covering map

CHAPTER 4

Covering Spaces

Oh! marvellous, O stupendous Necessity—by thy laws thou dost compel every effect to be the direct result of its cause, by the shortest path. These [indeed] are miracles...

— Leonardo da Vinci, *Codex Atlanticus* (c. 1500)
translated by Jean Paul Richter (1883)

Those who cannot remember the past are condemned to repeat it.

— George Santayana, *Reason in Common Sense* (1905)

A straight line may be the shortest distance between two points,
but it is by no means the most interesting.

— The Doctor [Jon Pertwee], *The Time Warrior* (1973)

Motivation, problem statement. **Merge into previous chapter?**

The important observation is missing: Polygons with holes look like graphs. The universal cover of a graph is the infinite tree that algorithms search when they forget to mark nodes as visited. So any algorithm to find X in a tree/polygon automatically finds "homotopic X" in any graph/holey polygon, as long as you forget to remember (or remember to forget) where you're going.

4.1 Definitions

A ***covering map*** is a continuous surjection $p\colon \widehat{X} \twoheadrightarrow X$ such that any point $x \in X$ has an open neighborhood U whose preimage $p^{-1}(U)$ can be written as the union of disjoint

sheets
hat X
covering space
simply connected
universal cover
tilde X

open sets $\bigsqcup_{i\in I} U_i$, and the restriction of p to each open set U_i is a homeomorphism to U. The open sets U_i are sometimes called ***sheets*** over U. If there is a covering map from a space $\widehat{X}$ to another space X, we call $\boldsymbol{\widehat{X}}$ a ***covering space*** of X. By convention, we require covering spaces to be connected. As a trivial example, every space is a covering space of itself, with the identity function (or any other homeomorphism) as the covering map.

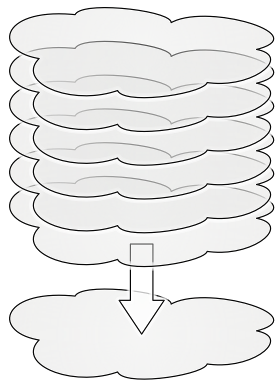

The local behavior of every covering map.

As an elementary example, suppose X is the "figure eight" graph with one vertex x and two edges r and b. Every covering space $\widehat{X}$ of X is a connected graph in which every vertex has degree 4. There are an infinite number of such graphs, some finite and others infinite. Suppose we arbitrarily orient the edges r and b, color edge r red, and color edge b blue. Then we can visualize any covering map from $\widehat{X}$ to X by orienting and coloring the edges of $\widehat{X}$ so that each vertex is incident to one incoming red edge, one incoming blue edge, one outgoing red edge, and one outgoing blue edge.

In fact, *every* connected graph where every vertex has degree 4 is a covering space of X.

Proof: Orient the edges of $\widehat{X}$ by following an arbitrary closed Euler tour. Color each oriented edge red if it is the first dart in the Euler tour leaving its tail and blue otherwise. Now every vertex in $\widehat{X}$ is incident to one incoming red edge, one incoming blue edge, one outgoing red edge, and one outgoing blue edge, as required. □

Universal Cover

Recall that a space is ***simply connected*** if every closed curve is contractible. A simply-connected covering space of X is called a ***universal cover*** of X. In fact, up to homeomorphism, every connected space X has a *unique* universal cover, denoted $\boldsymbol{\widetilde{X}}$. The name "univeral" is motivated by the fact that $\widetilde{X}$ is also a covering space (in fact, the universal cover) of every connected covering space of X.

basepoint
lift of a point
projection
lift of a point
lift of a homotopy

The universal covering space $\widetilde{X}$ can be described more directly as the set of all homotopy classes of paths from a fixed ***basepoint*** $x \in X$:

$$\widetilde{X}_x \;:=\; \{[\pi] \mid \pi \colon [0,1] \to X \text{ and } \pi(0) = x\}\,.$$

The associated covering map $\widetilde{p} \colon \widetilde{X}_x \to X$ simply maps any homotopy class to its final endpoint: $\widetilde{p}([\pi]) = \pi(1)$.[1] For any path α in X, the function $\phi_\alpha \colon \widetilde{X}_{\alpha(0)} \to \widetilde{X}_{\alpha(1)}$ defined by $\phi_\alpha([\pi]) = [\alpha \cdot \pi]$ is a homeomorphism. Thus, we can omit the basepoint from our notation with no loss of precision.

For example, the plane is its own universal covering space, as is the sphere. The universal cover of the closed annulus $\{(x,y) \mid 1 \le x^2 + y^2 \le 2\}$ is the infinite strip $\{(\theta, r) \mid 1 \le r^2 \le 2\}$ (where r and θ are *not* polar coordinates).

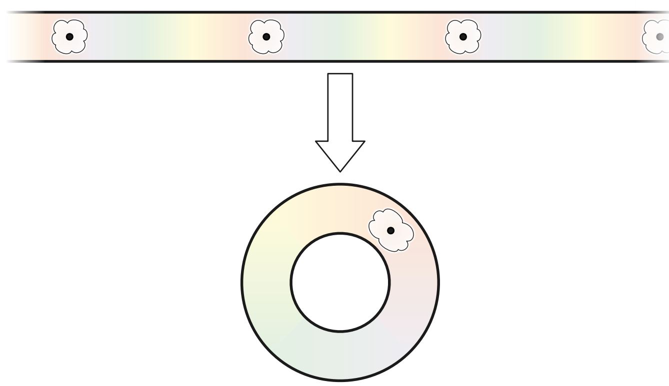

The infinite strip is the universal cover of the annulus.

Describe (universal) covering spaces of graphs, polygons with holes

Lifting Paths and Homotopies

Any point $\widehat{x} \in \widehat{X}$ is called a ***lift*** of its ***projection*** $p(\widehat{x}) \in X$; equivalently, the lifts of any point $x \in X$ are the points in the preimage $p^{-1}(x)$. A ***lift*** of a path $\pi \colon [0,1] \to X$ is any path $\widehat{\pi} \colon [0,1] \to \widehat{X}$ such that $\widehat{\pi} \circ p = \pi$. A lift of a path can be uniquely specified by choosing a list of one of its endpoints.

Lemma 4.1. *Let $p \colon \widehat{X} \to X$ be a covering map. For any path $\pi \colon [0,1] \to X$ and any point $\widehat{x} \in p^{-1}(\pi(0))$, there is a unique path $\Pi \colon [0,1] \to \widehat{X}$ such that $\Pi(0) = \widehat{x}$ and $\Pi \circ p = \pi$.*

Similarly, any homotopy $\widehat{h} \colon [0,1]^2 \to \widehat{X}$ is a ***lift*** of its projection $\widehat{h} \circ p$.

Lemma 4.2. *Let $p \colon \widehat{X} \to X$ be a covering map. For any homotopy $h \colon [0,1]^2 \to X$ and any point $\widehat{x} \in p^{-1}(h(0,0))$, there is a unique homotopy $H \colon [0,1]^2 \to \widehat{X}$ such that $H(0,0) = \widehat{x}$ and $H \circ p = h$.*

shortest homotopic path
metric space

This can be proved using crossing sequences for polygons with holes, or with a standard compactness argument for arbitrary domains.

Define lift of point, path, homotopy. Unique lift of paths and homotopies (or any function from a simply-connected domain). Lift construction for *canonical* paths/homotopies (and thus for crossing sequences) is straightforward by induction; construction for arbitrary paths/homotopies requires a standard compactness argument.

Covering spaces of graphs.

Lemma 4.3. *Let $p\colon \widehat{X} \to X$ be a covering map.*

- *For any contractible loop $\widehat{\ell}$ in $\widehat{X}$, the loop $p \circ \widehat{\ell}$ is contractible in X.*
- *Let ℓ be a contractible loop in X. For any point $\widehat{x} \in \widehat{X}$ such that $p(\widehat{x}) = \ell(0)$, there is a contractible loop $\widehat{\ell}$ in $\widehat{X}$ such that $\widehat{x} = \widehat{\ell}(0)$ and $p \circ \widehat{\ell} = \ell$.*

Why is universal cover unique? Prove simply-connected implies universal?

Corollary 4.4. *A loop ℓ in any space X is contractible if and only if ℓ lifts to a loop in the universal cover $\widetilde{X}$.*

Metric Spaces

Define metric, geodesic, isometry, complete, closure, metric completion. Canonical examples: The metric completion of an open polygon with holes is its closure. However, the metric completion of the plane minus points is *not* its closure (the plane).

Metrics lift to covering spaces. Any cover of the completion of X is the completion of a cover of X. Universal cover $\widetilde{P}$ of the plane minus points is *homeomorphic* to an open disk, but not *isometric* to an open disk. Thus, the metric completion of $\widetilde{P}$ is *not* a closed disk!

But all this is practically moot. If you don't think about it too hard, your algorithm will do the right thing. Sigh.

Corollary 4.5. *The shortest path in X homotopic to a given path α is the projection of the shortest path in $\widetilde{X}$ between the endpoints of any lift $\widetilde{\alpha}$ of α.*

4.2 Shortest (Homotopic) Paths in Polygons

The ***shortest homotopic path*** problem can be described as follows: Given a path π in some topological metric space X, find a path $\bar{\pi}$ of minimum length that is homotopic to π. In this chapter we consider an algorithm for a concrete special case of the shortest homotopic path problem, where the space X is a polygon with holes and the path π is a polygonal chain. This algorithm was first sketched by Leiserson and Maley [**?**], applied to more complex VLSI routing problems by Maley [**?**, **?**] and Gao *et al.* [**?**], and more formally described and analyzed by Hershberger and Snoeyink [**?**].

sleeve

Strip redundant definitions from earlier chapters: polygons (with holes), triangulation, homotopy, crossing sequences, etc. Define *metric* and *geodesic*!

The input to our algorithm is a simple polygon P and a polygonal path π (which may not be simple). The Jordan-Schönflies theorem implies that the shortest path $\bar{\pi}$ homotopic to π is just the shortest path in P from $\pi(0)$ to $\pi(1)$. Nevertheless, we will approach this special case as though its topology were nontrivial, because it illustrates important concepts that are useful in more general settings.

As in the previous chapter, we let n denote the number of edges of P, and let k denote the number of segments in the path π.

Move following parity discussion to Chapter 3.

We can also characterize reduced crossing sequences in terms of parity. For any diagonal edge e, the Jordan Curve Theorem implies that $P \setminus e$ has exactly two components. If an edge label e occurs an odd number of times in $X(\pi)$, then $\pi(0)$ and $\pi(1)$ lie in different components of $P \setminus e$; thus *any* path from $\pi(0)$ to $\pi(1)$ must cross e, and the shortest path must cross e exactly once. On the other hand, if an edge label e occurs an even number of times in $X(\pi)$, then $\pi(0)$ to $\pi(1)$ are in the same component of $P \setminus e$, so that shortest path $\bar{\pi}$ does not cross e at all. Thus, the reduced crossing sequence $\bar{X}$ contains precisely the edge labels that appear an odd number of times in $X(\pi)$; moreover, these labels are sorted by their first (or last) occurrence in $X(\pi)$.

Algorithm starts where our earlier contractibility algorithm left off. Compute a frugal triangulation of P in $O(n \log n)$ time. Compute the crossing sequence of π in $O(k+x) = O(nk)$ time. Reduce the crossing sequence in $O(x) = O(nk)$ time. If the reduced crossing sequence is empty, the shortest path is a line segment inside a single triangle. Otherwise, build the sleeve and then compute the funnel.

Sleeves and Funnels

Let $\bar{X}$ denote the reduced crossing sequence, and let $\bar{x}$ denote its length; because $\bar{X}$ contains each diagonal at most once, we have $\bar{x} \leq n-3$. With $\bar{X}$ in hand, we can restrict our attention to a subset of the triangles. The reduced crossing sequence defines a sequence of $\bar{x}+1$ triangles, starting with the triangle containing $\pi(0)$ and ending with the triangle containing $\pi(1)$. The ***sleeve*** of $\bar{X}$ is constructed by gluing together copies of the triangles in this sequence along their common edges.

Update figure; modify example from previous chapter.

funnel
tail of a funnel
apex!of a funnel
fan!of a funnel
wedge

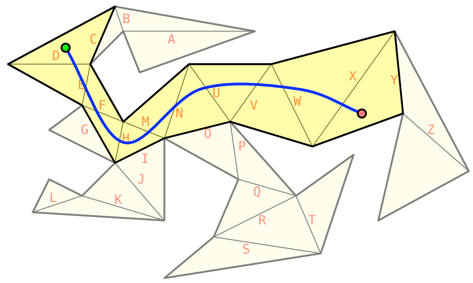

The sleeve for the reduced crossing word `DEFHMNVWX`.

If $\bar{x} = 0$, the sleeve consists of a single triangle, and the shortest path from $\pi(0)$ to $\pi(1)$ is a simple line segment. So assume from now on that $\bar{x} > 0$. We can clearly construct the sleeve in $O(\bar{x})$ time. Lemma **??** implies that $\bar{\pi}$ is the shortest path within the sleeve from $\pi(0)$ to $\pi(1)$.

Finally, we compute the shortest path through the sleeve using an algorithm independently discovered by Tompa [**?**], Chazelle [**?**], Lee and Preparata [**?**], and Leiserson and Maley [**?**]. The ***funnel*** of any diagonal e of the sleeve is the union of shortest paths from $\pi(0)$ to all points on e. The funnel consists of a polygonal path, called the ***tail***, from $\pi(0)$ to a point a called the ***apex***, plus a simple polygon called the ***fan***. The tail may be empty, in which case $\pi(0)$ is the apex. The fan is bounded by the edge e and two concave chains joining the apex to the endpoints of e. The shortest path from $\pi(0)$ to either endpoint of e consists of the tail plus one of the concave chains bounding the fan. Extending the edges of the concave chains to infinite rays defines a series of ***wedges***, which subdivide not only the fan but the triangle just beyond e.

CMYK-ify the funnel figures.

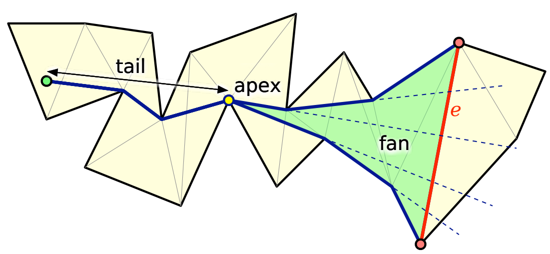

Anatomy of a typical funnel.

Beginning with a single triangle joining $\pi(0)$ to the first edge in $\bar{X}$, we extend the funnel through the entire sleeve one diagonal at a time. Each diagonal shares one endpoint with the previous diagonal; suppose we are extending the funnel from edge uv to edge vw. There are two cases to consider.

Let t be the predecessor of u on the shortest path from $\pi(0)$ to u. If the points v and w lie on opposite sides of the ray $\overrightarrow{ut}$, then the new endpoint w does not lie inside any wedge of the current fan. We can detect this case in $O(1)$ time with a single orientation test, and then *extend* the tunnel in $O(1)$ time by inserting w as a new fan vertex.

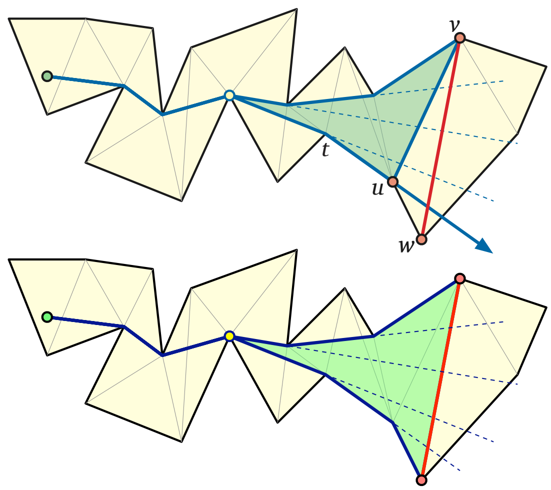

Extending and widening the funnel.

Otherwise, we *contract* the funnel, intuitively by moving u continuously along the boundary edge uw. Each time the point crosses the boundary of a wedge, we remove a vertex from the fan. If the removed vertex is the apex, its successor on the shortest path from $\pi(0)$ to v becomes the new apex, and the tail grows by an edge. We can detect whether the moving point will cross any wedge boundary in $O(1)$ time using our standard orientation test. Thus, the total time in this case is $O(d+1)$, where d is the number of vertices deleted from the fan. However, the total number of *deleted* vertices cannot exceed the total number of *previously inserted* vertices, so the *amortized* time to process this case is also $O(1)$.

Summing Up

When the funnel has reached the last edge in $\bar{X}$, we compute the shortest path from $\pi(0)$ to $\pi(1)$ in the sleeve by treating $\pi(1)$ as a triangle vertex and extending the funnel one last time. Thus, assuming the polygon is already triangulated, the overall time to compute the shortest path is $O(x + \bar{x}) = O(x) = O(nk)$.

In an actual implementation, it is not necessary to separate the algorithm into separate crossing sequence, reduction, sleeve, and funnel phases. Instead, it is possible to compute the tree of shortest paths from $\pi(0)$ to *every* vertex of *every* triangle crossed by the input path π in $O(x)$ time using single traversal of π, after which the shortest path to $\pi(1)$ can be extracted in $O(\bar{x})$ time.

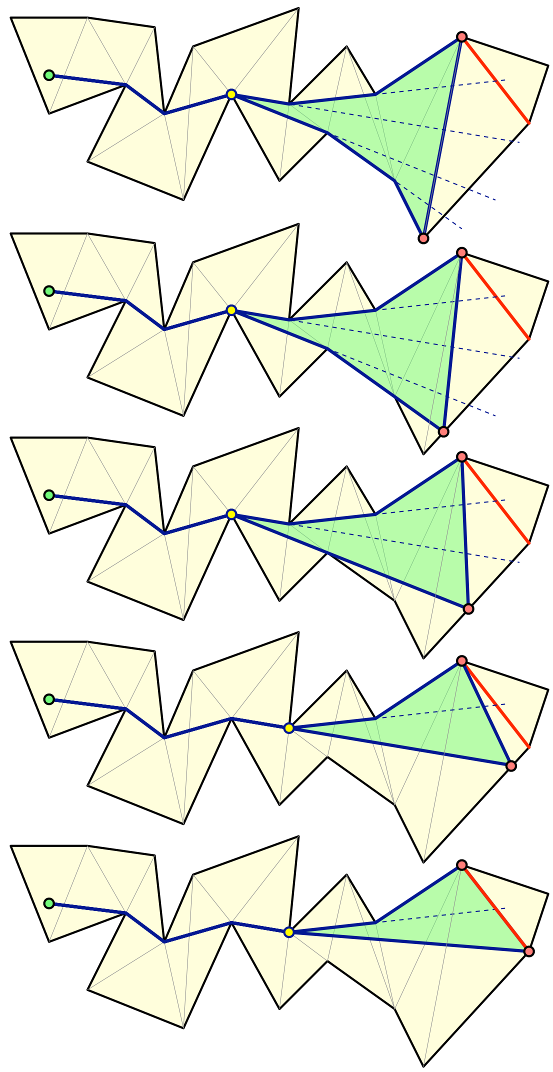

Extending and narrowing the funnel; the apex moves in the fourth step.

4.3 Polygons with Holes

Hershberger and Snoeyink [**?**] actually solve the shortest homotopic path problem for polygons with holes. Surprisingly, the algorithm we just described can be applied to polygons with holes with no significant modifications! As before, let n denote the number of edges in P, and let K denote the number of segments in π.

- In the preprocessing stage, compute a frugal triangulation of P in $O(n \log n)$ time.
- Compute the crossing sequence $X(\pi)$ in $O(x)$ time, where $x = O(kn)$ is the number of crossings.

- Reduce the crossing sequence in $O(x)$ time. We emphasize that the reduced crossing sequence $\bar{X}$ may contain the same edge label more than once, just not twice in succession.

- Compute the sleeve of $\bar{X}$ in $O(\bar{x})$ time, where $\bar{x}$ is the length of $\bar{X}$. Each time a path following $\bar{X}$ enters a triangle, we add ***a new copy*** of that triangle to the evolving sleeve. Thus, if a reduced path enters the same triangle five times, the resulting sleeve contains five different copies of that triangle. The sleeve is no longer a triangulated *simple* polygon; however, it is still homeomorphic to a disk. Moreover, if we represent the sleeve as a linked list of triangles, any self-overlaps are simply irrelevant.

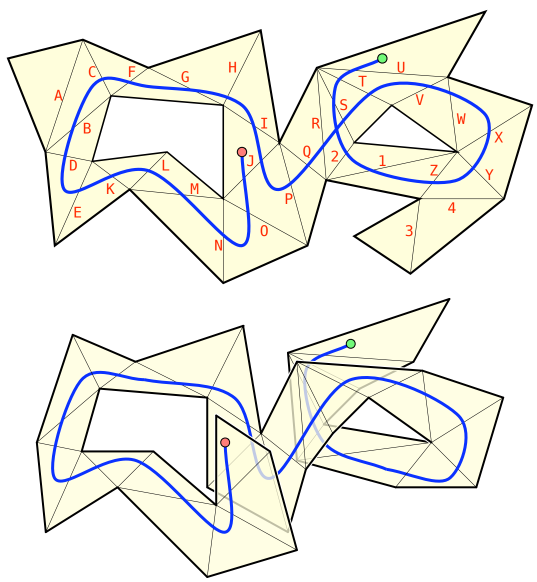

A reduced path in a polygon with two holes, with and the resulting non-simple sleeve.
The crossing sequence of the path is `UTS21ZYWVTSRQPJIHGFCBDEKLMNOJ`.

- Compute the shortest path in the sleeve using the funnel algorithm in $O(\bar{x})$ time. Even in this more general setting, the *fan* is always a simple polygon, so each extension step can be carried out exactly as described. The tail may intersect itself or the fan any number of times, but the algorithm won't notice.

- The overall running time of the algorithm is $O(n \log n + x + \bar{x}) = O(n \log n + nk)$.

Another useful way to think about the behavior of the algorithm is that it cannot distinguish between the original polygon with holes P and its universal cover $\widetilde{P}$. The correspondence between homotopy classes of paths and reduced crossing sequences implies the following description of $\widetilde{P}$ in terms of the triangulation of P. We describe an infinite triangulation of $\widetilde{P}$ by listing its constituent triangles and then declaring which pairs of edges should be identified.

basepoint
legal
equivalent!characters in crossing sequences
d-regular graph
bouquet of circles

Fix an arbitrary basepoint $p \in P$. Call a string X of edge labels legal if it is the reduced crossing sequence of a path π_X with $\pi_X(0) = p$. For each legal string X, let Δ_X denote ***a unique copy of*** the triangle containing $\pi_X(1)$. Thus, each triangle Δ lifts to an infinite number of triangles Δ_X, one for each reduced crossing sequence X ending in Δ. If X and Y are legal strings with $Y = Xe$, the triangles Δ_X and Δ_Y each contain a copy of edge e on their boundary; call these copies e_X and e_Y.

> The universal cover $\widetilde{P}$ is obtained from the disjoint union of all triangles Δ_X by identifying all pairs of edges e_X and e_Y such that $X = Ye$ for some edge e.

The choice of basepoint is unimportant; different basepoints induce different legal crossing sequences and therefore differently labeled triangles, but the resulting infinite triangulations are isomorphic.

That last sentence deserves a proof. Exercise?

The resulting triangulation of $\widetilde{P}$ is infinite. However, our shortest (homotopic) path algorithm only examines the *finite* set of triangles that intersect the path $\tilde{\pi}$.

We can also describe the transformation from P to $\widetilde{P}$ strictly in terms of crossing sequences. In any crossing sequence X, call two edge labels $X[i]$ and $X[j]$ ***equivalent*** if the substring $X[i\,..\,j]$ can be reduced to the empty string. For example, in the crossing sequence `ABCCCBCABCCBACBBC`, we can indicate equivalent labels with subscripts: $\mathtt{A_1B_1C_1C_1C_1B_2C_2A_2B_3C_3C_3B_3A_2C_2B_2B_2C_2}$. Then the reduced crossing sequence is simply the subsequence of distinct labels that occur an odd number of times, in order by their first occurrence, exactly as in the simply polygon case. For example, the crossing sequence $\mathtt{A_1B_1C_1C_1C_1B_2C_2A_2B_3C_3C_3B_3A_2C_2B_2B_2C_2}$ reduces to $\mathtt{A_1B_1C_1B_2C_2} = \mathtt{ABCBC}$. Nonequivalent occurrences of the same edge actually refer to two different lifts of that edge in the universal cover.

4.4 Closed Obstacles

Common practice to discuss shortest (homotopic) paths in the plane minus points, or minus line segments. Formal definition in terms of metric completion is a bit of a mess, but algorithmic extension is straightforward.

4.5 Exercises

1. A graph G is ***d-regular*** if every vertex of G has degree d. A ***bouquet of k circles*** is the unique graph with one vertex and k edges. Prove that for every positive integer k, every connected $2k$-regular graph is a covering space of the bouquet of k circles.
2. Prove that the graph of the cube is a covering space of the graph of the tetrahedron.
3. Prove that the underlying graph of any band decomposition $\Sigma^\square$ is a covering space of the bipartite graph with two vertices and three edges.

Notes

All spaces are nice!

1. (page **??**) Formally, we must also define the *topology* of the space $\widetilde{X}$. Let $\mathcal{U}$ be a basis of simply connected open sets for X. (All spaces are nice!) For any set $U \in \mathcal{U}$ and any path $\pi\colon [0,1] \to X$ with $\pi(1) \in U$, define

$$U_{[\pi]} = \Big\{ [\pi \cdot \eta] \;\Big|\; \eta\colon [0,1] \to U \text{ such that } \eta(0) = \pi(1) \Big\}.$$

Then the set $\widetilde{\mathcal{U}} := \{U_{[\pi]} \mid \pi(1) \in U \in \mathcal{U}\}$ is a basis for a topology on $\widetilde{X}$. For any two homotopic paths π and σ from x to a point in U, we have $U_{[\pi]} = U_{[\sigma]}$; on the other hand, if π and σ are not homotopic, then the sets $U_{[\pi]}$ and $U_{[\sigma]}$ are disjoint. The restriction of $\widetilde{p}$ to any set $U_{[\pi]}$ is a homeomorphism to the base neighborhood U; it follows that $\widetilde{p}\colon \widetilde{X} \to X$ is a covering map. Detailed proofs of these claims are given by Lee [**?**].

Bibliography

[1] Bernard Chazelle. A theorem on polygon cutting with applications. *Proc. 23rd Ann. IEEE Symp. Found. Comput. Sci.*, 339–349, 1982. (5)

[2] Shaodi Gao, Mark Jerrum, Michael Kaufmann, Kurt Mehlhorn, and Wolfgang Rülling. On continuous homotopic one layer routing. *Proc. 4th Ann. Symp. Comput. Geom.*, 392–402, 1988. (4)

[3] John Hershberger and Jack Snoeyink. Computing minimum length paths of a given homotopy class. *Comput. Geom. Theory Appl.* 4:63–98, 1994. (4, 7)

[4] Der-Tsai Lee and Franco P. Preparata. Euclidean shortest paths in the presence of rectilinear barriers. *Networks* 14:393–410, 1984. (5)

[5] John M. Lee. *Introduction to Topological Manifolds*. Graduate Texts in Mathematics 202. Springer, 2000. (11)

[6] Charles E. Leiserson and F. Miller Maley. Algorithms for routing and testing routability of planar VLSI layouts. *Proc. 17th Ann. ACM Symp. Theory Comput.*, 69–78, 1985. (4, 6)

[7] F. Miller Maley. *Single-Layer Wire Routing*. Ph.D. thesis, Massachusetts Inst. Tech., Cambridge, MA, August 1987. (4)

[8] F. Miller Maley. *Single-Layer Wire Routing and Compaction*. MIT Press, Cambridge, MA, 1990. (4)

[9] Martin Tompa. An optimal solution to a wire-routing problem. *J. Comput. Syst. Sci.* 23:127–150, 1981. (5)