

23 Shortest Interesting Cycles

In this lecture, I'll describe algorithms to find the shortest cycle in a surface map that is topologically interesting. I'll specifically consider two different definitions of “interesting”:

- A cycle is *noncontractible* if it cannot be continuously deformed to a single point; noncontractible cycles are *homotopically* nontrivial.
- A cycle is *nonseparating* if slicing along the cycle does not disconnect the surface; nonseparating cycles are *homologically* nontrivial.

The input to our algorithms is a surface map Σ with positively weighted edges.¹ I will assume throughout this presentation that there is a unique shortest path between any pair of vertices; this assumption can be enforced with standard perturbation schemes. Crucially, we will treat the underlying graph of Σ as a continuous topologically space; any edge with weight ℓ is isometric to the interval $[0, \ell]$ on the real line. Thus, we can reasonably consider shortest paths between points that lie in the interior of edges.

23.1 Properties of Shortest Nontrivial Cycles

It is unfortunately common for papers on surface-map algorithms to use the word “cycle” in two different senses. For topologists, a *cycle* is the continuous image of the circle S^1 , but for graph theorists, a *cycle* is a closed walk with no repeated vertices. That is, graph theorists assume that cycles are *simple* (or *injective*), but topologists do not. Fortunately for us, shortest nontrivial cycles are the same for both tribes.

Lemma: *The shortest noncontractible (or nonseparating) closed walk in a positively edge-weighted surface map is a simple cycle.*

Proof: For the sake of argument, suppose the shortest noncontractible closed walk W is *not* simple. We can decompose W into two closed walks $X \cdot Y$ at any vertex that W visits more than once. X and Y are both shorter than W , so they must both be contractible. The concatenation of two contractible closed walks is contractible. So W is contractible, contradicting its definition. Essentially the same argument applies to the shortest nonseparating closed walk. \square

3-Path Condition (Carsten Thomassen): *Let x and y be two points (either at vertices or in edge interiors) in a surface map Σ , and let α, β, γ be paths in Σ from x to y . If the cycles $\alpha \cdot \text{rev}(\beta)$ and $\beta \cdot \text{rev}(\gamma)$ are contractible (resp. separating), then the cycle $\alpha \cdot \text{rev}(\gamma)$ is also contractible (resp. separating).*

Proof: The concatenation of any two contractible closed walks is contractible.

The symmetric difference of any two separating cycles is separating. \square

Antipodality Condition (Carsten Thomassen): *Let σ be the shortest noncontractible (resp. nonseparating) cycle in a surface map Σ . Any pair of antipodal points partition σ into two equal-length shortest paths.*

Proof: Fix a pair x and \bar{x} of antipodal points in σ . These points clearly partition σ into two equal-length paths; call these paths α and β . Suppose α and β are *not* shortest paths, and let γ be a shortest path from x to \bar{x} . The cycles $\alpha \cdot \text{rev}(\gamma)$ and $\beta \cdot \text{rev}(\gamma)$ are both shorter than

¹With more effort, these algorithms can be generalized to *directed* surface maps with asymmetrically weighted darts.

$\alpha \cdot \text{rev}(\beta) = \sigma$ and thus are contractible (resp. separating). But then the 3-path property implies that σ is contractible (resp. separating), contradicting its definition. \square

Crossing Condition: *Any shortest path crosses the shortest noncontractible (resp. nonseparating) cycle at most once.*

Proof: Let π be a shortest path and let σ be a simple noncontractible (resp. separating) cycle, and suppose the intersection $\pi \cap \sigma$ is disconnected. Then we can decompose π into three subpaths $\pi^- \cdot \pi^\circ \cdot \pi^+$, where π° starts and ends at vertices of σ but is otherwise disjoint from σ . Because π is a shortest path, its subpath $\pi^c i r c$ is also a shortest path. Similarly, we can decompose σ into two paths α and β at the endpoints of π° . The 3-path condition implies that at least one of the cycles $\alpha \cdot \text{rev}(\pi^\circ)$ and $\beta \cdot \text{rev}(\pi^\circ)$ is noncontractible (resp. nonseparating). Because vertex-to-vertex shortest paths are unique, both of these cycles are shorter than σ . \square

23.2 A polynomial-time algorithm

Thomassen's 3-path condition implies the following algorithm to find the shortest noncontractible cycle σ . The antipodality condition implies that σ must consist of a shortest path from a vertex x to another vertex y , the edge yz , and the shortest path from z back to x . (Uniqueness of vertex-to-vertex shortest paths implies that the antipodal point \bar{x} lies in the interior of the edge yz .) There are only $O(n^2)$ loops of this form. We compute the distance between every pair of vertices in $O(n^2 \log n)$ time by running Dijkstra's algorithm n times, after which we can compute the length of each candidate cycle in $O(1)$ time. Finally, for each candidate loop γ , we test whether γ is contractible in $O(n)$ time, using Dehn's algorithm. Finally, we return the shortest candidate cycle that is noncontractible.

Theorem [Thomassen]: *Given any edge-weighted surface map Σ with complexity n , we can compute the shortest noncontractible cycle in Σ in $O(n^3)$ time.*

In fact, Dehn's algorithm is overkill for testing the contractibility of simple cycles. Instead we can rely on the classical lemma:

Lemma [David Epstein²]: *A nontrivial simple cycle σ on any surface S is contractible if and only if $S \setminus \sigma$ is disconnected and at least one component is an open disk.*

We can test whether a simple cycle σ is the boundary of a disk as follows. First, we perform a whatever-first search in the dual map Σ^* to determine whether the sliced map $\Sigma \setminus \sigma$ (or alternatively, the contracted map Σ/σ) is connected. If $\Sigma \setminus \sigma$ is connected, then σ is noncontractible. Otherwise, we compute the Euler characteristics and orientability of both components in $O(n)$ time. If either component is orientable with genus 1, that component is a disk and σ is contractible; otherwise, σ is non-contractible.

We can find the shortest *nonseparating* cycle in the same time bound, using a slightly *simpler* algorithm. A given simple cycle σ is non-separating if and only if $\Sigma \setminus \sigma$ is connected; we can test this condition in $O(n)$ time using whatever-first search in the dual map Σ^* . Otherwise, the algorithm is the same.

²David Epstein-with-one-p is a curly-haired English mathematician who was born in South Africa. David Eppstein-with-two-p's is a curly-haired American mathematician (and computer scientist) who was born in England. It's a minor miracle that they have never been coauthors.

Theorem [Thomassen]: *Given any edge-weighted surface map Σ with complexity n , we can compute the shortest nonseparating cycle in Σ in $O(n^3)$ time.*

23.3 Near-quadratic time

- Build a tree-cotree decomposition (T, L, C) where T is a shortest-path tree, in $O(n \log n)$ time, using Dijkstra's algorithm.
- Let X^* be the *dual* cut graph $(C \cup L)^*$. A fundamental loop $\text{loop}(x, yz)$ is separating if and only if $(yz)^*$ is a bridge of X^* , and contractible if and only if at least one component of $X^* \setminus (yz)^*$ is a tree.
- We can classify every edge of X^* as hair, bridge, or neither in $O(n)$ time. So we can find shortest noncontractible and nonseparating loops based at x in $O(n \log n)$ time.
- Try all basepoints $\implies O(n^2 \log n)$ time.

The high-level approach of this algorithm is due to Sarel Har-Peled and me, but we used a much more complicated algorithm to find shortest interesting loops, which interleaves Dijkstra steps and brute-force traversal of the dual map. Replacing Dijkstra's algorithm with a linear-time shortest-path algorithm reduces the overall running time to $O(n^2)$.

This is the fastest algorithm known for arbitrary genus. However, Martin Kutz proved that for any *constant* genus g , it is possible to compute shortest nontrivial cycles in only $O(n \log n)$ time, by a reduction to the planar minimum-cut problem (or more properly, to the dual problem of finding the shortest generating cycle in an annulus). The running time of Kutz's algorithm depends exponentially on the genus. Later work by Sergio Cabello and Erin Chambers reduced the dependence on g to a small polynomial.

23.4 Multiple-Source Shortest Paths

We can significantly improve Thomassen's algorithm using a generalization of either of the multiple-source shortest path algorithms that we saw earlier for planar maps.

[The rest of this section is only an outline.]

Parametric

Grove decomposition of the dual cut graph $T \cup L$ into $O(g)$ trees, each with a central cut path. Each pivot requires $O(g)$ dynamic-forest operations, and thus takes $O(g \log n)$ amortised time. Each dart pivots into the shortest path tree T at most $O(g)$ times. Total time is $O(g^2 n \log n)$, assuming generic edge weights. This can be improved to $O(g n \log n)$ time with more careful data structures and analysis.

Surface-Tree Lemma: *Let T be any tree embedded on a surface with exactly one boundary cycle B ; call any vertex in $T \cap B$ a boundary vertex. Let e be any edge of T , and let U and W be the components of $T \setminus e$. Either U contains every boundary vertex, or boundary vertices in U induce at most $g + 1$ paths in B .*

Proof(?): Let Σ be the surface map induced by $T \cup B$, and let o be the face of Σ bounded by B . Let (T, L, C) be any tree-cotree decomposition of Σ with the given spanning tree T .




Figure 1: Pivoting one edge into a grove decomposition.

Consider the subgraph $X^* = C^* \cup L^* \cup \{e^*\}$ of the dual map Σ^* . The induced embedding of X^* has exactly two faces, each containing (the faces of Σ^* dual to vertices of) one component of $T \setminus e$. Let Y^* be the boundary subgraph comprised of the non-isthmus edges of X^* . This boundary subgraph can be decomposed into at most $g + 1$ simple non-crossing cycles. (Recall the the *definition* of the genus G is the maximum number of disjoint cycles whose deletion does not disconnect the surface.) Each of these cycles crosses B at most twice. Thus, either X^* does not cross B at all, or X^* splits B into at most $2g + 2$ intervals, which alternate between vertices of U and vertices of W . \square

Recursive

Precompute homology annotations (via a system of cocycles) once at the start. Compute shortest-path trees T_i and T_j in $O(n \log n)$ time. Compute homology signatures of all shortest paths from s_i and s_k in $O(gn)$ time. Dart $u \rightarrow v$ is properly shared by T_i and T_k if it satisfies the following conditions:

- $\text{pred}_i(v) = \text{pred}_k(v) = u$
- $[\text{path}_i(u)] = [\text{path}_k(u)]$ — **This condition is new!**
- If $t = \text{pred}_i(u) = \text{pred}_k(u)$, then $t \rightarrow u$ is properly shared by T_i and T_k .
- Otherwise, darts $\text{pred}_i(u) \rightarrow u$ and $v \rightarrow u$ and $\text{pred}_k(u) \rightarrow u$ are oriented clockwise around u .

The second condition holds if and only if these two shortest paths define a separating arc from s_i through u to s_k .

Finding and contracting all properly shared darts takes $O(gn)$ time. Each vertex has $O(g \deg(v))$ interesting sources, so the total size of all maps at each level of recursion is $O(gn)$. The total preprocessing time is $O(gn \log n + g^2 n)$; the query time is still $O(\log n)$. It's unclear whether $O(gn \log n)$ time can be achieved with this approach.

23.5 Shortest Nonseparating Cycles in Near-Linear Time

Lemma [Cabello Mohar]: *The shortest nonseparating cycle in any surface map Σ crosses at least one cycle in a greedy system of cycles for Σ exactly once.*

Proof: *[[[to be written. The usual exchange argument implies 3 crossings is impossible, and every nonseparating cycle crosses some basis cycle an odd number of times]]]*

Lemma: *For any cycle γ , the shortest cycle that crosses γ exactly once can be computed in $O(\bar{g}^2 n \log n)$ time.*

Proof: Let $\Sigma' = \Sigma \setminus \gamma$. If γ is two-sided, then Σ' has two boundary cycles γ^- and γ^+ , each of which is a copy of γ . If γ is one-sided, then Σ' has a single boundary cycle γ^\pm that covers γ twice. In either case, Σ' contains two copies v^- and v^+ of every vertex v of γ .

Now let σ be the shortest cycle in Σ that crosses γ once. This cycle appears in Σ' as a shortest path from v^- to v^+ for some vertex v of Σ . We can compute all such shortest paths in $O(\bar{g}^2 n \log n)$ time using either MSSP algorithm, using either γ^+ or γ^\pm as the “outer” face. \square

Theorem: *The shortest nonseparating cycle in a surface map with Euler genus \bar{g} and complexity n can be computed in $O(\bar{g}^3 n \log n)$ time.*

Proof: Let (T, L, C) be a tree-cotree decomposition of Σ where T is a shortest-path tree. Let Q be the reduced cur graph obtained by removing degree-1 vertices from $T \cup L$, and let $\mathcal{C} = \{\gamma_1, \gamma_2, \dots, \gamma_{\bar{g}}\}$ be the system of cycles induced by T and L .

For each cycle γ_i , we compute the shortest cycle in Σ that crosses γ_i exactly once, in $O(\bar{g} n \log n)$ time, by running a multiple-source shortest-path algorithm in $\Sigma \setminus \gamma_i$. The shortest of these \bar{g} cycles is the shortest nonseparating cycle. \square

23.6 Shortest Noncontractible Cycles in Near-Linear Time (sketch)

This one is a bit more complicated.

Lemma: *Every noncontractible cycle in any surface map Σ crosses any cut graph of Σ at least once.*

The following lemma distills a more complex argument of Cabello, Chambers, and Erickson. **[[[Take this with a grain of salt until I wrote down a complete proof.]]]**

Lemma: *Let σ be the shortest noncontractible cycle in some surface map Σ . Let (T, L, C) be a tree-cotree decomposition of Σ where T is a shortest path tree and C is a maximum spanning tree. Let \mathcal{C} denote the corresponding system of cycles.*

- (a) σ crosses each cycle in \mathcal{C} at most once.
- (b) σ is a nonseparating cycle if and only if σ crosses some cycle in \mathcal{C} at least once.
- (c) If σ is a separating cycle, then σ is also the shortest non-contractible cycle in $\Sigma' = \Sigma \setminus \mathcal{C}$.

Proof: **[[[To be written. Notice that (a) is a stronger claim than the nonseparating crossing lemma!]]]**

Lemma: *Let σ be the shortest noncontractible cycle in some surface map Σ' with genus 0 and at least two boundaries. Let π be the shortest path between any two boundaries of Σ' . Either (a) σ crosses π exactly once, or (b) σ does not cross π and σ is the shortest noncontractible cycle in $\Sigma' \setminus \pi$.*

Proof: **[[[Mostly follows from the Crossing Condition.]]]**

Theorem: *The shortest noncontractible cycle in a surface map with Euler genus \bar{g} and complexity n can be computed in $O(\bar{g}^3 n \log n)$.*

Proof: Let (T, L, C) be a tree-cotree decomposition of Σ where T is a shortest-path tree and C is a maximum spanning tree. Let $\mathcal{C} = \{\gamma_1, \gamma_2, \dots, \gamma_{\bar{g}}\}$ be the system of cycles induced by T and L .

First, for each cycle γ_i , we compute the shortest cycle in Σ that crosses γ_i exactly once, in $O(\bar{g}^2 n \log n)$ time, by running a multiple-source shortest-path algorithm in $\Sigma \setminus \gamma_i$.

The shortest of these \bar{g} cycles is the shortest nonseparating cycle. Thus, if the shortest noncontractible cycle is nonseparating, then the shortest of these \bar{g} cycles is the shortest noncontractible cycle. Otherwise, we need to compute the shortest noncontractible cycle in the sliced surface $\Sigma' = \Sigma \setminus \mathcal{C}$.

The surface Σ' has genus 0 and b boundary cycles, for some $1 \leq b \leq \bar{g}$. If $b = 1$, then Σ' is a disk, so every cycle in Σ' is contractible. Otherwise, we compute a shortest path π between any two boundary cycles, compute the shortest cycle σ^\times in Σ' that crosses π exactly once via multiple-source shortest paths, and recursively search $\Sigma' \setminus \pi$. The algorithm halts after computing $b - 1$ cycles; the shortest of these is the shortest noncontractible cycle in Σ' .

Altogether the algorithm runs multiple-source shortest paths $\bar{g} + b = O(\bar{g})$ times, each on a surface map with complexity $O(n)$ and genus less than \bar{g} . So the overall running time is $O(\bar{g}^3 n \log n)$, as claimed. \square

23.7 References

1. Sergio Cabello, Erin W. Chambers, and Jeff Erickson. Multiple-source shortest paths in embedded graphs. *SIAM J. Comput.* 42(4):1542–1571, 2013. arXiv:1202.0314.
2. Sergio Cabello and Bojan Mohar. Finding shortest non-separating and non-contractible cycles for topologically embedded graphs. *Discrete Comput. Geom.* 37(2):213–235, 2007.
3. David B. A. Epstein. Curves on 2-manifolds and isotopies. *Acta Math.* 115:83–107, 1966.
4. Jeff Erickson. Shortest non-trivial cycles in directed surface graphs. *Proc. 27th Ann. Symp. Comput. Geom.*, 236–243, 2011.
5. Jeff Erickson, Emily Kyle Fox, and Luvsandondov Lkhamsuren. Holiest minimum-cost paths and flows in surface graphs. *Proc. 50th Ann. ACM Symp. Theory Comput.*, 1319–1332, 2018. arXiv:1804.01045.
6. Jeff Erickson and Sarel Har-Peled. Optimally cutting a surface into a disk. *Discrete Comput. Geom.* 31(1):37–59, 2004. arXiv:cs/0207004.
7. Emily Kyle Fox. Shortest non-trivial cycles in directed and undirected surface graphs. *Proc. 24th Ann. ACM-SIAM Symp. Discrete Algorithms*, 352–364, 2013. arXiv:1111.6990.
8. Martin Kutz. Computing shortest non-trivial cycles on orientable surfaces of bounded genus in almost linear time. *Proc. 22nd Ann. Symp. Comput. Geom.*, 430–438, 2006. arXiv:cs/0512064.
9. Carsten Thomassen. Embeddings of graphs with no short noncontractible cycles. *J. Comb. Theory Ser. B* 48(2):155–177, 1990.

23.8 Aptly Named Sir

- Enforcing shortest-path and optimal-cycle uniqueness
- Tight cycles
- Directed graphs
- Shortest systems of loops, arcs, or cycles
- Shortest *homotopic* paths, cycles, systems of loops, pants decompositions