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 surface without boundary. A cycle on $\mathcal{M}$ is *splitting* if it has no self-intersections and it partitions $\mathcal{M}$ into two components, neither 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. Specifically, we describe an algorithm to compute the shortest splitting cycle in $g^{O(g)}n \log n$ time.

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. Such problems are usually set in the *combinatorial surface* model. A combinatorial surface is a graph $G(\mathcal{M})$ embedded on a surface $\mathcal{M}$ that cuts $\mathcal{M}$ into topological disks. Curves on this surface are required to be walks on $G(\mathcal{M})$, and edges of $G(\mathcal{M})$ have positive weights, allowing to measure the length of a curve.

Many of these problems can be seen as the computation of 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 [11, Sect. 4.3]. For instance, the sets of non-separating and non-contractible cycles satisfy the 3-path condition.

In this paper, we study the following optimization problem: Given an orientable 2-manifold $\mathcal{M}$ with genus $g \geq 2$ without boundary, find a shortest simple non-contractible cycle that separates $\mathcal{M}$. 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 property. Removing a splitting cycle from any surface leaves two surfaces, each of genus at least one and with one boundary cycle.

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. Specifically, we describe an algorithm to compute the shortest splitting cycle in $g^{O(g)}n \log n$ time.

*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.

§CNRS, Laboratoire des Images et des Signaux, Grenoble, France, Francis.Lazarus@lis.inpg.fr

¶Department of Mathematics, University of Illinois, kwhittle@math.uiuc.edu

2 Topological Background

2.1 Curves on surfaces

We rely on several notions from combinatorial topology. In particular, we use the standard definitions for a *compact, orientable, and connected surface* $\mathcal{M}$, its *genus*, a *path*, a *loop*, or a *cycle* on $\mathcal{M}$, *homotopy with or without basepoint*. See also Hatcher [8] or previous papers [4, 2] for more details. We say that two loops are disjoint if they intersect only at their common basepoint. We say that a cycle *splits* a surface $\mathcal{M}$ if it is simple, non-contractible, and separating.

2.2 Systems of loops

If L is a set of pairwise disjoint simple loops, $\mathcal{M} \setminus L$ denotes the surface with boundary obtained by *cutting* $\mathcal{M}$ along the loops in L. A *system of loops* on $\mathcal{M}$ is a set of pairwise disjoint simple loops L 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 the loops appear in pairs on its boundary, and is called the *polygonal schema* associated to L.

2.3 Combinatorial and cross-metric surfaces

Like most earlier related results [1, 3, 4, 5, 6, 10], we state and prove our results in the *combinatorial surface* model. A combinatorial surface is an abstract surface $\mathcal{M}$ together with a weighted undirected graph $G = G(\mathcal{M})$, 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. A path is *simple* if it can be slightly deformed so as to become a simple path on the surface. 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 [2]. 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 [2] for further discussion of these two models.

3 NP-hardness

Theorem 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 reduction from the Hamilton cycle problem in grid graphs [9].

We describe a two-step reduction. Let H be a grid graph of size n. To begin the first reduction, we overlay n 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, top. 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.

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 $2n$ to the resulting surface so that each face of the final embedded graph is a disk. The resulting combinatorial surface $\mathcal{M}(H)$ has genus $2n - 2$ and complexity $O(n)$, and it can clearly be constructed in polynomial time. See Figure 1(c).

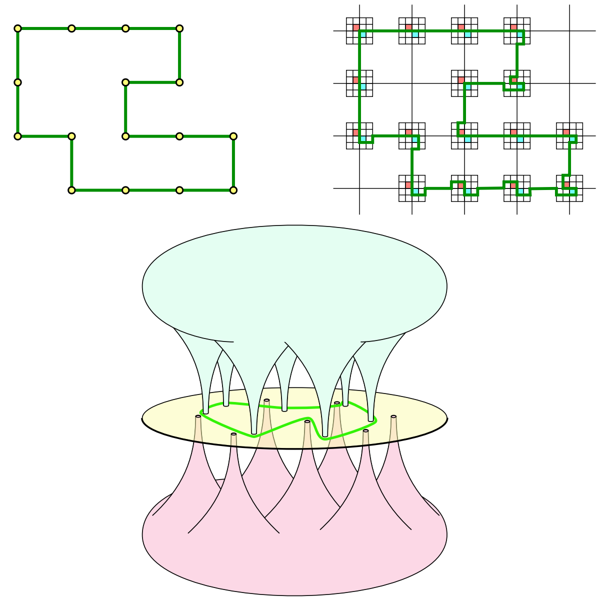

Figure 1: Top left: A Hamilton cycle of length n. Top right: The corresponding red/blue separating cycle (not to scale). Bottom: Separating heaven from hell (not to scale); the central disk is a small portion of Earth.

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. □

4 $O(g)$ Crossings with Any Shortest Path

Proposition 2 *Let P be a set of pairwise interior-disjoint shortest paths on a cross-metric surface $\mathcal{M}$. Some shortest splitting cycle crosses each path in P at most $O(g)$ times.*

Proof. 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.

Let γ be a shortest splitting cycle with 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. 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. Each arc corresponds to a loop in $\mathcal{M}/p$

with basepoint p/p. We say that two such arcs are *homotopic rel p* or *relatively homotopic* if the corresponding loops in $\mathcal{M}/p$ are homotopic.

For any two consecutive intersection points x and y along γ, the arc $\gamma[x,y]$ cannot be homotopic to $p[x,y]$, since otherwise we can obtain a no longer splitting cycle $\gamma[y,x]\cdot p[x,y]$ that has fewer crossings with the paths in P. It follows that none of the arcs of γ are contractible rel p. Since the arcs are disjoint except at their common endpoints, they correspond, in $\mathcal{M}/p$, to a set of simple, pairwise disjoint loops (except at their common basepoint) that are non-contractible and pairwise non-homotopic. Under these assumptions, the number of loops can be shown to be at most $12g-6$. Hence there are at most $12g-6$ relative homotopy classes of arcs.

We can partition the arcs into four types—LL, RR, LR, and RL—according to whether the arcs start on the left or right side of p, and whether they end on the left or right side of p. To complete the proof, we argue that there is at most one arc of each type in each relative homotopy class. It suffices to consider only types LL and RL; the other two cases follow from symmetric arguments.

Suppose for purposes of contradiction that there are two LL-arcs $u=\gamma[a,z]$ and $w=\gamma[c,x]$ that are homotopic rel p. Since the arcs are simple, pairwise disjoint, and homotopic, we may assume without loss of generality that the intersection points appear along p in the order a,c,x,z, and the cycle $u\cdot p[z,x]\cdot\bar{w}\cdot p[c,a]$ bounds a disk. Without loss of generality, we assume that no other arc homotopic rel p intersects this disk. Since γ is separating, there must be exactly one arc $v=\gamma[y,b]$ between u and w that is relatively homotopic to $\bar{u}$ and $\bar{w}$.

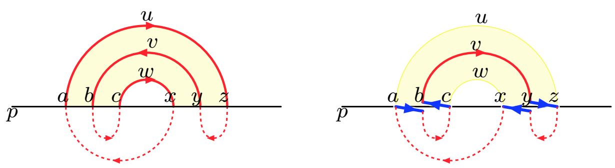

Figure 2: The exchange argument for LL arcs.

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 2. This cycle crosses p fewer times than γ, and crosses any other path in p no more than γ. We can prove that γ' is simple, separating, 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 the fact that γ is a shortest splitting cycle with the minimal number of crossings with paths in P. We conclude that any two LL-arcs must be in different relative homotopy classes.

The case of RL-arcs can be treated with a similar, though more complicated, analysis and exchange argument. □

5 Algorithm

We will prove the following result.

Theorem 3 *Let $\mathcal{M}$ be an orientable cross-metric surface without boundary; let g be its genus and n be its complexity. We can compute a shortest splitting cycle in $\mathcal{M}$ in $g^{O(g)}n\log n$ time.*

The algorithm proceeds in several stages, described in the following subsections. First, we compute the shortest system of loops from some arbitrary basepoint. Next, we enumerate all the plausible sequences of intersections of the splitting cycle with the loops in the system of loops. Note that a sequence of intersections determine the homotopy type for the splitting cycle. We discard any sequence that does not yield a valid splitting cycle. Finally, for each remaining sequence, we compute the shortest cycle with the same sequence of intersections. Out of all cycles constructed this way, the shortest one is the correct result.

5.1 Greedy Loops

Let v be any point of $\mathcal{M}$ in the interior of a face of $G^*(\mathcal{M})$. Let $\alpha_1,\ldots,\alpha_{2g}$ 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 [6].

Two key properties of this system of loops are that each loop α_i is as short as possible in its homotopy class, and is composed of two shortest paths β_i and β_i' in the primal graph $G(\mathcal{M})$. However, in general, these two paths meet at a point in the interior of some edge. We can handle this easily by a local modification of the primal or dual graph.

5.2 Enumeration of Plausible Sequences via Labeled Triangulations

Lemma 2 tells us that some shortest splitting cycle γ crosses each path β_i or β_i' at most $O(g)$ times, and thus crosses each loop α_i at most $O(g)$ times. Our algorithm therefore enumerates all sequences of intersections where each α_i is crossed at most $O(g)$ times. Since we are using the shortest system of loops, we may assume that the shortest splitting cycle does not cross any α_i consecutively in opposite directions.

Cut $\mathcal{M}$ along the loops α_i to obtain a polygonal schema. This operation also cuts the unknown cycle γ into many segments that go across the schema. Since γ is simple, no two of these segments cross. Without loss of generality, we can assume that γ does not pass through the basepoint v.

The segments of γ can be grouped into subsets according to which pair of edges they meet on the polygonal schema. We dualize the polygonal schema, replacing each edge with a vertex and connecting vertices that correspond to consecutive edges; now each subset of segments corresponds to an edge between two vertices of the dual $4g$-gon. Since no two segments cross, these edges cannot cross; in particular, all the edges belong to some triangulation of the dual polygon.

Thus the candidate sequences of intersections of a shortest splitting cycle are described by *labeled triangulations*, each of which consists of a triangulation of the dual polygon, in which every edge is labeled with an integer between 0 and $O(g)$. Intuitively, the label of an edge in the triangulation represents the number of times that the cycle runs along that edge. There are $C_{4g-2} = O(4^{4g})$ possible triangulations, where C_n is the nth Calatan number, which we can enumerate in $O(g)$ time each. There are $g^{O(g)}$ ways to label each triangulation, which we can enumerate in constant amortized time per labeling. We thus obtain a total of $g^{O(g)}$ potential labeled triangulations for γ.

5.3 Discarding Irrelevant Labeled Triangulations

Most of the labeled triangulations do not correspond to a splitting cycle, or to any cycle for that matter. We now explain how to throw away these possibilities. Namely, given a candidate labeled triangulation T, we want to decide if it (1) corresponds to a set of cycles, (2) is actually a single cycle, (3) is separating and (4) is non-contractible. (1) is satisfied if and only if, for each i, α_i and $\bar{\alpha}_i$ are crossed the same number of times. If this condition is satisfied, we can build a (set of) simple (pairwise disjoint) cycle(s) on $\mathcal{M}$ representing T, of complexity $O(g^2 n)$; for example, the segments can run along the boundary of the polygonal schema. If the second condition is satisfied, we can check conditions (3) and (4) using (simplified versions of) the algorithms by Erickson and Har-Peled [5]. This takes $O(g^2 n \log n)$ time. Actually, this whole step can be done more efficiently in $O(g^2)$ time.

5.4 Shortest Cycle for each Sequence of Intersections

In the last step, for each of the $g^{O(g)}$ non-discarded labeled triangulation, we compute the shortest cycle with the same sequence of intersections as defined by the labeled triangulation and keep the shortest one.

For this, we build a cylindrical surface by gluing copies of the polygonal schema. There is one copy per intersection in the sequence. If the ith intersection is α_j, then the ith copy is glued to the $i+1$th copy along α_j. (Each loop appears twice on the boundary of the polygonal schema and we must take orientation into account to make the proper gluing. Also, by construction, copies $i-1$ and $i+1$ are glued on different edges on the ith copy.) The last copy is glued to the first copy. The resulting cylinder is made of $O(g^2)$ copies of the polygonal schema, each of complexity $O(gn)$. By [7] (see also [2]) we can compute a shortest cycle homotopic to the boundaries of this cylinder in time $O(g^3 n \log n)$. Such a cycle, when projected back onto the surface $\mathcal{M}$, is a shortest cycle with the same prescribed sequence of intersections with the greedy system of loops. The total time spent is $g^{O(g)} n \log n$.

Finally, the output cycle may contain self-intersections. However, we are able to remove these intersections in the same amount of time. This concludes the proof of Theorem 3.

Acknowledgments. Support of a travel grant from UIUC/CNRS/INRIA is aknowledged. The authors also thank Martin Kutz for simplifying Section 5.4.

References

[1] S. Cabello and B. Mohar. Finding shortest non-separating and non-contractible cycles for topologically embedded graphs. In *Proc. 13th Annu. European Sympos. Algorithms*, volume 3669 of *LNCS*, pages 131–142. Springer-Verlag, 2005.

[2] É. Colin de Verdière and J. Erickson. Tightening non-simple paths and cycles on surfaces. In *Proc. 17th Annu. ACM-SIAM Sympos. Discrete Algorithms*, page to appear, 2006.

[3] É. Colin de Verdière and F. Lazarus. Optimal pants decompositions and shortest homotopic cycles on an orientable surface. In *Proc. 11th Sympos. Graph Drawing*, volume 2912 of *LNCS*, pages 478–490. Springer-Verlag, 2003.

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

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

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

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

[8] A. Hatcher. *Algebraic topology.* Cambridge University Press, 2002.

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

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

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