

COMPUTING SHORTEST CLOSED CURVES ON NON-ORIENTABLE SURFACES

Denys Bulavka,^{*} Éric Colin de Verdière,[†] and Niloufar Fuladi[‡]

ABSTRACT. We initiate the study of computing shortest non-separating simple closed curves with some given topological properties on non-orientable surfaces. While, for orientable surfaces, any two non-separating simple closed curves are related by a self-homeomorphism of the surface, and computing shortest such curves has been vastly studied, for non-orientable ones the classification of non-separating simple closed curves up to ambient homeomorphism is subtler, depending on whether the curve is one-sided or two-sided, and whether it is orienting or not (whether it cuts the surface into an orientable one).

We prove that computing a shortest orienting (weakly) simple closed curve on a non-orientable combinatorial surface is NP-hard but fixed-parameter tractable in the genus of the surface. In contrast, we can compute a shortest non-separating non-orienting (weakly) simple closed curve with given sidedness in $g^{O(1)} \cdot n \log n$ time, where g is the genus and n the size of the surface.

For these algorithms, we develop tools that can be of independent interest, to compute a variation on canonical systems of loops for non-orientable surfaces based on the computation of an orienting curve, and some covering spaces that are essentially quotients of homology covers.

1 Introduction

In computational topology of graphs on surfaces, much effort has been devoted to computing shortest closed curves with prescribed topological properties on a given surface. Most notably, the computation of shortest non-contractible, or shortest non-separating, closed curves on a combinatorial surface has been studied, under various scenarios, in at least a

^{*}*Einstein Institute of Mathematics, Hebrew University, Jerusalem 91904, Israel.* denys.bulavka@mail.huji.ac.il. Work done while this author was at Department of Applied Mathematics, Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic. This work was initiated while this author was visiting LIGM, funded by the French National Research Agency grant ANR-17-CE40-0033 (SoS), by the grant no. 21-32817S of the Czech Science Foundation (GACR) and by Charles University project PRIMUS/21/SCI/014.

[†]*LIGM, CNRS, Univ. Gustave Eiffel, F-77454 Marne-la-Vallée, France.* eric.colin-de-verdiere@univ-eiffel.fr. The work of this author is partly funded by the French National Research Agency grants ANR-17-CE40-0033 (SoS) and ANR-19-CE40-0014 (MIN-MAX).

[‡]*LORIA, CNRS, INRIA, Université de Lorraine, F-54000 Nancy, France.* niloufar.fuladi@aol.com. Work done while this author was at LIGM, CNRS, Univ. Gustave Eiffel, F-77454 Marne-la-Vallée, France. The work of this author is partly funded by the French National Research Agency grants ANR-17-CE40-0033 (SoS) and ANR-19-CE40-0014 (MIN-MAX).

dozen papers in the last twenty years [11, Table 23.2]. Also, algorithms have been given to compute shortest splitting closed curves [7], shortest essential closed curves [20], shortest closed curves within some non-trivial homotopy class [6], and shortest closed curves within a given homotopy class [12]. In all these cases, the purpose is to compute a shortest closed curve in a given equivalence class, for various notions of equivalence.

Identifying two closed curves on a given surface whenever there is a self-homeomorphism of the surface mapping one to the other is certainly one of the most natural equivalence relations. In particular, this is the most refined relation if we are only given the input surface, and it is relevant in particular in the context of mapping class groups [21, Section 1.3.1]. Under this notion, on an orientable surface, any two simple non-separating closed curves are equivalent: Any non-separating simple closed curve cuts the surface into an orientable surface that has (oriented) genus one less than the original surface and with two boundary components. However, for non-orientable surfaces, it turns out that the classification is subtler: Excluding some low-genus surfaces, a non-separating simple closed curve can be two-sided (have a neighborhood homeomorphic to an annulus) or one-sided (in which case it has a neighborhood homeomorphic to a Möbius band); furthermore it can be orienting (when cutting along it yields an orientable surface with boundary) or not.

In this paper, we study the complexity of computing shortest non-separating simple closed curves in non-orientable surfaces, under the constraint of being either one-sided or two-sided, either orienting or not, developing, in passing, tools to handle non-orientable surfaces algorithmically. Before describing our results in detail, we survey previous works.

1.1 Previous Works

One of the most basic and studied questions in topological algorithms for graphs on surfaces is that of computing a shortest non-contractible or non-separating closed curve (the length of such a curve is called *edge-width* in topological graph theory [1] or *systole* in Riemannian geometry [25]). Algorithmically, the simplest setup for graphs on surfaces is that of *combinatorial surfaces*: On a surface \mathcal{S} , one is given an embedding of a graph G that is cellular (all faces are homeomorphic to open disks); each edge of the graph has a positive weight. The goal is to compute a shortest closed walk in G that is non-trivial on \mathcal{S} either in homotopy or in homology. It turns out that such closed walks are simple, so they are, respectively, a shortest cycle that does not bound a disk on \mathcal{S} , or that does not separate \mathcal{S} .

Algorithms for computing shortest non-contractible or non-separating closed curves on surfaces have been developed since the early 1990s [35]. The current fastest algorithm in terms of the size of the input, the number n of vertices, edges, and faces of G , due to Erickson and Har-Peled [17], runs in $O(n^2 \log n)$ time. However, it is typical to view the genus g of the surface as a small parameter, and under this perspective it is worth mentioning the algorithms by Cabello, Chambers, and Erickson [5], which runs in $O(g^2 n \log n)$ for generic weights, and Fox [22], which runs in $2^{O(g)} \cdot n \log \log n$, but for orientable surfaces only. We refer to a survey [11, Table 23.2] for many other results.

Other topological properties have also been considered. Chambers, Colin de Verdière, Erickson, Lazarus, and Whittlesey [7] study the complexity of computing a shortest “simple”

closed curve that splits the (orientable) surface into two pieces, neither of which is a disk. Erickson and Worah [20] give a near-quadratic time algorithm to compute a shortest essential “simple” closed curve on an orientable surface with boundary. Cabello, DeVos, Erickson, and Mohar [6] provide a near-linear time algorithm (for fixed genus) to compute a shortest “simple” closed curve within some (unspecified) non-trivial homotopy class. In all these problems, one cannot expect in general the output closed curve to be a simple cycle in the input graph G : It sometimes has to repeat vertices and edges of G , but it is *weakly simple* [9] in the sense that it can be made simple by an arbitrary perturbation of the curve on the surface. In order to store weakly simple curves, both for the output and at intermediate steps of the algorithm, it is convenient to use the dual framework of *cross-metric surface* setting [12], in which these curves are really simple.

Very few works devote tools specifically to non-orientable surfaces. On the combinatorial side, Matoušek, Sedgwick, Tancer, and Wagner [30] carefully describe the result of cutting a non-orientable surface along a simple arc or closed curve, and in particular emphasize that there are various flavors of non-separating simple closed curves. Very recently, Fuladi, Hubard, and de Mesmay [23] prove the existence of a canonical system of loops on a non-orientable surface in which each loop has multiplicity $O(1)$, and show that such a system of loops can be computed in polynomial time. There are some good reasons not to neglect non-orientable surfaces [23, Introduction]: Random surfaces are almost surely non-orientable; graphs embeddable on an orientable surface of Euler genus $2g$ are embeddable on a non-orientable surface of genus $2g + 1$, while graphs embeddable on the projective plane can have arbitrarily large orientable genus; non-orientable surfaces appear naturally, e.g., in the graph structure theorem of Robertson and Seymour [32].

1.2 Our Results

We obtain the following results on orienting (simple) closed curves on non-orientable surfaces:

Theorem 1. *It is NP-hard to decide, given a cross-metric surface (\mathcal{S}, G^*) and an integer k , whether a shortest orienting closed curve on (\mathcal{S}, G^*) has length at most k .*

Theorem 2. *Given a non-orientable cross-metric surface (\mathcal{S}, G^*) of genus g and size n , we can compute a shortest orienting closed curve in (\mathcal{S}, G^*) in $2^{O(g)} \cdot n \log n$ time. Such a shortest closed curve has multiplicity at most two.*

It turns out that orienting curves are always non-separating, and that their sidedness is prescribed by the genus of the surface (see Lemma 3 below). In contrast, computing shortest non-orienting (simple) closed curves can be done in polynomial time:

Theorem 3. *Given a non-orientable cross-metric surface (\mathcal{S}, G^*) of genus g and size n , we can compute a shortest non-separating non-orienting one-sided (respectively, two-sided) closed curve in (\mathcal{S}, G^*) in $O(\text{poly}(g)n \log n)$ time. More precisely:*

- *We can compute a shortest non-separating non-orienting one-sided closed curve in $O(g^3 n \log n)$ time if g is even, and in $O(g^4 n \log n)$ time if g is odd;*

- we can compute a shortest non-separating non-orienting two-sided closed curve in $O(g^4 n \log n)$ time if g is odd, and in $O(g^5 n \log n)$ time if g is even.

Such shortest closed curves have multiplicity at most two.

These results implicitly assume that such curves exist (equivalently, $g \geq 2$ in the first case and $g \geq 3$ in the second one; see Lemma 3). They are stated in the cross-metric surface model, but their outputs immediately translate to shortest weakly simple closed curves in the dual combinatorial surface (\mathcal{S}, G) .

In passing, we develop tools of independent interest for non-orientable surfaces. First, we give an algorithm to compute an orienting curve in linear time (Proposition 1), refining a construction given by Matoušek, Sedgwick, Tancer, and Wagner [30]. Second, we introduce an analog of canonical systems of loops for non-orientable surfaces with a new combinatorial pattern, the *standard systems of loops*, and show that we can compute one in asymptotically the same amount of time as canonical systems of loops on orientable surfaces (Proposition 2); we remark that Fuladi, Hubard, and de Mesmay [23] compute a canonical system of loops in polynomial time, but they do not provide a precise estimate on the running time, which is likely higher than ours. Such standard systems of loops can be used to compute homeomorphisms between non-orientable surfaces; moreover, some algorithms for surface-embedded graphs rely on canonical systems of loops in the orientable case [24, 30], and for non-orientable surfaces they could use standard systems of loops instead. Third, we introduce *subhomology covers*, more compact than the *homology covers* of Chambers, Erickson, Fox, and Nayyeri [8]¹, but capturing exactly the topological information that we need to classify orienting/non-orienting, one-sided/two-sided closed curves (Section 7).

1.3 Techniques and Organization of the Paper

After the preliminaries (Section 2), we first show that our two algorithmic results (Theorems 2 and 3) follow from a single common statement (Theorem 4). In short, the topological properties that we are considering are homological: Knowing the homology class of a closed curve is enough to decide whether it is separating or not, one-sided or two-sided, orienting or not. Moreover, the non-separating non-orienting curves with given sidedness are characterized by the fact that they belong to the union of $\text{poly}(g)$ affine subspaces of small codimension in the homology group. The rest of the paper is devoted to the proof of Theorem 4. In Section 4, we give a linear-time algorithm to compute an orienting curve. This serves as a first step to compute a standard system of loops in Section 5. Section 6 provides a way to convert homology from a canonical basis to a standard one. At this point, given a closed curve, we are able to decide whether it has the desired topological properties efficiently. In Section 7, we introduce *subhomology covers*, and then show how to compute a shortest path in the subhomology cover that, when projected, will become the desired closed curve on the surface (Section 8); this uses a third kind of system of loops, the shortest system of loops [19, 10]. Section 9 concludes the algorithm. The NP-hardness proof (Theorem 1)

¹The article by Chambers, Erickson, Fox, and Nayyeri [8] combines several conference abstracts, including one by Erickson and Nayyeri [18]; the material that we use from [8] appeared first in [18].

is a reduction from the Hamiltonian cycle problem in grid graphs, in the same spirit as the NP-hardness proof of computing a shortest splitting closed curve by Chambers, Colin de Verdière, Erickson, Lazarus, and Whittlesey [7]; it is presented in Section 10.

2 Preliminaries

2.1 Curves and Graphs on Surfaces

We use standard terminology of topology of surfaces and graphs drawn on them; see, e.g., Armstrong [2] or, for a more algorithmic perspective, Colin de Verdière [11]. In this paper, the *genus* of a surface denotes its Euler genus (which equals twice the standard genus for orientable surfaces). Unless specified otherwise, surfaces are without boundary.

A simple closed curve on a surface \mathcal{S} is *two-sided* if it has a neighborhood homeomorphic to the annulus. Otherwise, it has a closed neighborhood homeomorphic to the Möbius band and it is called *one-sided*. A simple closed curve is called *non-separating* if the surface (with boundary) we obtain by cutting along it is connected; otherwise the curve is *separating*. Note that a separating curve is two-sided. A simple closed curve on a non-orientable surface is called *orienting* if by cutting along it, we obtain an orientable surface (with boundary).

A *homotopy* between two closed curves is a continuous family of closed curves between them. A closed curve is *contractible* if it is homotopic to a constant closed curve.

Finally, a *covering space* of \mathcal{S} is a topological space $\tilde{\mathcal{S}}$ with a continuous map $\pi: \tilde{\mathcal{S}} \rightarrow \mathcal{S}$ satisfying the local homeomorphism property, i.e., for every point x in \mathcal{S} there exists an open neighborhood U of x in \mathcal{S} and pairwise disjoint open sets U_1, \dots, U_d in $\tilde{\mathcal{S}}$ such that $\bigcup_{i=1}^d U_i = \pi^{-1}(U)$ and π restricted to each U_i is a homeomorphism with U . A *lift* of a path p on \mathcal{S} is a path \tilde{p} on $\tilde{\mathcal{S}}$ such that $\pi \circ \tilde{p} = p$.

2.2 Combinatorial and Cross-Metric Surfaces

A *combinatorial surface* (\mathcal{S}, G) is the data of a surface \mathcal{S} together with a positively edge-weighted graph G cellularly embedded on \mathcal{S} ; in this model, the curves (or the edges of the graphs) considered later are restricted to be walks in G . The length of a curve c is the sum of the weights of the edges of G used by c , counted with multiplicity. A closed curve (or a graph) on a combinatorial surface is *simple* if it is actually *weakly simple*; namely, if it admits an arbitrarily small perturbation on \mathcal{S} that turns it into a simple closed curve (or graph) on \mathcal{S} .

Weakly simple closed curves and graphs can traverse the same edge of G or visit the same vertex of G more than once. To keep track of these multiplicities, it is often useful to use the concept of *cross-metric surface* [12, Section 1.2], which refines the notion of combinatorial surface in dual form. A *cross-metric surface* (\mathcal{S}, G^*) is the data of a surface \mathcal{S} together with a positively edge-weighted graph G^* cellularly embedded on \mathcal{S} ; in this model, the graphs and curves considered later are in general position with respect to G^* . The length of a path p is the sum of the weights of the edges of G^* crossed by p , with multiplicity. The

multiplicity of a path or closed curve on (\mathcal{S}, G^*) is the maximum number of times it crosses a given edge of G^* . Every (weakly) simple closed curve on the combinatorial surface (\mathcal{S}, G) corresponds to a simple closed curve of the same length on the cross-metric surface (\mathcal{S}, G^*) , and conversely.

One can represent cellular graph embeddings on (possibly non-orientable) surfaces using, e.g., graph-encoded maps [29, 14]. We represent a graph embedded in the cross-metric surface (\mathcal{S}, G^*) by storing its *overlay* (or *arrangement*) with G^* , and our algorithms work on this representation. The *size* of a combinatorial or cross-metric surface is the number of vertices, edges, and faces of the underlying graph.

2.3 Shortest and Canonical Systems of Loops

A *system of loops* of a surface \mathcal{S} is a one-vertex graph L embedded on \mathcal{S} that has a single face, which is homeomorphic to an open disk. By Euler's formula, L has exactly g loops, where g is the (Euler) genus of \mathcal{S} . Cutting \mathcal{S} along L results in a $2g$ -gon, its *polygonal schema*, with edges of its boundary identified in pairs. The following result follows from earlier works.

Lemma 1 ([19, 10]). *Let (\mathcal{S}, G) be a combinatorial surface, orientable or not, of Euler genus g and size n . We can compute a set of $2g$ shortest paths on G such that each non-contractible closed walk in G intersects at least one of these shortest paths, in $O(gn + n \log n)$ time.*

Proof. We use the following result proved by Erickson and Whittlesey [19] and later extended by Colin de Verdière [10]: Let (\mathcal{S}, G^*) be a cross-metric surface of Euler genus g and let b be a point of \mathcal{S} not on G^* . We can compute, in $O(gn + n \log n)$ time, a shortest system of loops based at b in (\mathcal{S}, G^*) . Moreover, each loop is the concatenation of two shortest paths with endpoint b , together with a path that crosses a single edge of G^* .

Let (\mathcal{S}, G^*) be the cross-metric surface naturally associated to the combinatorial surface (\mathcal{S}, G) ; thus, G^* is the graph dual to G . Let b be an arbitrary point of \mathcal{S} not on the image of G^* . We apply the above result to (\mathcal{S}, G^*) and b , obtaining a set of g pairwise disjoint simple loops L based at b cutting \mathcal{S} into a disk. Moreover, the set of faces of G^* visited by the loops is included in the set of faces of G^* visited by $2g$ shortest paths in (\mathcal{S}, G^*) , or dually the set of vertices of G visited by the loops is included in the union of $2g$ shortest paths in G . Moreover, any non-contractible closed curve in G must cross at least one loop in L , since every closed curve in a disk is contractible. \square

One can describe a polygonal schema (see Figures 1 and 2) by assigning a symbol and an orientation to each loop in L , and listing the symbols corresponding to the loops encountered along the boundary of the polygonal schema, in clockwise order say, indicating a loop by a bar when encountered with the opposite orientation. A *canonical system of loops for an orientable surface* of Euler genus $2g$ is a system of loops L such that the polygonal schema associated to L has the form $a_1 b_1 \bar{a}_1 \bar{b}_1 \dots a_g b_g \bar{a}_g \bar{b}_g$ (Figure 1). A *canonical system of loops for a non-orientable surface* of genus g is a system of one-sided loops L such that




Figure 1: Two equivalent views of a canonical system of loops of an orientable surface of Euler genus 6. Left: A polygonal schema of the form $a_1b_1\bar{a}_1\bar{b}_1a_2b_2\bar{a}_2\bar{b}_2a_3b_3\bar{a}_3\bar{b}_3$. To reconstruct the surface, one identifies the edges of the polygon in pairs, respecting the orientations of the arrows. Right: By gluing together the corners of the polygonal schema, we see that all the vertices of the polygon get actually identified into a single vertex on the surface. The “half-arrows” indicate one side of each loop: In the present case, all loops are two-sided, which is reflected by the fact that the right side of a loop when leaving the vertex is still the right side of the loop when it comes back to the vertex.




Figure 2: Two equivalent views of a canonical system of loops of a non-orientable surface of Euler genus 3. Left: A polygonal schema of the form $a_1a_1a_2a_2a_3a_3$ with its top (visible) face blue and its bottom (hidden) face white. Right: The cyclic ordering of the edges around the vertex after identifying the edges of the polygonal schema. All loops are one-sided, which is reflected by the fact that the right side of each loop when leaving the vertex becomes the left side of that loop when it comes back to the vertex. The colors of the corners indicate the side of the polygon.

the polygonal schema associated to L has the form $a_1a_1a_2a_2\cdots a_ga_g$ (Figure 2). We will use the following result by Lazarus, Pocchiola, Vegter, and Verroust [28] to compute a canonical system of loops of an orientable surface:

Lemma 2 ([28]). *Let (\mathcal{S}, G^*) be an orientable cross-metric surface with genus g and size n , and let b be an arbitrary point in (\mathcal{S}, G^*) . We can compute a canonical system of loops of (\mathcal{S}, G^*) based at b in $O(gn)$ time, such that each loop has multiplicity at most four.*

2.4 Topological Characterizations via Signatures

Let $L = (\ell_1, \dots, \ell_g)$ be any system of loops and let p be a path on \mathcal{S} in general position with respect to L . The *signature* $\sigma(p, L)$ of p with respect to L is the vector in \mathbb{Z}_2^g whose i th component is the modulo-2 number of crossings of p with ℓ_i . (Although we will not use this, we remark that $\sigma(p, L)$ expresses the homology of p in the homology basis dual to L [8].)

The following lemma relates the topological type of a curve on a non-orientable surface to its signature with respect to a canonical system of loops.

Lemma 3 (Schaefer and Štefankovič [33]). *Let L be a canonical system of loops on a non-orientable surface and γ a simple closed curve in general position with respect to L .*

1. γ is one-sided if and only if $\sigma(\gamma, L)$ has an odd number of 1 elements;
2. γ is orienting if and only if all the elements in $\sigma(\gamma, L)$ are 1;
3. γ is separating if and only if all the elements in $\sigma(\gamma, L)$ are 0.

Proof. Given a non-orientable surface \mathcal{S} of genus g , a *crosscap decomposition* is a system of simple disjoint one-sided closed curves on \mathcal{S} such that when \mathcal{S} is cut along them we obtain the sphere with g boundary components. This was first introduced by Mohar and is called a *planarizing system of disjoint one-sided curves*, abbreviated as PD1S-system in [31]. One can turn a canonical decomposition L to a cross-cap decomposition by splitting the loops from the vertex; this can be done close to the vertex such that the crossings between γ and the curves in the system do not change after splitting.

Let P denote the polygonal schema associated to L .

1. By definition of a canonical system of loops (Figure 2), any loop crossing L exactly once is one-sided, in the sense that moving along it exactly once reverses the orientation (Figure 3). Thus, a simple closed curve γ in general position with respect to L is one-sided if and only if it crosses L an odd number of times.
2. Let Θ be a canonical decomposition we obtain from L as described above. It is sufficient to show that γ crosses each curve in Θ an odd number of times. By [33, Lemma 4], we know that a closed curve is orienting if and only if it crosses each curve in a crosscap decomposition an odd number of times. This finishes the proof.




Figure 3: Left: A two-sided curve crossing the canonical system of loops an even number of times. Right: A one-sided curve crossing the canonical system of loops an odd number of times.

3. The proof follows the same approach as in the previous case and using [33, Lemma 3] which states that a closed curve is separating if and only if it crosses each curve in a crosscap decomposition an even number of times.

As a side note, this statement is true if we replaced L by any system of loops. This is due to the fact that a curve that crosses each loop in such a system an even number of times is null-homologous and a simple closed curve is null-homologous if and only if it is separating. \square

Hence, orienting curves are non-separating, and their sidedness is prescribed by the genus.

3 Main Technical Result

All our algorithms will use the following theorem. As a motivation, we first show how Theorems 2 and 3 follow from it, while postponing its proof until the end of the paper.

Theorem 4. *Let (\mathcal{S}, G^*) be a cross-metric surface of Euler genus g and size n . Let k be an integer, $\rho : \mathbb{Z}_2^g \rightarrow \mathbb{Z}_2^k$ a linear map, and $A \subseteq \mathbb{Z}_2^k$.*

Then, for some (unspecified) canonical system of loops L that depends only on (\mathcal{S}, G^) , we can compute in $O(g^3 8^k n \log n)$ time a shortest closed curve c on (\mathcal{S}, G^*) such that $\rho(\sigma(c, L)) \in A$ (if such a curve exists). This curve is simple and has multiplicity at most two.*

We emphasize that the canonical system of loops L is not provided in the input of the algorithm, and is actually never computed explicitly.

Proof of Theorem 2, assuming Theorem 4. Recall that a curve is orienting if and only if it crosses every loop of a canonical system of loops an odd number of times (Lemma 3). Thus, to compute a shortest orienting closed curve, we apply Theorem 4 with $k = g$, ρ the identity, and A a single vector, the all-ones vector. \square

Proof of Theorem 3, assuming Theorem 4. Let $L = \{\ell_1, \dots, \ell_g\}$ be the canonical system of loops from Theorem 4 (which we know exists, even if we do not compute it). Let c be a closed curve. A loop ℓ_i is *even* (with respect to c) if ℓ_i and c cross an even number of times. Similarly, ℓ_i is *odd* if ℓ_i and c cross an odd number of times. The *oddtity number* of c is the number of odd loops with respect to c . We first consider the problem of computing shortest (non-separating) non-orienting, one-sided closed curves:

- If g is even, recall from Lemma 3 that a closed curve is non-separating, non-orienting, and one-sided if and only if its oddity number is odd. To compute a shortest such curve, it suffices to apply Theorem 4 with $k = 1$, $\rho(c_1, \dots, c_g) = \sum_i c_i$, and $A = \{1\}$.
- If g is odd, recall from Lemma 3 that a closed curve is non-separating, non-orienting, and one-sided if and only if its oddity number is odd, but different from g . To compute a shortest such curve, we do the following. For every $i = 1, \dots, g$, we apply Theorem 4 with $k = 2$, $\rho(c_1, \dots, c_g) = (c_1 + \dots + c_g, c_i)$, and $A = (1, 0)$. This computes a shortest closed curve with odd oddity number such that the i th loop of L is even. We return the shortest curve over all $i = 1, \dots, g$.

There remains to prove the theorem for non-separating, non-orienting, two-sided closed curves:

- If g is odd, recall from Lemma 3 that a closed curve is non-separating, non-orienting, and two-sided if and only if the oddity number is even and positive. To compute a shortest such curve, we do the following. For every $i = 1, \dots, g$, we apply Theorem 4 with $k = 2$, $\rho(c_1, \dots, c_g) = (c_1 + \dots + c_g, c_i)$, and $A = (0, 1)$. This computes a shortest closed curve with even oddity number such that the i th loop of L is odd. We return the shortest curve over all $i = 1, \dots, g$.
- If g is even, recall from Lemma 3 that a closed curve is non-separating, non-orienting, and two-sided if and only if the oddity number is even, positive, and different from g . To compute a shortest such curve, we do the following. For every $i \neq j \in \{1, \dots, g\}$, we apply Theorem 4 with $k = 3$, $\rho(c_1, \dots, c_g) = (c_1 + \dots + c_g, c_i, c_j)$, and $A = (0, 0, 1)$. This computes a shortest closed curve with an even oddity number such that the i th loop of L is even and the j th loop of L is odd. We return the shortest curve over all i, j . \square

4 Computing an Orienting Curve

Proposition 1. *Let (\mathcal{S}, G^*) be a non-orientable cross-metric surface of size n . We can compute an orienting curve on (\mathcal{S}, G^*) with multiplicity at most two in time $O(n)$.*

Matoušek, Sedgwick, Tancer, and Wagner [30, Proposition 5.5] proved the existence of such curve. We improve on their argument by providing a linear-time algorithm for its computation. We will use the following lemma (below and in Section 9), which is a variation of a very classical result: Every connected graph with even degrees has an Eulerian cycle.

Lemma 4. *Let (S, G) be a combinatorial surface of size n . Let E be the set of edges of G . Let $\mu : E \rightarrow \mathbb{Z}_+$ be a map such that (i) for each vertex v of G , the sum of the values of $\mu(e)$, for each e incident to v , is even, and (ii) the subgraph of G induced by the edges e such that $\mu(e) \geq 1$ is connected. Then one can compute a simple cycle in the dual cross-metric surface (S, G^*) such that, for each edge e of G , the cycle crosses the dual edge e^* exactly $\mu(e)$ times, in time linear in n plus the sum of the values of μ .*

We emphasize that, in Condition (ii), the considered subgraph cannot have isolated vertices, because it is induced by a set of edges, even though G itself may have vertices incident only to edges e such that $\mu(e) = 0$.

Proof. In the following proof we will use interchangeably the μ -value of an edge and its dual. The μ -degree of a vertex, as well as of a face, is the sum of the μ -values of its incident edges.

First, without loss of generality we can assume that $\mu(e) \in \{0, 1\}$: Indeed, otherwise, we subdivide each dual edge e^* whose μ -value is higher than 1 into $\mu(e)$ edges, each with μ -value equal to 1. Moreover, we can assume that each face of G^* has μ -degree 0, 2 or 4: Indeed, otherwise, we iteratively add a diagonal in a face of the dual graph of μ -degree at least six separating a part with μ -degree three from the rest, and set μ to 1 on the newly added diagonal; the hypotheses of the lemma still hold.

Within each face in G^* of μ -degree 0, 2 or 4, we draw 0, 1 or 2 disjoint simple paths, respectively, connecting the middle of each boundary edge with μ -value 1 in an arbitrary way. All these paths together form simple, pairwise disjoint cycles on the surface, but not necessarily a single cycle. To remedy this, we need to merge the cycles. At a high level, we perform a search in the *cycle graph*, the graph whose vertices are the cycles, and in which two vertices are adjacent if the corresponding cycles have a face in the dual graph (of the overlay of the cycles and G^*) incident to both of them. Then, in a second step, we merge the cycles together, by reconnecting each non-root cycle c to its parent cycle at the location where c was discovered.

To see that the cycle graph is connected, note that the edges of G^* cut the cycles into *pieces*. Let two pieces be *adjacent* if either they are consecutive pieces of the same cycle, or they share a face in the dual graph of the overlay of the cycles and G^* . By (ii), this graph is connected.

In more detail, we perform, e.g., a depth-first search rooted at an arbitrary vertex. To perform the depth-first search, two ingredients are needed: (1) we need to be able to compute the neighbors of a given cycle c (this is doable in time linear in the number of crossings of c with G^*), (2) we need to be able to mark cycles as explored or unexplored (for this purpose, with a linear-time preprocessing step, we make each edge that is a piece of a cycle c point to a separate data structure representing the cycle). Initially all cycles are unexplored. We start by exploring the root cycle, and as soon as we discover an unexplored cycle, we mark it as explored and recursively explore it. This takes linear time. Whenever a cycle c is discovered, we take note of the dual face of the overlay of G^* and of the cycles where this discovery happens.

In the second step, for each of these faces, we reroute c and its parent cycle to merge




Figure 4: Merging of two disjoint cycles inside a face without increasing the number of intersections with the graph.

them into a single cycle, see Figure 4. At the end of this process, all cycles are merged into a single one, which crosses every edge of G^* exactly once, as desired. \square

Proof of Proposition 1. Let G be the graph dual to G^* . We first choose an arbitrary orientation of every face of G . Let I be the set of *inconsistent* edges of G , which bound two faces of G with inconsistent orientations. If we start traversing the faces incident to a vertex of G by going around the vertex, in a complete turn, we must change orientations an even number of times to get back to the orientation of the initial face; this implies that each vertex of G is incident to an even number of edges in I . Thus the subgraph of G made of the edges in I has all its vertices of even degree, and cutting the surface along it yields an orientable surface. Moreover, we can compute I in linear time. See Figure 5. Now, for each edge e of G , let $\mu(e) = 1$ if $e \in I$ and $\mu(e) = 2$ otherwise. We apply Lemma 4 in the combinatorial surface (\mathcal{S}, G) (the hypotheses are obviously satisfied) to compute, in linear time, a simple cycle c in the cross-metric surface (\mathcal{S}, G^*) that crosses each edge e of G^* exactly $\mu(e)$ times.

There remains to prove that c is orienting. For this purpose, it suffices to exhibit an orientation of each face of the overlay of G^* and c in such a way that the inconsistent edges are exactly those arising from c . We do this as follows. We first orient the faces of the overlay that touch a vertex of G^* in the same way as the corresponding face of G was oriented in the beginning of the proof. Let f be a face of G^* . The *subfaces* of f are the faces of the overlay of G^* and c that lie inside it. Starting from an arbitrary subface of f that is already oriented, we propagate the orientation to all subfaces of f , in such a way that the subedges of c are inconsistent (see Figure 5, right). There is a unique way to do this, because the dual of the subfaces of f is a tree, and moreover these orientations are compatible with the already selected orientations of the subfaces of f touching a vertex of G , because each edge e^* of G^* is crossed an odd number of times by c if and only if e lies in I . By construction, the subedges of G^* are consistent. \square




Figure 5: Left: G and G^* with arbitrary orientations of the faces of G and the inconsistent edges of G (in red, dotted lines); the dual graph G^* is denoted in black lines. Right: an example of a cycle passing through all the edges and the derived orientation.

5 Computing a Standard System of Loops

In this section, we introduce standard systems of loops and show that we can compute one efficiently. We believe that this result can be of independent interest: It would perhaps be more natural to compute a canonical system of loops; but on non-orientable surfaces it is only known how to compute one in polynomial time [23] (the precise worst-case running time is certainly larger than $O(gn)$). We thus propose this alternate notion of *standard* systems of loops, which can be computed as quickly as the canonical systems of loops for orientable surfaces. In the next section, we show that for our purposes we can convert from one representation (in terms of parity of crossings) to the other.

A *standard system of loops* of a non-orientable surface \mathcal{S} of genus g is a system of loops such that the loops appear in the following order around the boundary of the corresponding polygonal schema (where bar denotes reversal and $p = \lfloor \frac{g-1}{2} \rfloor$): $zza_1b_1\bar{a}_1\bar{b}_1 \dots a_pb_p\bar{a}_p\bar{b}_p$ if g is odd, and $ywywa_1b_1\bar{a}_1\bar{b}_1 \dots a_pb_p\bar{a}_p\bar{b}_p$ if g is even.

Proposition 2. *Let (\mathcal{S}, G^*) be a non-orientable cross-metric surface of genus g and size n . In $O(gn)$ time, we can compute a standard system of loops on (\mathcal{S}, G^*) such that each loop has multiplicity at most 100.*

Proof. We first compute an orienting simple closed curve c with multiplicity at most two (Proposition 1). Let \mathcal{S}' be the orientable surface with boundary obtained by cutting \mathcal{S} along c . We remark that each edge of G^* corresponds to at most three edges in \mathcal{S}' , and these edges induce naturally a cross-metric structure on \mathcal{S}' .

Let us first assume that g is odd. Then c is one-sided, and \mathcal{S}' has a single boundary component B . Let u be an arbitrary point of B . Let \mathcal{S}' be the surface obtained by attaching a disk to the boundary component of \mathcal{S}' . In a first step, in $O(gn)$ time, we compute a canonical




Figure 6: The proof of Proposition 2, in the case where g is odd. Left: A pictorial view of the orientable surface \mathcal{S}' . The one-sided curve c is repeated twice on the boundary B of \mathcal{S}' . Right: On \mathcal{S} , the construction of the loop z , which is one-sided because c is one-sided.

system of loops L of $\bar{\mathcal{S}}'$, with basepoint v in the face incident with u , that has multiplicity at most four and does not cross any edge of B . We do this as follows. First, starting from \mathcal{S}' , we shrink the boundary component B to a point. Then, we apply the algorithm by Lazarus, Pocchiola, Vegter, and Verroust [28] (Lemma 2), obtaining a canonical system of loops of the resulting surface based at v . Finally, we expand back the point to the boundary component B .

Let us now connect u to the basepoint v with a path p that arrives at v at a suitable corner around v (Figure 6). By this we mean the following: The cyclic ordering of the edges at v is $a_1\bar{b}_1\bar{a}_1b_1\dots a_g\bar{b}_g\bar{a}_gb_g$ (where bar indicates the origin of an edge, and no bar indicates the target of an edge), see Figure 1; we make p arrive between two consecutive groups of the form $a_i\bar{b}_i\bar{a}_ib_i$. For this purpose, let f be the face containing the basepoint v of L ; it is incident with u . The loops cut f into *subfaces*; let f' be the subface incident with u . Starting from u in f' , we draw a path p that goes to a portion of L that lies on the boundary of f' , and then runs along L (possibly exiting f) until we get to the basepoint v at a suitable corner. Because, when running along the boundary of the polygonal schema, the suitable corners appear regularly, every four corners (see Figure 1), this can be done by running along at most two loops; each of these loops has multiplicity at most four on \mathcal{S}' , and thus p has multiplicity at most eight on \mathcal{S}' . This path p is computed in $O(n)$ time.

Finally, the set of loops L , together with the loop z that is the concatenation of the reversal of p , c , and (a slightly translated copy of) p , is a standard system of loops of \mathcal{S} , which is computed in $O(gn)$ time. Moreover, every edge of G^* is crossed at most twice by c and at most 24 times by p , so the multiplicity of each loop is at most 50 with respect to G^* .

There remains to consider the case where g is even; see Figure 7. In that case, \mathcal{S}' has exactly two boundary components. We choose an arbitrary basepoint u on one of the boundary components. Let u' be the unique point on the other boundary component such that u and u' are identified on \mathcal{S} . We first compute a path q of multiplicity one on \mathcal{S}' between u and u' . Let \mathcal{S}'' be the surface obtained from \mathcal{S}' by cutting along q ; it has a single boundary component. Let u'' be a point of \mathcal{S}'' mapping to u or u' after regluing to obtain \mathcal{S}' . Then we proceed as in the case where g is odd, using u'' in place of u . The final standard system of loops is made (after a small perturbation) of the canonical system of loops of \mathcal{S}'' , the loop y that is the concatenation of the reversal of p , q , and p (which is one-sided because the closed curve corresponding to q is one-sided; indeed, because c is




Figure 7: The proof of Proposition 2, in the case where g is even. Left: A pictorial view of the orientable surface \mathcal{S}' . The two-sided curve c corresponds to two boundary components of \mathcal{S}' . Right: On \mathcal{S} , the construction of the loops w and y .

two-sided and cuts \mathcal{S} into an orientable surface, any closed curve crossing c once must be one-sided, since otherwise \mathcal{S} would be orientable), and the loop w that is the concatenation of the reversal of p , c , and p (which is two-sided because c is two-sided). The multiplicity of each loop is at most 100 (because now we have to take into account the path q , which cuts some edges of \mathcal{S}' into two subedges). \square

6 Converting Between Canonical and Standard Signatures

Our next lemma describes a matrix allowing to change the coordinates in homology, from the standard basis to some canonical one. The motivation is that the topological properties we need are best expressed in terms of crossings with a canonical system of loops, while we can more easily compute a standard system of loops.

Lemma 5. *Assume L' is a standard system of loops on a non-orientable surface. There exists a canonical system of loops L such that, for each path p , we have $\sigma(p, L) = \varphi(\sigma(p, L'))$, where $\varphi : \mathbb{Z}_2^g \rightarrow \mathbb{Z}_2^g$ is the invertible linear map described by the following matrices depending on the parity of g .*




Figure 8: The cut-and-paste process to go from L' to L_1 in the case where g is odd. At each polygon, the blue diagonal introduces the cut and the red edge denotes the edge along which we paste. Here, X denotes the loops $\bar{a}_2\bar{b}_2a_3b_3\bar{a}_3\bar{b}_3 \dots a_p b_p \bar{a}_p \bar{b}_p$.

$$A_{\text{odd}} = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 1 & 0 & \cdots & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 1 & \cdots & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 1 & \cdots & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 1 & 0 & 1 & 0 & 1 & 0 & \cdots & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 & \cdots & 1 & 0 & 1 \end{bmatrix} \quad A_{\text{even}} = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 1 & \cdots & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & \cdots & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 1 & 0 & 1 & 0 & 1 & \cdots & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & \cdots & 1 & 0 & 1 \end{bmatrix}$$

We remark that we actually never compute the canonical system of loops L .

Proof. The proof uses the cut-and-paste technique used in the proof of the classification of surfaces. We start with the polygonal schema associated to L' and at each step we cut along a diagonal and paste along one of the sides of the polygon. Each diagonal can be seen as a concatenation of the edges in the polygonal schema and becomes a side in the new polygonal schema that we obtain after pasting.

This will prove that φ^{-1} is given by the following matrices, each being the inverse of the one stated in the lemma, as a simple computation shows.

$$A_{\text{odd}}^{-1} = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & \cdots & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 & 1 & \cdots & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & \cdots & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & \cdots & 1 & 1 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 1 & 1 & 0 \end{bmatrix} \quad A_{\text{even}}^{-1} = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & \cdots & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 & 1 & \cdots & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 & \cdots & 1 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & \cdots & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 1 & 1 & 0 \end{bmatrix}$$

First let us consider the case where $g = 2p + 1$ is odd. We first show that using 4 cut-and-paste moves we can transform the polygonal schema $L' : z z a_1 b_1 \bar{a}_1 \bar{b}_1 \dots a_p b_p \bar{a}_p \bar{b}_p$ to




Figure 9: The cut-and-paste process to go from L_{i-1} to L_i in the case where g is odd. Here, $X : \bar{a}_{i+1}\bar{b}_{i+1} \dots a_p b_p \bar{a}_p \bar{b}_p$ and $X' : c_1 c_1 \dots c_{2i-2} c_{2i-2}$.




Figure 10: The additional cut-and-paste move in the case where g is even.

$L_1 : c_1 c_1 c_2 c_2 z_1 z_1 a_2 b_2 \bar{a}_2 \bar{b}_2 \dots a_p b_p \bar{a}_p \bar{b}_p$, see Figure 8. We can see that $c_1 = e + a_1 = z + b_1$ in which $u + v$ denotes the concatenation of two loops u and v , ignoring ordering and orientation (which is irrelevant as far as crossing numbers are concerned); in other words, for any path p , we have $\sigma(p, c_1) = \sigma(p, z) + \sigma(p, b_1)$. Similarly, $c_2 = e = z + a_1 + b_1$ and $z_1 = e + b_1 = z + a_1$. We can see that these cut-and-paste scenarios can be described by the following matrix: let M_1 be a $g \times g$ matrix with columns m_i for $1 \leq i \leq g$ such that $m_1 = (1, 0, 1, 0, \dots, 0)$ (corresponding to c_1), $m_2 = (1, 1, 1, 0, \dots, 0)$ (corresponding to c_2), and $m_3 = (1, 1, 0, 0, \dots, 0)$ (corresponding to z_1). For $4 \leq i \leq g$, let m_i be the $g \times 1$ matrix with every element 0 except its i th element that is 1. Note that the entries of this matrix come from \mathbb{Z}_2 , the rows correspond to the elements in L' and the columns correspond to the elements in L_1 .

Let L_{i-1} denote the polygonal schema $c_1 c_1 \dots c_{2i-2} c_{2i-2} z_{i-1} z_{i-1} a_i b_i \bar{a}_i \bar{b}_i \dots a_p b_p$. Mimicking the same cut-and-paste moves depicted in Figure 8 for the sides z_{i-1} , a_i and b_i for $i > 1$ transforms L_{i-1} to the polygonal schema $L_i : c_1 c_1 \dots c_{2i} c_{2i} z_i z_i a_{i+1} b_{i+1} \bar{a}_{i+1} \bar{b}_{i+1} \dots a_p b_p \bar{a}_p \bar{b}_p$, see Figure 9. We denote the corresponding matrix by M_i ; this matrix changes the basis from L_i to L_{i-1} , keeping all the elements unchanged except z_{i-1} , a_i , and b_i : $c_{2i-1} = z_{i-1} + b_i$, $c_{2i} = z_{i-1} + a_i + b_i$, and $z_i = z_{i-1} + a_i$. Continuing this process, we obtain the canonical polygonal schema L with the final matrix being given by $A_{\text{odd}}^{-1} = M_1 M_2 \dots M_p$.

In the case where $g = 2p$ is even, we need one additional type of cut to transform $L' : y w \bar{y} w a_1 b_1 \bar{a}_1 \bar{b}_1 \dots a_{p-1} b_{p-1} \bar{a}_{p-1} \bar{b}_{p-1}$ to $L_1 : y y z z a_1 b_1 \bar{a}_1 \bar{b}_1 \dots a_{p-1} b_{p-1} \bar{a}_{p-1} \bar{b}_{p-1}$, see Figure 10. Then we can use the four types of cuts introduced for the case where the genus is odd to turn L_1 to L .

One can check that in this case the corresponding matrix is the matrix A_{even}^{-1} . \square

7 Subhomology Covers

In this section, we introduce subhomology covers, which are covering spaces related to quotients of the homology group. They generalize *cyclic double covers* introduced by Erickson [15] and used by Borradaile, Chambers, Fox, and Nayyeri [3] and are inspired from *homology covers* by Chambers, Erickson, Fox, and Nayyeri [8], but there are important differences with the latter: They can capture not necessarily the homology group, but arbitrary quotients of the homology group (which allows for faster algorithms), and they are defined on surfaces without boundary. This tool is not restricted to non-orientable surfaces and we present it for arbitrary surfaces. Our construction is inspired from the *voltage construction* by Gross and Tucker [26, Chapter 4].

Throughout this section, let (\mathcal{S}, G) be a combinatorial surface (orientable or not) of Euler genus g and size n , L a system of loops in general position with respect to G (we will only need the case where L is the standard system of loops, computed in Proposition 2, but we do not assume this for the construction). Furthermore, let k be an integer and $\rho : \mathbb{Z}_2^g \rightarrow \mathbb{Z}_2^k$ a linear map. We define $\alpha : E \rightarrow \mathbb{Z}_2^k$ by $\alpha(e) = \rho(\sigma(e, L))$.

Lemma 6. *The map α satisfies the Kirchhoff voltage law on every face: For every face of G with boundary edges e_1, \dots, e_m , we have that $\sum_{i=1}^m \alpha(e_i) = 0$.*

Proof. The boundary of a face bounds a disk and consequently is separating. Thus every loop in L intersects it an even number of times. \square

We define the graph \tilde{G} as a graph with vertices (v, ν) , one for every vertex v in G and ν in \mathbb{Z}_2^k . For each edge e in G , connecting the vertices u and v , and every ν in \mathbb{Z}_2^k , there is an edge in \tilde{G} connecting (u, ν) and $(v, \nu + \alpha(e))$. We now observe that each facial cycle in G (which bounds a face of G in \mathcal{S}) lifts to a cycle in \tilde{G} , by the Kirchhoff voltage law. By attaching a disk to the lift of each such cycle, we obtain a combinatorial surface $(\tilde{\mathcal{S}}, \tilde{G})$, which is naturally a (possibly non-connected) covering space of \mathcal{S} and is called the *subhomology cover associated to α* . An alternate way to build $(\tilde{\mathcal{S}}, \tilde{G})$ is to cut \mathcal{S} along L , obtaining a disk D , and to glue together 2^k copies of these disks, in such a way that the copy ν gets attached via a lift of loop $\ell_i \in L$ to the copy $\nu + \rho(s_i)$, where $s_i \in \mathbb{Z}_2^g$ has a single non-zero entry, the i th one (Figure 11); however, we will not need this equivalence.

Lemma 7. *Assume that one is given the combinatorial surface (\mathcal{S}, G) together with the map α . Then, in $O(2^k kn)$ time, we can compute the combinatorial surface $(\tilde{\mathcal{S}}, \tilde{G})$ that is the subhomology cover of (\mathcal{S}, G) associated to α . Moreover, $\chi(\tilde{\mathcal{S}}) = 2^k \chi(\mathcal{S})$.*

Again, the system of loops L is not part of the input; only the map α is specified.

Proof. The vertices of \tilde{G} are given by the lifts of the vertices of G , that is, by pairs (v, ν) for every vertex v of G and every $\nu \in \mathbb{Z}_2^k$. Every edge e from v to w in G lifts, given $\nu \in \mathbb{Z}_2^k$,




Figure 11: Left: A non-orientable surface of genus 3 cut along a canonical system of loops. Right: A non-orientable 4-sheeted covering space of the surface at left with genus 6 given by $\rho : \mathbb{Z}_3^k \rightarrow \mathbb{Z}_2^k$; $\rho(s_1) = (1, 0)$, $\rho(s_2) = (0, 0)$ and $\rho(s_3) = (0, 1)$. The notation l_i^j indicates the j th lift of the loop l_i .

to an edge from (v, ν) to $(w, \nu + \alpha(e))$. Finally, each facial cycle in G , incident to vertex v of G , lifts to 2^k faces of \tilde{G} in $\tilde{\mathcal{S}}$, each incident to (v, ν) for $\nu \in \mathbb{Z}_2^k$, as explained above.


The combinatorial map of \tilde{G} can thus be computed in $O(2^k kn)$ time. A more explicit construction would depend on the data structure used. For example, in the graph-encoded map data structure [29, 14], each flag f of G incident to vertex v corresponds to 2^k flags of \tilde{G} , denoted $\{(f, \nu) \mid \nu \in \mathbb{Z}_2^k\}$ where (f, ν) has associated vertex (v, ν) , and we can easily connect the flags via their three involutions, in overall $O(2^k kn)$ time.

The last claim follows from the fact that every vertex, edge, and face of G on \mathcal{S} corresponds to 2^k vertices, edges, and faces of \tilde{G} on $\tilde{\mathcal{S}}$, respectively. \square

Lemma 8. *Let c be a closed walk in G , and \tilde{c} be a lift of c in \tilde{G} , with endpoints (v, ν) and (v, ν') . Then $\rho(\sigma(c, L)) = \nu' - \nu$.*

Proof. By construction, if e connects w with w' , then each lift of e connects vertices (w, η) to $(w', \eta + \alpha(e))$ for some $\eta \in \mathbb{Z}_2^k$. Thus $\nu' - \nu$ equals the sum, over all edges e of c , of $\alpha(e) = \rho(\sigma(e, L))$. This, in turn, equals $\rho(\sigma(c, L))$. \square

Although we do not make an explicit use of it here, it is worth mentioning that our techniques are (co)homological in nature. Indeed, we could have defined the signature map ρ as a map from the first (co)homology group, over \mathbb{Z}_2 , instead of being from \mathbb{Z}_2^g . The difference between these two being that the later assumes a choice of basis, i.e., a system of loops. Different choices of basis will indeed produce isomorphic covering spaces since this

Figure 12: Proof of Lemma 9: the construction of c' .

one is uniquely determined by the normal subgroup of the fundamental group induced by the kernel of the map ρ .

8 Computing Shortest Closed Walks with Restriction on the Signature

Our algorithm will use the subroutine given in the following proposition, which is a variation on the strategy used earlier by Chambers, Erickson, Fox, and Nayyeri [8, Section 5.2]. It is a first step towards the proof of Theorem 4; however, instead of computing a simple closed curve in the cross-metric surface, we only compute a closed walk in the dual combinatorial surface.

Proposition 3. *Let (\mathcal{S}, G) be a combinatorial surface, orientable or not, of Euler genus g and size n . Moreover, let k be an integer, $\rho : \mathbb{Z}_2^g \rightarrow \mathbb{Z}_2^k$ a linear map, and $A \subseteq \mathbb{Z}_2^k$. Assume that for each edge e of G , the value of $\rho(\sigma(e, L))$ is given, for a fixed system of loops L of \mathcal{S} in general position with respect to G . Given this, we can compute in $O(g^3 8^k k n \log n)$ time a shortest closed curve c in G (a shortest closed walk) such that $\rho(\sigma(c, L)) \in A$.*

For the proof, we start by computing the subhomology cover $(\tilde{\mathcal{S}}, \tilde{G})$ associated to $\alpha(e) = \rho(\sigma(e, L))$ using Lemma 7. We also apply Lemma 1 to compute a set of $2g$ shortest paths p_1, \dots, p_{2g} in G that intersects every non-contractible closed walk in G .

The following lemma is inspired by Chambers, Erickson, Fox, and Nayyeri [8, Lemma 5.4].

Lemma 9. *Some shortest closed curve c in G such that $\rho(\sigma(c, L)) \in A$ has the following form: It is the projection of a shortest path in $(\tilde{\mathcal{S}}, \tilde{G})$ that starts with a subpath of a lift of some p_i and is otherwise disjoint from that lift.*

Proof. If A contains the zero vector, then any trivial closed curve (reduced to a single vertex) satisfies the desired property, so the lemma is trivially true. So, henceforth, we assume that A does not contain the zero vector.

Let c be a shortest closed curve in G such that $\rho(\sigma(c, L)) \in A$. Because A does not contain the zero vector, c is non-contractible and must thus meet some path p_i . We turn c into a loop ℓ by basing it at an arbitrary vertex of $c \cap p_i$. Let \tilde{p}_i be a path in \tilde{G} that is a lift of p_i in the subhomology cover $\tilde{\mathcal{S}}$, and let $\tilde{\ell}$ be a lift of ℓ that starts on \tilde{p}_i . See Figure 12. Let \tilde{w} be the last vertex of $\tilde{\ell} \cap \tilde{p}_i$ encountered when traversing $\tilde{\ell}$. Let $\tilde{\ell}'$ be the path in \tilde{G} with the same endpoints as $\tilde{\ell}$ obtained from $\tilde{\ell}$ by replacing its part before \tilde{w} with the subpath of \tilde{p}_i with the same endpoints. Finally, let ℓ' be the projection of $\tilde{\ell}'$ on \mathcal{S} , and c' be the associated closed curve.

Since $\tilde{\ell}$ and $\tilde{\ell}'$ have the same endpoints, Lemma 8 implies that $\rho(\sigma(c', L)) = \rho(\sigma(c, L))$, the latter being in A by hypothesis. Moreover, $\tilde{\ell}'$ is no longer than $\tilde{\ell}$, because \tilde{p}_i is a shortest path, since any lift of a shortest path is itself a shortest path. Thus, c' is no longer than c , and satisfies the desired properties. \square

We will also need the following immediate consequence of results by Erickson, Fox, and Lkhamsuren [16] and Cabello, Chambers, and Erickson [5].

Lemma 10. *Given a combinatorial surface (\mathcal{S}, G) , orientable or not, of Euler genus g and size n , with a distinguished face b , one can, after an $O(g^2 n \log n)$ -time preprocessing, compute the distance from any given vertex incident with b to any other vertex in $O(\log n)$ time.*

Proof. The case where the surface is orientable is precisely the result by Erickson, Fox, and Lkhamsuren [16, Theorem 5.1]. Moreover, as explained by Cabello, Chambers, and Erickson [5, Section 4.4], the non-orientable case reduces to the orientable one, by passing to the orientable double cover. \square

Proof of Proposition 3. As indicated above, we first compute the subhomology cover $(\tilde{\mathcal{S}}, \tilde{G})$ in $O(2^k kn)$ time (Lemma 7), and the paths p_1, \dots, p_{2g} in $O(gn + n \log n)$ time (Lemma 1).

Fix $i = 1, \dots, 2g$. We show below how to compute a shortest path in $(\tilde{\mathcal{S}}, \tilde{G})$ that starts with a subpath of a lift of p_i , is otherwise disjoint from that lift, and projects to a closed curve c such that $\rho(\sigma(c, L)) \in A$.

Let \tilde{p}_i be a lift of p_i and let $(v_0, \nu_0), (v_1, \nu_1), \dots, (v_m, \nu_m)$ be the sequence of vertices on \tilde{p}_i . We consider the combinatorial surface $\tilde{\mathcal{S}}'$ obtained by cutting $\tilde{\mathcal{S}}$ along the path \tilde{p}_i , thus forming a boundary, and then attaching a disk b to the resulting boundary component; each interior vertex (v_t, ν_t) of \tilde{p}_i , $1 \leq t \leq m - 1$, now corresponds to two vertices, $(v_t, \nu_t)^+$ and $(v_t, \nu_t)^-$; each duplicated edge has the same weight as its original. The vertices $(v_t, \nu_t)^\pm$ ($0 \leq t \leq m$, where for convenience $(v_0, \nu_0)^\pm$ and $(v_m, \nu_m)^\pm$ denote (v_0, ν_0) and (v_m, ν_m) , respectively) all lie on the boundary of the face b of $\tilde{\mathcal{S}}'$. By Lemma 8, it suffices to compute a shortest path, in this combinatorial surface, among all paths from some vertex $(v_t, \nu_t)^\pm$ to some corresponding vertex in the set $\{(v_t, \nu_t + a) \mid a \in A\}$. Lemma 10 allows us to do this: $(\tilde{\mathcal{S}}, \tilde{G})$ is a combinatorial surface of size $O(2^k n)$ and genus $O(2^k g)$ (by Lemma 7), and the same holds for the combinatorial surface in which we perform the computation; moreover, there are $O(2^k n)$ pairs of vertices between which we need to compute the distance. Thus, the preprocessing step takes $O(g^2 2^{3k} kn \log n)$ time, and the distance computations take $O(2^k kn \log n)$. Once we have computed a shortest distance between these pairs of points, we can compute an actual shortest path using Dijkstra's algorithm, without overhead.

Applying this for each $i = 1, \dots, 2g$, and returning the projection of the overall shortest path, we obtain the result. \square

9 Proof of Theorem 4

We will need the following intuitive lemma.

Lemma 11. *On a cross-metric surface (\mathcal{S}, G^*) , orientable or not, let L be a system of loops and c_1 and c_2 be two closed curves that cross each edge of G^* with the same parity. Then $\sigma(c_1, L) = \sigma(c_2, L)$.*

Proof. The lemma is obvious with homology tools: The assumption implies that c_1 and c_2 are homologous, so the conclusion holds. Below, we present an alternate, somewhat more self-contained, proof.

Let G be the dual graph of G^* . For $i = 1, 2$, we push c_i homotopically to the closed walk c'_i in G defined as follows: If c_i crosses edges e_1^*, \dots, e_k^* of G^* in this order, then c'_i traverses the corresponding edges e_1, \dots, e_k of G in the same order. We have (1) $\sigma(c_i, L) = \sigma(c'_i, L)$; indeed, it is a folklore result that the parity of the number of crossings between two closed curves depends only on the homotopy of these closed curves (e.g., because every homotopy between two multicurves can be realized by Reidemeister moves [13]). Moreover, for each edge e of G , the curves c'_1 and c'_2 traverse e with the same parity, which implies (2) $\sigma(c'_1, L) = \sigma(c'_2, L)$. The result follows from (1) and (2). \square

Proof of Theorem 4. Let (\mathcal{S}, G^*) be our input cross-metric surface, and let G be the dual graph of G^* . We start by computing, in $O(gn)$ time, a standard system of loops L' in general position with respect to G such that each loop crosses each edge of G at most 100 times; for this purpose, somewhat counterintuitively, we apply Proposition 2 in the cross-metric surface (\mathcal{S}, G) . For each edge e of G , we can now compute $\sigma(e, L')$, also in $O(gn)$ time.

Let $\rho' := \rho \circ \varphi$, where φ is the map from Lemma 5. For each edge e of G , we compute $\rho'(\sigma(e, L'))$, which by Lemma 5 equals $\rho(\sigma(e, L))$, for some fixed canonical system of loops L . This takes $O(kgn)$ time (there is no need to precompute the matrix of ρ' since its computation takes $O(kg^2)$ time). We then apply Proposition 3 in the combinatorial surface (\mathcal{S}, G) with the map ρ' and the standard system of loops L' : In $O(g^3 8^k kn \log n)$ time, we compute a shortest closed walk c in G such that $\rho'(\sigma(c, L')) \in A$, or equivalently $\rho(\sigma(c, L)) \in A$.

The remaining part of the proof is to turn c into a simple cycle. First, we build a map μ from the edges of G^* to $\{0, 1, 2\}$ based on c , as follows. Let e be an edge of G . If c does not traverse e , then we set $\mu(e) = 0$. If c traverses e a positive, even number of times, then we set $\mu(e) = 2$. Otherwise, we set $\mu(e) = 1$. Then, we apply Lemma 4 to (\mathcal{S}, G) and μ . This computes in linear time a simple closed curve c' in the cross-metric surface (\mathcal{S}, G^*) among those that cross each edge e^* of G^* exactly $\mu(e)$ times. We return c' . Indeed, by definition of μ , it has multiplicity at most two and it is no longer than c , which by Lemma 11 implies that it is a shortest closed curve such that $\rho(\sigma(c', L)) \in A$. \square

10 NP-Hardness of Computing a Shortest Orienting Closed Curve

In this section, we prove that it is NP-hard to decide, given a cross-metric surface and an integer k , whether a shortest orienting closed curve has length at most k (Theorem 1). The proof is a variation on the proof by Chambers, Colin de Verdière, Erickson, Lazarus, and




Figure 13: The reduction for the proof of Theorem 1. From left to right: (a) A part of the original grid graph G , where each vertex is overlaid with cycles of length eight, each denoted v_1, \dots, v_8 ; (b) the attached Möbius band (vertices marked a, b, c, v_1 , and v_5 are identified in pairs; the boundary of the Möbius band is thus cycle v_1, \dots, v_8). (c) Part of the system of loops inside a Möbius band.

Whittlesey [7] that computing the shortest splitting closed curve is NP-hard. For this proof, it is easier to reason in the realm of combinatorial surfaces.

Theorem 1. *It is NP-hard to decide, given a cross-metric surface (\mathcal{S}, G^*) and an integer k , whether a shortest orienting closed curve on (\mathcal{S}, G^*) has length at most k .*

Proof. A grid graph of size n is a graph induced by a set of n points in the two-dimensional integer grid. We reduce the problem of deciding whether there exists a Hamiltonian cycle in grid graphs, which is NP-hard [27], to that of computing a shortest orienting closed curve.

Let G be a grid graph of size n embedded in the sphere and without loss of generality assume that $n \geq 3$. First, we overlay G with n small cycles of length 8 each centered on a vertex of G ; see Figure 13(a). We then remove the interior of these cycles, obtaining a surface of genus zero with n boundary components. For each of these boundary components we associate a distinct Möbius band coming from a 4×4 -grid as in Figure 13(b). We identify the boundary of each Möbius band with the associated boundary component². Let G' be the resulting graph, which is embedded in a non-orientable surface \mathcal{S} of genus n . We assign weights to each edge of G' : Each edge in a Möbius band (including those edges in the cycles of length 8) has weight $\varepsilon < 1/(12n)$, and every other edge has weight one. Thus, (\mathcal{S}, G') is a combinatorial surface with size $O(n)$. We claim that G has a Hamiltonian cycle if and only if there exists an orienting closed curve of length at most $n + 1/2$ in the combinatorial surface (\mathcal{S}, G') . We prove this below, which in turn will conclude the proof.

As an auxiliary tool, we introduce a canonical system of loops L based at an arbitrary basepoint of \mathcal{S} , each loop associated to a given vertex of G . More precisely, for a given vertex of G let v_1, \dots, v_8 be the vertices of its corresponding 8-cycle, as in Figure 13(a), then its

²To illustrate this identification, let v_1, \dots, v_8 denote the vertices of a single boundary component of the surface as in Figure 13(a) and of the boundary of the associated Möbius band as in Figure 13(b). We proceed by identifying the edge $v_i v_{i+1}$ of the boundary component with the edge $v_i v_{i+1}$ of the associated Möbius band.




Figure 14: Local modifications at every vertex of a Hamiltonian cycle of G to turn it into a simple orienting cycle in the combinatorial surface (\mathcal{S}, G) .

associated loop is a small perturbation of the concatenation of a loop that goes from the basepoint to v_5 , goes “through the crosscap” from one copy of v_5 to the other copy, and finally goes back from v_5 to the basepoint via the same path, see Figure 13(c).

If G has a Hamiltonian cycle, thus of length n in G , we can easily modify that cycle (Figure 14) in such a way that it becomes a simple cycle of length at most $n + 6n\varepsilon < n + 1/2$ in G' crossing each loop in L exactly once, and thus orienting (Lemma 3).

Conversely, any orienting cycle c in (\mathcal{S}, G') of length at most $n + 1/2$ must cross each loop in L an odd number of times (Lemma 3), and thus pass through each of the n Möbius bands. Moreover, it uses at most n edges of weight 1 in G' . Thus, it corresponds, in G , to a Hamiltonian cycle. Indeed, since $n \geq 3$, any closed walk without vertex repetition in G is a Hamiltonian cycle in G . \square

11 Conclusion

We conclude with several remarks:

- Our main tool, Theorem 4, also holds for orientable surfaces. In that case, the proof is simpler, bypassing the detour with standard systems of loops: We compute a canonical system of loops using Lemma 2, compute the signature of every edge with respect to that system of loops, apply Proposition 3, and conclude as in the non-orientable case.
- There is a variant of our algorithm for non-orientable surfaces in which we can completely bypass the computation of a standard system of loops. Let us sketch the algorithm here. First, we compute an arbitrary system of loops L' , in $O(gn)$ time. Then, we compute the “change of coordinates” matrices between L' and some canonical system of loops L , as in the statement of Lemma 5; because L' is not fixed, we cannot make such matrices explicit, but since our proof is algorithmic (and indeed follows the proof of the classification theorem by Brahana [4, 34]), we can actually do this using the same principles, in time polynomial in g (without dependence on n). We emphasize that, after each cut-and-paste step, it is enough to remember only the signature vector $\sigma(\ell, L')$ for each loop ℓ in the current system of loops. More precisely,

L' is a polygonal schema of size $O(g)$; one can, in $O(g^2)$ time, obtain a sequence of $O(g)$ cut-and-paste operations to transform L' to L . Each such cut-and-paste operation replaces a loop, say ℓ_1 , with a new one, say ℓ_2 ; computing $\sigma(\ell_2, L')$ takes $O(g^2)$ time, because ℓ_2 is homotopic to the concatenation of $O(g)$ loops in the current system of loops, for which we know the signatures with respect to L' . Thus, in $O(g^3)$ total time, one can compute $\sigma(\ell, L')$ for all the loops $\ell \in L$, and then using matrix inversion on $g \times g$ -matrices one can compute $\sigma(\ell', L)$ for each loop $\ell' \in L'$, which are the values we need in the proof of Theorem 4.

The version presented here, while perhaps a bit longer, has some benefits compared to the route sketched above: (1) Our Propositions 1 and 2 are of independent interest and are likely to be reused in future work; (2) having an explicit relation on the homology bases of the standard and canonical systems of loops is somewhat nicer conceptually.

- Our techniques allow to compute a shortest overall one-sided closed curve in $O(g^3 n \log n)$ time (without controlling whether it is orienting or not); indeed, just apply Theorem 4 with $\rho(c_1, \dots, c_g) = \sum_i c_i$ and $A = \{1\}$.
- Our algorithms run in $O(n \log n)$ for fixed genus. It might be possible, in the case of two-sided curves, to obtain an algorithm with running time $O(n \log \log n)$, using the techniques by Chambers, Erickson, Fox, and Nayyeri [8, Section 4.3], though with a hidden dependence on the genus that is at least exponential.
- Finally, we have only considered non-separating closed curves, and leave open the complexity of computing a shortest separating closed curve with a given topological type (specifying the topology of the two surfaces with boundary resulting from cutting along it). Even on orientable surfaces, the following problem is fixed-parameter tractable in the genus [7, Theorem 6.1] but apparently neither known to be NP-hard nor polynomial-time solvable: Given an orientable surface \mathcal{S} of (orientable) genus g , compute a shortest simple closed curve that separates \mathcal{S} into a surface of (orientable) genus one and a surface of (orientable) genus $g - 1$.

Acknowledgments

We would like to thank Arnaud de Mesmay for stimulating discussions, and the anonymous reviewers for their useful comments.

References

- [1] Michael O. Albertson and Joan P. Hutchinson. The independence ratio and genus of a graph. *Trans. Amer. Math. Soc.*, 226:161–173, 1977. doi:10.2307/1997946.
- [2] Mark A. Armstrong. *Basic topology*. Undergraduate Texts in Mathematics. Springer-Verlag, New York-Berlin, 1983. doi:10.1007/978-1-4757-1793-8.

- [3] Glencora Borradaile, Erin W. Chambers, Kyle Fox, and Amir Nayyeri. Minimum cycle and homology bases of surface-embedded graphs. *J. Comput. Geom.*, 8(2):58–79, 2017. doi:[10.20382/jocg.v8i2a4](https://doi.org/10.20382/jocg.v8i2a4).
- [4] Henry R. Brahana. Systems of circuits on two-dimensional manifolds. *Ann. of Math.* (2), 23(2):144–168, 1921. doi:[10.2307/1968030](https://doi.org/10.2307/1968030).
- [5] Sergio Cabello, Erin W. Chambers, and Jeff Erickson. Multiple-source shortest paths in embedded graphs. *SIAM J. Comput.*, 42(4):1542–1571, 2013. doi:[10.1137/120864271](https://doi.org/10.1137/120864271).
- [6] Sergio Cabello, Matt Devos, Jeff Erickson, and Bojan Mohar. Finding one tight cycle. *ACM Trans. Algorithms*, 6(4):Article 61, 2010. doi:[10.1145/1824777.1824781](https://doi.org/10.1145/1824777.1824781).
- [7] Erin W. Chambers, Éric Colin de Verdière, Jeff Erickson, Francis Lazarus, and Kim Whittlesey. Splitting (complicated) surfaces is hard. *Comput. Geom.*, 41(1-2):94–110, 2008. doi:[10.1016/j.comgeo.2007.10.010](https://doi.org/10.1016/j.comgeo.2007.10.010).
- [8] Erin W. Chambers, Jeff Erickson, Kyle Fox, and Amir Nayyeri. Minimum cuts in surface graphs. *SIAM J. Comput.*, 52(1):156–195, 2023. doi:[10.1137/19M1291820](https://doi.org/10.1137/19M1291820).
- [9] Hsien-Chih Chang, Jeff Erickson, and Chao Xu. Detecting weakly simple polygons. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1655–1670. Society for Industrial and Applied Mathematics, 2015. doi:[10.1137/1.9781611973730.110](https://doi.org/10.1137/1.9781611973730.110).
- [10] Éric Colin de Verdière. Shortest cut graph of a surface with prescribed vertex set. In *Proceedings of the 18th European Symposium on Algorithms (ESA), part 2*, volume 6347 of *Lecture Notes in Computer Science*, pages 100–111. Springer-Verlag, Berlin, 2010. doi:[10.1007/978-3-642-15781-3_9](https://doi.org/10.1007/978-3-642-15781-3_9).
- [11] Éric Colin de Verdière. Computational topology of graphs on surfaces. In Jacob E. Goodman, Joseph O’Rourke, and Csaba Toth, editors, *Handbook of Discrete and Computational Geometry*, chapter 23, pages 605–636. CRC Press LLC, third edition, 2018. doi:[10.1201/9781315119601](https://doi.org/10.1201/9781315119601).
- [12] Éric Colin de Verdière and Jeff Erickson. Tightening nonsimple paths and cycles on surfaces. *SIAM J. Comput.*, 39(8):3784–3813, 2010. doi:[10.1137/090761653](https://doi.org/10.1137/090761653).
- [13] Maurits de Graaf and Alexander Schrijver. Making curves minimally crossing by Reidemeister moves. *J. Combin. Theory Ser. B*, 70(1):134–156, 1997. doi:[10.1006/jctb.1997.1754](https://doi.org/10.1006/jctb.1997.1754).
- [14] David Eppstein. Dynamic generators of topologically embedded graphs. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 599–608. Society for Industrial and Applied Mathematics, 2003.
- [15] Jeff Erickson. Shortest non-trivial cycles in directed surface graphs. In *Proceedings of the 27th Annual Symposium on Computational Geometry (SOCG)*, pages 236–243. Association for Computing Machinery, 2011. doi:[10.1145/1998196.1998231](https://doi.org/10.1145/1998196.1998231).

- [16] Jeff Erickson, Kyle Fox, and Luvsandondov Lkhamsuren. Holiest minimum-cost paths and flows in surface graphs. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 1319–1332. Association for Computing Machinery, 2018. doi:[10.1145/3188745.3188904](https://doi.org/10.1145/3188745.3188904).
- [17] Jeff Erickson and Sarel Har-Peled. Optimally cutting a surface into a disk. *Discrete Comput. Geom.*, 31(1):37–59, 2004. doi:[10.1007/s00454-003-2948-z](https://doi.org/10.1007/s00454-003-2948-z).
- [18] Jeff Erickson and Amir Nayyeri. Minimum cuts and shortest non-separating cycles via homology covers. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1166–1176. Society for Industrial and Applied Mathematics, 2011.
- [19] Jeff Erickson and Kim Whittlesey. Greedy optimal homotopy and homology generators. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1038–1046. Association for Computing Machinery, 2005.
- [20] Jeff Erickson and Pratik Worah. Computing the shortest essential cycle. *Discrete Comput. Geom.*, 44(4):912–930, 2010. doi:[10.1007/s00454-010-9241-8](https://doi.org/10.1007/s00454-010-9241-8).
- [21] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.
- [22] Kyle Fox. Shortest non-trivial cycles in directed and undirected surface graphs. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 352–364. Society for Industrial and Applied Mathematics, 2013.
- [23] Niloufar Fuladi, Alfredo Hubard, and Arnaud de Mesmay. Short topological decompositions of non-orientable surfaces. *Discrete Comput. Geom.*, 72(2):783–830, 2024. doi:[10.1007/s00454-023-00580-3](https://doi.org/10.1007/s00454-023-00580-3).
- [24] Cyril Gavoille and Claire Hilaire. Minor-universal graph for graphs on surfaces. *arXiv preprint arXiv:2305.06673*, 2023. doi:[10.48550/arXiv.2305.06673](https://doi.org/10.48550/arXiv.2305.06673).
- [25] Mikhael Gromov. Filling Riemannian manifolds. *J. Differential Geom.*, 18(1):1–147, 1983. doi:[10.4310/jdg/1214509283](https://doi.org/10.4310/jdg/1214509283).
- [26] Jonathan L. Gross and Thomas W. Tucker. *Topological graph theory*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons, Inc., New York, 1987.
- [27] Alon Itai, Christos H. Papadimitriou, and Jayme Luiz Szwarcfiter. Hamilton paths in grid graphs. *SIAM J. Comput.*, 11(4):676–686, 1982. doi:[10.1137/0211056](https://doi.org/10.1137/0211056).
- [28] Francis Lazarus, Michel Pocchiola, Gert Vegter, and Anne Verroust. Computing a canonical polygonal schema of an orientable triangulated surface. In *Proceedings of the 17th Annual Symposium on Computational Geometry (SOCG)*, pages 80–89. Association for Computing Machinery, 2001. doi:[10.1145/378583.378630](https://doi.org/10.1145/378583.378630).

- [29] Sóstenes Lins. Graph-encoded maps. *J. Combin. Theory Ser. B*, 32(2):171–181, 1982. doi:[10.1016/0095-8956\(82\)90033-8](https://doi.org/10.1016/0095-8956(82)90033-8).
- [30] Jiří Matoušek, Eric Sedgwick, Martin Tancer, and Uli Wagner. Untangling two systems of noncrossing curves. *Israel J. Math.*, 212(1):37–79, 2016. doi:[10.1007/s11856-016-1294-9](https://doi.org/10.1007/s11856-016-1294-9).
- [31] Bojan Mohar. The genus crossing number. *Ars Math. Contemp.*, 2(2):157–162, 2009. doi:[10.26493/1855-3974.21.157](https://doi.org/10.26493/1855-3974.21.157).
- [32] Neil Robertson and P. D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *J. Combin. Theory Ser. B*, 89(1):43–76, 2003. doi:[10.1016/S0095-8956\(03\)00042-X](https://doi.org/10.1016/S0095-8956(03)00042-X).
- [33] Marcus Schaefer and Daniel Štefankovič. The degenerate crossing number and higher-genus embeddings. *J. Graph Algorithms Appl.*, 26(1):35–58, 2022. doi:[10.7155/jgaa.00579](https://doi.org/10.7155/jgaa.00579).
- [34] John Stillwell. *Classical topology and combinatorial group theory*, volume 72 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1993. doi:[10.1007/978-1-4612-4372-4](https://doi.org/10.1007/978-1-4612-4372-4).
- [35] Carsten Thomassen. Embeddings of graphs with no short noncontractible cycles. *J. Combin. Theory Ser. B*, 48(2):155–177, 1990. doi:[10.1016/0095-8956\(90\)90115-G](https://doi.org/10.1016/0095-8956(90)90115-G).