

Finding shortest contractible and shortest separating cycles in embedded graphs

Sergio Cabello*

Abstract

We give a polynomial-time algorithm to find a shortest contractible cycle (i.e. a closed walk without repeated vertices) in a graph embedded in a surface. This answers a question posed by Hutchinson. In contrast, we show that finding a shortest contractible cycle through a given vertex is NP-hard. We also show that finding a shortest separating cycle in an embedded graph is NP-hard. This answers a question posed by Mohar and Thomassen.

1 Introduction

Let G be a graph embedded in a surface Σ . Closed walks in G correspond to closed curves in Σ , and hence we may ask for shortest closed walks with certain topological properties, such as for example being non-contractible. For the rest of this paper the term *cycle* will be used for a *closed walk without repeated vertices*. We need to clarify this because 'cycle' has been used for different concepts in (combinatorial) surfaces [5, 3] and in graph theory [19].

Thomassen [22] (c.f. Mohar and Thomassen [19, Section 4.3]) introduced the concept of *3-path condition*, which can be phrased as follows: a family of cycles \mathcal{C} satisfies the 3-path condition if the cycles not in \mathcal{C} generate a subspace of the cycle space that is disjoint from \mathcal{C} . When \mathcal{C} satisfies the 3-path condition, then we can find the shortest cycle in \mathcal{C} in polynomial time using the so-called *fundamental cycle method*. As particular cases, one obtains polynomial-time algorithms to find a shortest non-contractible cycle and a shortest (surface) non-separating cycle in an embedded graph.

The family of contractible cycles does not satisfy the 3-path condition. This naturally leads to the following problem, attributed to Hutchinson in [23] (c.f. [19, Problem 4.3.3.(a)]): is there a polynomial-time algorithm that finds a shortest contractible cycle? Note that the shortest contractible closed walk may




Figure 1: A graph embedded in an annulus. A shortest contractible cycle is strictly longer than a shortest contractible closed walk.

use some vertices twice, and hence perhaps it is not a cycle; Figure 1 shows an example in the annulus. The impossibility of visiting the same vertex twice is not a topological concept and gives to the problem a combinatorial flavour.

We show that a shortest contractible cycle in an embedded graph can be found in $O(E^2 \log E)$ time, where E is the number of edges in the graph. In contrast, we show that it is NP-hard to find the shortest contractible cycle through a given vertex. In the context of graphs on surfaces, it is the first problem for which we cannot find efficiently the shortest object through a given vertex but we can find the shortest object overall. This may have been the main obstacle for tackling the problem. A slight modification of the NP-hardness reduction also shows that it is NP-hard to find a shortest separating cycle. This answers a question by Mohar and Thomassen [19, Problem 4.3.3.(b)].

Our approach consists of using surgery, a standard technique to deal with topological surfaces that consists of cutting the surface along suitable curves. However, surgery introduces copies of vertices, and hence one has to be careful not to use in the resulting surface copies arising from the same vertex. This naturally leads to a closely related problem, which we call *shortest cycle*

*Department of Mathematics, IMFM, and Department of Mathematics, FMF, University of Ljubljana, Slovenia. E-mail: sergio.cabello@fmf.uni-lj.si. Research supported by the Slovenian Research Agency, project J1-7218 and program P1-0297.

with *forbidden pairs*: given a graph G and a family $\mathcal{F} \subseteq \binom{V(G)}{2}$ of *forbidden pairs* of vertices, the task is to find a shortest cycle that does not contain any pair of vertices from \mathcal{F} . This problem is closely related to the shortest path with forbidden pairs, which is NP-hard [12, 16]. Our main contribution is showing that the shortest cycle with forbidden pairs can be solved in polynomial time when G is a plane graph (a planar graph together with an embedding) and all the vertices in \mathcal{F} are cofacial. In some sense, the result is tight: the problem becomes NP-hard for plane graphs when all the vertices in \mathcal{F} can be covered by two faces, instead of just one.

Related work. The abstract and the introduction of the paper by Ren and Deng [20] claims to settle the problem of finding a contractible cycle in polynomial time. However, that paper only considers the problem for so-called LEW embeddings, which is a very restricted class of embeddings in which the shortest non-contractible cycle is longer than any facial walk. Closely related is the work by Cabello et al. [3], where the problem of finding a shortest contractible closed walk that encloses an unspecified face is reduced to the problem of finding a minimum cut in an edge-weighted planar graph. This results in a near-linear time algorithm for finding the shortest contractible closed walk.

It turns out that the shortest non-contractible and the shortest non-separating closed walks in an embedded graph are actually cycles, and hence the distinction between cycle and closed walk becomes irrelevant. Thomassen [22] provided the first polynomial-time algorithm to find a shortest non-contractible and a shortest non-separating cycle, running in roughly cubic time. Much work has been done in these problems since then [1, 2, 4, 10, 17]; the current best running time is $O(\min\{E, g^3\}E \log E)$ for graphs with E edges embedded in surfaces of genus g . Chambers et al. [5] considered the problem of finding a shortest splitting closed walk, that is, a non-contractible but separating closed walk that can be infinitesimally perturbed into a simple (injective) curve.

Other shortest ‘objects’ in an embedded graph have also been considered. Erickson and Har-Peled [10] considered the problem of finding a shortest subgraph with the property that cutting along it leaves a topological disk. Erickson and Whittlesey [11] studied shortest homotopy or homology generators. Colin de Verdière and Lazarus have studied shortest system of loops [7] and pants decompositions [8]. Colin de Verdière and Erickson [6] have studied the problem of finding the shortest walk (closed or not) homotopic to a given one.

Organization of the paper. In Section 2 we present some concepts used through the paper. In

Section 3 we provide a polynomial-time algorithm to solve the problem of finding a shortest cycle with forbidden pairs when all the vertices in the forbidden pairs are cofacial. In Section 4 we provide a polynomial-time algorithm to find a shortest contractible cycle in an embedded graph. In Section 5 we show NP-hardness of several related problems.

2 Background

Our presentation will be using terminology mostly used for graphs embedded in surfaces [19]. An alternative option would be using the terminology of combinatorial surfaces [6].

Topology. We assume that the reader is familiar with topology of surfaces; see any of the books [14, 18, 21]. Any surface Σ considered in the paper is compact, may have boundary, and may be non-orientable. The term *closed curve* is used for continuous mappings from \mathbb{S}^1 to Σ ; the term *arc* is used for paths whose endpoints are on the boundary of Σ . Through the paper, homotopy of paths and arcs is considered with fixed endpoints, while homotopy of closed curves is considered free homotopy. A closed curve α is *non-separating* if cutting the surface along (the image of) α gives rise to a unique connected component; the concept is related to \mathbb{Z}_2 -homology. A contractible closed curve is a separating curve. Two curves α, β cross c times if: (i) there exist infinitesimal continuous perturbations of α, β that cross transversally c times; and (ii) any infinitesimal continuous perturbations of α, β have at least c points in common.

Embeddings. We assume that the reader is familiar with graphs embedded in surfaces without boundaries. We only consider cellular embeddings, where the removal of the image of the graph leaves a set of topological disks. For graphs embedded in surfaces with boundary, we will only consider embeddings where each boundary component is contained in (the image of) the edge set. Alternatively, we may see these embeddings as those obtained from an embedding in a surface without boundary after removing the interior of certain faces. Formally, removing the interior of two faces adjacent to the same edge or removing the interior of a face whose facial walk is not a cycle does not give rise to a surface. However, this is a technicality that does not bring problems. An *embedded graph* consists of an abstract graph, a surface, and a combinatorial embedding of the graph in the surface. A *plane graph* is a graph embedded in the sphere or in a topological disk. We assume that the embedding is represented in a suitable way, like for example the one discussed by Eppstein [9].

Walks and surgery. A *walk* in a graph G is a sequence of vertices where two consecutive vertices are

connected by an edge in G . As mentioned before, the term *cycle* refers to a closed walk without repeated vertices. For embedded graphs, we will use the term *arc* for a walk whose end vertices are on the boundary of the surface. For a walk α through vertices x, y , we use $\alpha[x, y]$ to denote the subwalk of α between x, y ; if there are multiple appearances of x, y along α it will be clear from the context which ones are used. Homotopy of closed walks and cycles is the same as homotopy of closed curves. Homotopy of walks and arcs is the same as homotopy of paths.

Let G be an embedded graph and let α be a walk in G . We use the notation $G\mathcal{A}\alpha$ to denote the embedded graph obtained after cutting the surface that receives the embedding along the edges of α . This can be seen as ungluing the two faces adjacent to an edge e for each edge e in α . We denote by $G\mathcal{A}(\alpha_1, \dots, \alpha_k)$ the embedded graph obtained inductively as $(G\mathcal{A}(\alpha_1, \dots, \alpha_{k-1}))\mathcal{A}\alpha_k$.

Lengths. We assume that a graph G has positive edge-weights. The length $|\alpha|$ of a walk α is defined as the sum of the weights of the edges along α , counted with multiplicity. The distance $d_H(u, v)$ between vertices u, v is defined as the minimum length over all walks connecting u to v . A walk is *tight* if it is shortest in its homotopy class. The following result will be used in our algorithms.

LEMMA 1. *Let G be an embedded graph with E edges.*

- (a) *A shortest non-separating cycle in G can be found in $O(E^2 \log E)$ time.*
- (b) *If the surface where G is embedded has boundary, we can find in $O(E \log E)$ time a family of tight walks $\alpha_1 \dots, \alpha_k$ such that $G\mathcal{A}(\alpha_1, \dots, \alpha_k)$ is a plane graph with $O(E)$ edges.*

Proof. Item (a) is from [10]. Item (b) follows from Lemmas 4 and 5 of [3]; see Theorem 5 of [3].

Shortest cycle with forbidden pairs. Let G be a graph, not necessarily embedded. A family of forbidden pairs is a collection $\mathcal{F} \subseteq \binom{V(G)}{2}$, where $\binom{A}{2} \subset 2^A$ denotes the collection of subsets of A with two elements. A walk in G is *in compliance* with \mathcal{F} when u or v are not in the walk for each $\{u, v\} \in \mathcal{F}$. When G is a plane graph, the family \mathcal{F} is *cofacial* if all the vertices appearing in \mathcal{F} are cofacial.

3 Shortest cycle with cofacial forbidden pairs

In this section we consider the problem of finding a shortest cycle in a plane graph G with forbidden pairs \mathcal{F} , when the vertices in \mathcal{F} are cofacial. We will assume that G is 2-connected, as otherwise we can decompose the graph into maximal 2-connected subgraphs (usually

called *blocks*) in linear time, and treat each block separately. In a 2-connected plane graph all facial walks are cycles.

Let f_0 be the face adjacent to all vertices in \mathcal{F} , and let C_0 be the facial cycle defining f_0 . We may assume that G is embedded in a disk such that C_0 defines the boundary of the disk. Let $E(C_0)$ and $V(C_0)$ denote the edges and vertices of C_0 , respectively.

We next describe a recursive algorithm. The base case is when $G = C_0$ is a cycle, where we have two options: if \mathcal{F} is empty, then C_0 is the shortest cycle; if \mathcal{F} is nonempty, then G does not contain any cycle in compliance with \mathcal{F} . For the general case, we would like to cut along a tight arc connecting vertices of C_0 , and obtain two smaller problems. However, it may be that all such tight arcs are contained in C_0 , and hence no such cutting is possible. We fix this considering the graph $H = G - E(C_0)$, as follows. First note that some pair of distinct vertices of $V(C_0)$ are connected through a path in H because G is 2-connected. Take the distance $d_H(u, v)$ in H for each pair of vertices $u, v \in V(C_0)$, $u \neq v$, and let u_0, v_0 be a pair minimizing it; that is

$$(u_0, v_0) = \arg \min_{(u, v) \in (V(C_0))^2, u \neq v} d_H(u, v).$$

Let π_0 be a shortest path in H between u_0 and v_0 . Note that π_0 is disjoint from $E(C_0)$ by the definition of H . We have the following observation.

LEMMA 2. *If C_0 is not a shortest cycle in G in compliance with \mathcal{F} , then there exists a shortest cycle in G in compliance with \mathcal{F} that does not cross π_0 .*

Proof. Note that π_0 is not a shortest path in G , but in H , and therefore the result does not follow immediately. Through the proof we will assume that all cycles have the same orientation, which is clockwise in our figures.

Let \mathcal{Q} be the class of cycles of G in compliance with \mathcal{F} that are shortest. Each cycle of \mathcal{Q} bounds a topological disk. Let C_{short} be a cycle from \mathcal{Q} that bounds a disk $Disk_{\text{short}}$ whose closure does not contain any other cycle from \mathcal{Q} . We will show that π_0 is disjoint from the interior of $Disk_{\text{short}}$, and therefore C_{short} does not cross π_0 .

Let us assume, for the sake of contradiction, that π_0 is not disjoint from the interior of $Disk_{\text{short}}$. Consider a maximal subpath $\pi_0[x, y]$ with endpoints x, y in C_{short} that is contained in the interior of $Disk_{\text{short}}$, except for its endpoints. Let Q_1 be the closed walk obtained by concatenating $\pi_0[x, y]$ with $C_{\text{short}}[y, x]$ and let Q_2 be the closed walk obtained by concatenating $\pi_0[y, x]$ with $C_{\text{short}}[x, y]$. Note that Q_1, Q_2 are cycles, and they are in compliance with \mathcal{F} because π_0 does not contain edges from C_0 . Moreover, Q_1, Q_2 are contained in $Disk_{\text{short}}$.

We will show that Q_1 or Q_2 is no longer than C_{short} , leading to a contradiction with the choice of C_{short} . We distinguish cases depending on the overlap of C_{short} and C_0 . Note that $C_{\text{short}} \neq C_0$ by hypothesis, so we have the following three scenarios.

- (i) If C_{short} and C_0 do not have any common edge, then C_{short} is a shortest cycle in H . See Figure 2, left. In this case $|\pi_0[x, y]| \leq |C_{\text{short}}[x, y]|$ because π_0 is a shortest path in H . Therefore

$$\begin{aligned} |Q_1| &= |\pi_0[x, y]| + |C_{\text{short}}[y, x]| \\ &\leq |C_{\text{short}}[x, y]| + |C_{\text{short}}[y, x]| \\ &= |C_{\text{short}}|, \end{aligned}$$

and Q_1 is no longer than C_{short} .

- (ii) If C_{short} and C_0 agree in a single maximal subpath $C_{\text{short}}[v_1, u_1]$ of C_{short} , then $v_1, u_1 \in V(C_0)$ and $C_{\text{short}}[u_1, v_1]$ is contained in H . We have three subcases:

- x, y are vertices along $C_{\text{short}}[u_1, v_1]$. See Figure 2, center. Renaming the vertices x, y if necessary, we may assume that the ordering of vertices along $C_{\text{short}}[u_1, v_1]$ is u_1, x, y, v_1 , possibly with $u_1 = x$ or $y = v_1$. Since $C_{\text{short}}[u_1, v_1]$ is contained in H , also $C_{\text{short}}[x, y]$ is contained in H . We then have $|\pi_0[x, y]| \leq |C_{\text{short}}[x, y]|$ because π_0 is a shortest path in H , and therefore Q_1 is no longer than C_{short} .
- x, y are vertices along $C_{\text{short}}[v_1, u_1]$. See Figure 2, right. In this case $\{x, y\} = \{u_0, v_0\}$. Renaming the vertices x, y if necessary, we may assume that the ordering of vertices is such that $C_{\text{short}}[u_1, v_1] \subseteq C_{\text{short}}[x, y]$. By the choice of u_0, v_0 we have

$$\begin{aligned} |\pi_0[x, y]| &= |\pi_0[u_0, v_0]| \leq |C_{\text{short}}[u_1, v_1]| \\ &\leq |C_{\text{short}}[x, y]|, \end{aligned}$$

and therefore Q_1 is no longer than C_{short} .

- x, y are not both along $C_{\text{short}}[u_1, v_1]$ or both along $C_{\text{short}}[v_1, u_1]$. See Figure 3, left and center. In this case, $\{x, y\} \cap \{u_1, v_1\} = \emptyset$. Renaming the vertices u_0, v_0 if necessary, we may assume that v_0 is along the path $C_{\text{short}}[v_1, u_1]$. Renaming the vertices x, y if necessary, we may assume that x is along the path $C_{\text{short}}[u_1, v_1]$, and hence $y = v_0$. When $u_1 \neq u_0$ (Figure 3, left), the concatenation of $C_{\text{short}}[u_1, x]$ and $\pi_0[x, u_0]$ is in H and connects

two different vertices of $V(C_0)$. Therefore

$$\begin{aligned} |\pi_0[u_0, x]| + |\pi_0[x, v_0]| &= |\pi_0[u_0, v_0]| \\ &\leq |C_{\text{short}}[u_1, x]| + |\pi_0[x, u_0]|, \end{aligned}$$

and hence

$$\begin{aligned} |\pi_0[x, y]| &= |\pi_0[x, v_0]| \leq |C_{\text{short}}[u_1, x]| \\ &\leq |C_{\text{short}}[y, x]|. \end{aligned}$$

This implies that Q_2 is no longer than C_{short} . When $u_1 = u_0$ (Figure 3, center), then $v_1 \neq u_0$, and the same argument used with the concatenation of $\pi_0[v_0, x]$ and $C_{\text{short}}[x, v_1]$ shows that $|\pi_0[x, y]| \leq |C_{\text{short}}[x, y]|$, and hence Q_1 is no longer than C_{short} .

- (iii) If C_{short} and C_0 agree in at least two (disjoint) maximal subpaths of C_{short} , then there exist vertices $u_1, v_1, u_2, v_2 \in V(C_0)$ such that the paths $C_{\text{short}}[u_1, v_1], C_{\text{short}}[u_2, v_2]$ are disjoint from C_0 , except at its endpoints. See Figure 3, right. Since the paths $C_{\text{short}}[u_1, v_1], C_{\text{short}}[u_2, v_2]$ are contained in H , the choice of u_0, v_0 implies

$$|C_{\text{short}}[u_1, v_1]| \geq |\pi[u_0, v_0]|$$

and

$$|C_{\text{short}}[u_2, v_2]| \geq |\pi[u_0, v_0]|.$$

Since the paths $C_{\text{short}}[u_1, v_1], C_{\text{short}}[u_2, v_2]$ are disjoint, except possibly at its endpoints, we have

$$\begin{aligned} |C_{\text{short}}| &\geq |C_{\text{short}}[u_1, v_1]| + |C_{\text{short}}[u_2, v_2]| \\ &\geq 2 \cdot |\pi[u_0, v_0]|, \end{aligned}$$

and hence

$$\begin{aligned} |C_{\text{short}}[x, y]| + |C_{\text{short}}[y, x]| &= |C_{\text{short}}| \\ &\geq 2 \cdot |\pi[u_0, v_0]|. \end{aligned}$$

This means that $\max\{|C_{\text{short}}[x, y]|, |C_{\text{short}}[y, x]|\} \geq |\pi[u_0, v_0]| \geq |\pi[x, y]|$, and Q_1 or Q_2 is no longer than C_{short} .

THEOREM 1. *Given a plane graph G with n vertices and a family \mathcal{F} of forbidden pairs whose vertices are all adjacent to one face of G , we can find a shortest cycle in compliance with \mathcal{F} in $O(n^2 + nF)$ time.*

Proof. Like before, we assume that G is 2-connected, that C_0 is the facial cycle of G containing all vertices in \mathcal{F} , and that G is embedded in a disk whose boundary is defined by C_0 . Consider the following recursive algorithm:




Figure 2: Left: an example for case (i) of Lemma 2. Center: an example for the first item in case (ii) of Lemma 2. Right: an example for the second item in case (ii) of Lemma 2.




Figure 3: Left and center: examples for the third item in case (ii) of Lemma 2. Right: an example for case (iii) of Lemma 2.

- (i) If G is a cycle and \mathcal{F} is empty, we return the cycle G .
- (ii) If G is a cycle and \mathcal{F} is nonempty, we return that no cycle is in compliance with \mathcal{F} .
- (iii) If G is not a cycle, we construct $H = G - E(C_0)$, find a shortest path π_0 between vertices

$$(u_0, v_0) = \arg \min_{(u,v) \in (V(C_0))^2, u \neq v} d_H(u, v),$$

and cut G along π_0 to obtain two 2-connected plane graphs, denoted G_1, G_2 . We regard G_1, G_2 as embedded in topological disks. We also construct two families of forbidden sets $\mathcal{F}_1, \mathcal{F}_2$ from \mathcal{F} in the natural way: \mathcal{F}_i contains each forbidden pair of \mathcal{F} whose two vertices (or copies of them) are in G_i . If $\{u_0, v_0\} \in \mathcal{F}$, then $\{u_0, v_0\}$ goes to both $\mathcal{F}_1, \mathcal{F}_2$. All other forbidden pairs of \mathcal{F} go only to \mathcal{F}_1 or to \mathcal{F}_2 . For $i = 1, 2$, we *recursively* find a shortest cycle C_i in G_i in compliance with \mathcal{F}_i , or find out that no cycle of G_i is in compliance with \mathcal{F}_i . If \mathcal{F} is empty, we return the shortest among C_0, C_1, C_2 , or the subset of them that is defined. If \mathcal{F} is nonempty and for $i = 1, 2$ no cycle of G_i is in compliance with \mathcal{F}_i , we then return that no cycle of G is in compliance with \mathcal{F} . Otherwise, we return

the shortest cycle among C_1, C_2 , or the subset of them that is defined.

Correctness of the algorithm is shown by induction on $|E(H)|$. If $|E(H)| = 0$, then G is a cycle, and correctness follows from cases (i) and (ii) of the algorithm. If $|E(H)| \geq 1$, then G is not a cycle, and G_1 and G_2 are also 2-connected plane graphs. For $i = 1, 2$, let H_i be the subgraph of G_i containing the edges not in the boundary of the disk holding the embedding of G_i . It clearly holds $|E(H_i)| < |E(H)|$, and by the induction hypothesis the algorithm correctly finds a shortest cycle C_i in G_i that is in compliance with \mathcal{F}_i , and hence with \mathcal{F} , or tells that no such cycle exists. If C_0 is a shortest cycle in compliance with \mathcal{F} , then the algorithm returns C_0 , and hence the algorithm is correct. Otherwise, C_0 is not a shortest cycle in compliance with \mathcal{F} , and because of Lemma 2, the shortest cycle between C_1, C_2 is the shortest cycle in compliance with \mathcal{F} .

Let us now bound the time complexity of the algorithm. Finding the path π_0 can be done in linear time, as follows. Firstly, construct the graph $H = G - E(C_0)$ and add an extra vertex v_{new} with edges of length zero to all vertices $V(C_0)$. Let H' be the resulting planar graph. Secondly, construct in H' a shortest path tree T_{sp} from the source v_{new} in linear time [15]. Note that a shortest path in H consists of a subpath of T_{sp} ,

an edge from $E(H) \setminus E(T_{sp})$, and another subpath of T_{sp} . Therefore, we can take u_0, v_0 as the ancestors in $V(C_0)$ of the vertices in the edge

$$\arg \min_{uv \in E(H) \setminus E(T_{sp})} \{d_{T_{sp}}(u, v_{\text{new}}) + d_{T_{sp}}(v, v_{\text{new}}) + |uv|\}.$$

Finally, we compute the shortest path π_0 in H between u_0 and v_0 in linear time.

Consider a rooted tree where each node corresponds to a subproblem considered through the recursive algorithm; the root corresponds to the original problem. The two children of a node contain disjoint subsets of the faces from G . This tree has at most one leaf-node per face of G , and therefore a total of $O(n)$ subproblems are considered. In each subproblem we spend $O(n)$ time to find the path π_0 and $O(F + n)$ time to check if the boundary cycle C_0 is in compliance with \mathcal{F} . The running time follows.

In our forthcoming application, the forbidden pairs \mathcal{F} have some additional structure. A partition $\mathcal{V} = \{V_1, V_2, \dots\}$ of $V(G)$ defines the set of forbidden pairs $\mathcal{F}(\mathcal{V}) = \bigcup_{V_i \in \mathcal{V}} \binom{V_i}{2}$. Here we assume that $\binom{V_i}{2} = \emptyset$ whenever V_i has one element. Let us say that a partition \mathcal{V} is cofacial when $\mathcal{F}(\mathcal{V})$ is cofacial.

COROLLARY 1. *Given a plane graph G with n vertices and a cofacial partition \mathcal{V} of $V(G)$, we can find a shortest cycle in compliance with $\mathcal{F}(\mathcal{V})$ in $O(n^2)$ time.*

Proof. Given the partition \mathcal{V} of $V(G)$, we can check in $O(n)$ time if a cycle is in compliance with $\mathcal{F}(\mathcal{V})$ or not. Hence, in each of the $O(n)$ subproblems of the recursive algorithm of Theorem 1 we spend $O(n)$ time, instead of $O(n + F)$. The result follows.

4 Shortest cycle in embedded graphs

Let G be a graph embedded in a surface, possibly with boundary, and let \mathcal{F} be a family of forbidden pairs. Here we state the basic tool; see [3, Lemma 3] for a similar statement. The key observation is to consider the case when the vertices appearing in \mathcal{F} are on the boundary of the surface.

LEMMA 3. *Let α be a tight arc or closed walk in G . If all the vertices of \mathcal{F} are on the boundary of the surface then there exists a shortest contractible cycle in compliance with \mathcal{F} that does not cross α .*

Proof. Let \mathcal{Q} be the class of shortest contractible cycles of G in compliance with \mathcal{F} . Each cycle of \mathcal{Q} bounds a topological disk. Let C_{short} be a cycle from \mathcal{Q} bounding a disk $Disk_{\text{short}}$ whose closure does not contain any other cycle from \mathcal{Q} . We will show that α is disjoint

from the interior of $Disk_{\text{short}}$, and therefore C_{short} does not cross α .

Assume for the sake of contradiction, that α enters the interior of $Disk_{\text{short}}$. Consider a maximal subpath $\alpha[x, y]$ with the property that $\alpha[x, y] \setminus \{x, y\}$ is contained in the interior of $Disk_{\text{short}}$. Note that the endpoints x, y are vertices of C_{short} . It must be $x \neq y$, as otherwise $\alpha[x, y]$ would be contractible because it would be contained in a topological disk. Similarly, $\alpha[x, y]$ cannot have any repeated vertices, because the subpath between two repeated vertices would be contractible. It results that the closed walk C' , defined as the concatenation of $\alpha[x, y]$ and $C_{\text{short}}[y, x]$ has no repeated vertices, and hence is a cycle.

Note that the interior of $Disk_{\text{short}}$ cannot contain any vertex appearing in \mathcal{F} because $Disk_{\text{short}}$ does not contain any boundary component. Therefore, the cycle C' is also in compliance with \mathcal{F} because the ‘new’ subpath $\alpha[x, y] \setminus \{x, y\}$ does not contain vertices appearing in \mathcal{F} . Moreover, the paths $\alpha[x, y]$ and $C_{\text{short}}[y, x]$ are homotopic because $\alpha[x, y]$ is inside the disk $Disk_{\text{short}}$, and hence C' is also contractible. Using that $|\alpha[x, y]| \leq |C_{\text{short}}[y, x]|$ because α is tight, we see also that in fact C' is no longer than C_{short} . We conclude that $C' \in \mathcal{Q}$. Moreover, C' is contained in the closure of $Disk_{\text{short}}$ by construction. Therefore the properties of C' contradict the choice of C_{short} .

THEOREM 2. *Given a graph G with E edges embedded in a surface, possibly with boundary, we can find in $O(E^2 \log E)$ time a shortest contractible cycle in G .*

Proof. We first give an algorithm, then discuss its correctness, and finally derive its running time. We compute a shortest non-separating cycle α in G , and construct the embedded graph $G' = G \# \alpha$. Note that G' has at least one boundary component. We then construct the embedded graph $G'' = G' \# (\alpha_1, \alpha_2, \dots, \alpha_k)$, where $\alpha_1, \dots, \alpha_k$ is the tight system of arcs from Lemma 1(b). Note that G'' is a plane graph. We then group the vertices in the boundary of G'' as follows: for each vertex $v \in V(G)$, let W_v be the vertices in G'' that arose as copies of the original vertex v . Therefore, the groups $W_v, v \in V(G)$, form a partition of $V(G'')$. Note that the only groups in $W_v, v \in V(G)$, with cardinality larger than one have vertices in the boundary of G'' because copies of vertices are only introduced along the boundary. Therefore, $\mathcal{F}'' = \mathcal{F}(\{W_v \mid v \in V(G)\}) = \bigcup_{v \in V(G)} \binom{W_v}{2}$ is a cofacial family of forbidden pairs, and we can find the shortest cycle C_{short} in G'' in compliance with \mathcal{F}'' using Corollary 1. Finally, we return C_{short} . This finishes the description of the algorithm.

We next show the correctness of the algorithm. The embedded graph G'' is obtained iteratively from G by

cutting along a tight walk, until the plane graph G'' is obtained. Let $H_0, H_1, H_2, \dots, H_t$, where $H_0 = G$ and $H_t = G''$, be the sequence of embedded graphs constructed through the algorithm, where H_{i+1} is obtained from H_i by cutting along one tight walk. For any embedded graph H_i , let \mathcal{F}_i be the family of forbidden pairs given by pairs $\{u, v\} \in \binom{V(H_i)}{2}$ where u, v are copies of the same vertex in the original graph G . Note that by construction all the vertices in \mathcal{F}_i are in the boundary of H_i . A shortest contractible cycle in compliance with \mathcal{F}_{i+1} in H_{i+1} is a shortest contractible cycle in compliance with \mathcal{F}_i in H_i because of Lemma 3. By induction, this means that a shortest contractible cycle in compliance with \mathcal{F}_i in H_i is a shortest contractible cycle in $H_0 = G$. In particular, the cycle C_{short} , which is the shortest contractible cycle in compliance with $\mathcal{F}_t = \mathcal{F}''$ in $H_t = G''$, is a shortest contractible cycle in G .

We next show that the running time of the algorithm is $O(E^2 \log E)$. Finding the shortest non-separating cycle¹ takes $O(E^2 \log E)$ because of Lemma 1(a). Finding the arcs $\alpha_1, \dots, \alpha_k$ used to cut G' takes $O(n \log n)$ because we apply Lemma 1(b) to G' , which has $O(E)$ edges. We hence obtain a plane graph G'' with $O(E)$ edges. Finally, finding the shortest cycle in G'' in compliance with \mathcal{F}'' can be done in $O(E^2)$ time because of Corollary 1.

5 Hardness of related problems

Our NP-hardness results hold for unweighted graphs, so we will restrict our discussion in this section to unweighted graphs. All reductions are from the problem INDEPENDENTSET, which is NP-hard [13]: Given a graph H , find a largest independent set in H .

Consider an input graph H for INDEPENDENTSET, which we assume for simplicity that has $V(H) = \{1, 2, \dots, n\}$. Our reductions use an auxiliary graph G_H embedded in the sphere, that depends on H . The construction of G_H starts with an empty graph, and proceeds as follows:

1. For each vertex $v \in V(H)$ with neighbours $u_1 < \dots < u_{\deg(v)}$, create a path P_v with $\deg(v) + 1$ new vertices, denoted $v, v(u_1), v(u_2), \dots, v(u_{\deg(v)})$, and with edges between v and $v(u_1)$ and between $v(u_i)$ and $v(u_{i+1})$ for $i = 1, \dots, \deg(v) - 1$. Add the path P_v to G_H .
2. For each vertex $v \in V(H)$, make a path Q_v with $\deg(v) + 2$ vertices, whose second vertex is denoted v' . Add the path Q_v to G_H .

¹For orientable surfaces, we could use the tight closed walk of [3], which can be found faster. However, this would reduce the final running time only to $O(E^2)$.

3. For each vertex $v \in V(H)$, add new vertices s_v, t_v to G_H , add edges from s_v to the first vertex in P_v and the first vertex in Q_v , and add edges from t_v to the other ends of P_v and Q_v . This defines a cycle for each vertex $v \in V(H)$; see Figure 4 left.
4. Add a new vertex s to G_H , add the edges ss_1, st_n , and add the edges $t_i s_{i+1}$ for $i = 1, \dots, n - 1$; see Figure 4 right. This finishes the construction of G_H .

This finishes the description of the vertices and edges in G_H . The embedding of G_H in the sphere is fixed to be the one shown in Figure 4, right. In this embedding, the paths adjacent to s_i are ordered as $s_i t_{i-1}, s_i Q_i, s_i P_i$ and the paths adjacent to t_i are ordered as $t_i P_i, t_i Q_i, t_i s_{i+1}$, where we use $s = t_0 = s_{n+1}$.

Consider the family \mathcal{F}_{G_H} of forbidden pairs given by

$$\mathcal{F}_{G_H} = \{\{u(v), v(u)\} \mid uv \in E(H)\}.$$

This family \mathcal{F}_{G_H} of forbidden pairs has the following property: a cycle in G_H is in compliance with \mathcal{F} if and only if does not use both P_v and P_u for any $uv \in E(H)$. Let \mathcal{Q}_{G_H} be the set of cycles from G_H that are in compliance with \mathcal{F}_{G_H} and are of the form $ss_1 R_1 t_1 s_2 R_2 t_2 \dots s_n R_n t_n s$, where each R_v is either P_v or Q_v . The connection between H and G_H is summarized in the following observation.

LEMMA 4. *The graph H has an independent set of size k if and only if the shortest cycle in \mathcal{Q}_{G_H} has length at most $1 + 4n + 2 \cdot |E(H)| - k$.*

THEOREM 3. *The following two problems are NP-hard for unweighted graphs:*

- (a) *Given a plane graph G , a family of forbidden pairs \mathcal{F} , which may be cofacial, and a vertex s of G , find a shortest cycle in compliance with \mathcal{F} that passes through s .*
- (b) *Given a plane graph and a family of forbidden pairs \mathcal{F} that can be covered by 2 or more faces, find a shortest cycle in compliance with \mathcal{F} .*

Proof. To show the NP-hardness of problem (a), consider the plane graph G_H , the family \mathcal{F}_{G_H} of forbidden pairs, and take $s \in V(G_H)$ as the given vertex. The graph G_H is constructed such that any cycle through s is of the form $ss_1 R_1 t_1 s_2 R_2 t_2 \dots s_n R_n t_n s$, where each R_v is either P_v or Q_v . Hence, the family of cycles \mathcal{Q}_{G_H} is precisely the family of cycles in G_H through s that are in compliance with \mathcal{F}_{G_H} . Finding the shortest cycle in \mathcal{Q}_{G_H} would also solve the problem INDEPENDENTSET




Figure 4: Reduction in Theorem 3.

because of Lemma 4. This shows that problem (a) is NP-hard.

To show the NP-hardness of problem (b), one considers the plane graph G_H , and the family \mathcal{F} of forbidden pairs $\mathcal{F}' = \mathcal{F}_{G_H} \cup \{\{1, 1'\}, \{2, 2'\}, \dots, \{n, n'\}\}$. It is easy to see that the family of cycles in G_H in compliance with \mathcal{F}' is precisely \mathcal{Q}_{G_H} , and the NP-hardness follows by Lemma 4.

For studying contractible and separating cycles, the following observation will be useful.

LEMMA 5. *Let G be a graph embedded in a surface Σ and let α be a closed curve in Σ through a vertex v of G . If α does not touch any other edge or vertex of G besides v , a separating cycle in G cannot cross α .*

THEOREM 4. *The following problem is NP-hard for unweighted graphs: given a graph embedded in a surface and a vertex $s \in V(G)$, find a shortest contractible cycle through s .*

Proof. Consider the graph G_H embedded in the sphere, as defined above. Let f_0 be the face bounded by the cycle $ss_1P_1t_1s_2P_2t_2 \cdots s_nP_nt_ns$. We convert G_H into a graph embedded in an orientable surface of genus $|E(H)|$, as follows. For each vertex $v(u)$ of $V(G_H)$ take a small simple curve $\alpha_{v(u)}$ passing through the vertex $v(u)$ and that is otherwise contained in the interior of f_0 . Moreover, the curves $\alpha_{v(u)}$ should be pairwise disjoint for all $v(u) \in V(G_H)$. Let $\tilde{\Sigma}$ be the surface obtained by removing the interior of the two disks bounded by $\alpha_{v(u)}, \alpha_{u(v)}$, for all $uv \in E(H)$. Then, for each edge $uv \in E(H)$ we glue the two boundaries defined by $\alpha_{v(u)}$ and $\alpha_{u(v)}$, in such a way that the vertices $u(v), v(u)$ are identified and the resulting surface is orientable; see Figure 5. This finishes the transformation of G_H . Let G'_H be the resulting embedded graph, which is embedded in a certain surface Σ' . Since we add $|E(H)|$ ‘independent handles’, the surface Σ' has genus $|E(H)|$ and no boundary.

First, note that a separating cycle in G'_H cannot cross $\alpha_{u(v)} = \alpha_{v(u)}$ because of Lemma 5 applied to the

curve $\alpha_{u(v)}$ in the surface Σ' . Since cutting the surface Σ' along the curves $\alpha_{u(v)}$, for all $uv \in E(H)$, give rise to $\tilde{\Sigma}$, it follows that a cycle in G'_H that is separating must be a cycle in G_H that is in compliance with \mathcal{F}_{G_H} . As a particular case, any contractible cycle in G'_H must be a cycle in G_H that is in compliance with \mathcal{F}_{G_H} . Note also that any cycle in G_H becomes a contractible cycle in G'_H because f_0 is the only face affected by the topological surgery.

It follows that finding a shortest contractible cycle in G'_H through s is equivalent to finding a shortest cycle in G_H through s that is in compliance with \mathcal{F}_{G_H} . As seen in the proof of Theorem 3(a), this is equivalent to finding a shortest cycles in \mathcal{Q}_{G_H} , which in turn is equivalent to solving the problem INDEPENDENTSET in the graph H because of Lemma 4.

A similar proof where one also identifies each pair of vertices u, u' of G_H gives the following result.

THEOREM 5. *The following problem is NP-hard for unweighted graphs: given a graph G embedded in a surface, find a shortest surface-separating cycle.*

References

- [1] S. Cabello. Many distances in planar graphs. In *SODA '06: Proc. 17th Symp. Discrete Algorithms*, pages 1213–1220, 2006.
- [2] S. Cabello and E. W. Chambers. Multiple source shortest paths in a genus g graph. In *SODA '07: Proc. 18th Symp. Discrete Algorithms*, pages 89–97, 2007.
- [3] S. Cabello, M. DeVos, J. Erickson, and B. Mohar. Finding one tight cycle. Manuscript available at <http://www.imfm.si/preprinti/PDF/01047.pdf>; preliminary version in *SODA '08: Proc. 19th Symp. Discrete Algorithms*.
- [4] S. Cabello and B. Mohar. Finding shortest non-separating and non-contractible cycles for topologically embedded graphs. *Discrete Comput. Geom.*, 37:213–235, 2007. Preliminary version in *ESA '05*.




Figure 5: Illustration for the local transformation in the proof of Theorem 4. Left: part of the original graph G_H and two curves $\alpha_{v(u)}$ and $\alpha_{u(v)}$. Right: part of the resulting surface Σ' after adding the handle corresponding to the edge uv of H .

- [5] E. W. Chambers, É. Colin de Verdière, J. Erickson, F. Lazarus, and K. Whittlesey. Splitting (complicated) surfaces is hard. *Computational Geometry: Theory & Applications*, 41:94–110, 2008. Preliminary version in SOCG '06.
- [6] É. Colin de Verdière and J. Erickson. Tightening non-simple paths and cycles on surfaces. In *SODA '06: Proc. 17th Symp. Discrete Algorithms*, pages 192–201, 2006.
- [7] É. Colin de Verdière and F. Lazarus. Optimal system of loops on an orientable surface. *Discrete Comput. Geom.*, 33:507–534, 2005. Preliminary version in FOCS '02.
- [8] É. Colin De Verdière and F. Lazarus. Optimal pants decompositions and shortest homotopic cycles on an orientable surface. *J. ACM*, 54(4):18, 2007. Preliminary version in Graph Drawing 2004.
- [9] D. Eppstein. Dynamic generators of topologically embedded graphs. In *SODA '03: Proc. 14th Symp. Discrete Algorithms*, pages 599–608, 2003.
- [10] J. Erickson and S. Har-Peled. Optimally cutting a surface into a disk. *Discrete Comput. Geom.*, 31:37–59, 2004. Preliminary version in SOCG '02.
- [11] J. Erickson and K. Whittlesey. Greedy optimal homotopy and homology generators. In *SODA '03: Proc. 16th Symp. Discrete Algorithms*, pages 1038–1046, 2005.
- [12] H. N. Gabow, S. N. Maheswari, and L. J. Osterweil. On two problems in the generation of program test paths. *IEEE Trans. Software Eng.*, 2(3):227–231, 1976.
- [13] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.
- [14] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2001. Available at <http://www.math.cornell.edu/~hatcher/>.
- [15] M. Henzinger, P. Klein, S. Rao, and S. Subramanian. Faster shortest-path algorithms for planar graphs. *Journal of Computer and System Sciences*, 55(1):3–23, 1997.
- [16] V. Kann. Polynomially bounded minimization problems that are hard to approximate. *Nordic J. Comp.*, 1:317–331, 1994. Preliminary version in ICALP '93.
- [17] M. Kutz. Computing shortest non-trivial cycles on orientable surfaces of bounded genus in almost linear time. In *SOCG '06: Proc. 22nd Symp. Comput. Geom.*, pages 430–438, 2006.
- [18] W. S. Massey. *Algebraic Topology: An Introduction*. Springer Verlag, 1967.
- [19] B. Mohar and C. Thomassen. *Graphs on Surfaces*. Johns Hopkins University Press, Baltimore, 2001.
- [20] H. Ren and M. Deng. Short cycle structures for graphs on surfaces and an open problem of Mohar and Thomassen. *Science in China Series A: Mathematics*, 49(2):212–224, 2006.
- [21] J. Stillwell. *Classical Topology and Combinatorial Group Theory*. Springer-Verlag, New York, 1993.
- [22] C. Thomassen. Embeddings of graphs with no short noncontractible cycles. *J. Combin. Theory, Ser. B*, 48:155–177, 1990.
- [23] C. Thomassen. Recent results on graph embeddings. In Y. Alavi, G. Chartrand, O. R. Ollerman, and A. J. Schwenk, editors, *Graph Theory, Combinatorics, and Applications*, pages 1093–1103. John Wiley and Sons, 1991.