

A Discrete Analog of Tutte’s Barycentric Embeddings on Surfaces

Éric Colin de Verdière*

Vincent Despré†

Loïc Dubois‡

Abstract

Tutte’s celebrated barycentric embedding theorem describes a natural way to build straight-line embeddings (crossing-free drawings) of a (3-connected) planar graph: map the vertices of the outer face to the vertices of a convex polygon, and ensure that each remaining vertex is in *convex position*, namely, a barycenter with positive coefficients of its neighbors. Actually computing an embedding then boils down to solving a system of linear equations. A particularly appealing feature of this method is the flexibility given by the choice of the barycentric weights. Generalizations of Tutte’s theorem to surfaces of nonpositive curvature are known, but due to their inherently continuous nature, they do not lead to an algorithm.

In this paper, we propose a purely discrete analog of Tutte’s theorem for surfaces (with or without boundary) of nonpositive curvature, based on the recently introduced notion of *reducing triangulations*. We prove a Tutte theorem in this setting: every drawing homotopic to an embedding such that each vertex is *harmonious* (a discrete analog of being in convex position) is a *weak embedding* (arbitrarily close to an embedding). We also provide a polynomial-time algorithm to make an input drawing harmonious without increasing the length of any edge, in a similar way as a drawing can be put in convex position without increasing the edge lengths.

1 Introduction

Tutte’s barycentric method and its generalizations. One of the most basic problems in graph drawing is that of constructing straight-line embeddings (crossing-free drawings) of planar graphs [33]. Tutte’s celebrated barycentric embedding theorem [32] (1963) lies at the root of the wide class of force-directed drawing algorithms [26]. It provides a remarkably simple method to build straight-line embeddings of a (simple, 3-connected) planar graph, assuming the knowledge of a facial cycle in some (topological) planar embedding: map the vertices of that cycle to the vertices of a convex polygon, in the same order, and locate the other vertices in such a way that, if the inner edges are replaced by springs, then this physical system is at its equilibrium.

Tutte’s initial proof [32] has been revisited many times; see, in particular, Richter-Gebert [29, Section 12.2], Thomassen [31], or Edelsbrunner [11, Section I.4], thus leading to more insight. As it turns out, it suffices that each inner vertex be a barycenter of its neighbors, with some arbitrary positive coefficients [18]. (The coefficients need not be symmetric, so this is more general than requiring an equilibrium in a spring system, even if the springs may have different rigidities.) For this reason, Tutte drawings are called *barycentric*, and they can be computed by solving a system of linear equations. Equivalently, it suffices that each vertex be in *convex position* with respect to its neighbors: disregarding some degenerate cases, every straight line containing an inner vertex v sees edges incident to v on both sides. (The 3-connectivity assumption avoids overlaps.)

Tutte’s method has been seminal not only for graph drawing. Floater and Gotsman [19] and Gotsman and Surazhsky [22] use it to build morphings between two straight-line embeddings of the same planar graph. In a nutshell, one can interpolate between two embeddings in convex position by moving the barycentric coefficients from those of the initial embedding to those of the final embedding; Tutte’s theorem guarantees that the drawing will be an embedding at every step. Other applications include surface parameterization and approximation [18].

Of special interest to us is a particular generalization of Tutte’s theorem to graphs drawn on a surface S without boundary and of non-positive curvature (in particular, such surfaces are not homeomorphic to a sphere). Consider a graph G embedded as (the 1-skeleton of) a triangulation of S (the embedding needs not be geodesic). Y. Colin de Verdière [5, Théorème 2] (see also Hass and Scott [25, Lemma 10.12] and Luo, Wu, and Zhu [28, Theorem 1.6]) proves that if this embedding is deformed by a homotopy (a continuous motion) in such a way that each edge is geodesic and each vertex is in convex position, then it is actually an embedding. A similar result

*LIGM, CNRS, Univ Gustave Eiffel, F-77454 Marne-la-Vallée, France.

†Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France. This author was partially supported by grant ANR-23-CE48-0017 of the French National Research Agency ANR (project Abysm).

‡LIGM, CNRS, Univ Gustave Eiffel, F-77454 Marne-la-Vallée, France.

holds for surfaces with boundary [5, Théorème 3], fixing part of the triangulation to the boundary. Without the requirement to have a triangulation, degenerate cases could occur, so one would not always get an embedding. The assumption on non-positive curvature seems necessary because it ensures that geodesics (locally shortest paths) cross minimally.

As a side note, for the special case of the torus, other graph embedding methods based on Schnyder woods exist [20], but they lack the flexibility of barycentric embeddings, and to our knowledge they have not been extended to higher genus.

Towards a discrete version. In this paper, we introduce a combinatorial analog of Tutte’s barycentric method, in the plane and on surfaces, leading to a purely discrete algorithm. A first motivation comes from the instability of embeddings under small perturbations. Even in the plane, the resolution of Tutte embeddings (the maximum distance between two vertices, or between a vertex and a non-incident edge, assuming the minimum edge length is at least one) can be exponential in the number of vertices of the graph [8, 10], so rounding may easily create crossings. The situation is worse on surfaces. On piecewise-linear surfaces, computing shortest paths exactly requires exact arithmetic to compare sums of square roots, which is not known to be doable in polynomial time [13]. Obviously, the situation is even worse on smooth surfaces of negative curvature, e.g., hyperbolic surfaces.

A second motivation comes from the existing applications of Tutte’s theorem to morphing. In this setting, one cannot expect closed-form descriptions of the morph. In contrast, in a discrete world, one builds a sequence of local moves from an embedding to the other, in which each intermediate step is an embedding, thus computing a “discrete isotopy” between two given embeddings. Our result is a first step towards building such a sequence in which the edge lengths do not vary too much.

More generally, when devising algorithms for topological problems on surfaces, one often faces the following challenge. On the one hand, these problems are, in principle, easily solved by endowing the surface with a metric of negative curvature, which then enjoys useful properties (this is not possible for sporadic surfaces such as the sphere or the torus, but those can usually be dealt with separately). On the other hand, such smooth surfaces are not suitable for discrete algorithms. One thus has to design discrete models that retain the nice properties of smooth surfaces as much as possible. Such a discretization generally takes the following form: one considers a fixed “host” graph (cellularly) embedded on a surface S , and restricts all curves to lie in that graph. The properties of the host graph, and the definition of a length or of a “canonical”/“reduced”/shortest path, vary greatly, but a large body of the literature solves topological problems by following this pattern, and we provide a partial review now.

The earliest instance in this vein is probably Dehn’s algorithm [6] to determine whether a closed curve on a surface is contractible; this algorithm lies at the root of the entire field of small cancellation theory. Much more recent works study the same problem and extensions, using finer discrete models such as *octagonal decompositions* by É. Colin de Verdière and Erickson [4] to compute shortest homotopic curves, *systems of quads*, introduced by Lazarus and Rivaud [27] to test homotopy between curves, refined by Erickson and Whittlesey [17] for the same problem, and reused by Despré and Lazarus [7] to compute the geometric intersection number of curves. Arguably more powerful, and more related to our work, the model of *reducing triangulations*, very recently introduced by É. Colin de Verdière, Despré, and Dubois [3] has been used to decide whether a given drawing of a graph is homotopic to an embedding (and to compute one if it is the case). The host graph is (the 1-skeleton of) a triangulation of a surface (homeomorphic to neither the sphere nor the torus); in this model, curves enjoy the following properties, similar to surfaces of nonpositive curvature: each walk has a unique *reduced* walk homotopic to it, and the property of being reduced is stable by reversal and taking subwalks. Intuitively, each triangle is equilateral, but there are certain consistent rules to “break ties” among equal-length shortest paths.

Under this viewpoint, not every graph that embeds on S will have an embedding in the host graph, because there may be “not enough room” in it. It is thus natural to consider *weak embeddings*, namely, drawings of graphs in the host graph that can be made embeddings on S by an arbitrarily small perturbation. In this spirit, Atikaya, Fulek, and Tóth [1] provide a near-linear time algorithm to decide if a given drawing of a graph is a weak embedding, and if so to produce an actual embedding in a neighborhood of the host graph, encoded combinatorially.

Our contributions. Our discrete model of non-positively curved surfaces is that of reducing triangulations [3]. However, the original model consists of a triangulation with all vertices of degree at least eight, similar to surfaces of negative curvature. We relax this assumption by allowing degree-six vertices, allowing the discrete analog of “flat regions” on surfaces. The original model is quite constrained; for example, for a surface of fixed

genus, there exists only finitely many reducing triangulations homeomorphic to it. In contrast, our extension allows for finer and finer triangulations of various shapes. More importantly, we define *harmonious drawings*, which are a natural discrete analog of drawings in which each vertex is in convex position with respect to its neighbors.

Our central result is a purely discrete analog of Tutte’s theorem and its generalization to surfaces. The version for surfaces without boundary is as follows:

THEOREM 1.1. *Let S be an orientable surface without boundary homeomorphic to neither the sphere nor the torus. Let T be a reducing triangulation of S . Let G be a graph, and let $f : G \rightarrow T^1$ be a harmonious drawing. There is an embedding homotopic to f in S if and only if f is a weak embedding.*

We emphasize surfaces without boundary as they constitute the hardest cases, but we obtain similar results on all the surfaces with boundary, attaching or not vertices of the drawing to the boundary; this includes the case of the disk, similar to Tutte’s original result. In contrast, the case of the sphere is not relevant in this context (it does not admit reducing triangulations, because it cannot be endowed with a metric of nonpositive curvature). Moreover, our results are not valid on the torus; this case is also very particular, since it admits a flat metric, but it turns out that we need nonpositive curvature *and* at least one point of negative curvature; see Lemma 3.2 below. Tutte’s theorem on the special case of the torus has been studied recently by Erickson and Lin [16, Appendix A and Section 1.1.2], who overcome intrinsic difficulties in this case and apply it to morphing; see also Gortler, Gotsman, and Thurston [21].

Last but not least, we provide an efficient polynomial-time algorithm to make a drawing harmonious. Ideally, one would like not only to build a single harmonious drawing, but to have some flexibility similar to the choice of the barycentric weights. Tutte’s original method amounts to minimizing the energy of the physical system, namely, the weighted sum of the squares of the lengths of the edges. Our algorithm actually possesses a stronger property: it proceeds by “local” moves, which never increase the length of any edge of the drawing. In this way, it allows to build many harmonious drawings. In detail:

THEOREM 1.2. *Let S be an orientable surface without boundary homeomorphic to neither the sphere nor the torus. Let T be a reducing triangulation of S , with m edges. Let G be a graph, and let $f : G \rightarrow T^1$ be a drawing of size n . We can compute in $O((m+n)^2n^2)$ time a drawing $f' : G \rightarrow T^1$, harmonious, homotopic to f in S , such that for every edge e of G , the image of e under f' is not longer than under f .*

In contrast to the aforementioned works, our input graph G is arbitrary; we do not need any connectivity requirement, nor require G to be the 1-skeleton of a triangulation. As a byproduct, we remark that together, these two theorems allow us to decide whether an input drawing of a graph is homotopic to an embedding, in polynomial time, although the running time is not as good as the recent algorithm by É. Colin de Verdière, Despré, and Dubois [3]: simply turn it into a homotopic harmonious drawing by Theorem 1.2 and return whether it is a weak embedding using the result of Akitaya, Fulek, and Tóth [1]; correctness follows from Theorem 1.1.

The paper by É. Colin de Verdière et al. is restricted to reducing triangulations with minimum degree at least eight, and does not provide flexibility for computing an embedding if it exists [3, Lemma 4.1 and Proposition 4.2].

Overview of the techniques. After reviewing some preliminaries in Section 2, we define (our generalized version of) reducing triangulations, and introduce the notion of harmonious drawings, in Section 3.

Section 4 is devoted to the proof of our discrete analog of Tutte’s theorem (Theorem 1.1). In spirit, the proof follows some of the steps of previous proofs of Tutte’s theorem [5, 11]: we reduce to the case where f is homotopic to an embedding of a triangulation of S . We then have a continuous map φ from a triangulated copy of S to S itself. We prove that φ is orientation-preserving (or degenerate) on each triangle. We deduce that f can be turned into an embedding f' by homotopy not only in S , but even in the neighborhood Σ of the 1-skeleton of the reducing triangulation. In Σ , we transform f' by isotopy into an embedding arbitrarily close to f , thereby proving that f is a weak embedding. We remark that, since our goal is to prove that f is a weak embedding, we do not have to worry about degenerate cases, which is one of the difficulties in the continuous case. Rather, we *allow* such degeneracies, which leads to challenges with a different flavor.

In Section 5, we prove Theorem 1.2. We provide homotopy moves that can be applied to a drawing whenever it is not harmonious, and we describe a procedure that applies those moves in a specific order so that we can prove termination. Substantial technicalities appear since one of these moves does not decrease the length of the drawing strictly, nor, to our knowledge, any kind of potential that would immediately ensure termination of the algorithm.

Finally, in Section 6, we extend our results to surfaces with boundary; we reduce to the case of surfaces without boundary by filling the boundary components with surfaces (with genus), thereby extending the reducing triangulation.

2 Preliminaries

2.1 Graphs, drawings, and maps. In this paper *graphs* are undirected, but may have loops and multiple edges. We consider both finite and infinite graphs, all locally finite (each vertex has a finite degree), and we use standard notions of graph theory, such as *walks* (finite, semi-infinite, or bi-infinite). Writing the vertices of a bi-infinite walk W as $(\dots w_{-1}, w_0, w_1, \dots)$, we say that w_0 is the *central vertex* of W . The *non-negative part* of W is the semi-infinite walk $(w_0, w_1, \dots)$. We denote by $W_0 \cdot W_1$ the concatenation of two walks W_0 and W_1 , and by W_0^{-1} the reversal of W_0 . Given a closed walk C , and $n \geq 0$, we denote by C^n the n times concatenation of C by itself.

A *drawing* of a graph G on a graph H is a map $f : G \rightarrow H$ that sends every vertex of G to a vertex of H , and every edge of G to a walk (possibly a single vertex) in H . A drawing f is *simplicial* if it sends each edge of G to a vertex or an edge of H . A (graph) *homomorphism* is a drawing that sends every edge of G to an edge of H . Every drawing $f : G \rightarrow H$ *factors* uniquely as a *simplicial* map $\tilde{f} : \tilde{G} \rightarrow H$, where $\tilde{G}$ is a subdivision of G , where every edge e of G whose image walk $f \circ e$ has length $n \geq 2$ is subdivided into a path of length n in $\tilde{G}$, and where e is not subdivided otherwise. Also a simplicial map f *factors* uniquely as a *homomorphism* $\hat{f} : \hat{G} \rightarrow H$ for some graph $\hat{G}$: the graph $\hat{G}$ is obtained from G by contracting the edges mapped to single vertices, then f corresponds naturally to a homomorphism from $\hat{G}$ to H .

Given a map $f : A \rightarrow B$, and $A' \subseteq A$, $B' \subseteq B$ such that $f(A') \subseteq B'$, we occasionally use the notation $f|_{A'}^{B'}$ to denote the restriction of f to A' , corestricted to B' , thus yielding a map from A' to B' .

2.2 Surfaces. We use standard notions of topology, see any textbook [2, 30] for details. All the surfaces we consider but the plane are connected, compact, and orientable, so we omit these adjectives in the sequel. Every surface is determined up to homeomorphism by its *genus* and number of boundary components. A surface is *closed* if its boundary is empty.

On a surface S , a *path* is a map $p : [0, 1] \rightarrow S$; then $p(0)$ and $p(1)$ are the *end-points* of p . Its *reversal* is the map $p^{-1} : [0, 1] \rightarrow S$ that satisfies $p^{-1}(t) = p(1 - t)$. Similarly to walks, we denote by $p_0 \cdot p_1$ the concatenation of two paths p_0 and p_1 . A *loop* with *basepoint* b is a path whose both endpoints equal b . A path is *simple* if it is injective, except, of course, that $p(0) = p(1)$ if p is a loop. The *relative interior* of a simple path p is the image of $(0, 1)$ under p . A *closed curve* is a map from the circle to S ; it is simple if it is injective. On a surface S with boundary, an *arc* is a path that intersects the boundary precisely at its end-points. Paths and closed curves that differ only by their parameterizations are regarded as equal.

A *homotopy* between two paths p_0 and p_1 with the same endpoints is a continuous family of paths with the same endpoints between p_0 and p_1 . A (free) homotopy between two closed curves c_0 and c_1 is a continuous family of closed curves between them; this time, no point is required to be fixed. A loop or closed curve is *primitive* if it is not homotopic to another closed curve concatenated $n \geq 2$ times with itself. Given a graph G and a surface S , a homotopy between two maps $G \rightarrow S$ is a continuous family of maps between them. A homotopy that fixes a subset $X \subset G$ is *relative* to X . A map f is *contractible* if there exists a homotopy of f that turns the image of f into a single point.

An *embedding* of a graph G on a surface S is a “crossing-free drawing”, a continuous injective map from G to S . The *rotation system* of an embedding of G is the cyclic ordering of the edges of G incident to each vertex in the embedding. The *faces* of an embedding of G are the connected components of the complement of its image. A *triangulation* T is a surface together with a graph embedding in which every face is homeomorphic to an open disk and is bounded by three sides of edges. The *1-skeleton* T^1 of T is the graph embedded on T .

The *universal cover* $\tilde{S}$ of a closed surface S distinct from the sphere is the plane equipped with a *covering map* $\pi : \tilde{S} \rightarrow S$ that is a local homeomorphism. A *lift* of a point $x \in S$ is a point in the pre-image $\pi^{-1}(x) \subset \tilde{S}$. Similarly, a lift of a path $p : [0, 1] \rightarrow S$ is a path $\tilde{p} : [0, 1] \rightarrow \tilde{S}$ such that $\pi \circ \tilde{p} = p$. More generally, we can lift drawings in S and triangulations in S in a similar manner.

2.3 Weak embeddings. Let H be a graph embedded in a surface S , and let G be an abstract graph. A drawing $f : G \rightarrow H$ is a *weak embedding* [1] if there are embeddings of G in S that are arbitrarily close to f .




Figure 1: Subdividing the faces of a reducing triangulation.

As noticed by Akitaya, Fulek, and Tóth [1], the property for f to be a weak embedding can be reformulated, revealing that it does not depend on the actual embedding of H on S , but only on G , f , and the abstract graph H together with its rotation system. For this reformulation, we sketch the slightly different, equivalent presentation of $\hat{E}$. Colin de Verdière, Despré, and Dubois [3, Section 2.2]. To ease the reading, we now assume that H has no loop and no multiple edges; the following extends readily to general graphs H . The **patch system** of H is a surface with boundary Σ obtained by “thickening H ”. (This is similar to the *strip system* introduced by Akitaya et al. [1], and to the concept of fat graph [23] and ribbon graph [14]; the only difference is that, in our patch systems, each strip corresponding to a “fat edge” is contracted to a single path.) Each vertex v of H is assigned a closed disk D_v ; the disks D_w of the vertices w adjacent to v are attached to D_v , along pairwise disjoint segments on the boundary of D_v , in the order prescribed by the rotation system of H , and respecting the orientations of the disks. Every segment along which two disks D_v and D_w are glued becomes an arc in Σ **dual to** the edge of H between v and w . A drawing $f : G \rightarrow H \subset S$ is a **weak embedding** if and only if there exists an embedding $\psi : G \rightarrow \Sigma$, in general position, that **approximates** f in the following sense: for every edge e of G , the sequence of arcs of Σ crossed by the path $\psi \circ e$ is dual to the sequence of edges of H taken by the walk $f \circ e$.

2.4 Reducing triangulations. Our discrete model for surfaces is a slight extension of the notion of **reducing triangulation** recently introduced by $\hat{E}$. Colin de Verdière, Despré, and Dubois [3].

We emphasize that the model is restricted to orientable surfaces. A triangulation T of an (orientable) surface is **reducing** if each vertex in the interior of T has degree at least six, and the dual of T is bipartite (the faces of T can be colored red and blue so that adjacent faces receive distinct colors). Every closed surface other than the sphere admits a reducing triangulation: a reducing triangulation for the torus is made of two triangles, one red and one blue, glued together to form a quadrilateral, whose opposite sides are then identified (Figure 23); for higher-genus surfaces we refer to [3, Figure 17]. Every surface with boundary admits a reducing triangulation, as can be seen by considering subcomplexes of reducing triangulations of closed surfaces. (Compared to the original model, our reducing triangulations may have degree-six vertices, which allows in particular to refine an existing reducing triangulation by subdivision; see Figure 1. Also, they may have non-empty boundary.)

Reducing triangulations on surfaces with non-empty boundary will only be considered in Section 6, and the following definitions are relevant only for closed surfaces. A walk (e_1, e_2) of length two in the 1-skeleton T^1 of a reducing triangulation T makes a **turn** at its middle vertex v (see [3, Section 3.1 and Figure 3] for details). This turn is a k -turn, $k \geq 0$, if e_2 results from e_1 by k clockwise rotations of e_1 around v . It is a $-k$ -turn if e_2 results from k counter-clockwise rotations of e_1 around v . To be more precise, (e_1, e_2) makes a k_r -turn (or $-k_r$ -turn) if e_1 sees red on its left, and a k_b -turn (or $-k_b$ -turn) otherwise. The 0-turns, 1-turns, -1 -turns, 2_r -turns, and -2_r -turns are **bad**. All other turns are **good**. (We note that these notions depend on an orientation of the surface, which is why we require surfaces to be orientable.) A walk in T is **reduced** if it makes no bad turn. Then:

LEMMA 2.1. *Let T be a reducing triangulation of the plane. If x and y are vertices of T , then there is a unique reduced walk from x to y in T .*

Note that reducing triangulations of the plane cannot have loops or multiple edges (this would otherwise contradict Lemma 2.1), but reducing triangulations of more complex surfaces may. Lemma 2.1 is proved in [3, Proposition 3.1] on reducing triangulations that do not have degree six vertices. The proof in that paper extends verbatim to our case, only replacing the word “eight” by “six” in the second paragraph of the proof of [3, Lemma 3.2].

LEMMA 2.2. *Let T be a reducing triangulation of the plane. Let C be a simple closed walk in T , not a single vertex. Orient C so that the bounded side of C lies on its left. Then at least three of the turns of C are a 1-turn or a 2_r -turn.*

Lemma 2.2 can be seen as a consequence of Lemma 2.1:




Figure 2: (Left) The closed surface S of genus two, equipped with a reducing triangulation T , and a left line L in T . (Top Right) The universal covering space $\tilde{S}$ of S , i.e. the plane, and a left line $\tilde{L}$ that lifts L in $\tilde{S}$. (Bottom Right) A right line $\tilde{L}'$ in $\tilde{S}$.




Figure 3: (Left) The vertex v and the walk W in the definition of strong harmony. (Top right) The lifts $\tilde{L}$ and $\tilde{p}$, when L makes only 3_r -turns. (Bottom right) The path $\tilde{p}$ slightly unpacked to illustrate that $\tilde{p}$ can go back and forth, and stagnate.

Proof. By contradiction, assume C has two vertices x and y such that C never makes a 1-turn or a 2_r -turn except possibly at x and y . Consider the disk D bounded by C , and let T' be the reducing triangulation of D that is the restriction of T to D . Delete all of T outside D . Then extend T' to a new (infinite) reducing triangulation T'' of the plane (by adding triangles one by one), in which C does not make any -1 -turn, nor any -2_r -turn. Then the two sub-walks of C between x and y are distinct reduced walks in T'' , contradicting Lemma 2.1. $\square$

In this paper we say that a *closed* walk in T is reduced if it makes no bad turn. Note that this definition of reduced closed walks departs from that in [3] since, in our setting, a closed walk that makes only 3_r -turns is reduced. A consequence is that reduced closed walks may not be unique in their free homotopy class: we see the walks depicted in [3, Figure 9] (left and right) as two distinct freely homotopic reduced closed walks.

3 Harmonious drawings

In this section, we provide the key definition of *harmonious drawings* of graphs in a reducing triangulation, used in the statements of our theorems. We actually start with the notion of strongly harmonious drawings, which are the discrete analog of barycentric drawings, in which each inner vertex is drawn in *convex position*: every straight line I containing v sees edges incident to v on both sides, here understood in the *strong sense* that some edges incident to v enter the two *open* half-planes separated by I . Harmonious drawings are a slightly relaxed notion that is suitable for our results.

3.1 Preliminary definitions. Let S be a closed surface not homeomorphic to the sphere, and let T be a reducing triangulation of S . Let $\tilde{T}$ be the (infinite) reducing triangulation that lifts T in the universal cover $\tilde{S}$ of S . In $\tilde{T}$, we consider a left (resp. right) **line** to be a bi-infinite walk $\tilde{L}$ that makes only 3_r -turns (resp. -3_r -turns). See Figures 2 and 3. Note that $\tilde{L}$ is reduced, and is thus simple by Lemma 2.1. A path $\tilde{p}$ starting from the central vertex (Section 2.1) of $\tilde{L}$ **escapes** $\tilde{L}$ if $\tilde{p}$ enters the right (resp. left) side of $\tilde{L}$ at some point, and if the prefix of $\tilde{p}$ before this point is contained in the non-negative part of $\tilde{L}$.

On S , the **lines** are again the bi-infinite walks (not simple this time, since T is finite) that make only 3_r -turns or only -3_r -turns; note that the lines on S lift to the lines on $\tilde{S}$. On S , a path p **escapes** a line L if there are a lift $\tilde{L}$ of L , and a lift $\tilde{p}$ of p starting from the central vertex of $\tilde{L}$, such that $\tilde{p}$ escapes $\tilde{L}$.




Figure 4: A spur.




Figure 5: A graph G , a cycle graph O , and a surface S equipped with a reducing triangulation T (not represented). A simplicial map $f' : G \rightarrow O$ without spur, and a reduced closed walk $C : O \rightarrow T^1$. The composed drawing $C \circ f'$ is harmonious.

Let G and M be graphs, and let $f : G \rightarrow M$ be simplicial. A **cluster** of f is a connected subgraph G_0 of G whose edges are all mapped by f to a single vertex of M , maximal under this condition. A cluster G_0 of f is a **spur** if G_0 is not a connected component of G , and if the directed edges from G_0 to $G \setminus G_0$ are all mapped by f to the same directed edge of M ; see Figure 4. (Note that the direction of the edges matters when M has loops.)

3.2 Strongly harmonious drawings. A simplicial map $f : G \rightarrow T^1$ is **strongly harmonious** if for every vertex v of G , and for every line (left or right) L whose central vertex is $f(v)$ in T , there is a walk W based at v in G whose image path $f \circ W$ escapes L . See Figure 3.

Intuitively, f is strongly harmonious if for every vertex v of G , and for every bi-infinite walk L in T^1 starting at $f(v)$ making only 3_r -turns (resp. -3_r -turns), there is a walk W in G whose image by f may go “forward” on L , sometimes “backward” on L , but not to the point it goes before its central point, and then leaves L to its right (resp. left). But L is non-simple, and periodic since T is finite; W is allowed to wrap, say, 2.3 times forward along L , then 1.8 times backward, and then to leave L to its right (resp. left). Formalizing this phenomenon seems to be more easily captured using the universal covering space, as above.

3.3 Harmonious drawings. Harmony is a slightly relaxed version of strong harmony. On a connected graph G , a simplicial map $f : G \rightarrow T^1$ is **harmonious** if f is strongly harmonious, or if $f = C \circ f'$ for some cycle graph O , some reduced closed walk $C : O \rightarrow T^1$, and some simplicial map $f' : G \rightarrow O$, without spur. See Figure 5. In general, a simplicial map $f : G \rightarrow T^1$ is harmonious if f is harmonious on every connected component of G .

Finally, recall that every drawing $f : G \rightarrow T^1$ factors into a unique simplicial map $\bar{f} : \bar{G} \rightarrow T^1$. We say that f is harmonious (resp. strongly harmonious) if $\bar{f}$ is.

3.4 Remarks. We conclude this section with two remarks. First, we illustrate why strong harmony is not sufficient for our purposes. The reason is that some drawings of graphs in reducing triangulations cannot be made strongly harmonious by homotopy:

LEMMA 3.1. *Let S be a closed surface, not the sphere nor the torus. There are a reducing triangulation T of S , and a closed walk C in T , such that every closed walk freely homotopic to C is not strongly harmonious.*

Second, we give one reason (among others) why our results do not extend to the torus:

LEMMA 3.2. *There are a reducing triangulation T of the torus, and a closed walk C in T , such that every closed walk freely homotopic to C is not reduced.*

The proofs of Lemmas 3.2 and 3.1 are deferred to Appendix A.

4 A Tutte theorem for harmonious drawings: proof of Theorem 1.1

In this section we prove Theorem 1.1, which we restate for convenience:

THEOREM 1.1. *Let S be an orientable surface without boundary homeomorphic to neither the sphere nor the torus. Let T be a reducing triangulation of S . Let G be a graph, and let $f : G \rightarrow T^1$ be a harmonious drawing. There is an embedding homotopic to f in S if and only if f is a weak embedding.*

The “if” part is trivial, so we focus on the “only if” part. If Theorem 1.1 holds on simplicial drawings, then it holds on general drawings by definition. Thus, we assume in this section that drawings are simplicial, with the noticeable exception of some part of Section 4.1 where general drawings will be momentarily required: this will be clarified in due course.

4.1 Harmonious drawings on patch systems. In this section all graphs and drawings are finite, without further mention. We prove the following, which is roughly a Tutte result for patch systems:

PROPOSITION 4.1. *Let M be a graph embedded on a surface. Let G be a graph, and let $f : G \rightarrow M$ be simplicial and without spur. If there is an embedding homotopic to f in the patch system of M , then f is a weak embedding.*

The strategy to prove Proposition 4.1 is to find a sequence of moves (swaps) between the drawing f and a weak embedding, and to prove that applying a move to a weak embedding results in a weak embedding.

In a graph M , we say that a (closed) walk W is **canonical** if W does not use an edge of M and its reversal consecutively. We shall use the three following classical facts, see, e.g., Stillwell [30, Chapter 2]. (1) If two canonical walks are homotopic in M relatively to their end-vertices, then they are equal. (2) If two canonical closed walks are freely homotopic, then they differ by a cyclic permutation. (3) If two loops a and b based at the same point of M commute, i.e., if there is a loop c such that a is homotopic to $c \cdot b \cdot c^{-1}$ (where the homotopy fixes the basepoint), then a and b are homotopic to powers of a common loop.

LEMMA 4.1. *Let G and M be graphs, and let $f : G \rightarrow M$ be simplicial. If G is connected, if f is contractible, and if f has no spur, then $f(G)$ is a single vertex of M .*

Proof. We may assume without loss of generality that f is a homomorphism by contracting the edges that belong to clusters of f . Assume by contradiction that $f(G)$ is not a single vertex of M . Since G is connected, some edge of G is mapped by f to an edge of M . Since f has no spur, there is a semi-infinite walk W in G such that $f \circ W$ is canonical. Since G is finite, there is a subwalk W' of W , not a single vertex, that starts and ends at the same vertex of G . And the loop $f \circ W'$ is non-contractible in M by (1), contradicting the assumption that f is contractible. $\square$

Until the end of this section, we need to consider general drawings (instead of simplicial ones). Recall from Section 2 that every drawing $f : G \rightarrow M$ factors as a simplicial map $\tilde{f} : \tilde{G} \rightarrow M$, where $\tilde{G}$ is a subdivision of G . The clusters and spurs of f are those of $\tilde{f}$. The vertices of $\tilde{G}$ inserted in the edges of G are clusters of $\tilde{f}$. All other clusters of $\tilde{f}$ are subgraphs of G .

LEMMA 4.2. *Let M be a graph embedded on a surface. Let G be a graph, and let $f : G \rightarrow M$ be a drawing. If there is an embedding homotopic to f in the patch system of M , then there is a weak embedding $f' : G \rightarrow M$, homotopic to f , that has no spur.*




Figure 6: In the proof of Lemma 4.2, if f' had a spur, then the number of crossings of g with the arcs of Σ could be decreased.




Figure 7: Swap of the subgraph G_0 along the directed edge s .

Proof. Let Σ be the patch system of M . We regard f as a map from G to Σ . There is by assumption an embedding $g : G \rightarrow \Sigma$ homotopic to f . We can assume that g crosses the arcs of Σ as few times as possible subject to the constraint that it is an embedding homotopic to f . Let $f' : G \rightarrow M$ be the drawing of which g is an approximation. Then f' is a weak embedding homotopic to f .

There remains to prove that f' has no spur. By contradiction, assume that it does. See Figure 6. Let G_0 be the cluster of the spur, and let a be the arc of Σ dual to the directed edge of the spur. Let $E \subset \Sigma$ contain, for every edge e directed from G_0 to $G \setminus G_0$, the prefix of the image path $g(e)$ that leaves $g(G_0)$ to reach its first crossing with a . Let a_1 be the subpath of a that starts just before its first crossing with E and ends just after its last crossing with E . Because g is an embedding, there exists a simple path a_2 in Σ with the same endpoints as a_1 , otherwise disjoint from the arcs of Σ and from $g(G)$, such that the disk bounded by a_1 and a_2 contains $g(G_0)$. We consider a self-homeomorphism h of Σ that affects only a neighborhood of the disk bounded by a_1 and a_2 , and pushes a_2 to a_1 . Then $h \circ g$ is an embedding of G homotopic to g , with less crossings with the arcs of Σ , a contradiction. $\square$

A drawing $f : G \rightarrow M$ is **essential** if there is no connected component of G on which f is contractible. Let $f : G \rightarrow M$ be an essential drawing. Let G_0 be an induced subgraph of G , connected and mapped to a single vertex of T by f . Let s be a directed edge of T based at the vertex $f(G_0)$. Consider the following operation that transforms f homotopically into another essential drawing $f' : G \rightarrow M$ (Figure 7). First slide $f(G_0)$ along s . Then, among the image walks of the edges directed from G_0 to $G \setminus G_0$, those that initially admitted s as a prefix now start with the concatenation of s^{-1} and s : shorten those walks by removing this prefix. In the particular case where f and f' both have no spur (in addition of being essential), we call this operation a **swap**. A key but trivial observation is that, if f' results from a swap of f , then f results from a swap of f' .

Note that if f has no spur, if G_0 is a cluster of f , and if s is the image by f of some edge directed from G_0 to $G \setminus G_0$, then f' has no spur, so the operation is indeed a swap in this case. Every swap constructed in this section will either be such a swap, or the “inverse” of such a swap.

LEMMA 4.3. *Let G and M be graphs, and let $f, f' : G \rightarrow M$ be drawings. If f and f' are essential and without spur, and if they are homotopic, then there is a sequence of swaps between f and f' .*

Proof. We shall perform swaps on f and f' so that in the end $f = f'$. Assume without loss of generality that G is connected, and fix a spanning tree Y of G . We claim that there is a sequence of swaps that modify f so that in the end $f(Y)$ is a single vertex of M . To prove the claim, we say that an edge e of Y is contracted if $f(e)$ is a single vertex of M . We choose an arbitrary root for Y , and given any edge e of Y , we refer to the edges of Y distinct from e and separated from the root by e as the descendants of e . We prove the claim by considering an edge e of Y , by assuming that e is not contracted and that every descendant of e is contracted, and by exhibiting a swap that shortens $f \circ e$ while keeping the descendants of e contracted. Direct e so that the tail vertex of e is the closest to the root of Y . Let G_0 be the cluster containing the head vertex of e . Since the descendants of e are all contracted, they all belong to G_0 . Since e is not contracted, we may consider the last directed edge s of the walk $f \circ e$. The swap of $f(G_0)$ along the reversal of s shortens $f \circ e$, and keeps the descendants of e contracted. This proves the claim.

We use the claim immediately on both f and f' , then contract the spanning tree Y in both drawings. Every edge e of G not in Y becomes a loop. The image loop $f \circ e$ is contractible if and only if $f' \circ e$ is contractible, and in that case $f(e) = f(Y)$ and $f'(e) = f'(Y)$ are single vertices of M by (1). Also given any two edges e_1 and e_2 of G not in Y , we have $f \circ e_1 \simeq f \circ e_2$ if and only if $f' \circ e_1 \simeq f' \circ e_2$, where $\simeq$ denotes the homotopy of loops relatively to their basepoint. In that case $f \circ e_1 = f \circ e_2$ and $f' \circ e_1 = f' \circ e_2$ by (1). By contracting the contractible loops and identifying the homotopic loops in both f and f' , we may assume that every connected component of G has a single vertex (all its edges are loops), and that each of f and f' maps the edges of G to pairwise distinct non-trivial walks in M .

The end of the proof is adapted from the proof of [3, Lemma 8.6], to which we refer for details. Let v be the vertex of G . Since f and f' are essential, the graph G is not a single vertex, so G has a loop e . At this point, $f \circ e$ and $f' \circ e$ are canonical walks, but not necessarily canonical closed walks; however, by performing swaps on f and f' at v a few times if needed, we enforce that $f \circ e$ and $f' \circ e$ are canonical closed walks. Then by (2), and since $f \circ e$ and $f' \circ e$ are freely homotopic, $f \circ e$ is a cyclic permutation of $f' \circ e$. By performing swaps on f again a few times, we enforce that $f \circ e$ and $f' \circ e$ are actually equal (not up to cyclic permutation). At this point, we claim that we can slide $f(v)$ around $f \circ e$ by performing swaps on f , so that in the end f and f' are homotopic relatively to v . This claim implies $f = f'$ by (1), which proves the lemma. There remains to prove the claim. For this, consider the path followed by $f(v)$ during some *free* homotopy from f to f' . This path is a loop based at $f(v)$. Let W be the canonical walk homotopic to it. We prove the claim by showing that $f \circ e$ and W are equal to powers of the same walk. Indeed let C be a primitive canonical closed walk such that $f \circ e$ is freely homotopic to a power of C . Without loss of generality $f \circ e$ is equal to a power of C by (2). Also W commutes with $f \circ e$ since $f \circ e = f' \circ e$, and since $f \circ e$ is homotopic to the walk $W \cdot (f' \circ e) \cdot W^{-1}$. Thus the walk W is homotopic to a power of C by (3). And so W is equal to this power of C by (1). $\square$

LEMMA 4.4. *Let M be a graph embedded on a surface. Let G be a graph, and let $f, f' : G \rightarrow M$ be drawings. If f and f' are essential, if f' results from a swap of f , and if f is a weak embedding, then f' is a weak embedding.*

Proof. Let Σ be the patch system of M , and let $F : G \rightarrow \Sigma$ be an embedding that approximates f in Σ . Let G_0 be the subgraph of the swap, and let a be the arc of Σ dual to the directed edge of the swap. Let $E \subset \Sigma$ contain, for every edge e directed from G_0 to $G \setminus G_0$, the prefix of the image path $F(e)$ that leaves $F(G_0)$ to reach either a vertex of $F(G \setminus G_0)$ in the same face of Σ , or its first crossing with an arc of Σ . Let $E_1 \subset E$ contain the paths that reach a . Then $E_1 \neq \emptyset$, for otherwise G_0 would either be a spur in f' , or it would be a cluster mapped to a single vertex in f' . Let a_1 be the subpath of a that starts just before its first crossing with E_1 and ends just after its last crossing with E_1 . Because F is an embedding, there exists a simple path a_2 with the same endpoints as a_1 , otherwise disjoint from the arcs of Σ , and disjoint from $F(G)$ except for each path in $E \setminus E_1$ that a_2 may cross at most once, such that the disk D bounded by a_1 and a_2 contains $F(G_0)$.

We claim that D does not contain any part of $F(G)$ other than $F(G_0)$ and E . By contradiction assume that it does. Then D contains the image $F(v)$ of a vertex v of $G \setminus G_0$, since G_0 is an induced subgraph of G , and since any edge of $F(G \setminus G_0)$ intersecting D must have an endpoint in D . If a path based at $F(v)$ in $F(G)$ does not intersect a_1 nor $F(G_0)$, then this path stays in the interior of D , and so it does not intersect any arc of Σ . Therefore, in f' , the cluster containing v is either mapped to a single vertex, or is a spur, which is a contradiction.

Now consider a self-homeomorphism h of Σ that affects only a neighborhood of D , and pushes a_2 to a_1 . Then $h \circ F$ is an embedding of G , that approximates f' by the preceding claim. $\square$

The following lemma is easy and might be folklore, but we could not find a reference, so we provide a proof for completeness:

LEMMA 4.5. *Let S be a surface. Let G be a graph, and let $f : G \rightarrow S$ be a map. If f is contractible, and if there is an embedding homotopic to f , then G is planar.*

Proof. Without loss of generality G is connected. Fix a vertex r and a spanning tree T of G . There is an embedding $f' : G \rightarrow S$ homotopic to f . In f' , almost contract the image of T by isotopy, without changing the image of r , in order to push the image of T inside a small neighborhood N of $f(r)$. Every edge e of G not in T is mapped by f' to a simple contractible loop. Those loops can be pushed inside N by isotopy since, by a result of Epstein [15, Theorem 1.7], each of them bounds a disk with only (possibly) contractible loops inside it. $\square$

Finally, we prove Proposition 4.1:

Proof of Proposition 4.1. If G' is a connected component of G on which f is contractible, then $f(G')$ is a single vertex of M by Lemma 4.1, and $f|_{G'}$ can be made an embedding in an arbitrarily small disk in the patch system of Σ by Lemma 4.5. So we can assume that f is essential. By Lemma 4.2, there is a weak embedding $g : G \rightarrow M$, homotopic to f , without spur. By Lemma 4.3, there is a sequence of swaps from g to f . By Lemma 4.4, all the maps from G to M in this sequence are weak embeddings, thus f is itself a weak embedding, as desired. $\square$

4.2 A property of the coherently oriented maps homotopic to the identity. Let S be a surface, and let $Y \subset S$ be finite. A map $\varphi : S \rightarrow S$ is **coherently oriented** at Y if $\varphi^{-1}(Y)$ is finite and if φ is, locally, an orientation-preserving homeomorphism around every point of $\varphi^{-1}(Y)$, or an orientation-reversing homeomorphism around every point of $\varphi^{-1}(Y)$.

The following proposition is proved using the topological notion of the degree of a self-map $\varphi : S \rightarrow S$. It will be used in the case where Y contains one point per face of T , and we will regard $S \setminus Y$ as the patch system of T^1 .

PROPOSITION 4.2. *Let S be a closed surface. Let $\varphi : S \rightarrow S$ be a map, homotopic to the identity map of S . Let $Y \subset S$ be finite. If φ is coherently oriented at Y , then $\#\varphi^{-1}(Y) = \#Y$ and $\varphi|_{S \setminus \varphi^{-1}(Y)}^{S \setminus Y}$ is homotopic to a homeomorphism $S \setminus \varphi^{-1}(Y) \rightarrow S \setminus Y$.*

Proof. Let $y \in Y$. We claim that $\varphi^{-1}(y)$ contains only one point, and that φ is orientation-preserving around this point. To prove this claim, let n^+ and n^- denote the number of points of $\varphi^{-1}(y)$ around which φ is respectively orientation-preserving and orientation-reversing. The difference $n^+ - n^-$ does not depend on the choice of y (as long as y is chosen so that φ is locally a homeomorphism around every point of $\varphi^{-1}(y)$) and is known as the degree of φ .

The degree of a map is invariant by homotopy, and φ is homotopic to the identity of S , so $n^+ - n^- = 1$. We assumed $n^- = 0$ or $n^+ = 0$. Thus $n^- = 0$ and $n^+ = 1$.

Using our claim, for every $y \in Y$, there is a closed disk $B_y \subset S$ containing y in its interior, such that $\varphi^{-1}(B_y)$ is a closed disk A_y , and such that $\varphi|_{A_y}^{B_y}$ is an orientation-preserving homeomorphism. Without loss of generality, the disks $\{B_y\}_{y \in Y}$ are pairwise disjoint. Let N be the surface obtained from S by removing the interiors of the disks $\{B_y\}_{y \in Y}$, and M be obtained by removing the interiors of the disks $\{A_y\}_{y \in Y}$. The map $\varphi' := \varphi|_M^N$ is defined. Since φ is a degree one map, $\varphi' : M \rightarrow N$ is a degree one map. By construction, φ' maps the boundary of M to the boundary of N and the interior of M to the interior of N , and the restriction and corestriction of φ' to the boundaries of M and N is a homeomorphism. It follows from a result by Edmonds [12, Theorem 4.1] that φ' is homotopic to a homeomorphism $M \rightarrow N$, where the homotopy is relative to ∂M . (More precisely, this result follows from [12, Theorem 3.1] by noting that (1) each branch covering of degree ± 1 is a homeomorphism, and that (2) each pinch map from M to M is homotopic to the identity, because the simple closed curve defining the pinch must bound a disk.) $\square$

4.3 From graphs to triangulations. Let T be a reducing triangulation, and let Z be a triangulation. A map $\varphi : Z \rightarrow T$ is **simplicial** if φ maps Z^1 to T^1 simplicially, and if φ sends every face of Z to a vertex, an edge, or a face of T . It is **strongly harmonious** if $\varphi|_{Z^1}^{T^1}$ is strongly harmonious.

In this section, we prove the following proposition, which (essentially) allows us to consider simplicial drawings of entire triangulations, instead of graphs.

PROPOSITION 4.3. *Let S be a closed surface distinct from the sphere. Let T be a reducing triangulation of S . Let G be a finite graph embedded in S , and let $f : G \rightarrow T^1$ be simplicial. Assume that f is strongly harmonious, and that f is homotopic to the inclusion map $G \hookrightarrow S$. There are a triangulation Z of S whose 1-skeleton contains a subdivision of G as a subgraph, and a simplicial map $\varphi : Z \rightarrow T$ with $\varphi|_G^{T^1} = f$, such that φ is strongly harmonious and homotopic to the identity map of S .*

The proof of Proposition 4.3 relies on two lemmas:

LEMMA 4.6. *Let S be a closed surface distinct from the sphere. Let T be a reducing triangulation of S . Let G be a graph obtained from another graph G' by inserting a path graph P between (possibly equal) vertices of G' . Let $f : G \rightarrow T^1$ be simplicial. If $f|_{G'}$ is strongly harmonious, and if f maps P to a reduced walk in T , then f is strongly harmonious.*

Proof. If P is a single edge, then it does not affect strong harmony. So assume that P has an interior vertex v , and let L be a left line (the right line case being similar) in T whose central vertex is $f(v)$. We will exhibit a walk based at v in G whose image path escapes L . This will prove the lemma. Consider the universal covering triangulation $\tilde{T}$ of T , and a lift $\tilde{L}$ of L in $\tilde{T}^1$. Consider also the reduced walk $X := f|_P$. There are a lift $\tilde{P}$ of P , a lift $\tilde{v}$ of v in $\tilde{P}$, and a lift $\tilde{X} : \tilde{P} \rightarrow \tilde{T}^1$ of X , such that $\tilde{X}(\tilde{v})$ is the central vertex of $\tilde{L}$. Let $\tilde{w}_0$ and $\tilde{w}_1$ be the two end-vertices of $\tilde{P}$, and let $\tilde{P}_0$ and $\tilde{P}_1$ be the two sub-walks of $\tilde{P}$ that go from $\tilde{v}$ to respectively $\tilde{w}_0$ and $\tilde{w}_1$. Since $\tilde{X}$ is a reduced walk, and since $\tilde{L}$ makes only 3_r -turns, one of $\tilde{P}_0$ and $\tilde{P}_1$, say $\tilde{P}_0$ without loss of generality, is such that $\tilde{X} \circ \tilde{P}_0$ either escapes $\tilde{L}$, or stays in the non-negative part of $\tilde{L}$. In the first case, considering the vertex w_0 of G lifted by $\tilde{w}_0$, the portion P_0 of P from v to w_0 is such that $f \circ P_0$ escapes L . In the latter case, replace the central vertex of $\tilde{L}$ by $\tilde{X}(\tilde{w}_0)$, and project the resulting line onto the surface, obtaining a line L' which is a shift of L . Since w_0 belongs to G' , and since $f|_{G'}$ is strongly harmonious, there is a walk Q based at w_0 in G such that $f \circ Q$ escapes L' . Then the concatenation P'_0 of P_0 and Q is such that $f \circ P'_0$ escapes L . $\square$

LEMMA 4.7. *Let w, v_0, v_1 be pairwise distinct vertices of a plane reducing triangulation T . If v_0 and v_1 are adjacent in T , then there is $i \in \{0, 1\}$ such that, along the reduced walk from v_i to w , the vertex consecutive to v_i is adjacent to or equal to v_{1-i} .*

Proof. Let e be the edge of T between v_0 and v_1 ; direct e so that it sees blue at its left, and assume without loss of generality that it is directed from v_0 to v_1 . Let W_1 be the reduced walk from v_1 to w . Assume that the vertex consecutive to v_1 in W_1 is neither equal nor adjacent to v_0 , for otherwise there is nothing to do. Let W_0 be the concatenation of e with W_1 . By assumption, W_0 does not make a 0-turn, 1-turn, or -1 -turn at v_1 , and by our choice of direction of e , it also does not make a 2_r -turn, so it is reduced. And W_0 starts with edge e , as desired. $\square$

Proof of Proposition 4.3. In this proof, we will use the following standard topological fact: Every graph embedded on a surface can be extended to a triangulation by adding edges (not vertices). This can be done by repeatedly inserting an edge e inside a face F , where e “differs” from every edge e' on the boundary of F in the sense that the concatenation of e and e' does not bound a disk in F .

At any time we denote by $G^\diamond$ the graph derived from G by removing its degree two vertices. The proof is in three steps; see Figure 8. In the first step, as long as some face of $G^\diamond$ is not a triangle, we insert an edge e in this face, as described in the preceding paragraph. Doing so, we consider the homotopy from the inclusion map $G \hookrightarrow S$ to the map f , and apply this homotopy to the two end-vertices of e , thus extending e to a path e' between vertices of T . Let p be the unique reduced path homotopic to e' (Lemma 2.1). If the length n of p is greater than one, then we insert $n - 1$ vertices in G along e . We also extend f to e by mapping e to p . In this way, f remains




Figure 8: In the proof of Proposition 4.3, the embedded graph G is extended to a triangulation.

homotopic to the inclusion map $G \hookrightarrow S$, and f remains strongly harmonious by Lemma 4.6. Now, every face of $G^\diamond$ is a triangle.

In the second step (see Figure 8, bottom center), for every (triangular) face m of $G^\diamond$, we consider an edge e incident to m in $G^\diamond$. If e is subdivided in G , then we insert a path in m from each interior vertex of e to the vertex of m opposite to e . As in the previous paragraph, we map each of the new paths to a reduced path in T , so that f remains homotopic to the inclusion map $G \hookrightarrow S$, and so that f remains strongly harmonious. Now, every (triangular) face m of $G^\diamond$ is incident to an edge that is not subdivided in G .

In the last step, Lemma 4.7 ensures that we can triangulate m by inserting new edges in m (see Figure 8, bottom right), and by mapping in f each such new edge to a *vertex or an edge* of T , keeping f strongly harmonious and homotopic to the inclusion map $G \hookrightarrow S$. Now G is the 1-skeleton of a triangulation Z , and f trivially extends to a simplicial map $\varphi : Z \rightarrow T$. $\square$

4.4 A property of the maps homotopic to the identity and strongly harmonious. Here is another key step towards the proof of Theorem 1.1: we prove that our simplicial drawings of triangulations orient the (non-degenerate) triangles coherently.

PROPOSITION 4.4. *Let S be a closed surface distinct from the sphere. Let T be a reducing triangulation of S . Let Z be a triangulation of S , and let $\varphi : Z \rightarrow T$ be simplicial. If φ is strongly harmonious, and if φ is homotopic to the identity map of S , then there cannot be two faces z_+ and z_- of Z for which $\varphi|_{z_+}$ is positive and $\varphi|_{z_-}$ is negative.*

In this section only, it is convenient to consider the plane P , and an (infinite) reducing triangulation T of P ; the notion of strong harmony immediately extends to that setting. And again, if Z is a triangulation of P , a simplicial map $\varphi : Z \rightarrow T$ is strongly harmonious if $\varphi|_{Z^1}$ is strongly harmonious.

In the 1-skeleton T^1 of T , the minimum number of edges of a path between two given vertices defines a distance on the vertex set of T . A map $\varphi : P \rightarrow P$ is **uniformly homotopic** to the identity map $1_{P \rightarrow P}$ if it is homotopic to the identity and there is $\kappa > 0$ such that every point $x \in P$ mapped to a vertex v of T by φ lies in a vertex, edge, or face of T whose incident vertices are at distance less than κ from v . We shall prove the following:

PROPOSITION 4.5. *Let P be the plane. Let T be a reducing triangulation of P . Let Z be a triangulation of P , and let $\varphi : Z \rightarrow T$ be simplicial. If φ is strongly harmonious, and if φ is uniformly homotopic to the identity map of P , then there cannot be two faces z_+ and z_- of Z for which $\varphi|_{z_+}$ is positive and $\varphi|_{z_-}$ is negative.*

Proposition 4.4 (concerning surfaces) easily follows from Proposition 4.5 (concerning the plane) by lifting:

Proof of Proposition 4.4, assuming Proposition 4.5. The universal covering space $\tilde{S}$ of S is the plane. Also T lifts to a reducing triangulation $\tilde{T}$ of $\tilde{S}$, Z lifts to a triangulation $\tilde{Z}$ of $\tilde{S}$, and φ lifts to a simplicial map $\tilde{\varphi} : \tilde{Z} \rightarrow \tilde{T}$. Since φ is strongly harmonious, $\tilde{\varphi}$ is strongly harmonious. We claim that $\tilde{\varphi}$ is uniformly homotopic to the identity map of $\tilde{S}$. This claim implies by Proposition 4.5 that there cannot be two faces z_+ and z_- of $\tilde{Z}$ for which $\tilde{\varphi}|_{z_+}$ is positive and $\tilde{\varphi}|_{z_-}$ is negative. Then the same holds for Z and φ , proving the proposition.




Figure 9: A cone.




Figure 10: The impossible cases in the proof of Lemma 4.8.

To prove the claim, we use the assumption that there is a homotopy H between φ and the identity map of S , and we lift H to a homotopy $\tilde{H}$ between $\tilde{\varphi}$ and the identity map of $\tilde{S}$. For every lift $\tilde{x} \in \tilde{S}$ of a point $x \in S$, the path $\tilde{p}$ followed by the image of $\tilde{x}$ in $\tilde{H}$, and the path p followed by the image of x in H are such that the number of vertices, edges, and faces of $\tilde{T}$ crossed by $\tilde{p}$ is equal to the number of vertices, edges, and faces of T crossed by p . The latter is uniformly bounded since S is compact. $\square$

The rest of this section is devoted to the proof of Proposition 4.5. We need some definitions and lemmas. See Figure 9. Let T be a reducing triangulation of the plane. The part of T on a given side of a bi-infinite reduced path ∂H , and not on ∂H , is a **half-plane** H of T , and ∂H is the boundary of H . We emphasize that half-planes are open. A half-plane H is **nested** in another half-plane H' if $H \subset H'$ and $\partial H \cap \partial H' = \emptyset$. Let e be a directed edge e of T that sees blue on its left, and let L be the bi-infinite reduced walk in T that contains the head vertex v of e , and makes only 3_b - and 3_r -turns, except at v where it makes a 4_b -turn whose middle edge is e . The **cone** of e is the half-plane H on the right of L , and v is the **tip** of H .

LEMMA 4.8. *Let T be a reducing triangulation of the plane. If H is a half-plane of T , and if $v \in H$ is a vertex of T , then v is the tip of a cone H' nested in H .*

Proof. There is a reduced path P between v and a vertex u of ∂H , such that P is internally included in H . Let g be the edge of P incident to v , directed toward v . First assume that g sees blue on its left. We claim that the cone H' of g is nested in H . First we prove $\partial H' \cap \partial H = \emptyset$ by contradiction. See Figure 10 Assuming that ∂H and $\partial H'$ share a vertex w , let Q' be the concatenation of P and of the subpath of $\partial H'$ between v and w , and let Q be the subpath of ∂H between u and w . Then Q and Q' are distinct reduced paths with the same end-vertices in T , contradicting Lemma 2.1. It then easily follows that $H' \subset H$, and we now provide the details. We have $\partial H' \subset H$ since $\partial H' \cap H \neq \emptyset$ and $\partial H' \cap \partial H = \emptyset$. The reduced path P is internally disjoint from ∂H and $\partial H'$ by Lemma 2.1, and since ∂H and $\partial H'$ are reduced. So the relative interior of P lies in the open region R between ∂H and $\partial H'$. Also, the relative interior of g is disjoint from H' , and so is R .

Now assume that g sees red on its left. We claim that there is a rotation of one turn of g around its head vertex v (either clockwise or counter-clockwise) after which the cone H' of g satisfies $\partial H' \cap \partial H = \emptyset$. As above, this claim implies that H' is nested in H . Let I be the boundary of H' after a clockwise-rotation around v , and let J be the boundary after a counter-clockwise rotation around v . Assume by contradiction that I and J both intersect ∂H . See Figure 10. Cut I into two semi-infinite walks at v , the right part (with respect to g) denoted by I_0 , and the left part denoted by I_1 . Cut J into two parts at v , the right part J_0 , the left part J_1 . The




Figure 11: Combs.

concatenation of P and I_1 is a reduced walk, so I_1 is disjoint from ∂H by Lemma 2.1, and so I_0 intersects ∂H in a vertex a . The concatenation of P and J_0 is a reduced walk, so J_0 is disjoint from ∂H by Proposition 2.1, and so J_1 intersects ∂H in a vertex b . Let Q' be the concatenation of the subpath of I_0 between a and v , and of the subpath of J_1 between v and b . Let Q be the subpath of ∂H between a and b . Then Q and Q' are distinct reduced paths with the same end-vertices in T , contradicting Lemma 2.1. $\square$

LEMMA 4.9. *Let T be a reducing triangulation of the plane. Let G be a graph, and let $f : G \rightarrow T^1$ be simplicial. Let H be a cone of T , and let v be a vertex of G . If f is strongly harmonious, and if $f(v)$ is the tip of H , then there is a walk W based at v in G that satisfies $f(W) \subset H \cup \partial H$ and $f(W) \not\subset \partial H$.*

Proof. Let $G_0 \subset G$ be the cluster of f containing v . Since f is strongly harmonious, there is a directed edge e from G_0 to $G \setminus G_0$ such that $f(e) \in H \cup \partial H$. If $f(e) \in H$, then we are done. So we assume $f(e) \in \partial H$. The boundary ∂H splits into two semi-infinite walks I_0 and I_1 based at $f(v)$, where H lies on the left of I_0 and on the right of I_1 . Assume $f(e) \in I_0$, the other case being similar. Let w be the head vertex of e . Let J be the suffix of I_0 obtained by removing its first edge ($f(w)$ is the first vertex of J). Let L be the bi-infinite walk that makes only 3_r -turns and contains J as a subwalk. Since f is strongly harmonious, there is a walk W based at w in G , such that $f \circ W$ may stay in J for a while, and then leaves J on its left. $\square$

In order to ease the reading of the proof of Proposition 4.5, we shall not use Lemma 4.8 and Lemma 4.9 as they are. Instead, we shall use two corollaries:

COROLLARY 4.1. *Let T be a reducing triangulation of the plane. Let G be a graph, and let $f : G \rightarrow T^1$ be simplicial. Let H be a half-plane of T , and let v be a vertex of G . Let $\kappa > 0$. If f is strongly harmonious, and if $f(v) \in H$, then there is a walk W from v to a vertex w in G , such that $f(W) \subset H$, and such that $f(w)$ is at distance greater than κ from ∂H .*

Proof. This follows immediately from repeated iterations of Lemmas 4.8 and 4.9. $\square$

The next corollary is easier to state with the following definition. See Figure 11. In a reducing triangulation T of the plane, a right (resp. left) **comb** is a pair (H, x) where H is a half-plane of T , and x is a vertex of ∂H , that satisfy the following: directing ∂H so that H lies on its right (resp. left), every turn of ∂H at x and after x is 3_r or 2_b (resp. -3_r or -2_b). The **glen** of (H, x) is the union of H and of the part of ∂H consecutive to x , including x .

COROLLARY 4.2. *Let T be a reducing triangulation of the plane. Let G be a graph, and let $f : G \rightarrow T^1$ be simplicial. Let (H, x) be a (left or right) comb of T , and let v be a vertex of G . Let $\kappa > 0$. If f is strongly harmonious, and if $f(v) = x$, then there is a walk W from v to a vertex w in G , such that $f(W)$ is included in the glen of (H, x) , and such that $f(w)$ is at distance greater than κ from ∂H .*

Proof. The vertex x is the tip of a cone H' such that $H' \cup \partial H'$ is included in the glen of (H, x) . Lemmas 4.8 and 4.9 conclude. $\square$

We need a few more technical lemmas.

LEMMA 4.10. *Let T be a plane reducing triangulation, and let T' be subgraph of T^1 . If T' is connected, finite, and not a single vertex, then T' admits at least two vertices x such that the edges of T' incident to x are all included in some bad turn of T .*

Proof. If T' has more than one vertex of degree one, then we are done. So assume that this is not the case. Then T' has at least three vertices since it has neither loops nor multiple edges (Section 2.4). First assume that T' has no degree one vertex. Then the outer closed walk W of T' has length greater than or equal to three, and never uses an edge of T^1 and its reversal consecutively. Orient W so that the outer-face of T' lies on the right of W . Walk along W until some vertex x_0 is visited for the second time, then cut the portion of W from x_0 to itself, thus obtaining a closed walk C based at x_0 . Then C is a simple closed walk, not a single vertex. Thus, since T^1 has no loop nor multiple edges, C has length greater than or equal to three. Also, every vertex $x \neq x_0$ of C is such that all the edges of T' incident to x either lie on C , or on the left of C . Lemma 2.2 ensures that at least three of the left turns of C are bad, and at least two of them do not occur at x_0 , proving the lemma in this case.

Now assume that T' has exactly one vertex of degree one. Removing degree one vertices from T' as long as there is one immediately gives the following: T' is the union of a path P , not single vertex, and of a connected graph Q , not a single vertex, such that the intersection of P and Q is one of the two end-vertices x of P , and such that Q has no degree one vertex. By the preceding, Q has a vertex distinct from x that suits our need. And the end-vertex of P distinct from x also suits our needs. $\square$

In the following, we denote by T^0 the set of vertices of a triangulation T .

LEMMA 4.11. *Let P be the plane. Let T be a reducing triangulation of P . Let Z be a triangulation of P , and let $\varphi : Z \rightarrow T$ be simplicial. If φ is strongly harmonious, and if φ is uniformly homotopic to the identity map of P , then $\varphi^{-1}(T^0)$ does not disconnect P .*

Proof. By contradiction, assume that it does. There is no simple bi-infinite walk W in Z^1 that belongs entirely to $\varphi^{-1}(T^0)$, for otherwise $\varphi(W)$ would be a single vertex of T (being a connected subset of T^0), contradicting the assumption that φ is uniformly homotopic to the identity map. Therefore, and since φ is simplicial, there are a simple cycle C in Z^1 , a vertex $v \in Z^0$ in the interior of the bounded side of C , and a vertex x of T , such that $\varphi(C) = x$ and $\varphi(v) \neq x$. Let G be the subgraph of Z^1 that contains C and the part of Z^1 lying in the bounded side of C . Then $\varphi(G)$ is a finite connected subgraph of T , not a single vertex. So by Lemma 4.10 there is a vertex $y \neq x$ in $\varphi(G)$ such that the set B of edges of $\varphi(G)$ incident to y is included in a bad turn of T . There is a cluster $G_0 \subset Z^1$ such that $\varphi(G_0) = y$, and such that the edges between G_0 and $Z^1 \setminus G_0$ all belong to G . Those edges are mapped to B by φ , contradicting the assumption that φ is strongly harmonious. $\square$

In the following, we say that 1-turns and 2_r -turns are **bad left turns**.

LEMMA 4.12. *In a plane reducing triangulation T , let I be a simple bi-infinite walk that does not make any bad left turn. The vertices of T on the left of I at distance one from I are the vertices of a simple bi-infinite walk I' that does not make any bad left turn.*

Proof. Consider the sequence E of directed edges of T emanating from vertices of I to the left side of I . If e_1 and e_2 are consecutive in E , then either e_1 and e_2 have the same target vertex, or they have the same source vertex. In the first case, e_2 is the counter-clockwise rotation of e_1 around their target vertex, in the second case e_2 is the clockwise-rotation of e_1 around their source vertex. Let I' be the bi-infinite walk corresponding to the target vertices of the directed edges in E . The vertices of I' are those on the left of I at distance one from I .

Then I' does not make any 0-turn, nor any bad left turn. For otherwise there are three consecutive $e_1, e_2, e_3 \in E$ that all have the same target vertex (at which I' makes a bad left turn), such that e_2 sees red on its left. But then I makes a 2_r -turn at the source vertex of e_2 , a contradiction.

Moreover I' is simple. For otherwise some portion W of I' is a non-trivial simple closed walk, and the bounded side of W lies on its right by Lemma 2.2, and since W makes at most one bad left turn. But then I is contained in the bounded side of W , a contradiction. $\square$




Figure 12: The construction in the proof of Proposition 4.5.

Proof of Proposition 4.5. The overall proof is illustrated in Figure 12. By contradiction, assume that there are two faces z_0 and z_- of Z such that $\varphi|_{z_+}$ is positive and $\varphi|_{z_-}$ is negative. By Lemma 4.11, there is a simple path p from z_+ to z_- in the dual of Z , that is disjoint from $\varphi^{-1}(T^0)$. Up to replacing z_+ and z_- by other faces of Z , we may assume that φ is null on every face of Z intermediately visited by p . Extend p to the respective vertices of z_+ and z_- not incident to the first and last edges crossed by p . Because φ is null on every face visited by p except the first and last ones, and because p is disjoint from $\varphi^{-1}(T^0)$, there is an edge $s \in T^1$ such that the image, by φ , of all the edges of Z crossed by p is exactly s . Moreover, one can direct s in such a way that all the edges of Z crossed by p from left to right are mapped to s . Consider the left-most bi-infinite reduced walk of T that contains s , and let H be the half-plane on its left. Because $\varphi|_{z_+}$ is positive and $\varphi|_{z_-}$ is negative, the interiors $\dot{z}_+$ and $\dot{z}_-$ of triangles z_+ and z_- are mapped inside H .

Using the assumption that φ is uniformly homotopic to the identity, there is $\kappa > 0$ such that each vertex of T is at distance smaller than κ from its pre-images under φ . Let p' be any of the directed edges crossed by p from left to right. By Corollary 4.1 there is, in Z^1 , a walk U_0 (resp. U_1) based at the source (resp. target) end-vertex of p such that, for $i = 1, 2$, $\varphi(U_i) \subset H$, and such that $\varphi(U_i)$ contains a point at distance greater than 3κ from ∂H . Cut p into two parts p_0 and p_1 at its intersection with p' , and reverse p_0 . Let q_i be the concatenation of p_i and U_i . Make q_i a simple path by shortening it if necessary. Then q_0 and q_1 are disjoint from p' except for their basepoint, since $\varphi(U_0), \varphi(U_1) \subset H$, $\varphi(p') = s$, and $\varphi(\dot{z}_+), \varphi(\dot{z}_-) \subset H$. We prove by contradiction that q_0 and q_1 are disjoint except for their basepoint. If not, then $q_0 \cup q_1$ contains a simple closed curve whose bounded side contains an endpoint y of p' . By Corollary 4.2 there is a walk X based at y in Z^1 such that $\varphi(X)$ is disjoint from $\dot{s} \cup H$, and such that X intersects $q_0 \cup q_1$. That contradicts the fact that $\varphi(q_0 \cup q_1) \subset \dot{s} \cup H$.

Now cut p' into two parts p'_0 and p'_1 at its intersection with p . Reverse p'_0 . By Corollary 4.2 there is a simple walk U'_0 (resp. U'_1) based at the source (resp. target) end-vertex of p' in Z^1 , such that $\varphi(U'_i)$ is disjoint from $\dot{s} \cup H$, and such that $\varphi(U'_i)$ contains a point at distance greater than 3κ from ∂H . Let q'_i be the concatenation of p'_i and U'_i . Since $\varphi(U'_0)$ and $\varphi(U'_1)$ are disjoint from $\dot{s} \cup H$, the paths q'_0 and q'_1 are simple, and each of them is disjoint from q_0 and q_1 except for their basepoint. An argument by contradiction similar to the one of the previous paragraph shows that q'_0 and q'_1 are disjoint except for their basepoint.

Let Q be the concatenation of q_0 and q_1 , and Q' be the concatenation of q'_0 and q'_1 . By construction, Q and Q' are simple, and cross exactly once. Until now, we have considered separately the situation in Z (before applying φ) and in T (after applying φ). But recall that Z and T are both triangulations of the plane, and that each vertex of T is at distance smaller than κ from its pre-images under φ . Because of this, and by the properties of $\varphi(Q)$ and $\varphi(Q')$, we have that Q has its endpoints inside H , at distance at least 2κ from ∂H , and may enter the complement of H , but only by a distance of at most κ . Similarly, Q' has its endpoints outside H , at distance at least 2κ from ∂H , and may enter H but only by a distance of at most κ . By (repeated applications of) Lemma 4.12, the vertices of T inside H at distance 2κ from ∂H are the vertices of a simple bi-infinite walk I . Then I separates the end-points of Q from every point of Q' . Join the end-points of Q by a path P , where P is separated from Q' by I , thus turning Q into a (not necessarily simple) closed curve C . Because Q and Q' cross exactly once, and because P and Q' are disjoint, then C and Q' cross exactly once. Similarly, there is a simple bi-infinite walk that separates the end-points of Q' from C , so Q' can be extended to a closed curve C' such that C and C' cross exactly once. But this is impossible, since any two closed curves in the plane cross an even number of times. $\square$

4.5 Proof of Theorem 1.1. We now have almost all the material to prove Theorem 1.1. We first prove the theorem for strongly harmonious drawings, in the following proposition.

PROPOSITION 4.6. *Let S be a closed surface distinct from the sphere. Let T be a reducing triangulation of S . Let G be a finite graph embedded in S , and let $f : G \rightarrow T^1$ be simplicial. If f is strongly harmonious, and if f is homotopic to the inclusion map $G \hookrightarrow S$, then f is a weak embedding.*

(As a side note, observe that Proposition 4.6 considers also the torus.)

Proof. By Proposition 4.3 there are a triangulation Z of S whose 1-skeleton contains a subdivision of G as a subgraph, and a simplicial map $\varphi : Z \rightarrow T$ with $\varphi|_G^{T^1} = f$, such that φ is strongly harmonious and homotopic to the identity map of S . By Proposition 4.4 there cannot be two faces z_+ and z_- of Z for which $\varphi|_{z_+}$ is positive and $\varphi|_{z_-}$ is negative. Thus, letting $Y \subset S$ contain one point in the interior of each face of T , φ is coherently oriented at

Y . By Proposition 4.2 the map $\varphi|_{S \setminus \varphi^{-1}(Y)}^{S \setminus Y}$ is homotopic to a homeomorphism $S \setminus \varphi^{-1}(Y) \rightarrow S \setminus Y$. In particular, and since $G \cap \varphi^{-1}(Y) = \emptyset$, the map $f|^{S \setminus Y}$ is homotopic to an embedding $G \rightarrow S \setminus Y$. Also $S \setminus Y$ is the patch system of T , and f has no spur since f is strongly harmonious. So f is a weak embedding by Proposition 4.1. $\square$

The proof of Theorem 1.1 relies on a few additional lemmas.

LEMMA 4.13. *Let S be a closed surface distinct from the sphere and the torus. Let T be a reducing triangulation of S . Let C be a collection of closed walks in T . If the walks in C are reduced, and if C is homotopic to an embedding, then C is a weak embedding.*

The proof of Lemma 4.13 adapts arguments from the proofs of [3, Proposition 4.2] and [9, Proposition 1.5], though it is considerably simpler due to the fact we consider a collection of closed curves, homotopic to an embedding (instead of general graph drawings, and instead of curves with arbitrary geometric intersection number).

Proof. Without loss of generality the walks in C are not single vertices, and are thus not contractible by Lemma 2.1. Let Σ be the patch system of T . We claim that there is a simple collection of closed curves homotopic to C in Σ . This claim implies the lemma by Proposition 4.1 since C has no spur. To prove the claim consider a collection of closed curves Γ homotopic to C in Σ , and self-crossing as few times as possible in Σ . We prove the claim by contradiction so assume that Γ self-intersects. Let $\tilde{S}$ be the universal covering space of S . Let $\tilde{\Sigma}$ and $\tilde{T}$ be the respective lifts of Σ and T in $\tilde{S}$. Since Γ self-intersects while being homotopic to a simple collection of closed curves, there are either a non-simple closed curve in Γ that could be made simple by homotopy, or there are two simple closed curves in Γ that cross while they could be made simple altogether by homotopy. There are two cases.

First assume that in $\tilde{S}$ some lift of a curve from Γ self-intersects. Then some portion of this lift is a simple loop ℓ in $\tilde{S}$, based at the intersection point. If the bounded side D of ℓ does not contain any face of $\tilde{T}$, equivalently if $D \subset \tilde{\Sigma}$, then the self-intersection of γ can be removed by homotopy in Σ , contradicting the assumption that Γ self-crosses as few times as possible. Otherwise consider the walk W in $\tilde{T}$ that encodes the sequence of crossings of ℓ with the arcs of $\tilde{\Sigma}$. Then W is not a single vertex, the two end-vertices of W are the same vertex, and W is reduced, contradicting Lemma 2.1.

Now assume that in $\tilde{S}$ every lift of every curve in Γ is simple. There are two simple lifts $\tilde{\gamma}_0$ and $\tilde{\gamma}_1$ of curves from Γ that intersect at least twice (see e.g. [24, Theorem 3.5]). Then some portion ℓ_0 of $\tilde{\gamma}_0$ and some portion ℓ_1 of $\tilde{\gamma}_1$ are such that ℓ_0 and ℓ_1 are the two sides of a bigon D embedded in $\tilde{S}$, between two lifted self-intersections of γ . If $D \subset \tilde{\Sigma}$, then the two self-intersections of γ can be removed by homotopy in Σ , contradicting the assumption that γ self-crosses as few times as possible. Otherwise consider the walks W_0 and W_1 in $\tilde{T}$ that encode the sequence of crossings of ℓ_0 and ℓ_1 with the arcs of $\tilde{\Sigma}$. Then W_0 and W_1 are distinct, have the same end-vertices, and are reduced, contradicting Lemma 2.1. $\square$

Preparing for the next lemma, observe that if M is a finite graph embedded on a surface, then the universal cover of M is a tree $\tilde{M}$ naturally equipped with a rotation system. Moreover, if C_0 is a closed walk without spur in M , then every lift $\tilde{C}_0$ of C_0 partitions $\tilde{M}$ into three parts: the left of $\tilde{C}_0$, the right of $\tilde{C}_0$, and the image of $\tilde{C}_0$.

LEMMA 4.14. *Let M be a finite graph embedded on a surface. Let G be a finite graph simplicially mapped to M , and let C be a collection of closed walks in M . Assume that G and C are weak embeddings without spur, and that their union is not a weak embedding. Then in the universal cover of M there is a lift of G that has vertices on both sides of a lift of a walk C_0 from C .*

Proof. Without loss of generality every edge of G is mapped to an edge of M (contract the clusters of G otherwise). Let Σ be the patch system of M , and let G' and C' be embeddings approximating G and C in Σ . Without loss of generality G' , C' , and the arcs of Σ are in general position. Also, G' and C' cross minimally. By assumption they cross at a point x in the interior of some face of Σ . Let $\tilde{\Sigma}$ be the universal covering space of Σ , and let $\tilde{x}$ be a lift of x in $\tilde{\Sigma}$. Let $\tilde{C}'_0$ be the lift of a walk C'_0 from C' that contains $\tilde{x}$. Cut the lifts of G' at every intersection with $\tilde{C}'_0$, and let $\tilde{G}'$ be one of the two cuts that meet $\tilde{x}$. Since the number of crossings between G' and C' is minimal, $\tilde{G}'$ is not contained in the union of the faces and arcs of $\tilde{\Sigma}$ used by $\tilde{C}'_0$. $\square$

The following is analogous to Lemma 4.5.

LEMMA 4.15. *Let S be a surface. Let G be a graph, and let $f : G \rightarrow S$ be a map. If $f(G)$ is a simple circle $S_0 \subset S$, if S_0 does not bound a disk in S , and if f is homotopic to an embedding in S , then f is homotopic to an embedding in a tubular neighborhood of S_0 .*

Proof. Without loss of generality G is connected. Fix a vertex r of G and a spanning tree T of G rooted at r . Let ℓ be the simple loop on S based at $x := f(r)$ whose image is S_0 . We first claim that we can assume without loss of generality that each edge of T is mapped to x by f , and that each edge not in T is mapped, under f , to a power of ℓ . To see this, contract, in S_0 , the edges in T ; each remaining edge is homotopic to some power of ℓ in S_0 . We second claim that each edge of T is actually mapped to either ℓ , to ℓ^{-1} , or to the constant loop in S_0 . This is due to a result of Epstein [15, Theorem 4.2], and since the edges not in T are mapped to loops that can be made simple by homotopy in S .

On the other hand, there is an embedding $f' : G \rightarrow S$ homotopic, in S , to f . Without loss of generality, we can assume that this homotopy between f and f' holds the image of r fixed. Indeed consider a homotopy H from f to f' , and the path $p : [0, 1] \rightarrow S$ followed by the image of r under H . There is an ambient isotopy H' of S that “counteracts” H in the sense that it maps $p(t)$ to x for all $t \in [0, 1]$. Composing the maps in H by the maps in H' gives a homotopy from f to an embedding (not f') that holds the image of r fixed.

We contract the edges of T in f' to a small neighborhood of x , this time preserving the fact that we have an embedding. The remaining edges are loops that, under f' , are homotopic to their counterparts under f (if T would be really contracted to x). The contractible ones can be pushed by isotopy into a small neighborhood of x , as each of them bounds a disk with only (possibly) contractible loops inside it [15, Theorem 1.7]. The other loops can be bundled together parallel, since any pair of them bounds a disk with only (possibly) contractible or homotopic loops inside it. Then they can be pushed altogether in a neighborhood of ℓ . $\square$

Proof of Theorem 1.1. Clearly if f is a weak embedding, then there is an embedding homotopic to f in S . For the other direction assume that there is an embedding homotopic to f in S . We shall prove that f is a weak embedding.

Partition G into two subgraphs A and B such that $f|_B$ is strongly harmonious, and f is not strongly harmonious on any of the connected components of A . Then $f|_A = C \circ f'$ for some disjoint union O of cycle graphs, mapped to reduced closed walks by $C : O \rightarrow T^1$, and some simplicial map $f' : A \rightarrow O$, without spur.

Our first claim is that the collection of closed walks C is homotopic to an embedding in S . Indeed C can be realized as the restriction of f to a collection of disjoint cycles in A , as follows. For every cycle O_0 in O , since f' has no spur, there is a simple closed walk W in A mapped by f' to a non-trivial power of O_0 . Then W is mapped by f to a non-trivial power of the closed walk $C_0 := C|_{O_0}$, homotopic to an embedding (since f is), and so it is actually mapped to $C|_0$ or its reversal [15, Theorem 4.2].

Thus C is a weak embedding by Lemma 4.13, and since the walks in C are reduced. Now f' is a weak embedding by Lemma 4.15 and Proposition 4.1. Our second claim is that $f|_B \cup C$ is a weak embedding. This claim proves the theorem as $f|_{A \cup B}$ is then a weak embedding. We prove the claim by contradiction so assume that $f|_B \cup C$ is not a weak embedding. Let U_T be the universal cover of the 1-skeleton of T . There is by Lemma 4.14 some lift of $f|_B$ in U_T that contains vertices on both sides of some lift of a walk C_0 from C . Now let $\tilde{S}$ be the universal cover of S . In $\tilde{S}$, there are lifts $\tilde{B}$ and $\tilde{C}_0$ of $f|_B$ and C_0 such that $\tilde{B}$ contains vertices on both sides of $\tilde{C}_0$. Then $\tilde{C}_0$ is a semi-infinite reduced walk in the lift of T , and $\tilde{B} \cup \tilde{C}_0$ is uniformly homotopic to an embedding. That contradicts Corollary 4.1. $\square$

5 Harmonizing a drawing monotonically: Proof of Theorem 1.2

In this section we prove Theorem 1.2, which we restate for convenience:

THEOREM 1.2. *Let S be an orientable surface without boundary homeomorphic to neither the sphere nor the torus. Let T be a reducing triangulation of S , with m edges. Let G be a graph, and let $f : G \rightarrow T^1$ be a drawing of size n . We can compute in $O((m+n)^2 n^2)$ time a drawing $f' : G \rightarrow T^1$, harmonious, homotopic to f in S , such that for every edge e of G , the image of e under f' is not longer than under f .*

The strategy is to transform a drawing f iteratively by some “local” moves satisfying the property that, if no move can be applied, then the current drawing is harmonious. Two of these moves, the shortening and balancing




Figure 13: The moves. (Left) Each disk represents a single cluster. (Right) Each disk may represent several clusters. The balancing is slightly counter-clockwise (clockwise would not decrease any edge length here).

moves, decrease strictly the length of the drawing. On the other hand, the flip move does not affect the length of the drawing. Roughly but not exactly, the algorithm performs moves iteratively, giving priority to the shortening and balancing moves over the flip moves. To bound the complexity of the algorithm, it then suffices to bound the length of a flip sequence, assuming that at each step of this sequence, no shortening or balancing move is possible. The details are actually more complicated; in particular, the flip sequence is chosen carefully.

5.1 Reduction to simplicial maps. We have the following preliminary observation:

LEMMA 5.1. *To prove Theorem 1.2, we can assume that f is simplicial.*

Proof. Let $f : G \rightarrow T^1$ be as in the statement of Theorem 1.2. Let $\bar{f} : \bar{G} \rightarrow T^1$ be the associated simplicial map that factorizes f . If Theorem 1.2 holds for simplicial maps, then we obtain a map $\bar{f}' : \bar{G} \rightarrow T^1$ that is harmonious, is homotopic to $\bar{f}$, and does not increase the length of the edges (compared to $\bar{f}'$). In particular, $\bar{f}'$ is simplicial. It naturally corresponds to a drawing f' of G on T^1 that satisfies the desired properties. $\square$

5.2 Flips, shortenings, and balancings. Throughout this section, T is a reducing triangulation of a surface S distinct from the sphere and the torus, G is a graph, and $f : G \rightarrow T^1$ is a *simplicial* map. Recall that f factors uniquely into a homomorphism $\hat{f} : \hat{G} \rightarrow T^1$. We now define the three moves bringing $\hat{f}$ (and thus f) closer to harmony; see Figure 13.

- First, if the edges of $\hat{G}$ incident with v leave v via two edges around $\hat{f}(v)$, which together form a 2_r -turn around $\hat{f}(v)$, then we can perform a **flip move** to $\hat{f}$, which transforms $\hat{f}$ into a homotopic map $\hat{f}' : \hat{G} \rightarrow T^1$ (Figure 13, top left), which is actually also a homomorphism. From $\hat{f}'$, we immediately deduce a simplicial map $f' : G \rightarrow T^1$.

Since $\hat{f}, \hat{f}' : \hat{G} \rightarrow T^1$ are both homomorphisms, we can also view a flip as a specific operation that turns a homomorphism into another one (we will use this point of view later).

- Second, let v be a vertex of $\hat{G}$. If the edges of $\hat{G}$ incident with v leave v via one, two, or three consecutive edges of T^1 around $\hat{f}(v)$, then we can perform a **shortening move** to $\hat{f}$, which transforms $\hat{f}$ into a homotopic, simplicial map $\hat{f}'$ of $\hat{G}$ in which no edge of $\hat{G}$ is longer than in $\hat{f}$ (Figure 13, bottom left). From $\hat{f}'$, we immediately deduce a simplicial map $f' : G \rightarrow T^1$.
- Third, consider a simple directed cycle C in $\hat{G}$ that makes only 3-turns under f . We say that a walk of $\hat{G}$, identified by its sequence of directed edges $(e_0, \dots, e_k)$, *follows* C if there is a walk on C , its *following walk*,




Figure 14: In the proof of Lemma 5.2, the first two and the last two edges of J cannot be mapped to two edge-disjoint walks.

(which may go back and forth on C), identified by its directed edges $(c_0, \dots, c_k)$ such that $f(e_i) = f(c_i)$ for each i . A walk $(e_0, \dots, e_k, e_{k+1})$ of $\hat{G}$ *pulls C to the left* if $(e_0, \dots, e_k)$ follows C , with following walk $(c_0, c_1, \dots, c_k)$, and if moreover edge $f(e_{k+1})$ leaves the image of the directed cycle C , at the end-vertex of c_k , to its left. The notion of being pulled to the right is defined analogously.

If C makes only 3-turns under f , is pulled left, and is not pulled right, we can define a **balancing move** as follows (Figure 13, right). We consider all the vertices of $\hat{G}$ that are parts of walks following C (including the vertices of C themselves) and move them to the left of C , in such a way that we still have a homomorphism when restricted to the vertices following C . There are exactly two possibilities to do this, depending on whether the cycle is rotated “slightly clockwise” or “slightly counterclockwise” (as in Figure 13). In any case, the length of the image of an edge of $\hat{G}$ does not increase, and we choose a possibility that strictly decreases the length of the image of at least one edge, based on how a walk pulls C to the left. As before, from $\hat{f}'$, we immediately deduce a simplicial map $f' : G \rightarrow T^1$.

These three moves are useful in the following sense:

LEMMA 5.2. *If f cannot be flipped, shortened, or balanced, then f is harmonious.*

Proof. Without loss of generality G is connected. Also f is a homomorphism (for otherwise f factors into a homomorphism $\hat{f} : \hat{G} \rightarrow T^1$ such that no move applies to $\hat{f}$ by definition, and such that if $\hat{f}$ is harmonious, then f also). If f is strongly harmonious, we are done, so we can assume that f is not strongly harmonious. So let v be a vertex of G , and let L be a (left or right) line in T^1 whose central vertex is $f(v)$ such that, for each walk W in G based at v , the image of W under f does not escape L . In the universal cover, this means the following. There are a vertex $\tilde{x}$ of $\tilde{T}^1$, projecting to $f(v)$, and a (left or right) line $\tilde{L}$ in $\tilde{T}^1$ whose central vertex is $\tilde{x}$, such that, for each walk W in G based at v , the lift of $f \circ W$ based at $\tilde{x}$ does not escape $\tilde{L}$. We assume that $\tilde{L}$ is a left line (equivalently, that it makes only 3_r -turns), the other case being similar. We claim that, for each walk W in G based at v , the lift of $f \circ W$ based at $\tilde{x}$ is actually *contained* in $\tilde{L}$.

First, we explain why the claim implies the lemma. Because T^1 is finite, there are a cycle graph O and a reduced closed walk $C : O \rightarrow T^1$ such that $\tilde{L}$ is a lift of C . Using the claim, and since $\tilde{L}$ is simple by Lemma 2.1, there is a simplicial map $f' : G \rightarrow O$ such that $f = C \circ f'$. Moreover, f' has no spur since f cannot be shortened. So f is harmonious, proving the lemma. There remains to prove the claim, which we do in the remaining part of the proof.

Recall that $\tilde{L}$ makes only 3_r -turns. In G , one can build a semi-infinite walk I based at v such that the lift of $f \circ I$ based at $\tilde{x}$ is equal to the non-negative part of $\tilde{L}$. (Indeed, at a given step on $\tilde{L}$, there is no edge that goes strictly to the right of $\tilde{L}$ since no walk based at v can escape L under f ; if all edges go strictly to the left of $\tilde{L}$ or backward on $\tilde{L}$, a shortening or a flip could be applied to f .) Since I is a semi-infinite walk in G , it contains a subwalk J that, after removing its first and last edge, becomes a simple loop Q in G , based at some vertex w of G . Let P be the prefix of I leading to Q . The lift of $f \circ P$ based at $\tilde{x}$ is a portion of the non-negative part of $\tilde{L}$, ending at a vertex $\tilde{y}$ of $\tilde{T}^1$. Also $f \circ Q$, regarded as a closed walk by concatenating it with itself, makes a 3_r -turn also at the middle occurrence of w ; Otherwise, the first two and the last two edges of J would map, under f , to two edge-disjoint walks of length two making 3_r -turns at $f(w)$ (Figure 14), and so one could stop I at this point and escape from $\tilde{L}$, a contradiction. In particular the lift of $f \circ Q$ based at $\tilde{y}$ is (a translate of) $\tilde{L}$.

We conclude in two steps. First we prove that for every walk W based at v in G , the lift of $f \circ W$ based at $\tilde{x}$ cannot enter the right side of $\tilde{L}$, even after staying in $\tilde{L}$ for a while. By contradiction, assume that it does. For every $n \geq 1$ the walk $P \cdot Q^n \cdot P^{-1} \cdot W$ is based at v in G , and is such that the lift of $f \circ (P \cdot Q^n \cdot P^{-1} \cdot W)$ based at $\tilde{x}$ enters the right side of $\tilde{L}$ after staying in $\tilde{L}$. There is n such that this lift does not intersect $\tilde{L}$ outside of its non-negative part. We obtained a walk based at v that escapes L under f , a contradiction.

Now we prove that the lift of $f \circ W$ based at $\tilde{x}$ cannot enter the left side of $\tilde{L}$, even after staying in $\tilde{L}$ for a while. By contradiction, assume that it does. Without loss of generality, removing the last edge from W gives a walk W' such that the lift of $f \circ W'$ based at $\tilde{x}$ is contained in $\tilde{L}$. Then the lift of $f \circ (P^{-1} \cdot W')$ based at $\tilde{y}$ is contained in $\tilde{L}$, which lifts $f \circ Q$, and so $P^{-1} \cdot W'$ follows Q . Then $P^{-1} \cdot W$ pulls Q to the left. Since Q is pulled to the left, and since no balancing applies, Q is pulled to the right. So there is a walk U based at w in G such that the lift of $f \circ U$ based at $\tilde{y}$ is contained in $\tilde{L}$, except for its last edge that enters the right side of $\tilde{L}$. Then the walk $P \cdot U$, based at v , is such that the lift of $f \circ (P \cdot U)$ based at $\tilde{x}$ enters the right side of $\tilde{L}$ after staying in $\tilde{L}$ for a while. That contradicts the previous paragraph. $\square$

5.3 Preliminaries on flip sequences. By the preceding lemma, a natural strategy is to apply flips, shortenings, and balancings as much as possible until it is not possible any more. Shortenings and balancings strictly decrease the length of the map, so only finitely many such moves can be applied. Most of the argument thus focuses on sequences of flips in which no shortening or balancing can be applied at any step. Recall that flips transform a homomorphism into another one, and thus *henceforth we consider maps from G to T^1 that are homomorphisms*.

Formally, a **flip sequence** is a sequence of homomorphisms $f_0, \dots, f_p : G \rightarrow T^1$ such that f_{i+1} results from a flip of f_i for every $0 \leq i < p$, and no shortening or balancing can be applied to any of $f_0, \dots, f_{p-1}$. We use the following conventions. Given $0 \leq i < p$ we call flip i and abusively denote by i the flip from f_i to f_{i+1} . Given $0 \leq i \leq j \leq p$, we denote by $F_{i \rightarrow j}$ the flip sequence $f_i, \dots, f_j$. Given a vertex v of G , we denote by $F|v$ the walk performed in T by the image of v through the flips of v in F .

The map f induces a **left-blue direction** of G , obtained by directing each edge e of G in such a way that $f(e)$ has a blue triangle on its left. Thus G becomes a digraph. A **source** in a digraph is a vertex that has no incoming edge. We will use the following trivial observation repeatedly, without mentioning it explicitly: Each flippable vertex v is a source of its left-blue direction, and each flip reverses the direction of the edges incident to v . We need a series of easy lemmas.

LEMMA 5.3. *Let F be a flip sequence. If v and w are two adjacent vertices in G , then v is flipped in F in-between any two flips of w .*

Proof. When w is flipped, it is a source in its left-blue direction, and it can only become a source again after v is flipped. $\square$

LEMMA 5.4. *Let F be a flip sequence. If, before the first flip of F , vertex v is a source in its left-blue direction, and v cannot be flipped, then v is not flipped in F .*

Proof. Vertex v cannot be flipped before at least one of its neighbors is flipped, but no neighbor w of v can be flipped before v is flipped, because w is not a source. $\square$

LEMMA 5.5. *Let F be a flip sequence. If C is a cycle in G (not reduced to a single vertex) that is a directed cycle in its left-blue direction, then no vertex of C is flipped in F .*

Proof. The first vertex of C that would be flipped would not be a source (in its left-blue direction) before the flip. $\square$

Let e be an edge of G . In a flip sequence, assume that flip i flips an end-vertex of e , that flip j flips the other end-vertex of e , and that the end-vertices of e are not flipped between i and j . We say that flip j **counteracts** flip i if the image of e is rotated clockwise by i and counter-clockwise by j , or if it is rotated counter-clockwise by i and clockwise by j . See Figure 15.

LEMMA 5.6. *Let F be a flip sequence, and let $i < j$ be two flips of the same vertex v of G , such that no flip of v appears between i and j . Then we have the following properties:*




Figure 15: (Left) The second flip of the edge e does not counteract the first flip of e . (Right) The second flip of e counteracts the first one.




Figure 16: In a flip sequence, between two contiguous flips of a vertex v , the neighbors of v are flipped once.

- Every neighbor of v is flipped exactly once between i and j ;
- if every such flip counteracts flip i , then $F_{i \rightarrow j+1}|v$ is a 3_r -turn; otherwise, it is either a 1_r -turn or a -1_r -turn.

Proof. By Lemma 5.3, every neighbor of v is flipped exactly once between i and j , so there are three cases depicted in Figure 16. $\square$

LEMMA 5.7. Let $f_0, \dots, f_p$ be a flip sequence. If $i < j$ are flips of distinct adjacent vertices v and w respectively, and if no flip between i and j flips a neighbor of w , then flip j counteracts flip i .

Proof. Assume, for the sake of a contradiction, that flip j does not counteract flip i . See Figure 17. Let e be the edge of G between v and w , directed from v to w . Assume that flips i and j rotate the image of e clockwise, the other case being similar. Let N be the directed edges of G with source w .

We look at the situation in f_j , and thus just before flip j . We have that f_j maps N to two directed edges a and b of T^1 such that the reversal of a , followed by b , make a 2_r -turn in T . Since j rotates e clockwise, $f_j(e)$ is the reversal of b .

Note that w is not flipped between flips i and j (because otherwise v , a neighbor of w , would also be flipped between flips i and j , by Lemma 5.3). Since also no neighbor of w is flipped between flips i and j , we have $f_{i+1}(N) = f_j(N) = \{a, b\}$. Since i rotates the image of e clockwise, $f_i(e)$ is the edge in the middle of the 2_r -turn formed by a and b (directed towards the image of w). Thus $f_i(N)$ is included in a set of three consecutive directed edges with source $f_i(w)$, and contains the middle directed edge, and so f_i can be shortened, contradicting the fact that we have a flip sequence. $\square$

5.4 Proof of Theorem 1.2. In this section we prove Theorem 1.2. The proof follows from a few definitions and lemmas.




Figure 17: (Left) In the proof of Lemma 5.7, the vertices v and w , the edge e , and the set N of edges incident to w . (Right) If flips i and j both rotate the image of e clockwise, then f_i can be shortened.




Figure 18: (Left) Flat zone of size six, and its central vertex. (Right) In the proof of Lemma 5.9, if there is a flat zone $\tilde{T}_0$ of size nine in the universal covering triangulation of T , then the central vertex $\tilde{v}$ of $\tilde{T}_0$ projects to a vertex v on T such that every vertex of T at distance two or less from v has degree six.

LEMMA 5.8. Assume that G is connected and has q vertices. Let r be a vertex of G , and let F be a flip sequence of G that never flips r . Then F is composed of $O(q^2)$ flips.

Proof. In G every vertex is at distance less than q from r . By Lemma 5.3 every vertex at distance $i \geq 0$ from r is flipped at most i times. $\square$

In a reducing triangulation T , a **flat zone** of size $m' \geq 1$ is a subtriangulation T_0 of T isomorphic to the subdivision of a triangle depicted in Figure 18, in which the sides of T_0 have length m' .

LEMMA 5.9. Every flat zone of the universal covering triangulation of T has size less than $3(m+1)$, where m is the number of edges of T^1 .

The key property used in the proof is that, as we have assumed, S is not a torus. This is actually the only place where this assumption is used. (Since T is a reducing triangulation, S cannot be a sphere anyway.)

Proof. Assume the existence of a flat zone $\tilde{T}_0$ of size $3(m+1)$ in the universal covering triangulation of T . See Figure 18. The central vertex of $\tilde{T}_0$ is at distance $m+1$ from the boundary of $\tilde{T}_0$. So every vertex of T admits a lift in the interior of $\tilde{T}_0$. Then every vertex of T has degree six, so S is a torus by Euler's formula, contradicting our assumption. $\square$

A flip sequence F is **k -forward**, $k \geq 1$, if some vertex v of G is such that $F|v$ contains a subwalk of length k that makes only 3_r -turns.




Figure 19: In the proof of Lemma 5.11, if W_i makes a 1_r -turn, then W_j makes a 1_r -turn.

LEMMA 5.10. *Let F be a flip sequence of homomorphisms $G \rightarrow T^1$. If F is k -forward for some $k \geq 1$, then the universal covering triangulation of T has a flat zone of size k .*

Proof. Let $\tilde{T}$ be the universal covering triangulation of T . Lift F to a sequence $\tilde{F}$ of homomorphisms $\tilde{G} \rightarrow \tilde{T}^1$, where $\tilde{G}$ is a covering space of G . Assume without loss of generality (up to restricting to a smaller flip sequence) that some vertex v of $\tilde{G}$ is such that $\tilde{F}|v$ has length k and makes only 3_r -turns, and that the first and last flips of $\tilde{F}$ are flips of v . We shall prove by induction that $\tilde{F}|v$ is a side of a flat zone in $\tilde{T}$ lying to the left of $\tilde{F}|v$. The case $k = 1$ is clear since a flat zone of size one is just a triangle of $\tilde{T}$. So assume $k \geq 2$. Consider the first flip of v in $\tilde{F}$, and let $f : \tilde{G} \rightarrow \tilde{T}^1$ be the first map in $\tilde{F}$. Let e be the directed edge of $\tilde{T}$ along which the flip is performed (e is the first directed edge of $\tilde{F}|v$). Consider the triangle t of $\tilde{T}$ on the left of e , and the vertex x of t that is not incident to e .

There is a neighbor w of v in $\tilde{G}$ such that $\tilde{f}(w) = x$. By Lemma 5.3 there is in $\tilde{F}$ a flip of w in-between any two flips of v . Recall that v makes only 3_r -turns; by Lemma 5.6, this implies that every flip of w counteracts the previous flip of v . It follows that $\tilde{F}|w$ is a walk of length $k - 1$ that makes only -3_r -turns, running parallel to $\tilde{F}|v$; by Lemma 5.6 again, these walks can only be 3_r , 1_r , or -1_r -turns, and because the degree of each vertex of T^1 is at least six, the only possibility is that $\tilde{F}|w$ makes only 3_r -turns. By induction, there is a flat zone of size $k - 1$ on the left of $\tilde{F}|w$, so there is a flat zone of side k on the left of $\tilde{F}|v$. $\square$

We now need some definitions. Consider an ordering $v_0, \dots, v_{q-1}$ of the vertices of a graph. We say that v_i is a **lowpoint** if every neighbor of v_i , except perhaps v_0 , is higher than v_i in the ordering. The ordering is **proper** if (1) each vertex but v_0 is adjacent to a lower vertex, and (2) the set of lowpoints has the form $\{0, \dots, k\}$ for some k . A flip sequence $f_0, \dots, f_p$ is **proper** if there is a proper ordering $v_0, \dots, v_{q-1}$ of the vertices of G such that for every $0 \leq i < p$ the vertex flipped from f_i to f_{i+1} is $v_{i \bmod q}$.

LEMMA 5.11. *Assume that G is connected and has q vertices. Let F be a proper flip sequence of $kq + 2$ homomorphisms $G \rightarrow T^1$, for some integer k . Then F is k -forward.*

Proof. Let $v_0, \dots, v_{q-1}$ be the corresponding proper ordering of the vertices of G . Since F is composed of $kq + 1$ flips, the lowest vertex v_0 is both the first and the last vertex flipped in F , and v_0 is flipped $k + 1$ times in F , so $F|v_0$ has length $k + 1$. Let w_0 be the highest neighbor of v_0 . There are two cases. First assume that w_0 is a lowpoint. Then, consider any neighbor x of v_0 . Because w_0 is a lowpoint, x is a lowpoint as well. Thus v_0 is the last neighbor of x to be flipped in F before a flip of x , and so every flip of x counteracts the last flip of v_0 by Lemma 5.7. That being true for every neighbor x of v_0 , Lemma 5.6 implies that $F|v_0$ makes only 3_r -turns, and thus that F is $(k + 1)$ -forward.




Figure 20: In the proof of Lemma 5.11, W_i cannot make a 1_r -turn.

Now assume that there is a neighbor of w_0 distinct from v_0 that is lower than w_0 in L . Among the neighbors of w_0 lower than w_0 , let w_1 be the highest. Among the neighbors of w_1 lower than w_1 , let w_2 be the highest $\dots$, and so on until $w_{m-1} = v_0$ for some $m \geq 3$. Then $w_{m-1}, \dots, w_0$ are the vertices of a directed cycle C in G , in order, such that for every i the vertex w_{i+1} is both the last neighbor of w_i flipped before w_i , and the last vertex of C flipped before w_i , indices being taken modulo m .

Let I contain the flips $i \in \{0, \dots, (k-1)q\}$ whose vertex flipped, here $v_{i \bmod q}$, belongs to C . For every $i \in I$, the image of $v_{i \bmod q}$ makes a walk W_i of length two between f_i and f_{i+q+1} . Indeed $v_{i \bmod q}$ is here flipped twice: in flips i and $i+q$.

We claim that for every two consecutive $i < j \in I$, if W_i makes a 1_r -turn, then W_j also makes a 1_r -turn. See Figure 19. Consider the edge e between $v_{i \bmod q}$ and $v_{j \bmod q}$, and consider the flips $i, j, i+q$, and $j+q$. By definition no neighbour of $v_{j \bmod q}$ is flipped between i and j , nor between $i+q$ and $j+q$, so by Lemma 5.7 flip j counteracts flip i , and flip $j+q$ counteracts flip $i+q$. Then flip i cannot rotate the image of e clockwise (as in the bottom part of Figure 19), for otherwise flip j would rotate the image of e counter-clockwise since it counteracts flip i , and then flip $i+q$ could not be such that W_i makes a 1_r -turn. So flip i rotates the image of e counter-clockwise (as in the top part of Figure 19). Then flip j rotates the image of e clockwise, since it counteracts flip i . Flip $i+q$ is also clockwise since W_i makes a 1_r -turn, and flip $j+q$ is counter-clockwise since it counteracts flip $i+q$. We proved that W_j makes a 1_r -turn.

We use the claim immediately to prove by contradiction that for every $i \in I$, if $i \leq (k-2)q$, then W_i does not make a 1_r -turn. For otherwise, by our claim, every $l > i \in I$ is such that W_l makes a 1_r -turn. In particular, the smallest $j > i \in I$ is such that W_j makes a 1_r -turn, so the flips $i, j, i+q$, and $j+q$ are such as depicted in Figure 20. But then flip $i+2q$ cannot be such that W_{i+q} makes a 1_r -turn, which is a contradiction.

The same arguments show that for every $i \in I$, if $i \leq (k-2)q$, then W_i does not make a -1_r -turn. So W_i makes a 3_r -turn by Lemma 5.6. In particular $F_{0 \rightarrow (k-1)q+1}|v_0$ makes only 3_r -turns, so F is k -forward. $\square$

An ordering $v_0, v_1, \dots$ of the vertices of a digraph D is **monotonic** if every edge of D is directed from a vertex v_i to a vertex v_j such that $i < j$.

LEMMA 5.12. *Let D be a digraph of size n . If D has no directed cycle and has a single source, then we can compute in $O(n)$ time an ordering of the vertices of D that is both proper and monotonic.*

Proof. Without loss of generality, D has at least two vertices. Let v_0 be the unique source of D . Since D has no directed cycle, the vertex set of D can be partitioned into sets $I_0, \dots, I_m$, $m \geq 1$, where $I_0 := \{v_0\}$, and where I_k contains the sources of $D \setminus \bigcup_{i < k} I_i$ for every $1 \leq k \leq m$. Order arbitrarily each of the sets $I_0, \dots, I_m$, and concatenate them into an ordering L of the vertices of D . Then L is monotonic. Moreover, each vertex but v_0 is adjacent to a lower vertex, and $I_0 \cup I_1$ is the set of lowpoints of L , so L is proper. $\square$

Before proving Theorem 1.2 we detail how to detect, given a *homomorphism* f , if a balancing move applies to f :

LEMMA 5.13. *Let T be a reducing triangulation of size m . Let G be a graph of size n , and let $f : G \rightarrow T^1$ be a homomorphism. In $O(m+n)$ time we can determine if there is a simple closed walk C in G such that $f \circ C$ makes only 3-turns, is pulled left, and not right. In that case, both C and the subgraph G_0 of G spanned by the walks that follow C are computed at the same time.*

Proof. Without loss of generality, assume that we look for C such that $f \circ C$ makes only 3_r -turns (the 3_b -turns case being symmetric). Consider the following algorithm. As a preliminary, build a graph G' from G by detaching every vertex v of G from its incident edges, and by re-attaching those edges to copies of v as follows. Let N contain the directed edges emanating from v in G , whose basepoints have thus been detached from v . For every two directed edges a and b based at $f(v)$ in T such that (a^{-1}, b) makes a 3_r -turn, consider all the directed edges in N that are mapped to a or b (if any), and attach all their basepoints to a common copy v' of v . Mark v' with red if $f(N)$ contains a directed edge of T on the right of (a^{-1}, b) . Mark v' green if $f(N)$ does not contain any such directed edge, and if it contains a directed edge on the left of (a^{-1}, b) .

The graph G' *projects* to G in the sense that the edges of G' are those of G , and every vertex of G' corresponds to a vertex of G (though not in a one-to-one manner). We say that a vertex v' of G' *lifts* a vertex v of G if v is the projection of v' .

Build the graph G' in $O(m + n)$ time. Direct every edge e' of G' so that $f(e')$ sees red on its left. Then determine in $O(n)$ time if there is a connected component G'_0 of G' that does not contain any red vertex, that contains a green vertex, and that contains a simple directed cycle C' . If there is none, then return that nothing was found. Otherwise, return the closed walk C in G to which C' projects, and the subgraph G_0 of G' to which G'_0 projects.

Let us now prove that this algorithm is correct. Every closed walk C in G whose image walk makes only 3_r -turns is lifted by a directed cycle C' in G' , and if C is simple, then C' is simple. Conversely, every simple directed cycle C' in G' projects to a closed walk C that makes only 3_r -turns. And if the connected component G'_0 of C' contains no red vertex, then C is simple; Indeed any self-intersection vertex of C would either correspond to a self-intersection vertex of C' , which is impossible since C' is simple, or otherwise it would correspond to red vertices in C' . We conclude with the claim that G_0 is the subgraph of G spanned by the walks following C , and that C is pulled left (resp. right) if and only if G'_0 contains a green (resp. red) vertex. Indeed every walk W that follows C in G lifts to a walk W' based at some vertex of C' in G'_0 . Also, if W can be extended by one edge into a walk that pulls C to the left (resp. right), then the end-vertex w' of W' is marked green (resp. red). Reciprocally, any walk in G'_0 from C' to a vertex w' projects to a walk W in G that follows C . And if w' is green (resp. red), then W can be extended by one edge to pull C to the left (resp. right). $\square$

Proof of Theorem 1.2. By Lemma 5.1, we can assume that f is simplicial. Without loss of generality, we assume that G is connected, for otherwise we could apply the algorithm to each connected component separately. Making f harmonious is done using balancings, shortenings, and flips. Since balancings and shortenings decrease the length of the drawing, we give them priority over flips. More precisely, the algorithm consists in applying the routine given below, which only performs flips, with the following important twist, left implicit in the description: *whenever a balancing or a shortening is possible, we apply it and resume the routine from scratch.* Recall that the simplicial map $f : G \rightarrow T^1$ factors as a homomorphism $\hat{f} : \hat{G} \rightarrow T^1$; it is convenient to express the routine in terms of $\hat{f}$, since flips can be described at the homomorphism level. Here is the routine:

1. Choose an arbitrary vertex r of $\hat{G}$, and flip any vertex of $\hat{G}$ other than r , in any order, as long as possible.
2. Direct $\hat{G}$ with the left-blue direction. If $\hat{G}$ has no directed cycle and has a single source, then do the following: apply Lemma 5.12 to build in $O(n)$ time a proper and monotonic ordering $v_0, \dots, v_{q-1}$ of the $q \geq 1$ vertices of $\hat{G}$; initialize $i := 0$. Then, while it is possible to flip $v_{i \bmod q}$, flip it and increment i . (Some precisions: (a) we go to Step 3 as soon as $v_{i \bmod q}$ is not flippable; (b) The ordering is fixed during this entire step; if, after a flip, we update the direction of the edges to preserve the left-blue direction, it ceases to be monotonic.)
3. Flip any vertex of $\hat{G}$ (even possibly r), in any order, as long as possible.

If the algorithm terminates, then f is harmonious by Lemma 5.2. We now bound the number of flips of the routine, assuming that it is not interrupted by a balancing or a shortening. Step 1 does not flip r , so it consists of $O(n^2)$ flips by Lemma 5.8. Also the flip sequence of Step 2 is proper, so it has length $O(mn)$ by Lemmas 5.9, 5.10, and 5.11. Let F be the flip sequence of Step 3. For the sake of the analysis, we preserve the left-blue direction of the edges of $\hat{G}$ after each flip (equivalently, at each flip of a vertex v , we reverse the direction of the edges incident to v). We now prove that in all cases, some vertex of $\hat{G}$ is not flipped in F :

- If Step 2 was skipped because $\hat{G}$ has a directed cycle, then this cycle remains fixed by Lemma 5.5;

- if Step 2 was skipped because $\hat{G}$ has a source v distinct from r , then v is a source that cannot be flipped, so v is not flipped in Step 3 by Lemma 5.4;
- if Step 2 was executed, we claim that in Step 2, every attempt to flip a vertex v happens when v is a source. Indeed, at the first round of flips, the neighbours of v that have already been flipped correspond precisely to the edges that were directed towards v in the initial monotonic ordering, which have thus been reversed by the flips, making v a source.

After each vertex of $\hat{G}$ has been flipped, the directed graph is again monotonic. This proves the claim. It follows that v cannot be flipped in Step 3 by Lemma 5.4.

This implies that the flip sequence F of Step 3 has length $O(n^2)$ by Lemma 5.8, as desired. Thus, overall, if not interrupted, the routine terminates after $O((m+n)n)$ flips if not interrupted. Since there are $O(n)$ balancings or shortenings, the total number of flips, balancings, and shortenings is $O((m+n)n^2)$.

To prove the claimed running time of $O((m+n)^2n^2)$ time, it remains to note that finding and applying the next move, or correctly asserting that no move can be applied anymore, takes $O(m+n)$ time. Indeed, $\hat{G}$ and $\hat{f}$ can be computed in $O(n)$ time by constructing the clusters of f . On $\hat{G}$ finding a flip, or correctly asserting that there is none, takes $O(m+n)$ time. Same for the shortenings. Concerning the balancings, an application of Lemma 5.13 dermines in $O(m+n)$ time if there is a simple closed walk C in $\hat{G}$ such that $\hat{f} \circ C$ makes only 3-turns, is pulled left, and is not pulled right. If there is none, then no balancing is available. Otherwise C can be balanced, and Lemma 5.13 computes both C and the subgraph $\hat{G}_0$ of $\hat{G}$ that will move with C during the balancing. The modification brought to f by the balancing then takes $O(m+n)$ time. $\square$

6 Extensions to surfaces with boundary

In this section we extend the results of Sections 4 and 5 to reducing triangulations with boundary. To mimic the classical setting of Tutte's theorem, we consider the constraint of attaching vertices to the boundary of the surface. We formalize this constraint in the following definitions. Let S be a surface with non-empty boundary (such as the disk). Let G be a graph, and let $f : G \rightarrow S$ be a map. We use the following conventions: the boundary ∂S of S is directed so that the interior of S lies on its right, and the rotation system of a graph embedded on S records, for every vertex v embedded on S , the *clockwise* order of the edges meeting v .

Anchor. An *anchor* is a (possibly empty) set A of vertices of G mapped to ∂S by f , together with linear orderings of those vertices mapped to the same point of ∂S . Note that some vertices of G can be mapped to ∂S without belonging to A .

Untangling relatively to an anchor. Let $f' : G \rightarrow S$ be obtained from f by sliding infinitesimally along ∂S the images of the vertices in A , so as to separate them in the orders prescribed by A . We say that f can be *untangled relatively to A* if there is an embedding homotopic to f' , where the homotopy fixes the image of every vertex in A .

Weak embeddings relative to an anchor. Informally, we say that f is a weak embedding *relative to A* if there exist embeddings arbitrarily close to f in which the vertices in A are embedded along ∂S , and in the orders prescribed by A . We shall not use this definition as is, and instead we consider an equivalent formulation in the case where f is a drawing in the 1-skeleton T^1 of a reducing triangulation T of S . This formulation better suits our needs. It is based on the combinatorial formulation of Section 2.3. More precisely, we extend T^1 to a graph M^* equipped with a rotation system, we extend G and f to a graph G^* and a drawing $f^* : G^* \rightarrow M^*$ as follows. Consider every two consecutive edges e and f along the boundary of T , and let x be the vertex of T between e and f . Let $v_1, \dots, v_k$ be the $k \geq 0$ vertices in A that are mapped to x by f , in the order prescribed by A . In M^* , create k degree one vertices attached to x by edges $s_1, \dots, s_k$: place them in the rotation system of M^* so that $e, s_1, \dots, s_k, f$ are consecutive around x . Then for every $1 \leq i \leq k$, create an edge incident to v_i in G^* , and map this edge to s_i in f^* . We say that f is a *weak embedding relative to A* if f^* is a weak embedding.

Our result is the following :

THEOREM 6.1. *Let S be a surface with boundary. Let T be a reducing triangulation of S , with m edges. Let G be a graph of size n , and let $f : G \rightarrow T^1$ be simplicial. Let A be an anchor for f . If f can be untangled rel. A , then we can compute in $O((m+n)^2n^2)$ time a simplicial map $f' : G \rightarrow T^1$, homotopic to f relatively to A , weak embedding relative to A , not longer than f .*




Figure 21: Construction of the 1-gadget.




Figure 22: Construction of the triangulation $T_\bullet$ in the proof of Theorem 6.1.

The proof of Theorem 6.1 goes by extending T to a reducing triangulation without boundary (in order to apply Theorem 1.2). This extension step uses ad-hoc gadgets that we now define. Let S be the surface of genus one with one boundary component. The **1-gadget** is the reducing triangulation T of S depicted in Figure 21. The boundary of S is the union of two edges of T : one edge is incident to a blue face of T , and the other to a red face. The **3-gadget** is then the reducing triangulation T' obtained from disjoint 1-gadgets T_1, T_2, T_3 by identifying the blue-incident edge of T_i with the red-incident edge of T_{i+1} for every $i \in \{1, 2\}$. The boundary of T' is the union of the red-incident edge of T_1 and the blue-incident edge of T_3 .

We call **crown** any reducing triangulation T of the annulus obtained from a circular list of $k \geq 2$ triangles $t_0, \dots, t_{k-1}$ by identifying some side of t_i with some side of t_{i+1} for every i , where indices are modulo k .

Proof of Theorem 6.1. We shall extend T to a reducing triangulation *without boundary* $T_\bullet$, and then extend f to a drawing on $T_\bullet$. The construction of $T_\bullet$ is as follows. First extend T to a reducing triangulation T_0 by attaching a crown to each boundary component of T . Choose the crowns so that if x is a vertex of the boundary of T , and if $k \geq 0$ vertices of A are mapped to x by f , then x is incident to at least $k + 6$ edges in the interior of its crown; Crucially, we make the observation (O) that the first three, and the last three edges incident to x in the interior of the crown will not be used by the extended drawing. Now build $T_\bullet$ from T_0 as follows. See Figure 22. Consider a copy T_1 of T_0 . Reverse the direction of T_1 , and exchange the colors of its faces. Identify every edge of the boundary of T_0 with its copy in T_1 . At this point the faces of $T_\bullet$ are properly colored, but $T_\bullet$ may not be a reducing triangulation since vertices on the boundary of T_0 may have low degree. This is fixed by cutting open in $T_\bullet$ every edge of the boundary of T_0 , and by identifying the two cut paths with the two boundary paths of a 3-gadget (keeping the coloring of the faces proper).

Now extend G and f to a graph $G_\bullet$, and a simplicial map $f_\bullet : G_\bullet \rightarrow T_\bullet^1$; First extend G and f to some G_0 and $f_0 : G_0 \rightarrow T_0^1$, as follows. Consider every vertex x of the boundary of T . Let $v_1, \dots, v_k$ for some $k \geq 0$ be the vertices of A mapped to x by f , in order. Let $e_1, \dots, e_{k+6}$ be consecutive edges around x in the interior of the crown incident to x . For every $1 \leq i \leq k$, add in G_0 an edge s_i incident to v_i , and let $f_0(s_i) := e_{i+3}$. We say that the newly created end-vertex of s_i is a tip vertex of G_0 , and that s_i is a tip edge of G . Build $G_\bullet$ and $f_\bullet$ from G_0 and f_0 by considering a copy G_1 of G_0 , and the mirror map $f_1 : G_1 \rightarrow T_1^1$, and by identifying in $G_0 \cup G_1$ and $f_0 \cup f_1$ every tip vertex of G_0 with its mirror vertex in G_1 .

To conclude, transform $f_\bullet$ to a map $f'_\bullet : G_\bullet \rightarrow T_\bullet^1$ with the algorithm of Theorem 1.2, and return $f' := f'_\bullet|_G$.

Let us show why this is correct. Let $T^\flat$ be the sub-triangulation of $T_\bullet$ that is the union of T and the mirror of T . Let $G^\flat$ be the subgraph of $G_\bullet$ that is the union of the two copies of G . By the above observation (O), applying the moves described in Section 5 to $f_\bullet$ does not modify the images of the stem edges, and preserves the fact that $f_\bullet(G^\flat) \subset T^\flat$. In particular f' is homotopic to f rel. A . Moreover $f_\bullet$ is homotopic to an embedding in $T_\bullet$ since f can be untangled relatively to A in T . Thus $f'_\bullet$ is a weak embedding by Theorem 1.1, and $f'_\bullet$ is not longer than $f_\bullet$. And so f' is a weak embedding relative to A , not longer than f . $\square$

Acknowledgments. We thank the reviewers for their detailed comments that led to several improvements in the presentation and to a critical fix in the definition of a swap.

References

- [1] Hugo A. Akitaya, Radoslav Fulek, and Csaba D. Tóth. Recognizing weak embeddings of graphs. *ACM Transactions on Algorithms (TALG)*, 15(4):1–27, 2019.
- [2] Mark Anthony Armstrong. *Basic topology*. Springer Science & Business Media, 2013.
- [3] Éric Colin de Verdière, Vincent Despré, and Loïc Dubois. Untangling graphs on surfaces. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4909–4941. SIAM, 2024.
- [4] Éric Colin de Verdière and Jeff Erickson. Tightening nonsimple paths and cycles on surfaces. *SIAM Journal on Computing*, 39(8):3784–3813, 2010.
- [5] Yves Colin de Verdière. Comment rendre géodésique une triangulation d’une surface ? *L’Enseignement Mathématique*, 37:201–212, 1991.
- [6] Max Dehn. Transformation der Kurven auf zweiseitigen Flächen. *Mathematische Annalen*, 72:413–421, 1912.
- [7] Vincent Despré and Francis Lazarus. Computing the geometric intersection number of curves. *Journal of the ACM (JACM)*, 66(6):1–49, 2019.
- [8] Giuseppe Di Battista and Fabrizio Frati. From Tutte to Floater and Gotsman: on the resolution of planar straight-line drawings and morphs. In *Proceedings of the 30th International Symposium on Graph Drawing and Network Visualization (GD)*, pages 109–122, 2021.
- [9] Loïc Dubois. Making Multicurves Cross Minimally on Surfaces. In *32nd Annual European Symposium on Algorithms ESA*, pages 50:1–50–15, 2024.
- [10] Peter Eades and Patrick Garvan. Drawing stressed planar graphs in three dimensions. In *Proceedings of the 4th International Symposium on Graph Drawing (GD)*, pages 212–223. Springer, 1996.
- [11] Herbert Edelsbrunner and John Harer. *Computational topology. An introduction*. AMS, Providence, Rhode Island, 2009.
- [12] Allan L Edmonds. Deformation of maps to branched coverings in dimension two. *Annals of Mathematics*, 110(1):113–125, 1979.
- [13] Friedrich Eisenbrand, Matthieu Haeberle, and Neta Singer. An improved bound on sums of square roots via the subspace theorem. In *Proceedings of the 40th International Symposium on Computational Geometry (SOCG)*, pages 54:1–54:8, 2024.
- [14] Joanna A. Ellis-Monaghan and Iain Moffatt. *Graphs on surfaces: dualities, polynomials, and knots*, volume 84. Springer-Verlag, 2013.
- [15] David B. A. Epstein. Curves on 2-manifolds and isotopies. *Acta Mathematica*, 115:83–107, 1966.
- [16] Jeff Erickson and Patrick Lin. Planar and toroidal morphs made easier. *Journal of Graph Algorithms and Applications*, 27:95–118, 2023.
- [17] Jeff Erickson and Kim Whittlesey. Transforming curves on surfaces redux. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1646–1655, 2013.
- [18] Michael Floater. Parametrization and smooth approximation of surface triangulations. *Journal of Computer-Aided Geometric Design*, 14(3):231–250, 1997.
- [19] Michael S. Floater and Craig Gotsman. How to morph tilings injectively. *Journal of Computational and Applied Mathematics*, 101:117–129, 1999.
- [20] Daniel Gonçalves and Benjamin Lévêque. Toroidal maps: Schnyder woods, orthogonal surfaces and straight-line representations. *Discrete & Computational Geometry*, 51(1):67–131, 2014.
- [21] Steven J. Gortler, Craig Gotsman, and Dylan Thurston. Discrete one-forms on meshes and applications to 3D mesh parameterization. *Journal of Computer-Aided Geometric Design*, 33:83–112, 2006.
- [22] Craig Gotsman and Vitaly Surazhsky. Guaranteed intersection-free polygon morphing. *Computers and Graphics*, 25(1):67–75, 2001.
- [23] Jonathan L. Gross and Thomas W. Tucker. *Topological graph theory*. Wiley, 1987.

- [24] Joel Hass and Peter Scott. Intersections of curves on surfaces. *Israel Journal of mathematics*, 51:90–120, 1985.
- [25] Joel Hass and Peter Scott. Simplicial energy and simplicial harmonic maps. *Asian Journal of Mathematics*, 19(4):593–636, 2015.
- [26] Stephen G. Kobourov. Force-directed drawing algorithms. In Roberto Tamassia, editor, *Handbook of graph drawing and visualization*, chapter 12. Chapman and Hall, 2006.
- [27] Francis Lazarus and Julien Rivaud. On the homotopy test on surfaces. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 440–449. IEEE, 2012.
- [28] Yanwen Luo, Tianqi Wu, and Xiaoping Zhu. The deformation space of geodesic triangulations and generalized Tutte’s embedding theorem. arXiv:2105.00612, 2021.
- [29] Jürgen Richter-Gebert. *Realization spaces of polytopes*, volume 1643 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1996.
- [30] John Stillwell. *Classical topology and combinatorial group theory*. Springer-Verlag, New York, second edition, 1993.
- [31] Carsten Thomassen. Tutte’s spring theorem. *Journal of Graph Theory*, 45(4):275–280, 2004.
- [32] William Thomas Tutte. How to draw a graph. *Proceedings of the London Mathematical Society*, 3(1):743–767, 1963.
- [33] Luca Vismara. Planar straight-line drawing algorithms. In Roberto Tamassia, editor, *Handbook of graph drawing and visualization*, chapter 6. Chapman and Hall, 2006.




Figure 23: (Left) A reducing triangulation T of the torus, and a closed walk C of length two in T . (Right) A portion of the universal covering triangulation of T , and a portion of a lift of C . In T there is no reduced closed walk freely homotopic to C .

A Proof of Lemmas 3.2 and 3.1

In this section we prove Lemmas 3.2 and 3.1, which we restate for convenience.

LEMMA 3.2. *There are a reducing triangulation T of the torus, and a closed walk C in T , such that every closed walk freely homotopic to C is not reduced.*

Proof. Consider the reducing triangulation T , and the closed walk C of Figure 23. Assume the existence of a reduced closed walk C' in T^1 . If at some point C' takes a directed edge of T that sees red on its left, then C' makes only 3_r -turns. Otherwise C' makes only 3_b -turns. In both cases C' is not freely homotopic to C . $\square$

LEMMA 3.1. *Let S be a closed surface, not the sphere nor the torus. There are a reducing triangulation T of S , and a closed walk C in T , such that every closed walk freely homotopic to C is not strongly harmonious.*

Proof. There is a reducing triangulation T of S whose vertex degrees are all greater than or equal to eight, see [3, Figure 17]. There is a closed walk C_2 in T^1 that makes only 3_r -turns; Indeed every walk that makes only 3_r -turns will repeat itself since T is finite. Let C_1 be the closed walk obtained by pushing C_2 to the left in the way depicted in [3, Figure 9]: C_2 makes only -3_r -turns. Then C_1 and C_2 are not strongly harmonious. Let C' be a closed walk in T^1 , freely homotopic to C_1 or C_2 . If C' is strongly harmonious, then C' is reduced, and it follows from the work of É. Colin de Verdière, Despré, and Dubois [3] that C' is equal to C_1 or to C_2 (up to cyclic permutation and reversal). Indeed T fits their more restrictive definition of reducing triangulation. Also if C' does not make only 3_r -turns, then C' fits their more restrictive definition of reduced closed walk, and so $C' = C_1$ by [3, Proposition 3.4]. If C' makes only 3_r -turns, then reversing the colors of T makes C' reduced under their definition, and gives $C' = C_2$. $\square$