first and last vertex incident to $E_b[i]$.

3. $Q_a[i]$ to be the minimal subpath of Q_i that contains $x_a[i]$ and $y_a[i]$, and $Q_b[i]$ to be the minimal subpath of Q_i that contains $x_b[i]$ and $y_b[i]$.

4. $V_R[i] = V(Q_a[i]) \cup V(Q_b[i])$, note that $V_R[i]$ is composed of the vertices of at most 2 subpaths of Q_i.

5. $V_R = V_R[1] \cup V_R[2]$.

6. The *B-essential drawing* induced from φ, to be the V_R-extended drawing induced from φ.

Lemma 7.7 (Path splitting)**.** *Let G be a graph of genus g and α_1, α_2 and γ be three mutual vertex disjoint paths in G such that each edge in $E(G) \setminus E(\gamma)$ that is incident to γ is also incident to $\alpha_1 \cup \alpha_2$. Let H be a graph satisfying the following conditions.*

1. *H contains two copies of $V(\gamma)$, V_1 and V_2, and one copy of all vertices in $V(G) \setminus V(\gamma)$. For any vertex $x \in V(\gamma)$ we write $x[i]$ to denote the copy of x in V_i, $1 \le i \le 2$.*

2. *For any $\{u, v\} \in E(G)$ if $u, v \in V(\gamma)$ then $\{u[1], v[1]\}, \{u[2], v[2]\} \in E(H)$.*

3. *For any $\{u, v\} \in E(G)$ if $u, v \notin V(\gamma)$ then $\{u, v\} \in E(H)$.*

4. *For any $\{u, v\} \in E(G)$ and $1 \le i \le 2$ if $u \in V(\alpha_i)$ and $v \in V(\gamma)$ then $\{u, v[i]\} \in E(H)$.*

Then, H has genus $O(g)$.

Proof. Let φ be an embedding of G on a surface of genus g and consider the cyclic order of edges around γ in φ. We say two edges are homotopic if and only of they are homotopic after contracting α_1, α_2 and γ. For each edge e incident to γ we say that e has type one if it is incident to α_1 and has type two if it is incident to α_2. Observe that the set of type one edges (and similarly the set of type two edges) fall into $O(g)$ homotopy classes. Since two homotopic edges, α_i ($i \in \{1,2\}$) and γ form the boundary of a topological disc, different homotopy classes cannot interleave. It follows that in the cyclic order of edges around γ, switching between type one and type two edges happen $O(g)$ times.

Now, to build an embedding of H of genus $O(g)$, we add γ', a copy of γ, and for each maximal consecutive set of type one edges we disconnect them from γ and add one handle to reroute them and connect them to γ'. □

Lemma 7.8. *Let G be a graph of genus g, and let P be a Hamiltonian path in G. Let $B_1, B_2 \subseteq E(G)$ be bands with spine P, and with primary segments Q_1, and Q_2, respectively. Assume that $B = B_1 \cap B_2$. Let $H = Q_1 \cup Q_2 \cup L_1 \cup L_2 \cup B$. Then, there is an embedding φ of H on a surface of genus $O(g^3)$ such that the B-essential drawing induced from φ has all its $V(H)$-marginal edges on a collection of at most $O(g^3)$ disjoint intervals.*

Proof. Let M be the set of $V(H)$-marginal edges. Let M_{out} be the set of vertices such that each $v \in M_{\text{out}}$ is an endpoint of at least one marginal edge and $v \notin Q_1 \cup Q_2$. M_{out} can be decomposed into at most 3 disjoint segments, γ_1, γ_2 and γ_3, of P that are disjoint from Q_1 and Q_2.

Let H' be the graph obtained by adding the edges M to H. We use Lemma 7.7 to obtain the genus $O(g)$ graph H'' by splitting γ_1, γ_2 and γ_3. We obtain paths γ_i' and γ_i'', for $1 \le i \le 3$, so that there is no edge between γ_i' and Q_2 and there is no edge between γ_i'' and Q_1. It follows that we can extend Q_1 and Q_2 and treat marginal edges as local edges in H''.

Now we use Lemma 5.1 to obtain a drawing of H'' on a surface of genus $O(g^3)$. Since the marginal edges form a cut between $Q_1 \cup Q_2$ and the rest of the graph they are on at most $O(g^3)$ handles, and so $O(g^3)$ segments. □

Lemma 7.9 (Simple subband drawing)**.** *Let G be a graph of genus g, and let P be a Hamiltonian path in G. Let $B_1, B_2 \subseteq E(G)$ be bands with spine P, and with primary segments Q_1, and Q_2, respectively. Assume that $B = B_1 \cap B_2 \neq \emptyset$. Then, there exists $S \subseteq V(G)$, and an S-extended drawing ψ, satisfying the following conditions.*

1. *For any edge $\{u,v\} \in B$, both u and v are in S.*
2. *$P[S]$ is composed of at most four subpaths of $Q_1 \cup Q_2$.*
3. *ψ has genus $O(g^3)$.*
4. *Let $M \subseteq E(G)$ be the set of S-marginal edges. Then, there exists a collection of (ψ, P)-intervals $(M_1, P_1), \ldots, (M_k, P_k)$, for some $k = O(g^3)$, such that $M \setminus E(P) = \bigcup_{i=1}^{k} M_i$.*

Moreover, S and ψ can be computed in polynomial time.

Proof. We build an auxiliary graph X with genus $O(g)$ that is composed of two bands. We find an embedding of X, φ_X, and use it as a guide to build S and an S-extended embedding ψ.

Let P', Q_1, P'', Q_2 and P''' be vertex disjoint subpaths of P such that $P' \preceq Q_1 \preceq P'' \preceq Q_2 \preceq P'''$ and $V(P') \cup V(Q_1) \cup V(P'') \cup V(Q_2) \cup V(P''') = V(P)$.

We build X as follows. First we make two copies of P', P_1' and P_2', two copies of P'', P_1'' and P_2'', and two copies of P''', P_1''' and P_2'''. We build the path α_1 by connecting P_1', Q_1, P_1'' and P_1''' in this order and the path α_2 by connecting P_2', P_2'', Q_2 and P_2''' in this order; here the result of connecting two paths β and γ is the path obtained by identifying the last vertex of β with the first vertex of γ.

For each edge $(u,v) \in G$, (i) if $u,v \in V(Q_1) \cup V(Q_2)$ then $(u,v) \in E(X)$; observe that there is only one copy of u and one copy of v in $V(X)$, (ii) if $u \in V(Q_i)$ $(1 \le i \le 2)$ and $v \in V(P') \cup V(P'') \cup V(P''')$ then we connect u to the copy of v that is in $V(P_i') \cup V(P_i'') \cup V(P_i''')$ in X.

Since the genus of G is g there is an embedding φ^* of $G[B_1 \cup B_2 \cup L_1 \cup L_2 \cup Q_1 \cup Q_2]$ on a surface of genus $O(g)$. We use φ^* to build an embedding φ_X^* of X that has genus $O(g)$. We delete at most 4 edges from φ^* to disconnect P', Q_1, P'', Q_2 and P'''. Then, we split each of P', P'' and P''' using Lemma 7.7. Finally for $1 \le i \le 2$ we add edges to connect P_i' to Q_i, Q_i to P_i'' and P_i'' to P_i''' by adding a constant number of handles. We also add an edge between the first vertices of P_1' and P_2' to ensure that X is still Hamiltonian, which is useful when we want to apply Lemma 7.8. So the genus of X is $O(g)$.

As X is Hamiltonian, it has genus $O(g)$, and (B,α_1) and (B,α_2) are bands in X, we can use Lemma 7.8 to obtain a drawing φ of X on a surface of genus $O(g^3)$. Then we set $S = V_R$, the essential vertex set of B, and ψ to be the B essential drawing of φ. It is easy to check that S and ψ satisfy the required properties in the lemma. □

Lemma 7.10 (Subband drawing). *Let G be a graph of genus g, and let P be a Hamiltonian path in G. Let $B_1, B_2 \subseteq E(G)$ be bands with spine P, and with primary segments Q_1, and Q_2, respectively. Assume that $B = B_1 \cap B_2 \neq \emptyset$. Then, there exist $S_1, S_2, \ldots, S_r \subseteq V(G)$ and for each $1 \le i \le r$ an S_i-extended drawing ψ_i, satisfying the following conditions.*

1. *$r = O(1)$.*
2. *For any edge $\{u,v\} \in B$, there is an $1 \le i \le r$ such that both u and v are in S_i.*
3. *For all $1 \le i \le r$, $P[S_i]$ is composed of at most four subpaths of $Q_1 \cup Q_2$.*
4. *For all $1 \le i \le r$, ψ_i has genus $O(g^3)$.*
5. *For all $1 \le i \le r$, let $M_i \subseteq E(G)$ be the set of S_i-marginal edges. Then, there exists a collection of disjoint (ψ_i, P)-intervals $(M_{i,1}, P_{i,1}), \ldots, (M_{i,k}, P_{i,k})$, for some $k = O(g^3)$, such that $M_i = \bigcup_{j=1}^{k} M_{i,j}$.*

Moreover, S_i's and ψ_i's can be computed in polynomial time.

Proof. For $1 \le i \le 2$ the band (B_i, Q_i) has a primary segment and $3 \le t_i \le 4$ outlets $P_{i,1}, \ldots, P_{i,t_i}$. For each possible pair of $\alpha \in \{Q_1 \cap P_{2,i}\}_{1 \le i \le t_1}$ and $\beta \in \{Q_2 \cap P_{1,i}\}_{1 \le i \le t_2}$, we use the algorithm of Lemma 7.9 to obtain a $(V(\alpha) \cup V(\beta))$-drawing.

There are at most 16 possible pairs of α and β to consider, so $r = O(1)$. Property (2) holds because we are considering every possible intersections. The other properties are implied by Lemma 7.9. □

Lemma 7.11 (Decomposition). *Let G be a graph of genus g, and let P be a Hamiltonian path in G. Then, we can compute in polynomial time a collection $S_1, \ldots, S_k$ of subsets of $V(G)$, and for every $i \in \{1, \ldots, k\}$, an S_i-extended drawing ψ_i, so that the following conditions are satisfied.*

1. *$k = O(g^2)$.*
2. *For any $e \in E(G)$, there exists $i \in \{1, \ldots, k\}$ such that e is S_i-internal.*

3. *For any $i \in \{1,\dots,k\}$, $P[S_i]$ has at most four connected components.*

4. *For any $i \in \{1,\dots,k\}$, ψ_i has genus $O(g^3)$, and all the S_i-marginal edges are on $O(g)$ disjoint (ψ_i, P)-intervals. Moreover, for each $j \in \{1,\dots,j\}$ the set of S_i-marginal edges that are S_j-internal, are on $O(g)$ disjoint (ψ_i, P)-intervals.*

Proof. We compute a band covering $\mathcal{B} = \{(B_i, Q_i)\}_{i=1}^{t}$ of G using the algorithm of Lemma 2.5; here, $t = O(g)$.

Using $\mathcal{B}$ we build a collection that is composed of two types of vertex subsets, we call them global and local.

To make the global subsets, for each pair $1 \le i, j \le t$ with $B_i \cap B_j \neq \emptyset$ we use Lemma 7.10 to obtain an $O(1)$ number of subsets $\{S_1, S_2, \dots, S_r\}$ and their extended drawings, satisfying the properties stated in the lemma. Note that each global edge of $\mathcal{B}$ is an internal edge of some S_i. However, there may be local edges that are not internal (or even marginal) edges for any S_i. We build local sets to include such local edges.

For each $1 \le i \le t$, consider the planar embedding φ_i of $(B_i \cup L_i \cup (P \setminus \gamma_i))$. We color the edges of B_i so that two edges are assigned the same color if and only if they are homotopic (when the endpoints of Q_i are treated as punctures) and they are both in B_j for some $j \neq i$. Since there are a constant number of homotopy classes in each band, and there are $O(g)$ bands, it follows that the required number of colors is $O(g)$.

In φ_i and on each side of Q_i we define a set of *segments*, each to be the minimal subpath of Q_i that is incident to all edges of a certain color on that side of Q_i; a segment inherits the color of its incident global edges. By the definition of homotopy the segments on each side are edge disjoint. Observe that there may be vertices on Q_i that are not assigned any color on one or both sides. To cover each side of Q_i we introduce $O(g)$ segments of *void* color on each side to cover the gaps between already colored. We say that a subpath of Q_i is *colorful* on a certain side if and only if it does not intersect any segment of void color on that side.

To compute the local subsets we remove all maximal subpaths of Q_i that are colorful on both sides to obtain $O(g)$ candidate subpaths. We say that two candidate subpaths are connected if and only if there is an edge on $E(G) \setminus E(P)$ with one endpoint incident to each of them. Planarity of φ implies that each candidate subpath is connected to at most one other candidate subpath.

We start with the vertex set of $O(g)$ candidate subpaths and merge two such sets if their corresponding subpaths are connected. So, we acquire a collection of $O(g)$ subsets each composed of at most 2 subpaths of Q_i. By construction the genus of any local set is zero and all its marginal edges are on $O(g)$ disjoint intervals. The union of all global and local sets together with the embeddings is a collection of vertices and extended embeddings satisfying all of the required properties. $\square$

Lemma 7.12 (Conflict resolution)**.** *Let G be a graph, and let P be a Hamiltonian path of G. Let $S_1, S_2 \subseteq V(G)$, and M_i, I_i and ψ_i be the set of marginal edges, the set of internal edges and S_i-extended embeddings, $1 \le i \le 2$, all satisfying the following properties.*

1. *$P[S_i]$ is composed of at most four subpaths of P.*

2. *ψ_i has genus $\le \gamma$ and all marginal edges M_i are on $\le \alpha$ disjoint (ψ_i, P)-intervals.*

3. *In ψ_1 the marginal edges $M_1 \cap I_2$ can be covered by $\leq \alpha$ disjoint (ψ_1, P)-intervals. Similarly, in ψ_2 the marginal edges $M_2 \cap I_1$ can be covered by $\leq \alpha$ disjoint (ψ_2, P)-intervals.*

Then, in polynomial time the extended embedding ψ_1' of genus $O(\gamma+\alpha)$ that agrees with ψ_2 can be computed. Further, all edges in $M_1 \cap I_2$ can be covered by $\leq 2\alpha$ disjoint intervals in ψ_1'

Proof. Let $A = \{a_1, a_2, \cdots, a_p\}$ be the set of intervals of ψ_1 that covers $M_1 \cap I_2$ and $B = \{b_1, b_2, \cdots, b_q\}$ be the set of intervals of ψ_2 that covers $M_2 \cap I_1$. Then, we consider, $\Pi = \{\pi_1, \pi_2, \ldots, \pi_r\}$, the set of the mutual intersection of $a \in A$ and $b \in B$. Both A and B are composed of intervals and by assumption $p \leq \alpha$ and $q \leq \alpha$. It follows that $r \leq 2\alpha$.

Observe that π_i's are the only possible subpaths of P on which the embeddings $\psi_1[I_1 \cup I_2]$ and $\psi_2[I_1 \cup I_2]$ may not agree. Further, on each π_i edges in M_1 and M_2 are on two different sides in both ψ_1 and ψ_2.

Now we build a new extended embedding ψ_1' of S_1 that agrees with ψ_2. We start with ψ_1 and a copy ψ' of $\psi_2[S_1 \cap S_2]$ and modify ψ_1 by rerouting all the edges in $I_1 \cap M_2$ to go to ψ' and deleting $\psi_1[S_1 \cap S_2]$ at the end of the day. By assumption ψ_1 and ψ' both have genus at most γ, we use $\leq 2\alpha+4$ handles or Möbius bands for rerouting the edges, so ψ_1' will have genus $\leq 2\gamma+2\alpha+8$.

We first reroute the edges of $I_1 \cap M_2$ that are in P. Since, each of $P[S_1]$ and $P[S_2]$ are composed of at most 3 subpaths, $P[S_1 \cap S_2]$ is composed of at most 6 subpaths. Therefore the edges of P that are in $I_1 \cap M_2$ can be rerouted by adding at most 12 handles.

We reroute the edges of $I_1 \cap M_2$ that are not in P by considering π_i's in turn. For a fixed π_i we know that all the edges in $I_1 \cap M_2$ are connected to π_i from a single side in both ψ_1 and ψ_2 (and so ψ'). So we can reroute that set of edges by using a handle or a Möbius band. Thus, we need at most 2α handles or Möbius bands to reroute all such edges.

Since no subpath of Π is incident to any marginal edge of $M_1 \backslash I_2$, all other marginal edges are still on α disjoint intervals after the above surgery in ψ_1'.

It follows from the construction that all marginal edges $M_1 \cap I_2$ are on 2α intervals of ψ_1'. □

Theorem 7.13 (Main result)**.** *There exists a polynomial-time algorithm which given a graph G of orientable genus g, and a Hamiltonian path in G, outputs a drawing of G on a surface of either orientable, or non-orientable genus $O(g^7)$.*

Proof. First we use Lemma 7.11 to compute $\{S_1, S_2, \ldots, S_k\}$ and $\{\psi_1, \psi_2, \ldots, \psi_k\}$ satisfying the conditions of Lemma 7.11.

Then, for each pair of S_i and S_j we resolve the conflict using Lemma 7.12. Let the extended embedding we obtain for each S_i at the end of all conflict resolutions be ψ_i'.

The genus of the extended embedding of each S_i increases by $O(g)$ after each conflict resolution and $k = O(g^2)$, so the genus of each ψ_i' is $O(g^3)$. Further, for all $1 \leq i, j \leq k$ the set of edges $M_i \cap I_j$ are on $O(g^3)$ disjoint intervals of ψ_i'. It follows that M_i are on $O(g^5)$ disjoint intervals of ψ_i'.

Since $\{\psi_1', \psi_2', \ldots, \psi_k'\}$ mutually agree, they collectively describe an embedding of G. On the other hand, each ψ_i' has genus $O(g^3)$ and can be disconnected from the rest of G by cutting along $O(g^5)$ handles or Möbius bands. It follows that the genus of the entire embedding is $O(kg^5) = O(g^7)$. □

References

[1] CHANDRA CHEKURI AND ANASTASIOS SIDIROPOULOS: Approximation algorithms for Euler genus and related problems. In *Proc. 54th FOCS*, pp. 167–176. IEEE Comp. Soc. Press, 2013. [doi:10.1109/FOCS.2013.26, arXiv:1304.2416] 2, 3

[2] JIANER CHEN, SAROJA P. KANCHI, AND ARKADY KANEVSKY: A note on approximating graph genus. *Inform. Process. Lett.*, 61(6):317–322, 1997. [doi:10.1016/S0020-0190(97)00203-2] 2

[3] ION S. FILOTTI, GARY L. MILLER, AND JOHN H. REIF: On determining the genus of a graph in $O(v^{O(g)})$ steps. In *Proc. 11th STOC*, pp. 27–37. ACM Press, 1979. [doi:10.1145/800135.804395] 2

[4] JONATHAN L. GROSS AND THOMAS W. TUCKER: *Topological Graph Theory*. Courier Corporation, 1987. 2

[5] ALLEN HATCHER: *Algebraic Topology*. Cambridge Univ. Press, 2002. Available at the author's webpage. 2

[6] JOHN E. HOPCROFT AND ROBERT ENDRE TARJAN: Efficient planarity testing. *J. ACM*, 21(4):549–568, 1974. [doi:10.1145/321850.321852] 2

[7] KEN-ICHI KAWARABAYASHI, BOJAN MOHAR, AND BRUCE A. REED: A simpler linear time algorithm for embedding graphs into an arbitrary surface and the genus of graphs of bounded tree-width. In *Proc. 49th FOCS*, pp. 771–780. IEEE Comp. Soc. Press, 2008. [doi:10.1109/FOCS.2008.53] 2

[8] KEN-ICHI KAWARABAYASHI AND ANASTASIOS SIDIROPOULOS: Beyond the Euler characteristic: Approximating the genus of general graphs. In *Proc. 47th STOC*, pp. 675–682. ACM Press, 2015. [doi:10.1145/2746539.2746583, arXiv:1412.1792] 2, 3

[9] YURY MAKARYCHEV, AMIR NAYYERI, AND ANASTASIOS SIDIROPOULOS: A pseudo-approximation for the genus of Hamiltonian graphs. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pp. 244–259. Springer, 2013. [doi:10.1007/978-3-642-40328-6_18] 1

[10] ALEKSANDER MALNIČ AND BOJAN MOHAR: Generating locally cyclic triangulations of surfaces. *J. Combin. Theory Ser. B*, 56(2):147–164, 1992. [doi:10.1016/0095-8956(92)90015-P] 9

[11] BOJAN MOHAR: A linear time algorithm for embedding graphs in an arbitrary surface. *SIAM J. Discrete Math.*, 12(1):6–26, 1999. Preliminary version in STOC'96. [doi:10.1137/S089548019529248X] 2

[12] BOJAN MOHAR: Face covers and the genus problem for apex graphs. *J. Combin. Theory Ser. B*, 82(1):102–117, 2001. [doi:10.1006/jctb.2000.2026] 2

[13] BOJAN MOHAR AND CARSTEN THOMASSEN: *Graphs on Surfaces*. John Hopkins Univ. Press, 2001. 2

[14] NEIL ROBERTSON AND PAUL D. SEYMOUR: Graph minors. VIII. A Kuratowski theorem for general surfaces. *J. Combin. Theory Ser. B*, 48(2):255–288, 1990. [doi:10.1016/0095-8956(90)90121-F] 2

[15] CARSTEN THOMASSEN: The graph genus problem is NP-complete. *J. Algorithms*, 10(4):568–576, 1989. [doi:10.1016/0196-6774(89)90006-0] 2

[16] CARSTEN THOMASSEN: Triangulating a surface with a prescribed graph. *J. Combin. Theory Ser. B*, 57(2):196–206, 1993. [doi:10.1006/jctb.1993.1016] 2

[17] ARTHUR T. WHITE: *Graphs of Groups on Surfaces: Interactions and Models*. Volume 188. Elsevier, 2001. 2

AUTHORS

Yury Makarychev
Associate professor
Toyota Technological Institute – Chicago
yury@ttic.edu
http://ttic.uchicago.edu/~yury/

Amir Nayyeri
Assistant professor
Oregon State University
Corvallis, OR
nayyeria@eecs.oregonstate.edu
http://web.engr.oregonstate.edu/~nayyeria/

Anastasios Sidiropoulos
Assistant professor
The Ohio State University
sidiropoulos.1@osu.edu
http://sidiropoulos.org

ABOUT THE AUTHORS

YURY MAKARYCHEV is an associate professor of computer science at TTIC. He received an MS in Mathematics in 2001 from Moscow State University and a Ph. D. in Computer Science in 2008 from Princeton University, advised by Moses Charikar. Yury served as a postdoctoral researcher at Microsoft Research in Redmond, WA, and Cambridge, MA. Upon completion of the postdoc at Microsoft, Yury joined TTIC in 2009. Yury's research interests include combinatorial optimization, approximation algorithms, and metric geometry.

AMIR NAYYERI received his Ph. D. in computer science in 2012 from the University of Illinois at Urbana-Champaign, advised by Jeff Erickson. He is currently an assistant professor at Oregon State University. His research interest is in algorithm design, particularly for problems in geometry and topology.

ANASTASIOS SIDIROPOULOS, called "Tasos" by his friends and colleagues, is an assistant professor at the Computer Science & Engineering and the Mathematics Departments at The Ohio State University. He received his Ph. D. from the Massachusetts Institute of Technology in 2008, advised by Piotr Indyk. His research focuses on developing algorithms for the analysis of graphs and geometric data sets.