Splitting (Complicated) Surfaces Is Hard*

Erin W. Chambers[†] Éric Colin de Verdière[‡] Jeff Erickson[§]

Francis Lazarus[¶] Kim Whittlesey[||]

Abstract

Let $\mathcal{M}$ be an orientable combinatorial surface. A cycle on $\mathcal{M}$ is *splitting* if it has no self-intersections and it partitions $\mathcal{M}$ into two components, neither of which is homeomorphic to a disk. In other words, splitting cycles are simple, separating, and non-contractible. We prove that finding the shortest splitting cycle on a combinatorial surface is NP-hard but fixed-parameter tractable with respect to the surface genus g and the number of boundary components b of the surface. Specifically, we describe an algorithm to compute the shortest splitting cycle in $(g+b)^{O(g+b)} n \log n$ time, where n is the complexity of the combinatorial surface.

1 Introduction

Optimization problems on surfaces in the fields of computational topology and topological graph theory have received much attention in the past few years. Erickson and Har-Peled [9] prove that computing the shortest graph whose removal cuts a surface into a topological disk is NP-hard. Colin de Verdière and Lazarus consider the problem of finding the shortest *simple* loop [7] or cycle [6] within a given homotopy class. A polynomial-time algorithm for the generalization of this problem to non-simple curves was recently obtained by Colin de Verdière and Erickson [5]. Erickson and Whittlesey [10] provide simple polynomial-time algorithms to compute the shortest homology basis and the shortest fundamental system of loops on a given surface.

Several authors have considered the problem of computing a shortest cycle with some prescribed topological property, such as non-contractibility. When the set of cycles with the desired property satisfies the so-called *3-path condition*, a generic algorithm of Mohar and Thomassen finds a shortest such cycle in $O(n^3)$ time [17, Sect. 4.3]. For example, the sets of non-separating and non-contractible cycles on any surface satisfy the 3-path condition, but the set of separating cycles on a surface does not. Erickson and Har-Peled describe a faster algorithm to compute non-separating and non-contractible cycles in $O(n^2 \log n)$ time [9]. Cabello and Mohar [3], Cabello [2], and Kutz [14] develop

*See http://www.cs.uiuc.edu/~jeffe/pubs/splitting.html for the most recent version of this paper.

[†]Department of Computer Science, University of Illinois, erinwolf@uiuc.edu. Research partially supported by an NSF Graduate Research Fellowship and by NSF grant DMS-0528086.

[‡]CNRS, Laboratoire d'informatique de l'École normale supérieure, Paris, France, Eric.Colin.de.Verdiere@ens.fr

[§]Department of Computer Science, University of Illinois, jeffe@cs.uiuc.edu. Research partially supported by NSF grants CCR-0093348, CCR-0219594, and DMS-0528086.

[¶]GIPSA-Lab, CNRS, Grenoble, France, Francis.Lazarus@gipsa-lab.inpg.fr

[||]Department of Mathematics, University of Illinois, kwhittle@math.uiuc.edu

faster algorithms when the surface has small genus; the fastest of these runs in $(g+b)^{O(g+b)} n \log n$ time [14].

In this paper, we consider the following natural optimization problem: Given an orientable 2-manifold $\mathcal{M}$, compute a shortest simple cycle that separates $\mathcal{M}$ into two topologically non-trivial components. For simplicity, we will call a simple non-contractible separating cycle a *splitting* cycle. The set of splitting cycles does not satisfy the 3-path condition, so a different approach is required to compute the shortest one.

After reviewing a few necessary concepts from topology and proving some preliminary results, we prove in Section 3 that computing the shortest splitting cycle on a given surface is NP-hard. In Section 4, we prove that a shortest splitting cycle on a surface of genus g with b boundary components cuts any shortest path on the surface $O(g+b)$ times; this bound is tight if the surface has no boundary. This property leads to an algorithm to compute a shortest splitting cycle in $(g+b)^{O(g+b)} n \log n$ time, which we describe in Section 5. Thus, we show that the shortest splitting problem is fixed-parameter tractable with respect to the genus and the number of boundary components of the surface. This is the first result of this kind among the previously cited works. In particular, although Erickson and Har-Peled provide an algorithm to compute the minimum cut graph on any surface of constant genus and constant number of boundary components in polynomial time, the order of the polynomial depends on the genus and on the number of boundary components [9].

In Section 6, we consider the problem of finding a shortest cycle that splits the surface into two surfaces of prescribed topology. For example, given a surface of genus g with b boundary components, and given $g' \leq g$, $b' \leq b+1$, we may look for a shortest cycle that splits $\mathcal{M}$ into two surfaces, one of which has genus g' and b' boundary components. Or we may look for a shortest splitting cycle that is not homotopic to a boundary. We give an algorithm to solve these questions in $(g+b)^{O(g+b)} n \log n$ time.

Finally, iteratively cutting a surface of genus g without boundary along a shortest splitting cycle and attaching disks to the boundaries of the resulting surface provides a way to compute a decomposition into g punctured tori, but we show in Section 7 that one does not necessarily obtain the shortest such decomposition with this method.

2 Preliminaries

2.1 Topological Background

We recall several notions from combinatorial and computational topology. See also Hatcher [12], Stillwell [18], or Zomorodian [20] for more details.

Curves on surfaces.

A *surface* (or 2-manifold with boundary) $\mathcal{M}$ is a topological Hausdorff space where each point has a neighborhood homeomorphic either to the plane or to the closed half-plane. The points with no neighborhood homeomorphic to the plane comprise the *boundary* of $\mathcal{M}$. All the surfaces considered in this paper are *compact*, *connected*, and *orientable*. Such a surface is homeomorphic to a sphere with g handles attached and b open disks removed, for some unique non-negative integers g and b called respectively the *genus* and the *number of boundaries* of $\mathcal{M}$.

Let $\mathcal{M}$ be a surface. A *path* on $\mathcal{M}$ is a continuous map $p\colon [0,1] \to \mathcal{M}$. The *endpoints* of p are $p(0)$ and $p(1)$. A *loop* with *basepoint* x is a path with both endpoints equal to x. An *arc* is a path that intersects the boundary of the surface precisely at its endpoints. A *cycle* is a continuous map $\gamma\colon S^1 \to \mathcal{M}$, where S^1 denotes the unit circle. A path, loop, or cycle is *simple* if it is one-to-one, except, of course, for the endpoints of a loop. Two cycles are *disjoint* if they do not intersect. Two paths are *disjoint* if they do not intersect, except possibly at their endpoints; in particular, two disjoint loops can intersect only at their common basepoint.

If p is a path, its *reversal* is the path $\bar{p}(t) = p(1-t)$. The *concatenation* $p \cdot q$ of two paths p and q with $p(1) = q(0)$ is defined by setting $(p \cdot q)(t) = p(2t)$ for all $t \leq 1/2$, and $(p \cdot q)(t) = q(2t-1)$ for all $t \geq 1/2$.

Systems of loops, system of arcs, and homotopy.

If S is a set of pairwise disjoint simple loops or arcs, $\mathcal{M} \setminus S$ denotes the surface with boundary obtained by *cutting* $\mathcal{M}$ along the loops or arcs in S. Let $\mathcal{M}$ be a surface with genus g and with b boundary components.

If $b = 0$, a *system of loops* on $\mathcal{M}$ is a set of pairwise disjoint simple loops L with a common basepoint such that $\mathcal{M} \setminus L$ is a topological disk. Any system of loops contains exactly $2g$ loops. $\mathcal{M} \setminus L$ is a $4g$-gon where each loop appears as two boundary edges; this $4g$-gon is called the *polygonal schema* associated with L.

If $b \geq 1$, a *system of arcs* on $\mathcal{M}$ is a set of pairwise disjoint simple arcs A such that $\mathcal{M} \setminus A$ is a topological disk. Any system of arcs contains exactly $2g + b - 1$ arcs, by Euler's formula and standard double-counting argument. $\mathcal{M} \setminus A$ is a $(8g + 4b - 4)$-gon where each arc appears as two boundary edges, the remaining $4g + 2b - 2$ edges of the $(8g + 4b - 4)$-gon corresponding to pieces of the boundaries of $\mathcal{M}$; this $(8g + 4b - 4)$-gon is called the *polygonal schema* associated with A.

A *homotopy* between two paths p and q is a continuous map $h\colon [0,1] \times [0,1] \to \mathcal{M}$ such that $h(0,\cdot) = p$, $h(1,\cdot) = q$, $h(\cdot,0) = p(0) = q(0)$, and $h(\cdot,1) = p(1) = q(1)$. A (free) *homotopy* betweeen two cycles γ and δ is a continuous map $h\colon [0,1] \times S^1$ such that $h(0,\cdot) = \gamma$ and $h(1,\cdot) = \delta$. Two paths or cycles are *homotopic* if there is a homotopy between them. For any two loops ℓ and ℓ' with the same basepoint, their concatenations $\ell \cdot \ell'$ and $\ell' \cdot \ell$ are freely homotopic as cycles but not necessarily homotopic as loops. A loop or cycle is *contractible* if it is homotopic to a constant loop or cycle.

Two cycles are *homologous* (with $\mathbb{Z}_2$ coefficients) if one can be continuously deformed into the other via a deformation that may include splitting cycles at self-intersection points, merging intersecting pairs of cycles, or adding or deleting contractible cycles. Thus, if two cycles are homotopic, they are also homologous, but the converse is not necessarily true. A cycle or loop is *null-homologous* if it is homologous to a constant loop. A simple cycle γ is null-homologous if and only if it is *separating*, that is, if $\mathcal{M} \setminus \gamma$ has two components. Every contractible simple cycle is separating (in fact, it bounds a disk), but not all simple cycles are contractible.

We say that a cycle *splits* a surface $\mathcal{M}$ if it is simple, non-contractible, and separating. Every surface admits a splitting cycle, except those homeomorphic to a sphere, a disk, or a torus.

Combinatorial and cross-metric surfaces.

Like most earlier works [3, 6, 7, 9, 10, 15], we state and prove our algorithmic results in the *combinatorial surface* model. A combinatorial surface is an abstract surface $\mathcal{M}$ together with a weighted undirected graph $G(\mathcal{M})$, embedded on $\mathcal{M}$ so that each open face is a disk. In this model, the only allowed paths are walks in G; the length of a path is the sum of the weights of the edges traversed by the path, counted with multiplicity. Paths on a combinatorial surface may visit a vertex or an edge several times without crossing. The *complexity* of a combinatorial surface is the total number of vertices, edges, and faces of G.

It is often more convenient to work in an equivalent dual formulation of this model introduced by Colin de Verdière and Erickson [5]. A *cross-metric surface* is also an abstract surface $\mathcal{M}$ together with an undirected weighted graph $G^* = G^*(\mathcal{M})$, embedded so that every open face is a disk. However, now we consider only *regular* paths and cycles on $\mathcal{M}$, which intersect the edges of G^* only transversely and away from the vertices. The *length* of a regular curve p is defined to be the sum of the weights of the dual edges that p *crosses*, counted with multiplicity. See [5] for further discussion of these two models.

We emphasize that most of our structural results (Lemma 4.1, Proposition 4.2, and Theorem 4.3) are *not* limited to the combinatorial surface model, but rather apply to surfaces with a wide range of metrics, including piecewise-linear, piecewise-algebraic, and abstract Riemannian surfaces. Even the NP-hardness reduction in Section 3 can be easily modified to apply in these more general surface models. The restriction to the combinatorial surface model is only necessary for our algorithmic results in Section 5, largely because this is the only known surface model in which exact shortest paths can be computed efficiently without additional assumptions.*

2.2 Preliminary Lemma

Let L be a set of simple, pairwise disjoint loops with common basepoint, all in the interior of $\mathcal{M}$. Let $\mathcal{M}'$ be a connected component of $\mathcal{M} \setminus L$. If $\mathcal{M}'$ is a disk that has one, two, or three copies of the basepoint on its boundary, it is called respectively a *monogon*, a *bigon*, or a *trigon*. If it is an annulus such that one of its boundary components has exactly one copy of the basepoint and the other is a boundary of the original surface $\mathcal{M}$, it is called an *elementary annulus*.

Lemma 2.1. *Let $\mathcal{M}$ be a surface with genus g and b boundaries. Let $L \neq \emptyset$ be a set of simple, pairwise disjoint loops with common basepoint, all in the interior of $\mathcal{M}$. Assume that no component of $\mathcal{M} \setminus L$ is a monogon or a bigon. Then $|L| \leq 6g + 2b - 3$.*

Proof: We first prove that we can extend L to a set L' of simple, pairwise disjoint loops with the same basepoint x such that $\mathcal{M} \setminus L$ consists of trigons and elementary annuli only. Initially, let $L' = L$.

Let $\mathcal{M}'$ be any component of $\mathcal{M} \setminus L$; it has at least one copy of x on at least one of its boundaries. If $\mathcal{M}'$ contains a boundary component of $\mathcal{M}$, we add to L' a loop in $\mathcal{M}'$ that encloses this boundary component, splitting $\mathcal{M}'$ into an elementary annulus and another surface. Now we consider any component $\mathcal{M}'$ of $\mathcal{M} \setminus L'$ that is not an elementary annulus; such a component contains no boundary component of the original surface $\mathcal{M}$.

If $\mathcal{M}'$ has genus at least one, we add to L' a non-separating loop in $\mathcal{M}'$, based at some copy of x on its boundary. This does not create monogons or bigons and decreases the genus of $\mathcal{M}'$ by one. Iterating the procedure, we may assume that $\mathcal{M}'$ has genus zero. If $\mathcal{M}'$ has at least two boundaries (and is not an elementary annulus), then each boundary contains a copy of the basepoint x, and we can add a loop in $\mathcal{M}'$ that connects two copies of x on different boundaries. This does not separate $\mathcal{M}'$, creates no monogon or bigon, and decreases its number of boundary components by one. So we may assume that $\mathcal{M}'$ is a disk.

*Efficient shortest-path algorithms for piecewise-linear surfaces [4, 16] require exact real arithmetic. Even if the input coordinates are integers, shortest path lengths are sums of square roots of integers; it is an open question whether two such sums can be compared in polynomial time on an integer RAM [1]. The analysis of these algorithms also assumes that the shortest path between any two points in the same face of the surface is contained in that face. Although this condition is satisfied by (possibly non-convex) polyhedra in any Euclidean space, it is not true for arbitrary abstract PL surfaces.

For some $k \geq 3$, the boundary of the disk $\mathcal{M}'$ contains k copies of x. Such a disk may be triangulated by adding $k-2$ loops. So we obtain a set $L' \supseteq L$ of simple, pairwise disjoint loops with common basepoint that split $\mathcal{M}$ into b elementary annuli and t trigons, for some integer t.

Finally, counting the edge-face incidences in two different ways, we obtain $2|L'| = 3t + b$. Euler's formula, applied to the surface $\mathcal{M}$ with each boundary component filled with a disk, implies $2 - 2g = 1 - |L'| + t + b$. We conclude that $|L'| = 6g + 2b - 3$. Since $|L| \leq |L'|$, the result holds. □

2.3 Finding a Splitting Cycle

If length is not a consideration, we can compute a splitting cycle on a surface $\mathcal{M}$ in $O(n)$ time as follows.

If $\mathcal{M}$ is a sphere, a disk, or a torus, then no splitting cycle exists. Otherwise, if $\mathcal{M}$ has at least one boundary component, then any cycle enclosing one boundary component is splitting. So we may assume that $\mathcal{M}$ has genus at least two and has no boundary.

First we construct a simple non-contractible cycle, using a variant of an algorithm of Erickson and Har-Peled [9], which finds the shortest non-contractible loop with a given basepoint in $O(n \log n)$ time. The running time is dominated by the computation of a shortest-path tree rooted at x using Dijkstra's algorithm. If we ignore the edge lengths and use breadth-first search instead, the running time drops to $O(n)$; the modified algorithm still returns a simple non-contractible cycle α. If α is separating, then we are done. Otherwise, we compute another non-contractible cycle that crosses α exactly once, in $O(n)$ time. To do this, we choose an arbitrary vertex $x \in \alpha$ and compute a simple path β from one copy of x to the other copy in $\mathcal{M} \setminus \alpha$, using (for example) depth-first search. The cycle $\gamma = \alpha \cdot \beta \cdot \bar{\alpha} \cdot \bar{\beta}$ is simple and null-homologous, but not contractible; specifically, one of the components of $\mathcal{M} \setminus \gamma$ is a punctured torus. Thus γ is a splitting cycle.

3 NP-hardness

Theorem 3.1. *Finding the shortest splitting cycle on a combinatorial surface is NP-hard.*

Proof: A *grid graph of size* n is a graph induced by a set of n points on the two-dimensional integer grid. We describe a two-step reduction from the Hamilton cycle problem in grid graphs [13].

Let H be a grid graph of size n. To begin the first reduction, we overlay n many 4×4 square grids of width $\epsilon < 1/4n$, one centered on each vertex of H. In each small grid, we color the square in the second row and second column *red* and the square in the third row and third column *blue* (we fix the origin at the upper left corner). We now easily observe that the following question is NP-complete: *Does the modified grid contain a cycle of length at most* $n+1/2$ *that separates the red squares from the blue squares?* Any Hamilton cycle for H can be modified to produce a separating cycle of length at most $n + 1/2$ by locally modifying the Hamilton cycle within each small grid, as shown in Figure 1, left and middle. Conversely, any separating cycle must pass through the center points of all n small grids, which implies that any separating cycle of length at most $n + 1/2$ must contain n grid edges that comprise a Hamilton cycle for H. We note that this result holds also if the cycle is allowed to visit vertices and edges of the modified grid several times, as in the case of combinatorial surfaces.

In the second reduction, we reduce the problem to finding a minimum-length splitting cycle. We isometrically embed the modified grid on a sphere, which we call *Earth*. We remove the red and blue squares to create $2n$ punctures, which we attach to two new punctured spheres, called *heaven* and *hell*. We attach the n punctures in heaven to the n blue punctures on Earth; similarly, we attach the n punctures in hell to the n red punctures on Earth. We append edges of length $n+1$ to the resulting surface so that each face of the final embedded graph is a disk. The resulting

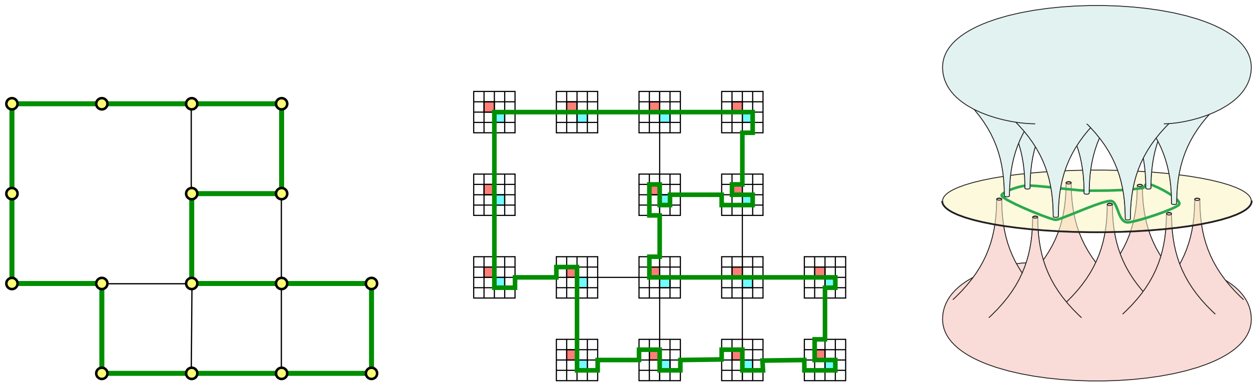

Figure 1. Left: A Hamilton cycle of length n. Middle: The corresponding red/blue separating cycle (not to scale). Right: Separating heaven from hell (not to scale); the central disk is a small portion of Earth.

combinatorial surface $\mathcal{M}(H)$ has no boundary, genus $2n-2$, and complexity $O(n)$; it can clearly be constructed in polynomial time. See Figure 1, right.

If the shortest cycle γ that splits $\mathcal{M}(H)$ has length less than $n+1/2$, then it must lie entirely on Earth. Moreover, γ must separate the blue punctures from the red punctures; otherwise, $\mathcal{M}(H)\setminus\gamma$ would be connected by a path through heaven or through hell. Thus, γ is precisely the shortest cycle that separates the red and blue squares in our intermediate problem. Testing whether γ has length less than $n+1/2$ is thus NP-complete from the first reduction. □

With a few trivial modifications, our reduction also implies that computing the shortest splitting cycle on a polyhedral or Riemannian surface is NP-hard, although it is an open question whether the corresponding decision problems are still in NP.

4 Structural Properties

4.1 Multiplicity Bound

For any two points x and y on a cycle α, we let $\alpha[x,y]$ denote the path from x to y along α, taking into account the orientation of α. For a path or a dual edge α, the same notation is used for the unique simple path between x and y on α.

Lemma 4.1. *Any shortest splitting cycle on an orientable cross-metric surface $\mathcal{M}$ crosses each edge e^* of $G^*(\mathcal{M})$ at most once in each direction.*

Proof: Assume for the purpose of contradiction that some shortest splitting cycle γ crosses some dual edge e^* twice in the same direction, say left to right. Let x and z be consecutive left-to-right intersection points along e^*; that is, γ does not cross $e^*[x,z]$ from left to right. We claim that γ must cross e^* (exactly once) from right to left at some point y between x and z. Indeed, because γ is separating and $\mathcal{M}$ is orientable, the orientation of the crossings of e^* with γ must alternate along e^*.

The cycles $\gamma[x,y]\cdot e^*[y,x]$ and $\gamma[y,z]\cdot e^*[z,y]$ are non-contractible; otherwise, we could shorten γ by removing two crossings with e^* without changing its homotopy class.

Define a new cycle $\gamma' = \gamma[x,y]\cdot e^*[y,z]\cdot\gamma[z,x]\cdot e^*[x,y]\cdot\gamma[y,z]\cdot e^*[z,x]$, as shown in Figure 2. The new cycle γ' is simple, because x, y, z are consecutive along e^*. The cycle γ' is in the same homology class as γ and is therefore null-homologous. By translating the non-contractible cycles $\gamma[x,y]\cdot e^*[y,x]$ and $\gamma[y,z]\cdot e^*[z,y]$ away from γ', we obtain two non-contractible cycles a' and b', one in each component of $\mathcal{M}\setminus\gamma'$; so γ' is non-contractible. Finally, γ' crosses e^* two times fewer

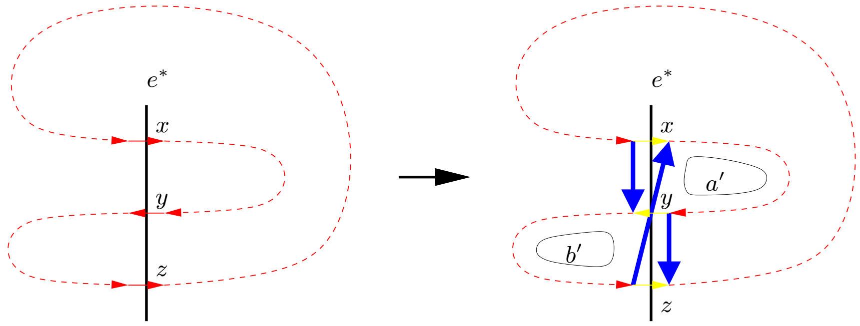

Figure 2. Lemma 4.1. If γ crosses e^* twice in the same direction, we can remove two crossings.

than γ and crosses every other edge in $G^*(\mathcal{M})$ the same number of times as γ. We conclude that γ' is a splitting cycle that is shorter than γ, which is impossible. This concludes the proof.

As a side note, let $\mathcal{M}_\ell$ and $\mathcal{M}_r$ be the components of $\mathcal{M} \setminus \gamma$ on the left and on the right of γ; define $\mathcal{M}'_\ell$ and $\mathcal{M}'_r$ similarly to be the components of $\mathcal{M} \setminus \gamma'$ on the left and on the right of γ'. We remark that the topology of $\mathcal{M}_\ell$ and $\mathcal{M}'_\ell$ is the same, and similarly for $\mathcal{M}_r$ and $\mathcal{M}'_r$: topologically, $\mathcal{M}'_\ell$ is obtained from $\mathcal{M}_\ell$ by splitting it into two pieces along an arc, and by regluing these two pieces along a subpath of their boundaries. A more formal way to see this result is to note that the Euler characteristic and the number of boundaries of $\mathcal{M}_\ell$ and $\mathcal{M}'_\ell$ are the same. $\square$

4.2 Shortest-Path Crossing Bound

Proposition 4.2. *Let $\mathcal{M}$ be an orientable cross-metric surface with genus g and b boundary components. Let P be a set of shortest paths on $\mathcal{M}$ such that the intersection of any two shortest paths is a (possibly empty) set of common endpoints. There is a shortest splitting cycle that crosses each path in P at most $12g + 4b - 6$ times.*

Proof: Amongst all shortest splitting cycles, let γ have the minimum number of crossings with paths in P. We can assume that γ does not pass through the endpoints of any path in P, since we are on a cross-metric surface and could simply perturb γ slightly. Consider any path p in P that intersects γ.

The intersection points $\gamma \cap p$ partition γ into *arcs*. These arcs may intersect other paths in P. Let $\mathcal{M}/p$ be the quotient surface obtained by contracting p to a point p/p. The set of all arcs corresponds in $\mathcal{M}/p$ to a set of simple, pairwise disjoint loops L with basepoint p/p.

We claim that no component of $(\mathcal{M}/p) \setminus L$ can be a monogon. Otherwise, there would be two intersection points x and y of γ and p such that $\gamma[x, y]$ and $p[y, x]$ bound a disk. But then we could obtain a splitting cycle $\gamma[y, x] \cdot p[x, y]$, no longer than γ, that has fewer crossings with the paths in P, a contradiction.

We prove below that each loop in L is, in $\mathcal{M}/p$, incident to at most one bigon of $(\mathcal{M}/p) \setminus L$. Now, whenever one component of $(\mathcal{M}/p) \setminus L$ is a bigon, remove one of the two incident loops, and iterate until there is no bigon. The previous remark implies that at most half of the loops have been removed; so, if L' denotes the new set of loops, we have $|L| \leq 2|L'|$. Furthermore, no component of $(\mathcal{M}/p) \setminus L'$ is a monogon or a bigon, hence, by Lemma 2.1, we have $|L'| \leq 6g + 2b - 3$. So $|L| \leq 12g + 4b - 6$, as claimed.

To complete the proof, it thus suffices to prove that no loop in L bounds in $\mathcal{M}/p$ two disks of $(\mathcal{M}/p) \setminus L$, themselves bounded by two loops. We assume the contrary: there are three arcs $u = \gamma[a, z]$, $v = \gamma[y, b]$, and $w = \gamma[c, x]$ such that u/p and v/p comprise the boundary of a component of $\mathcal{M}/p$ that is a disk, and similarly for v/p and w/p; see Figure 3, left. Since γ is separating, for a fixed orientation of $\mathcal{M}$, the clockwise boundaries of these disks are $u/p \cdot v/p$ and $\bar{v}/p \cdot \bar{w}/p$. Gluing these disks along v/p, we see that there is an open disk D_p in $\mathcal{M}/p$ bounded by $u/p \cdot \bar{w}/p$ such that v/p is the only loop in L intersecting this disk.

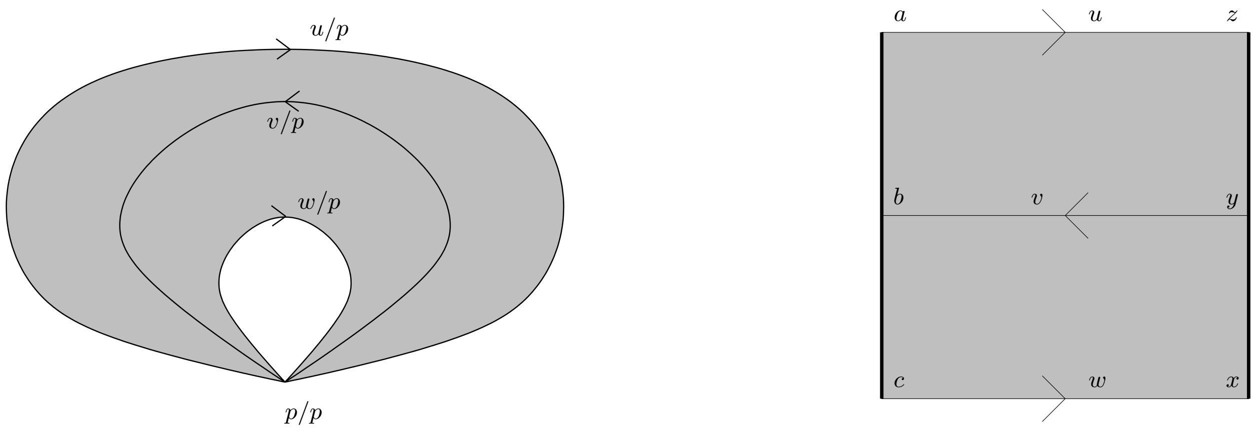

Figure 3. Left: Loop v/p is incident with two bigons. D_p is shaded. Right: A drawing obtained by expanding back the basepoint p/p. The two bold lines correspond to subsegments of p'. These two subsegments can overlap in the sense that a point on ac and a point on zx may correspond to the same point of p'.

Consider the minimal subpath p' of p that includes the endpoints a, c, z, x of u and w. In particular, the surfaces $(\mathcal{M}/p) \setminus (u \cup w)$ and $(\mathcal{M}/p') \setminus (u \cup w)$ are isomorphic. By our previous assumption, one component of $\mathcal{M} \setminus (p' \cup u \cup w)$ is an open disk D such that the portion of γ inside D is precisely arc v.

The boundary of D is a walk in the embedded graph $p' \cup u \cup w$ and is composed of u, $\bar{w}$, and pieces of p'. Since the embedded graph $p' \cup u \cup w$ has no vertex of degree one, this boundary must in fact be $u \cdot p'[z, x] \cdot \bar{w} \cdot p'[c, a]$. See Figure 3, right.

We claim that p cannot enter the open disk D. Otherwise, a component q of $p \setminus p'$ would be entirely contained in D. Call s the common endpoint of q and p', where $s \in \{a, c, z, x\}$. By symmetry, we may assume $s = a$. Then, since w does not cross q and because the crossings are transverse, the arc u' preceding u along γ must be in D, whence $u' = v$. See Figure 4. It follows that $p[y, z]$ and $v \cdot u = \gamma[y, z]$ bound a disk. But γ does not cross $p[y, z]$, because otherwise there would be, in D, an arc different from v. We can thus replace the part $\gamma[y, z]$ of γ by $p[y, z]$; this does not increase its length, removes at least three crossings with the paths in P, and the result is a simple cycle homotopic to γ; this contradicts the choice of γ.

From the previous claim, it follows that D is also a component of $\mathcal{M} \setminus (p \cup u \cup w)$ with boundary $u \cdot p[z, x] \cdot \bar{w} \cdot p[c, a]$. There are two cases to consider depending on the relative directions of $p[z, x]$ and $p[c, a]$ along p.

If $p[z, x]$ and $p[c, a]$ are directed the same way along p, then $p[z, x]$ and $p[c, a]$ cannot strictly overlap since otherwise $\mathcal{M}$ would contain a Möbius strip while $\mathcal{M}$ is orientable. Up to a change of direction of γ and to an exchange between u and w, we are in the situation depicted on top Figure 5 (note that by simplicity of γ, $c \neq x$).

Suppose on the contrary that $p[z, x]$ and $p[c, a]$ have opposite directions along p. Up to a change of direction of γ and to an exchange between u and w, we can assume a arises first along p. We now claim that $p[z, x]$ and $p[c, a]$ cannot stricly overlap. Otherwise z would occur strictly between a and c, so that the arc u' following u along γ would enter D at z, whence $u' = v$. It follows that $z = y$ on $\mathcal{M}$ (refer to Figure 3, right). Again, since the crossings between γ and p are transverse, this implies that p enters D, contradicting a previous claim.

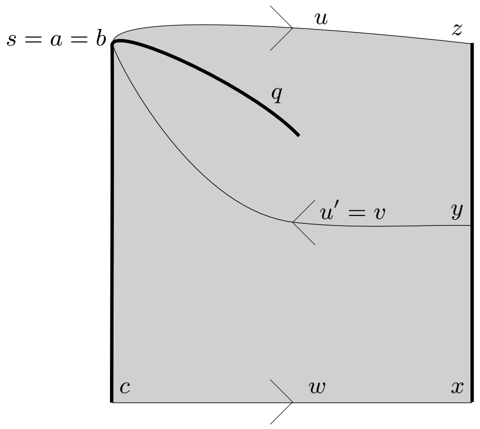

Figure 4. The two bold lines correspond to subsegments of p. Points s, y, and z must be pairwise distinct on $\mathcal{M}$ since $gamma$ is simple.

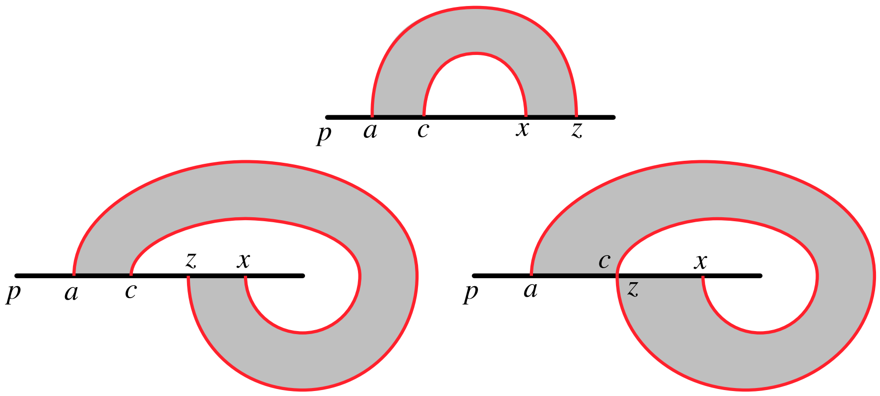

Figure 5. The three possible configurations for arcs $u = \gamma[a, z]$ and $w = \gamma[c, x]$; disk D is shaded.

We conclude the proof by showing that none of the top, bottom left or bottom right configurations on Figure 5 can actually happen.

Top case: Recall that γ intersects D along v.

Without loss of generality, suppose the path $\gamma[x, a]$ does not contain any of the arcs u, v, or w. Consider the cycle

$$\gamma' = p[a, b] \cdot \gamma[b, c] \cdot p[c, b] \cdot \bar{\gamma}[b, y] \cdot p[y, z] \cdot \gamma[z, y] \cdot p[y, x] \cdot \gamma[x, a]$$

obtained by removing u and w from γ, reversing v, and connecting the remaining pieces of γ with subpaths of p; see Figure 6. This cycle crosses p fewer times than γ, and crosses any other path in p no more than γ. An argument identical to the proof of Lemma 4.1 implies that γ' is simple, null-homologous, and non-contractible. Because p is a shortest path, u cannot be shorter than $p[a, z]$, which implies that γ' is no longer than γ, contradicting our assumption that γ is a shortest splitting cycle with the minimal number of crossings with paths in P.

Bottom left and right cases: As in the previous case, γ intersects D along v.

If $c = z$, there is only one way to connect these three arcs to form the cycle γ; if $c \neq z$, there are two possibilities. See Figure 7. In all three cases, by deleting arcs u and w, reversing arc v, and connecting the remaining pieces of γ with subpaths of p, we create a splitting cycle γ' that is no longer than γ and crosses fewer times the paths in P than γ, which is impossible. We omit the tedious details.

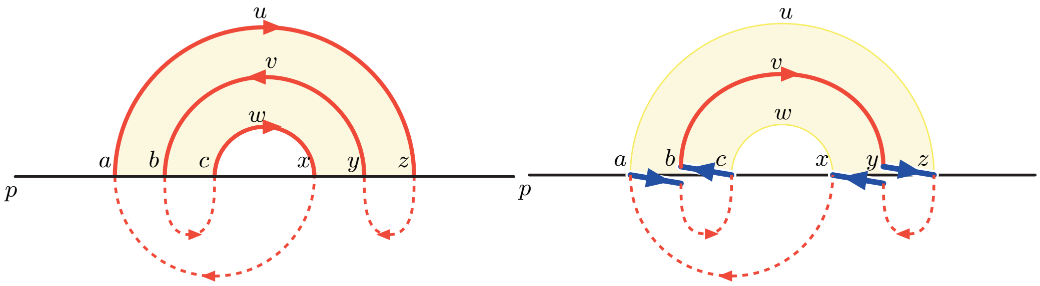

Figure 6. The exchange argument for the top configuration.

We note that, as in the proof of Lemma 4.1, the topology of the surface that is the part of $\mathcal{M}$ on either side of γ does not change during these exchanges: in all cases, topologically non-trivial components of this surface are cut along paths with endpoints on the boundary and reglued differently. $\square$

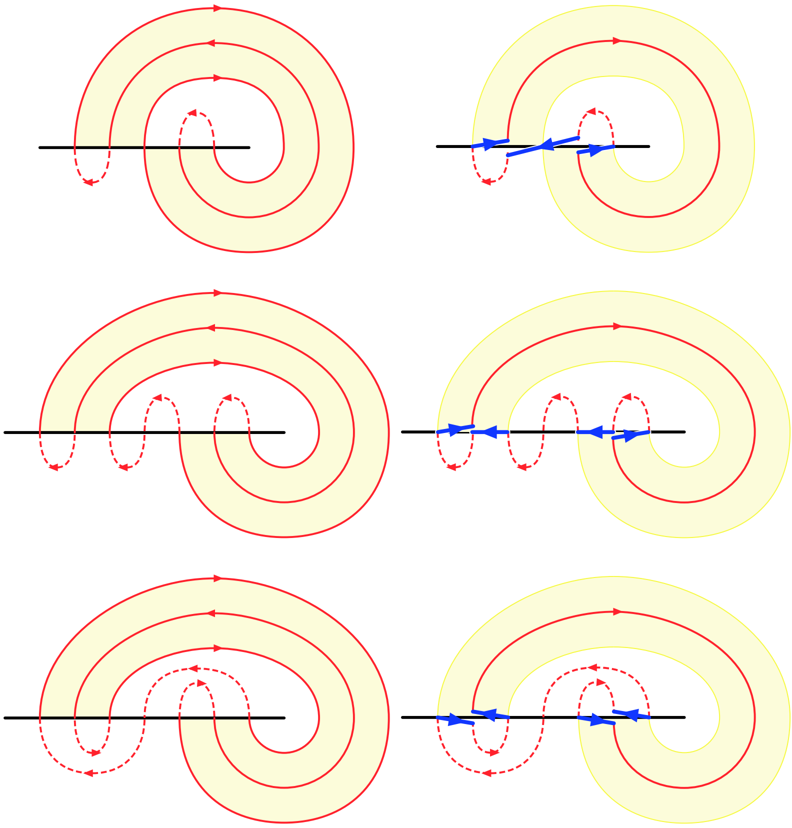

Figure 7. Three exchange arguments for the bottom configurations.

4.3 Crossing Lower Bound

Earlier algorithms for computing shortest non-contractible, non-separating, or essential cycles rely on the fact that the desired cycle is the concatenation of two equal-length shortest paths [17, 9].[†]

[†]In general, this characterization requires shortest paths that terminate in the interior of edges, but if we refine edges appropriately, the shortest cycle will indeed be the concatenation of two shortest vertex-to-vertex paths.

Cabello and Mohar [3], Cabello [2], and Kutz [14] exploit a slightly different property to compute shortest nontrivial cycles more quickly on surfaces of constant genus: The desired shortest cycle crosses any shortest path at most twice. As we prove next, neither of these properties holds for the shortest splitting cycle; in particular, the upper bound of Proposition 4.2 is tight up to constant factors for surfaces without boundary.

Theorem 4.3. *For any $g \geq 2$, there is an orientable combinatorial surface $\mathcal{M}_g$ of genus g, without boundary, whose unique shortest splitting cycle (up to orientation) crosses a shortest path g times and (therefore) cannot be decomposed into fewer than g shortest paths and edges.*

Proof: We consider only the case where g is even. We construct a combinatorial surface $\mathcal{M}_g$ of genus g such that the shortest non-contractible cycle γ and the shortest splitting cycle μ cross g times; see Figure 8. The shortest non-contractible cycle γ can be partitioned into two shortest paths, one of which will be disjoint from μ; it follows that μ crosses some shortest path g times.

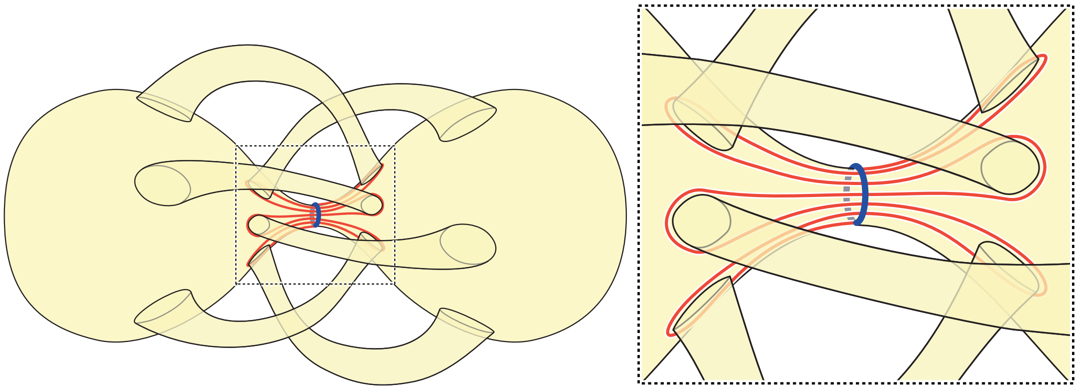

Figure 8. A surface whose shortest splitting cycle cuts a shortest path g times, and a closeup of the undulating shortest splitting cycle.

The base surface $\mathcal{M}_0$ is a sphere whose geometry approximates an hourglass. Let γ be the central cycle that partitions the two lobes of the hourglass. To construct $\mathcal{M}_g$, we attach g handles to this hourglass, each joining a small circle c_i on the neck of the hourglass, just to one side of γ, to a large circle C_i far away on the opposite lobe. The small circles c_i are arranged symmetrically around the neck of the hourglass, alternating between the two sides of γ; the large circles C_i are also partitioned evenly between the two lobes of the hourglass.

Let μ be a cycle that undulates around the small circles in order, crossing γ a total of g times, as shown in Figure 8. We easily verify that μ splits $\mathcal{M}_g$ into two surfaces of genus $g/2$. Let $x_1, x_2, \ldots, x_g = x_0$ denote the g intersection points of μ and γ. For each i, let μ_i and γ_i respectively denote the subpaths of μ and γ between x_{i-1} and x_i. Finally, we partition path γ_1 into three subpaths $\gamma_1^\flat$, $\gamma_1^\natural$, and $\gamma_1^\sharp$ at arbitrary points y and z. Let F denote the union of these $2g+2$ paths.

To obtain a combinatorial structure on $\mathcal{M}_g$, we embed a weighted graph G that contains F as a subgraph, such that every face of the embedding is a topological disk. We assign length 1 to edge γ_i for each $i \neq 1$, length 1 to edges $\gamma_1^\flat$ and $\gamma_1^\sharp$, length $g+1$ to edge $\gamma_1^\natural$, length $4g^2$ to each edge μ_i, and length at least $10g^5$ to every other edge in G. The cycle μ does not contain a single (even approximate!) vertex-to-vertex shortest path. Even if we allow shortest paths between points in the interior of edges, each such path contains at most one vertex x_i. The cycle γ can clearly be partitioned into two shortest paths of length $g+1$ at points y and z, and μ crosses one of these paths g times. Thus, to complete the proof, it suffices to show that μ is in fact the shortest splitting cycle in $\mathcal{M}_g$.

Let α be a shortest splitting cycle; the assigned edge weights guarantee that α is a cycle in F. By Lemma 4.1, α traverses each path γ_i and μ_i at most once in each direction. For any i, there is a path in $\mathcal{M}_g \setminus F$ from one side of γ_i to the other, so α must traverse each edge γ_i either twice (in opposite directions) or not at all.

Consider any simple cycle β in a tubular neighborhood $\mathcal{N}$ of F that traverses every edge in F either once in each direction or not at all. This cycle must be null-homologous, and therefore separating in $\mathcal{M}$ and $\mathcal{N}$. Furthermore, the boundary of $\mathcal{N}$ belongs entirely to some component of $\mathcal{M} \setminus \beta$. It follows that one component of $\mathcal{N} \setminus \beta$ is only bounded by β. This component must be a disk since $\mathcal{N}$ has genus zero. We conclude that β is contractible.

The splitting cycle α is not contractible, so it must traverse some edge μ_i exactly once. But α must traverse the edges adjacent to any vertex x_i an even number of times, which implies that α traverses *every* edge μ_i exactly once. Thus, every splitting cycle in F is at least as long as μ. We conclude that μ is indeed the unique shortest splitting cycle. □

5 Algorithm

In this section, we prove that computing the shortest splitting cycle is fixed-parameter tractable with respect to the genus and the number of boundary components of the surface.

Theorem 5.1. *Let $\mathcal{M}$ be an orientable cross-metric surface; let g be its genus, b be its number of boundary components, and n be its complexity. We can compute a shortest splitting cycle in $\mathcal{M}$ in $(g+b)^{O(g+b)} n \log n$ time.*

The algorithm proceeds in several stages, described in detail in the following subsections. First, we compute a set of $O(g + b)$ loops or arcs that cut the surface $\mathcal{M}$ into a disk, using (a variant of) the greedy algorithm of Erickson and Whittlesey [10]; each loop or arc is the concatenation of two shortest paths. Next, we enumerate all possible sequences of crossings of this system of loops or arcs by a simple cycle that crosses each loop or arc $O(g + b)$ times. Proposition 4.2 implies that the shortest splitting cycle must have one of these crossing sequences. We discard any crossing sequence that does not correspond to a splitting cycle. For each valid crossing sequence, we compute a shortest cycle with that crossing sequence using the recent algorithm of Kutz [14]. The shortest of these cycles, γ, corresponds to the desired crossing sequence; finally, we post-process γ to remove any self-intersections, without changing its length or its free homotopy type.

5.1 Greedy System of Loops or Arcs

We first compute a set of interior-disjoint shortest paths that split $\mathcal{M}$ into a disk.

If $\mathcal{M}$ has no boundary, let v be any point of $\mathcal{M}$ in the interior of a face of $G^*(\mathcal{M})$. Let $\alpha_1, \alpha_2, \ldots, \alpha_{2g}$ be the shortest system of loops of $\mathcal{M}$ with basepoint v; this system of loops can be computed in $O(n \log n + gn)$ time using a greedy algorithm of Erickson and Whittlesey [10].

A key property of this greedy system of loops is that each loop α_i is composed of two shortest paths in the primal graph $G(\mathcal{M})$. However, in general these two paths meet at a point m_i in the interior of some edge e_i. To simplify our algorithm, we split e_i into two edges at m_i—or equivalently, in the dual graph, we replace the dual edge e_i^* with two parallel edges—partitioning the length appropriately, so that α_i consists of two vertex-to-vertex shortest paths β_i and β_i' in $G(\mathcal{M})$. Another property that will be used later is that each α_i is a shortest loop in its homotopy class.

If $\mathcal{M}$ has at least one boundary component, then an easy variant on the aforementioned algorithm by Erickson and Whittlesey [10] allows to compute a greedy *system of arcs*, which bears properties similar to the greedy system of loops: it is made of $O(g+b)$ arcs α_i; it can be computed in $O(n \log n + (g+b)n)$; each arc α_i is composed of two shortest paths β_i and β'_i; and each arc is as short as possible in its homotopy class. To compute the greedy system of loops with basepoint v, the idea is to maintain a set L, initially empty, of simple, pairwise disjoint loops. The algorithm iteratively adds to L the shortest loop based at v such that $\mathcal{M} \setminus L$ is connected, relying on a shortest paths tree rooted at the basepoint. To build the greedy system of arcs, the idea is similar: we iteratively add to the set of arcs L the shortest arc such that $\mathcal{M} \setminus L$ is connected. To implement this efficiently, it suffices to compute simultaneously shortest paths trees rooted at every boundary component.

5.2 Simple Crossing Sequences

The *crossing sequence* of a cycle γ records the intersections of γ with the greedy loops or arcs α_i, in cyclic order along γ. Any two cycles with the same crossing sequence are homotopic, although two homotopic cycles can have different crossing sequences. A crossing sequence is *simple* if it can be generated by a simple cycle; non-simple cycles can have simple crossing sequences.

Proposition 4.2 implies that some shortest splitting cycle γ crosses each path β_i or β'_i $O(g+b)$ times, and thus crosses each loop or arc α_i $O(g+b)$ times. Our algorithm therefore enumerates a superset of all simple crossing sequences that cross each loop or arc α_i $O(g+b)$ times. Note that there are $(g+b)^{\Theta((g+b)^2)}$ crossing sequences with $O(g+b)$ occurrences of each loop or arc, but the vast majority of these are not simple. Thus, we cannot naively enumerate crossing sequences that satisfy Proposition 4.2 and then check each sequence for simplicity. Our enumeration algorithm enforces simplicity from the beginning.

We begin by cutting $\mathcal{M}$ along the loops or arcs α_i to obtain a polygonal schema $\mathcal{D}$; this is a cross-metric disk with complexity $O((g+b)n)$. This cutting operation also cuts the unknown splitting cycle γ into *segments* that cut across $\mathcal{D}$. Because γ is simple, no two of these segments cross. If $b = 0$, because the basepoint v does not lie on $G^*(\mathcal{M})$, we can slightly perturb γ without changing its length (in the crossing metric) or its homotopy class; thus, we assume that γ does not pass through the basepoint v. If $b \geq 1$, because no endpoint of a segment can be on an edge of the polygonal schema corresponding to a piece of boundary of $\mathcal{M}$, we contract these $4g+2b-2$ edges, working with a *contracted polygonal schema* with $4g+2b-2$ edges instead of $8g+4b-4$.

The segments of γ can be grouped into subsets according to which pair of greedy loops they meet on the boundary of $\mathcal{D}$ (Figure 9). We abstract and dualize the (contracted) polygonal schema by replacing each edge of the (contracted) polygonal schema with a vertex and connecting vertices that correspond to consecutive edges. Now each subset of segments corresponds to a diagonal between two vertices of the dual $4g$-gon (if $b = 0$) or $(4g+2b-2)$-gon (if $b \neq 0$). Since no two segments cross, these diagonals cannot cross. In particular, all the diagonals belong to some triangulation of the dual polygon.

Thus the candidate crossing sequences of a shortest splitting cycle are described by *weighted triangulations*, which consist of a triangulation of the dual polygon, each of whose edges is weighted with an integer between 0 and $O(g+b)$. The label of an edge in the triangulation represents the number of times that the cycle runs along that edge. There are C_{k-2} possible triangulations of a k-gon, where $C_k = O(4^k)$ is the kth Catalan number, which we can enumerate in $O(k)$ time each. (This is essentially identical to the enumeration of binary trees.) Here $k = O(g+b)$. There are $(g+b)^{O(g+b)}$ ways to label each triangulation, which we can enumerate in constant amortized time per labeling.

We thus obtain a total of $(g+b)^{O(g+b)}$ weighted triangulations. Most of these do not correspond to a splitting cycle, or indeed to *any* cycle. We now explain how to discard these possibilities.

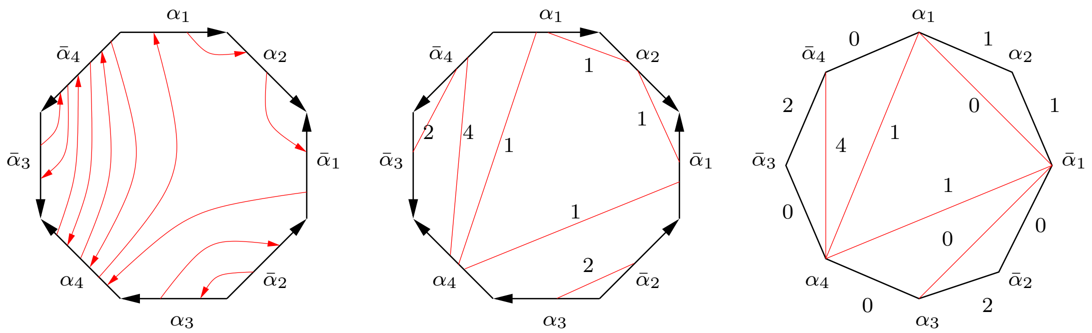

Figure 9. Left: A splitting cycle on a double-torus ($g = 2$, $b = 0$). Middle: The corresponding subsets of segments; each label indicates the number of segments contained in a subset. Right: The corresponding weighted triangulation.

5.3 Testing Weighted Triangulations

Given a candidate weighted triangulation T, we want to test whether T corresponds to a splitting cycle. We must check that (1) T corresponds to a set of cycles, (2) this set contains exactly one cycle, and (3) this cycle is separating but non-contractible. (We already know that the cycle is simple.) We describe how to perform these tests in $O((g+b)^2)$ time. If any of these properties is not satisfied, we simply discard T.

To check that T corresponds to a set of cycles, it suffices to check that the two edges of $\mathcal{D}$ that correspond to the same α_i are crossed the same number of times. (If $\mathcal{M}$ has a boundary, no edge of $\mathcal{D}$ corresponding to a part of the boundary of $\mathcal{M}$ can be crossed, since we worked in the contracted polygonal schema.)

For the remaining tests, we build a combinatorial surface $\mathcal{M}'$ homeomorphic to $\mathcal{M}$, whose graph G' is the arrangement of the greedy system of loops and the candidate cycle(s) defined by T. We cut the abstract (non-contracted) polygonal schema along the subsets of edges given by the triangulation and then identify corresponding subpaths on the boundary of the polygonal schema. If we have multiple edges running along the same edge of the triangulation, these define thin rectangular strips. The complexity of the resulting surface is $O(g+b)^2$. The $O(n)$ internal complexity of the original surface $\mathcal{M}$ is ignored.

Once we have constructed $\mathcal{M}'$, we can check whether T defines a single cycle γ by a simple depth-first search.

To test that γ is separating but non-contractible, we use a simplification of an algorithm of Erickson and Har-Peled [9]. To test separation, we perform a depth- or breadth-first search on the faces of $\mathcal{M}'$, starting from any initial face, but forbidding crossings of any segment of γ. The cycle γ is separating if and only if the search halts before visiting every face of the surface. (Alternately, γ is separating if and only if γ crosses each α_i the same number of times in both directions.) We also compute the Euler characteristic of the reachable portion of the surface during the search by counting vertices, edges, and faces. The cycle is non-contractible if and only if this Euler characteristic is neither 1 (a disk) nor $1 - 2g - b$ (the complement of a disk).

5.4 From Crossing Sequence to Cycle

For each valid weighted triangulation, we compute the shortest cycle with the corresponding crossing sequence in time $O((g+b)^3 n \log n)$ using the recent algorithm of Kutz [14]. For the sake of completeness, we sketch the algorithm here.

First we glue together a cycle of $O((g+b)^2)$ distinct copies of the polygonal schema $\mathcal{D}$, one per crossing in the sequence. Each successive pair of copies is glued along the edge specified by the corresponding entry in the crossing sequence. Because we consider only crossing sequences without *spurs*—the cycle does not cross a loop α_i and then immediately recross the same loop α_i in the opposite direction—the resulting combinatorial surface is an annulus, which we denote $\mathcal{D}^\circ$. (In Kutz's terms, we are considering only *curl-free* splitting cycles.) The polygonal schema $\mathcal{D}$ has complexity $O((g+b)n)$, so the complexity of $\mathcal{D}^\circ$ is $O((g+b)^3 n)$.

We then compute the shortest cycle γ° in $\mathcal{D}^\circ$ that is freely homotopic to the boundaries of $\mathcal{D}^\circ$, using an algorithm of Frederickson [11]; see also [5, Lemma 3.3(d)]. Given a combinatorial annulus of complexity N, Frederickson's algorithm finds a shortest generating cycle in $O(N \log N)$ time. Thus, we compute γ° in $O((g+b)^3 n \log((g+b)^3 n)) = O((g+b)^3 n \log n)$ time.

Finally, projecting γ° back to the original surface $\mathcal{M}$ gives us a shortest cycle γ with the given crossing sequence.

Since there are $(g+b)^{O(g+b)}$ valid crossing sequences, the total time spent in this phase of our algorithm is $(g+b)^{O(g+b)} O((g+b)^3 n \log n) = (g+b)^{O(g+b)} n \log n$.

5.5 Removing Self-intersections

Let γ be the shortest cycle computed in the previous phase, over all possible valid weighted triangulations. This cycle is null-homologous, non-contractible, as short as possible in its homotopy class, and freely homotopic to a simple cycle, but *not* necessarily simple. Although the cycle γ° is simple, projecting it back to the original surface may introduce self-intersections. We will prove that some simple cycle γ' is homotopic to γ and has the same length as γ (this also follows from [6]); hence, γ has the correct homotopy class of a shortest splitting cycle. We then need to compute such a γ'.

We consider the polygonal schema defined by the greedy loops or arcs. A *segment* is a path that intersects the polygonal schema precisely at its endpoints. A set of segments S *respects* a weighted triangulation T if, for each pair of edges of the polygonal schema, the number of segments of S between them equals the weight on the corresponding edge of T (or zero if there is no such edge). We need the following lemma.

Lemma 5.2. *Let T be a weighted triangulation and let S be a set of segments that respects T. There is a set of simple, pairwise disjoint segments of total length no larger than the total length of S that respects T and intersects the greedy loops or arcs at the same points as S.*

Proof: We prove the result by induction on the number of crossings in S. If S has no crossings, there is nothing to show. Otherwise, we will transform S into another set of segments that has one or two fewer crossings, still respects T, and intersects the greedy loops or arcs at the same points.

Consider a crossing in S in the polygonal schema. If this corresponds to a self-crossing of a single segment, we may remove the loop based at this crossing point and conclude by the induction hypothesis.

Otherwise, the crossing corresponds to two segments s and s'. If these segments cross at least twice, they form a bigon (two simple interior-disjoint paths with the same endpoints) inside the polygonal schema, which can be flipped to obtain a set of segments respecting T that is no longer and has two crossings fewer.

Otherwise, s and s' cross exactly once, at some crossing point p; they must thus correspond to (possibly identical) edges e and e' of the triangulation. In a triangulation, no two edges cross; if e and e' do not share any endpoint, then s and s' would cross an even number of times, which is not the case; so e and e' must share at least one endpoint, which is also an edge E of the polygonal schema. Hence swapping the parts of s and s' that are between E and p also results in a set of segments respecting T; this operation decreases the number of crossings by one and does not increase the length. □

Recall that γ is the cycle computed in the previous section. Let T be the weighted triangulation corresponding to γ; it was also computed above.

We apply the following algorithm. Consider the intersections of the cycle γ with the greedy loops or arcs. There is exactly one way to connect these points with simple, disjoint segments respecting T. We iteratively connect these pairs of points by shortest segments, forbidding, at each step, any crossing with the previously computed segments. By a simple exchange argument, the segments created are shortest paths in the polygonal schema; so this set of segments is a shortest set of segments, among all sets of simple, pairwise disjoint segments respecting T and intersecting the polygonal schema at the same points as γ. By Lemma 5.2, this set of segments is no longer than the set of segments of γ. Since it is a set of simple, pairwise disjoint segments respecting the valid triangulation T, it forms a *single* cycle γ', which has the same crossing sequence as γ by construction of γ. So γ' is the desired splitting cycle.

There are $O((g+b)^2)$ shortest paths to compute, in a planar graph of complexity $O((g+b)n)$. The cost of this uncrossing step is thus $O((g+b)^3 n)$. We note that, using a result by Takahashi, Suzuki, and Nishizeki [19, Section 3], we may improve the algorithm to obtain a time complexity of $O((g+b)n\log(g+b))$.

This concludes the proof of Theorem 5.1.

6 Splitting Surfaces into two Surfaces of Prescribed Topology

Let $\mathcal{M}$ be a surface with genus g and with b boundary components. A splitting cycle splits $\mathcal{M}$ into two surfaces $\mathcal{M}_1$ and $\mathcal{M}_2$, with respective genus g_1 and g_2 and with respective number of boundary components $b_1 \geq 1$ and $b_2 \geq 1$, where $g_1 + g_2 = g$ and $b_1 + b_2 = b + 2$. The techniques developed in the previous section allow to compute the shortest splitting cycle that splits the surface into two surfaces of prescribed genus and number of boundaries.

More precisely, let $S = \{(g_1, b_1), \ldots, (g_k, b_k)\}$ be a set of ordered pairs of integers. We say that a splitting cycle on $\mathcal{M}$ is *allowed* by S if it splits $\mathcal{M}$ into two surfaces, one of which has genus g' and b' boundary components, for some $(g', b') \in S$. In particular, if $S = (\{0, \ldots, g\} \times \{1, \ldots, b+1\}) \setminus \{(0,1), (0,2), (g,b), (g,b+1)\}$, the splitting cycles allowed by S are precisely the essential splitting cycles (the splitting cycles that do not bound an annulus).

Theorem 6.1. *Let $\mathcal{M}$ be an orientable cross-metric surface with genus g and b boundary components; let S be a set of ordered pairs. We can compute a shortest splitting cycle allowed by S in $(g+b)^{O(g+b)} n \log n$ time.*

Proof: If S contains $(0,0)$ or $(g, b+1)$, then a contractible cycle of length zero is a shortest splitting cycle allowed by S. Otherwise, the algorithm described in the previous sections, with very minor modifications, also works. Lemma 4.1 and Proposition 4.2 extend verbatim to the case where we are looking for a shortest splitting cycle allowed by S, because their proofs are based on exchange arguments that change the splitting cycle without modifying the topology of the surface on either side of the cycle. The algorithm has to be modified in Section 5.3 because we now have to check that a given weighted triangulation corresponds to a splitting cycle that is allowed by S; this is straightforward and does not increase the running-time of the algorithm. □

7 Decomposition into Punctured Tori

A *decomposition into punctured tori* of an orientable surface $\mathcal{M}$ with genus g and no boundary is a set of $g-1$ pairwise disjoint, pairwise non-homotopic, splitting cycles. Equivalently, it is a set of simple, pairwise disjoint cycles that split $\mathcal{M}$ into g punctured tori.

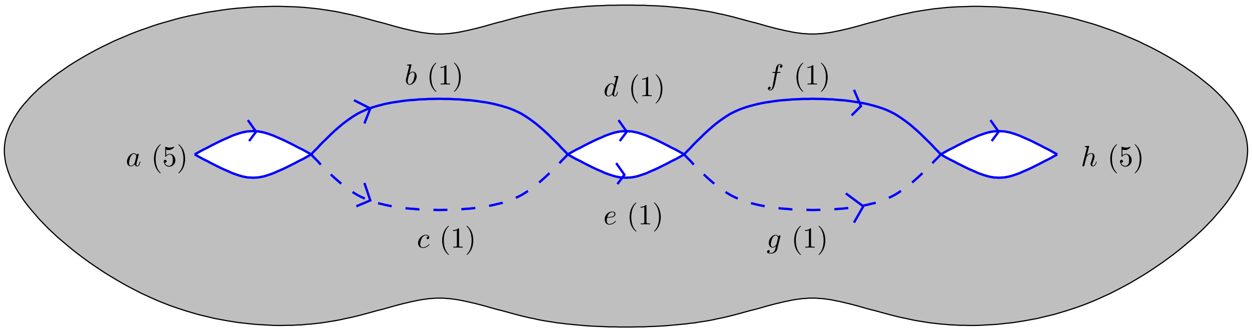

Figure 10. The graph G embedded on a triple-torus $\mathcal{M}$ in the proof of Proposition 7.1. The weights of the edges of G are indicated in parentheses.

There is a naïve greedy algorithm to compute a decomposition into punctured tori: compute the shortest splitting cycle of $\mathcal{M}$, cut along it, fill the boundaries obtained with a disk, and recurse in each of the two resulting surfaces.

Proposition 7.1. *The greedy algorithm does not necessarily provide the shortest decomposition into punctured tori.*

Proof: We use the graph G, embedded in a triple-torus, shown on Figure 10. The graph is extended, with edges of very large weight, to obtain a combinatorial surface $\mathcal{M}$. Any decomposition into punctured tori of $\mathcal{M}$ consists of precisely two cycles.

There is a decomposition into punctured tori of $\mathcal{M}$ of length 28, see Figure 11: one cycle is $\bar{b}ab\bar{c}\bar{a}c$, the other one is $fh\bar{f}g\bar{h}\bar{g}$. It is actually easy to prove that it is the shortest decomposition into punctured tori, but we won't need that. We will prove that the greedy algorithm gives the decomposition into punctured tori depicted in Figure 12, of length 30.

G can be decomposed into five edge-disjoint cycles that are homologically independent in $\mathcal{M}$ and form a cycle basis of G. It follows that the shortest splitting cycle uses two of these cycles at least twice. From this remark and from the assignment of the weights of G, we deduce that there are exactly two shortest splitting cycles, $bd\bar{e}\bar{b}ce\bar{d}\bar{c}$ and its symmetric, $\bar{f}\bar{d}ef\bar{g}\bar{e}dg$, on $\mathcal{M}$; they both have length 8. Without loss of generality, assume that the greedy algorithm chooses the first one, γ, as first cycle of the decomposition into punctured tori.

Now, cutting $\mathcal{M}$ along γ and filling its boundaries with disks yields two connected surfaces, one being a torus and the other one being a double-torus. The double-torus will be split with another splitting cycle; it is a combinatorial surface $\mathcal{M}'$, as shown on Figure 13, left.

Consider the surface $\mathcal{M}''$ on Figure 13, right. It is the same surface as $\mathcal{M}'$, except that the pairs of edges $b'd'$ and $b''e'$, and $c'd''$ and $c''e''$, have been merged and the leftmost endpoints of these edges have been identified. Any splitting cycle on $\mathcal{M}'$ gives a splitting cycle on $\mathcal{M}''$, of the same length. The graph on $\mathcal{M}''$ is made of three homologically independent cycles; arguing as above, the shortest splitting cycle on $\mathcal{M}''$ uses two cycles at least twice, and thus must have length at least 22. Hence the shortest splitting cycle on $\mathcal{M}'$ has length at least 22; but, as shown on Figure 12, there is such a cycle. Hence the greedy algorithm provides a decomposition into punctured tori of total length 30. $\square$

8 Conclusions

The results of this paper suggest several open problems. Most notably, can we approximate the shortest splitting cycle, or is that also NP-hard? The following high-level approach seems promising.

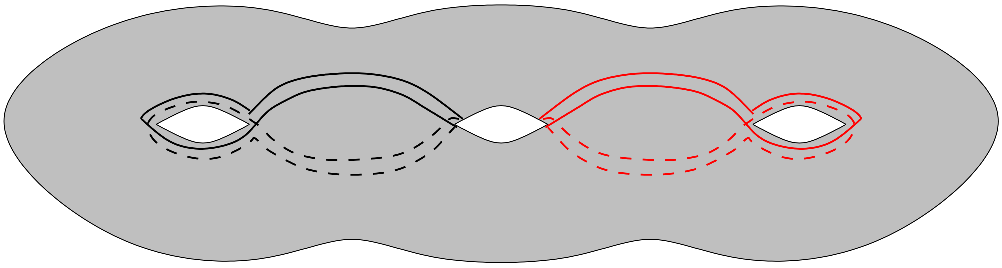

Figure 11. The shortest decomposition into punctured tori, of length 28.

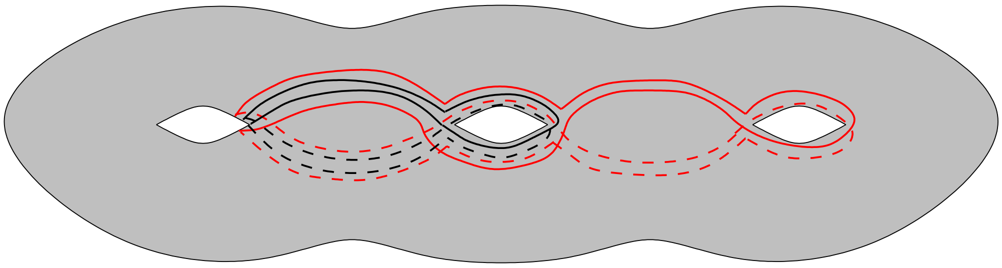

Figure 12. The greedy decomposition into punctured tori, of length 30. The first cycle has length 8, the second one has length 22.

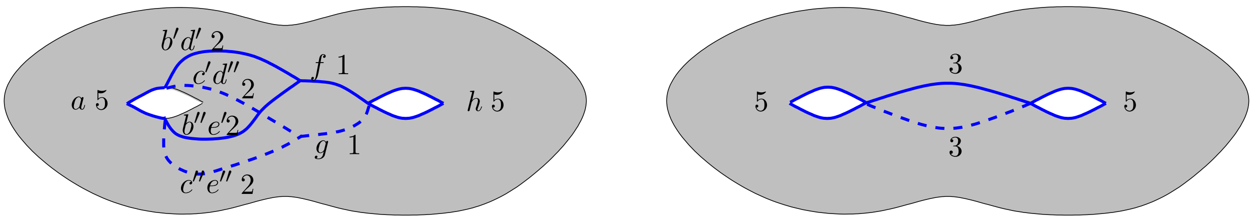

Figure 13. Left: the double-torus $\mathcal{M}'$ obtained after cutting $\mathcal{M}$ along the shortest splitting cycle and filling the boundary with a disk. Each edge b, c, d, and e on $\mathcal{M}$ corresponds to two edges on $\mathcal{M}'$. The weights of the edges are indicated. Right: the surface $\mathcal{M}''$.

Compute shortest simple cycles in each non-trivial homotopy class in order of increasing length, stopping either when we find a separating cycle, or when we find two cycles α and β that intersect an odd number of times. If we find a separating cycle, it is of course the shortest splitting cycle. If we find two cycles with odd intersection number, an exchange argument implies that the intersection number is 1. For some orientation of α and β, the cycle $\alpha \cdot \beta \cdot \bar{\alpha} \cdot \bar{\beta}$ is a splitting cycle whose length is at most four times the length of the shortest splitting cycle. Can this algorithm be implemented efficiently? How quickly can we enumerate the k shortest homotopy classes of (simple) cycles? Techniques of Eppstein [8] for enumerating k shortest paths may be useful here.

If we iterately cut along splitting cycles on the surface where each boundary is filled with a disk, we obtain a decomposition of the surface into punctured tori. However, we proved that repeatedly cutting the surface along its shortest splitting cycle does not necessarily yield the shortest such decomposition. Does this algorithm provide a constant-factor approximation for this problem? Is computing the shortest torus decomposition NP-hard? Similarly, a *pants decomposition* is a set of disjoint simple cycles decomposing a surface into pairs of pants, or spheres with three boundary components [6]. Is it NP-hard to compute the shortest pants decomposition? Are either of these problems fixed-parameter tractable?

Finally, Erickson and Har-Peled [9] prove that finding the minimum cut graph is NP-hard, by a reduction from the fixed-parameter tractable rectilinear Steiner tree problem. Is computing the shortest cut graph fixed-parameter tractable? The most serious bottleneck here seems to be computing the shortest cut graph in a given homotopy class, where the homotopy is allowed to move the vertices.

Acknowledgments. Portions of this work were done while all five authors were at the University of Illinois and later at École normale supérieure, thanks to the support of a UIUC/CNRS/INRIA travel grant. We thank the other participants of the 1st UIUC/CNRS/INRIA Workshop on Computational Topology for interesting discussions and meals. We also thank Therese Biedl, Sergio Cabello, Sariel Har-Peled, Martin Kutz, Kevin Milans, and Günter Rote for useful discussions while this work was in progress. We are especially grateful to Martin Kutz for sending us an advance copy of his paper [14].

References

[1] J. Blömer. Computing sums of radicals in polynomial time. *Proc. 32nd Annu. IEEE Sympos. Found. Comput. Sci.*, 670–677, 1991.

[2] S. Cabello. Many distances in planar graphs. *Proc. 17th Annu. ACM-SIAM Sympos. Discrete Algorithms*, 1213–1220, 2006.

[3] S. Cabello and B. Mohar. Finding shortest non-separating and non-contractible cycles for topologically embedded graphs. *Proc. 13th Annu. European Sympos. Algorithms*, 131–142, 2005. Lecture Notes in Computer Science 3669, Springer-Verlag.

[4] J. Chen and Y. Han. Shortest paths on a polyhedron, part I: Computing shortest paths. *Int. J. Comput. Geom. Appl.* 6(2):127–144, 1996.

[5] É. Colin de Verdière and J. Erickson. Tightening non-simple paths and cycles on surfaces. *Proc. 17th Annu. ACM-SIAM Sympos. Discrete Algorithms*, 192 – 201, 2006.

[6] É. Colin de Verdière and F. Lazarus. Optimal pants decompositions and shortest homotopic cycles on an orientable surface. *Proc. 11th Sympos. Graph Drawing*, 478–490, 2003. Lecture Notes Comput. Sci. 2912, Springer-Verlag.

[7] É. Colin de Verdière and F. Lazarus. Optimal systems of loops on an orientable surface. *Discrete Comput. Geom.* 33(3):507–534, 2005.

[8] D. Eppstein. Finding the k shortest paths. *SIAM J. Computing* 28(2):652–673, 1998.

[9] J. Erickson and S. Har-Peled. Optimally cutting a surface into a disk. *Discrete Comput. Geom.* 31(1):37–59, 2004.

[10] J. Erickson and K. Whittlesey. Greedy optimal homotopy and homology generators. *Proc. 16th Annu. ACM-SIAM Sympos. Discrete Algorithms*, 1038–1046, 2005.

[11] G. N. Frederickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM J. Comput.* 16(6):1004–1022, 1987.

[12] A. Hatcher. *Algebraic topology.* Cambridge University Press, 2002. ⟨http://www.math.cornell.edu/~hatcher/⟩.

[13] A. Itai, C. H. Papadimitriou, and J. L. Szwarcfiter. Hamilton paths in grid graphs. *SIAM J. Comput.* 11:676–686, 1982.

[14] M. Kutz. Computing shortest non-trivial cycles on orientable surfaces of bounded genus in almost linear time. *Proc. 22nd Annu. ACM Sympos. Comput. Geom.*, 430–437, 2006.

[15] F. Lazarus, M. Pocchiola, G. Vegter, and A. Verroust. Computing a canonical polygonal schema of an orientable triangulated surface. *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, 80–89, 2001.

[16] J. S. B. Mitchell, D. M. Mount, and C. H. Papadimitriou. The discrete geodesic problem. *SIAM J. Comput.* 16(4):647–668, 1987.

[17] B. Mohar and C. Thomassen. *Graphs on Surfaces.* Johns Hopkins University Press, 2001.

[18] J. Stillwell. *Classical Topology and Combinatorial Group Theory.* Springer-Verlag, New York, 1993.

[19] J. Takahashi, H. Suzuki, and T. Nishizeki. Shortest noncrossing paths in plane graphs. *Algorithmica* 16:339–357, 1996.

[20] A. Zomorodian. *Topology for Computing.* Cambridge University Press, 2005.