Finding one tight cycle*

Sergio Cabello† Matt DeVos‡ Jeff Erickson§ Bojan Mohar¶

Abstract

A cycle on a combinatorial surface is tight if it as short as possible in its (free) homotopy class. We describe an algorithm to compute a single tight, non-contractible, simple cycle on a given orientable combinatorial surface in $O(n \log n)$ time. The only method previously known for this problem was to compute the globally shortest non-contractible or non-separating cycle in $O(\min\{g^3, n\}\, n \log n)$ time, where g is the genus of the surface. As a consequence, we can compute the shortest cycle freely homotopic to a chosen boundary cycle in $O(n \log n)$ time and a tight octagonal decomposition in $O(gn \log n)$ time.

1 Introduction

Cutting along curves is the basic tool for working with topological surfaces. When the surface is equipped with a metric, the surgery is typically made along shortest non-trivial cycles, where non-trivial may mean non-contractible or (surface) non-separating, depending on the application. Here, we are interested in cycles with a different metric property: a cycle is *tight* if it is shortest in its free homotopy type. Note that a shortest non-trivial cycle is going to be tight, but the converse does not hold.

We are interested in the algorithmic aspects of finding a tight, non-trivial cycle. Like most previous algorithmical works concerning curves on surfaces [1, 2, 3, 4, 5, 6, 7, 8, 9, 14], we consider the *combinatorial surface* model. A combinatorial surface $\mathcal{M}$ is an edge-weighted multigraph G embedded on a surface, and only paths arising from walks in G are considered. The length of a path is the sum of the weights of its edges, counted with multiplicity. The complexity of a combinatorial surface, denoted by n, is the sum of the number of its vertices, edges, and faces.

The theory of graphs embedded on surfaces, a natural generalization of the theory of planar graphs, is a very active research area. See the monograph [15] for an introduction. Algorithmical aspects of topological graph theory are also playing an important role in several graph problems. See for example the recent linear-time algorithm of Kawarabayashi and Reed [13] for testing if a given graph has bounded crossing number.

The main result of this paper is an algorithm to compute a tight, surface non-separating cycle on an orientable combinatorial surface in $O(n \log n)$ time. The best previous solution to the problem of finding a tight, non-trivial cycle was to compute the globally shortest non-trivial cycle, which can be done in $O(n^2 \log n)$ time with an algorithm by Erickson and Har-Peled [8] or in $O(g^3 n \log n)$ time with an algorithm by Cabello and Chambers [1]. (See [2, 14] for other relevant results.)

This new algorithm has the following implications:

- In the approach of Colin de Verdière and Erickson [4] for finding shortest curves homotopic to a given one, the bottleneck of the preprocessing part was to find a tight, non-trivial cycle. With our result, we can speed up their preprocessing from $O(\min\{g^3, n\} n \log n)$ to $O(gn \log n)$.

- We can compute the shortest cycle homotopic to a given boundary component in $O(n \log n)$ time. The previous best algorithm [4] used $O(gn \log n)$ time.

- In topological graph theory, several of the proofs based on cutting along shortest non-trivial cycles carry out if instead we cut along a tight, non-trivial cycle. Thus, algorithmic counterparts of several basic theorems can be improved with our new result.

2 Background

Surfaces. We summarize some basic concepts of topology. See [12, 16] for a comprehensive treatment.

*S.C. was partially supported by the European Community Sixth Framework Programme under a Marie Curie Intra-European Fellowship, and by the Slovenian Research Agency, project J1-7218.

†Department of Mathematics, IMFM, and Department of Mathematics, FMF, University of Ljubljana, Slovenia, sergio.cabello@fmf.uni-lj.si

‡Department of Mathematics, Simon Fraser University, Burnaby, B.C., mdevos@sfu.ca

§Department of Computer Science, University of Illinois at Urbana-Champaign, jeffe@cs.uiuc.edu

¶Department of Mathematics, Simon Fraser University, and Department of Mathematics, IMFM, mohar@sfu.ca

A (topological) *surface* (or 2-manifold) Σ is a compact topological space where each point has a neighbourhood homeomorphic to the plane or to a closed halfplane. A boundary point in Σ is a point with the property that no neighbourhood is homeomorphic to the plane. The *boundary* of Σ is the union of all boundary points, and it is known to consist of a finite number (possibly 0) of connected components, each component homeomorphic to a circle. The surface is *non-orientable* if it contains a subset homeomorphic to the Möbius band, and *orientable* otherwise.

All surfaces considered in this paper are orientable, and we will regularly not mention the adjective 'orientable'. A (g, b)-surface is a sphere with a number $g \geq 0$ of handles attached to it, and a number $b \geq 0$ of open disks removed. Up to homeomorphism, any surface is a (g, b)-surface for a unique pair $g, b \geq 0$.

A *path* in Σ is a continuous mapping $p : [0, 1] \to \Sigma$, a *cycle* is a continuous mapping $\gamma : \mathbb{S}^1 \to \Sigma$, a *loop* with basepoint x is a path such that $x = p(0) = p(1)$, and an *arc* is a path whose endpoints are on the boundary. *Curve* is a generic term used for paths, cycles, arcs, and loops. A curve is *simple* when the mapping is injective, except for the common endpoint in the case of loops.

Two paths p, q with $p(0) = q(0)$ and $p(1) = q(1)$ are *homotopic* if there is a continuous function $h : [0, 1]^2 \to \Sigma$ such that $p(\cdot) = h(0, \cdot)$, $q(\cdot) = h(1, \cdot)$, $h(\cdot, 0) = p(0)$, and $h(\cdot, 1) = p(1)$. Two cycles α, β are *(freely) homotopic* if there is a continuous function $g : [0, 1] \times \mathbb{S}^1 :\to \Sigma$ such that $\alpha(\cdot) = g(0, \cdot)$ and $\beta(\cdot) = g(1, \cdot)$. Two arcs are *homotopic* if they are homotopic as paths after contracting the boundary components that contain its endpoints; the intuition is that we allow the endpoints to freely move along the same boundary component. Simple curves are typically identified with their image because, up to reversal of the parameterization, any two parameterizations with the same image correspond to homotopic curves.

A cycle is *contractible* if it is homotopic to the constant loop. Cutting along a simple contractible cycle gives two connected components, and one of them is a topological disk. A simple cycle α is *non-separating* if cutting the surface along (the image of) α gives rise to a unique connected component. Non-separating cycles are non-contractible, while contractible cycles are separating. Being contractible or separating is a property invariant under homotopy of cycles.

We use the notation Σ✄α to denote the surface obtained after cutting Σ along α. We denote by Σ✄$(\alpha_1, \ldots, \alpha_k)$ the surface obtained inductively as $(\Sigma$✄$(\alpha_1, \ldots, \alpha_{k-1})$✄$\alpha_k$.

Combinatorial surface. All our results will be phrased in the combinatorial surface model. This model is dual to the cross-metric surface model; see [4] for a discussion. A *combinatorial surface* $\mathcal{M}$ is a surface $\Sigma(\mathcal{M})$ together with a multigraph $G(\mathcal{M})$ embedded on Σ so that each face of G is a topological disk. The complexity of a combinatorial surface $\mathcal{M}$ is defined as the sum of the number of vertices, edges, and faces of $G(\mathcal{M})$. The genus and the number of boundary components of $\mathcal{M}$ are those of $\Sigma(\mathcal{M})$.

In the combinatorial surface model, we only consider curves that arise as paths in $G(\mathcal{M})$. A curve is said to be *homotopically simple* if there is an infinitesimal continuous perturbation that makes it simple. *All curves considered in this paper are homotopically simple*, and we will drop the adjective 'homotopically simple' in most cases. The *multiplicity* of a curve α is the maximum number of times that an edge appears in the graph-walk that defines α.

We assume that the graph $G(\mathcal{M})$ has edge-weights, which gives a "metric" to the model. The length $|\alpha|$ of a curve α is defined as the sum of the weights of the edges in the graph-walk that defines α, counted with multiplicity. A cycle or an arc is *tight* if it is shortest in its homotopy class.

Families of curves. We say that two cycles α, β *include a bigon* if there are simple subpaths $p_\alpha \subseteq \alpha$ and $p_\beta \subseteq \beta$ with common endpoints such that p_α and p_β bound a topological disk.

A *tight system of disjoint arcs* in a combinatorial surface with boundary is a family of simple curves $\alpha_1, \alpha_2, \ldots, \alpha_k$ such that

- no two distinct arcs α_i, α_j share an edge or cross;
- the arc α_i is a tight arc in $\mathcal{M}$✄$(\alpha_1, \ldots, \alpha_{i-1})$.

If $\mathcal{M}$ is a surface of complexity n and $\alpha_1, \alpha_2, \ldots, \alpha_k$ is a tight system of disjoint arcs, then the complexity of $\mathcal{M}$✄$(\alpha_1, \ldots, \alpha_k)$ is at most $2n$ because each edge gets doubled at most once during the cutting.

In the cross-metric model, one can also consider the concept of arrangements of families of curves. We refer the reader to [4], where also the concept of *tight octagonal decomposition* is introduced.

3 Toolbox

We next list results that will be used in our proofs and algorithms.

Lemma 1. ([11]) *Given two homotopic cycles α, β in an orientable surface, if they have some common point, then they include a bigon.*

Lemma 2. ([8]) *For any given basepoint x, we can find in $O(n \log n)$ time a shortest non-separating loop with basepoint x.*

Lemma 3. *Let $\mathcal{M}$ be a surface with at least two boundary components, let σ be one of its boundary components, and let $\alpha_1, \ldots, \alpha_k$ be a tight system of disjoint arcs where each α_i does not have endpoints in σ. Then every tight cycle homotopic to σ in $\mathcal{M}\text{✂}(\alpha_1, \cdots, \alpha_k)$ is also a tight cycle homotopic to σ in $\mathcal{M}$.*

Proof. The proof is by induction on k. There is nothing to prove if $k = 0$. The induction step is the same as the proof when $k = 1$, which we assume henceforth.

Let α be a tight cycle in $\mathcal{M}$ homotopic to σ. Then σ and α bound a cylinder D in $\mathcal{M}$. We choose α such that D is smallest possible, i.e., no other tight cycle bounds a cylinder which is contained in D. We will show that α_1 is disjoint from the interior of D, which will then imply that α is also a tight cycle in $\mathcal{M}' = \mathcal{M}\text{✂}\alpha_1$ homotopic to σ. To see this, suppose that α_1 enters D. If $\alpha_1 \cap D$ contains a simple arc α_1' whose endpoints x, y are on α, then α_1 and α include a bigon. Because of tightness of α and α_1, both segments of this bigon have the same length, and we can replace the segment of α with α_1'. The new curve is homotopic to σ and contradicts the minimality of D. This completes the proof.

The proof of the previous lemma also shows the following.

Lemma 4. *Let $\mathcal{M}$ be a surface with at least one boundary component, let σ be one of its boundary components, and let α be a simple tight cycle in $\mathcal{M}$. Then every tight cycle homotopic to σ in $\mathcal{M}\text{✂}\alpha$ is also a tight cycle homotopic to σ in $\mathcal{M}$.*

Lemma 5. *Let $\mathcal{M}$ be a combinatorial surface of complexity n, genus g, and exactly one boundary component σ. We can find in $O(n \log n)$ time a tight system of disjoint arcs $\alpha_1, \ldots, \alpha_{2g}$ such that $\mathcal{M}\text{✂}(\alpha_1, \ldots, \alpha_{2g})$ is a topological disk.*

Proof. Contract σ to a point p_σ, and construct a greedy system of loops $\alpha_1, \ldots, \alpha_{2g}$ at p_σ in $O(n \log n)$ time, as explained by Erickson and Whittlesey [9]. Unmaking the contraction, the curves $\alpha_1, \ldots, \alpha_{2g}$ become arcs in $\mathcal{M}$ with endpoints at σ, and cutting along them the surface becomes a topological disk. It follows from the greediness of the construction that each α_i is tight in $\mathcal{M}\text{✂}(\alpha_1, \ldots, \alpha_{i-1})$; see [9]. An edge could appear in several curves α_i, but we assign it to the one with smallest index i where it appears, and remove it from the rest. This can be done in $O(n)$ from the implicit representation of the greedy system of loops provided by Erickson and Whittlesey [9].

4 Finding one tight cycle

Let $\mathcal{M}$ be a combinatorial surface of complexity n.

Lemma 6. *Suppose that $\mathcal{M}$ has $b \geq 2$ boundary components and let σ be one of its boundary cycles. We can find in $O(n \log n)$ time a tight cycle homotopic to σ.*

Proof. Assume first that $b = 2$, and let σ' be the boundary component distinct from σ. Glue a disk over σ, and construct a tight system of disjoint arcs $\alpha_1, \ldots, \alpha_{2g}$ as described in Lemma 5. Cutting the surface $\mathcal{M}$ along $\alpha_1, \ldots, \alpha_{2g}$ leaves an annulus $\mathcal{A}$ whose boundary components are σ and σ'. Note that $\mathcal{A}$ has linear complexity because we cut $\mathcal{M}$ along a tight system of disjoint arcs. Furthermore, it follows from Lemma 3 that a tight cycle homotopic to σ in $\mathcal{A}$ is a tight cycle homotopic to σ in $\mathcal{M}$. Finally, the shortest generating cycle in $\mathcal{A}$ can be computed in $O(n \log n)$ time using the algorithm by Frederickson [10] because $\mathcal{A}$ has linear complexity. This concludes the case when $b = 2$.

The case when $b > 2$ can be reduced to $b = 2$ as follows. Let $\sigma_1, \ldots, \sigma_{b-1}$ be the boundary cycles different from σ. We contract each σ_i to a point p_i, and find a shortest path tree T from p_1. This can be done in $O(n \log n)$ time. Let π_i denote the shortest path from p_1 to p_i contained in T. Each edge can appear in several paths π_i, but we proceed like in the proof of Lemma 5: we assign each edge to the path p_i with smallest index that contains it, and delete it from the rest. Let $q_2, \ldots, q_{b-1}$ be the paths that are obtained. The curves $q_2, \ldots, q_{b-1}$, in the original surface $\mathcal{M}$, form a tight system of disjoint arcs. Therefore, a tight cycle homotopic to σ in $\mathcal{M}' = \mathcal{M}\text{✂}(q_2, \ldots, q_{b-1})$ is a tight cycle homotopic to σ in $\mathcal{M}$ because of Lemma 3. Note that $\mathcal{M}'$ has complexity $O(n)$ because it is obtained from $\mathcal{M}$ by cutting along a tight system of disjoint arcs. Since $\mathcal{M}'$ has two boundary components, the result follows from the previous case.

Lemma 7. *Let ℓ_x be a shortest non-separating loop based at x, and let γ_1 be a tight cycle homotopic to ℓ_x in $\mathcal{M}\text{✂}\ell_x$. Then the cycle γ_1 is tight in $\mathcal{M}$ as well.*

Proof. We will show that in $\mathcal{M}$ there is a cycle γ that is homotopic to ℓ_x, it is tight, and does not cross ℓ_x. Since γ does not cross ℓ_x, then γ is also homotopic to one of the copies of ℓ_x in $\mathcal{M}\text{✂}\ell_x$, and it is tight. Therefore $|\gamma_1| = |\gamma|$, so γ_1 is tight in $\mathcal{M}$.

Let γ be a shortest cycle that is homotopic to ℓ_x (in $\mathcal{M}$) and crosses ℓ_x as few times as possible. We want to show that γ and ℓ_x do not cross. Assume for contradiction that γ and ℓ_x cross. Then, by Lemma 1, they include a bigon. Let $\pi_\gamma \subseteq \gamma$ and π_{ℓ_x} be the

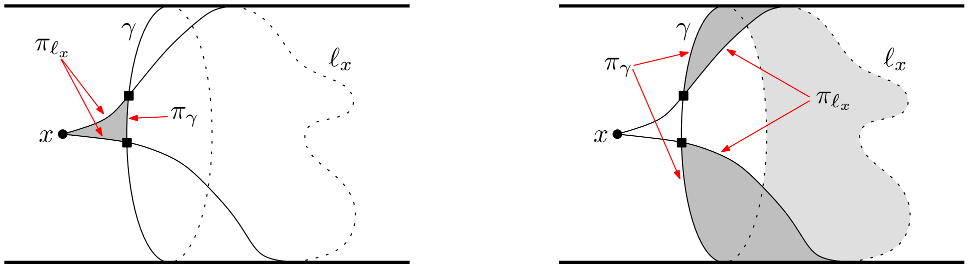

Figure 1: Figure for Lemma 7. The grey region represents a bigon between ℓ_x and γ. Left: the case $x \in \pi_{\ell_x}$. Right: the case $x \in \pi_\gamma$.

two subpaths that enclose the bigon; π_γ and π_{ℓ_x} are homotopic paths; see Figure 1. Let q_γ be the subpath $\gamma \setminus \pi_\gamma$ and let q_{ℓ_x} be the subpath $\ell_x \setminus \pi_{\ell_x}$. We distinguish two cases:

π_{ℓ_x} **contains** x**.** Let δ be the cycle π_γ concatenated with q_{ℓ_x}. Note that δ crosses ℓ_x twice less than γ does. Since π_γ and π_{ℓ_x} are homotopic, δ is homotopic to ℓ_x. Since π_{ℓ_x} concatenated with q_γ is a non-separating cycle through x, it holds that

$$|\ell_x| = |\pi_{\ell_x}| + |q_{\ell_x}| \leq |\pi_{\ell_x}| + |q_\gamma|,$$

which implies $|q_{\ell_x}| \leq |q_\gamma|$. We conclude that

$$|\delta| = |\pi_\gamma| + |q_{\ell_x}| \leq |\pi_\gamma| + |q_\gamma| = |\gamma|,$$

and since δ crosses ℓ_x twice less than γ, we get a contradiction.

π_{ℓ_x} **does not contain** x**.** Let δ be the cycle π_{ℓ_x} concatenated with q_γ. Note that δ crosses ℓ_x twice less than γ does. Since π_γ and π_{ℓ_x} are homotopic, δ is homotopic to γ and ℓ_x. Since q_{ℓ_x} concatenated with π_γ is a non-separating cycle through x, it holds that

$$|\ell_x| = |q_{\ell_x}| + |\pi_{\ell_x}| \leq |q_{\ell_x}| + |\pi_\gamma|,$$

which implies $|\pi_{\ell_x}| \leq |\pi_\gamma|$. We conclude that

$$|\delta| = |\pi_{\ell_x}| + |q_\gamma| \leq |\pi_\gamma| + |q_\gamma| \leq |\gamma|,$$

and since δ crosses ℓ_x twice less than γ, we get a contradiction.

Theorem 1. *Let $\mathcal{M}$ be an orientable combinatorial surface of complexity n. We can find in $O(n \log n)$ time a tight cycle that is (homotopically) simple and surface non-separating.*

Proof. Choose a point $x \in \mathcal{M}$, and construct a shortest non-separating loop ℓ_x with basepoint x. Since $\mathcal{M}$ is an orientable surface, $\mathcal{M}' = \mathcal{M} \,\text{✂}\, \ell_x$ has two boundary components ℓ'_x and ℓ''_x arising from ℓ_x. We find γ', a tight cycle homotopic to ℓ'_x in $\mathcal{M} \,\text{✂}\, \ell_x$, and γ'', a tight cycle homotopic to ℓ''_x in $\mathcal{M} \,\text{✂}\, \ell_x$, and return the shorter cycle among γ', γ''. This finishes the description of the algorithm.

The cycle γ_{min} returned by the algorithm is tight because of Lemma 7. Since the cycle γ_{min} is homotopic to the simple, non-separating loop ℓ_x in $\mathcal{M}$, it follows that γ_{min} is also non-separating and simple. As for the running time, note that ℓ_x can be found in $O(n \log n)$ time because of Lemma 2, and the cycles γ', γ'' can also be obtained in $O(n \log n)$ time using Lemma 6 because $\mathcal{M} \,\text{✂}\, \ell_x$ has at least two boundary components.

5 Consequences and conclusions

Theorem 2. *Let $\mathcal{M}$ be an orientable combinatorial surface of complexity n, and let σ be a given boundary cycle in $\mathcal{M}$. We can find in $O(n \log n)$ time a tight cycle that is homotopic to σ.*

Proof. If $\mathcal{M}$ has more than two boundary components, the result follows from Lemma 6. If σ is the only boundary component, we compute a tight non-separating, simple cycle γ using Theorem 1, and then find a tight cycle $\tilde{\sigma}$ homotopic to σ in $\mathcal{M} \,\text{✂}\, \gamma$. Finally, we return the cycle $\tilde{\sigma}$. This finishes the description of the algorithm.

The algorithm is correct because the returned cycle $\tilde{\sigma}$ is homotopic to σ and is tight because of Lemma 4. As for the running time, γ is obtained in $O(n \log n)$ time because of Theorem 1, and $\tilde{\sigma}$ is obtained in $O(n \log n)$ time using Lemma 6 because $\mathcal{M} \,\text{✂}\, \gamma$ has precisely three boundary components.

We can find a tight octagonal decomposition of a surface $\mathcal{M}$ without boundary in $O(gn \log n)$ time, improving the previous $O(n^2 \log n)$ time bound by Colin de Verdière and Erickson [4]. This also improves the preprocessing time in their results.

Theorem 3. *Let $\mathcal{M}$ be an orientable cross-metric surface with complexity n, genus $g \geq 2$, and no boundary.*

We can construct a tight octagonal decomposition of $\mathcal{M}$ in $O(gn \log n)$ time.

Proof. Consider the construction described in Theorem 4.1 of [4]. Their first step is to find a tight cycle in $\mathcal{M}$, which they implement finding a globally shortest non-separating cycle in $O(n^2 \log n)$ time. (Finding this cycle can be done in $O(g^3 n \log n)$ time using the more recent result of Cabello and Chambers [1].) Using Theorem 1, we can now perform this first step in $O(n \log n)$ time. After this, the rest of their construction takes $O(gn \log n)$ time, and the result follows.

THEOREM 4. *Let $\mathcal{M}$ be an orientable combinatorial surface with complexity n, genus $g \geq 2$, and no boundary. Let p be a path on $\mathcal{M}$, represented as a walk in $G(\mathcal{M})$ with complexity k. We can compute a shortest path p' homotopic to p with complexity k' in $O(gn \log n + gk + gn\bar{k})$ time, where $\bar{k} = \min\{k, k'\}$.*
For a cycle γ, we can do the same in $O(gn \log n + gk + gn\bar{k} \log(n\bar{k}))$ time.

Proof. The preprocessing time for constructing a tight octagonal decomposition has gone down from $O(n^2 \log n)$ to $O(gn \log n)$ because of the previous result. The result follows then from the algorithms in [4].

References

[1] 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.

[2] 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.

[3] E. W. Chambers, É. Colin de Verdière, J. Erickson, F. Lazarus, and K. Whittlesey. Splitting (complicated) surfaces is hard. In *SOCG '06: Proc. 22nd Symp. Comput. Geom.*, pages 421–429, 2006.

[4] 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.

[5] É. Colin de Verdière and F. Lazarus. Optimal pants decompositions and shortest homotopic cycles on an orientable surface. In *11th International Symposium on Graph Drawing, Perugia, Italy*, volume 2912 of *LNCS*, pages 478–490. Springer, 2004.

[6] 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.

[7] D. Eppstein. Dynamic generators of topologically embedded graphs. In *SODA '03: Proc. 14th Symp. Discrete Algorithms*, pages 599–608, 2003.

[8] 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.

[9] J. Erickson and K. Whittlesey. Greedy optimal homotopy and homology generators. In *SODA '03: Proc. 16th Symp. Discrete Algorithms*, pages 1038–1046, 2005.

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

[11] H. Hass and P. Scott. Intersections of curves on surfaces. *Israel J. Math.*, 51:90–120, 1985.

[12] A. Hatcher. *Algebraic Topology.* Cambridge University Press, 2001. Available at `http://www.math.cornell.edu/~hatcher/`.

[13] K. Kawarabayashi and B. Reed. Computing crossing number in linear time. In *STOC '07: Proc. 39th ACM Symp. Theory Comput.*, pages 382–390, 2007.

[14] 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.

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

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