Testing Contractibility in Planar Rips Complexes*

Erin W. Chambers[†] Jeff Erickson Pratik Worah

Department of Computer Science
University of Illinois, Urbana-Champaign
{erinwolf,jeffe,pworah2}@uiuc.edu

ABSTRACT

The *(Vietoris-)Rips complex* of a discrete point-set P is an abstract simplicial complex in which a subset of P defines a simplex if and only if the diameter of that subset is at most 1. We describe an efficient algorithm to determine whether a given cycle in a planar Rips complex is contractible. Our algorithm requires $O(m \log n)$ time to preprocess a set of n points in the plane in which m pairs have distance at most 1; after preprocessing, deciding whether a cycle of k Rips edges is contractible requires $O(k)$ time. We also describe an algorithm to compute the shortest non-contractible cycle in a planar Rips complex in $O(n^2 \log n + mn)$ time.

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations*

General Terms: Algorithms, Performance

Keywords: Computational topology, homotopy, Vietoris-Rips complex, Rips shadow

1. INTRODUCTION

A fundamental class of problems in computational topology deals with properties of paths and cycles in various topological spaces that are invariant under continuous deformation, or *homotopy*. For example, given two paths, can one be continuously deformed into the other? Given a topological metric space, what is the shortest cycle that cannot be continuously contracted to a single point? These and similar problems have been studied extensively for regions of the plane with holes [36, 31, 25, 3, 2, 25, 5] and graphs embedded on surfaces [26, 27, 15, 12, 6, 4, 11, 39, 44, 41, 16, 19]. Applications of these algorithms include problems in graph drawing [20], map simplification [5], simplification and parameterization of surface meshes [33, 52], and approximation algorithms [17] and fixed-parameter tractable algorithms [38] for generalizations of planar graphs.

Relatively little is known about these problems in more general topological spaces. Even determining whether two paths are homotopic is undecidable if the ambient space is an arbitrary 2-dimensional simplicial complex [42] or a 4-manifold, and decidability is open for arbitrary 3-manifolds [47]. Consequently, most topological algorithms for non-surfaces use *homology*, which provides a cruder classification of topological features than homotopy, but generalizes more easily [23, 24, 10, 29, 54].

In this paper, we develop efficient algorithms for basic homotopy problems in a special class of simplicial complexes called *(Vietoris-)Rips complexes*. The Rips complex of a set of points is a simplicial complex that contains a simplex for each subset with diameter less than 1. These complexes were introduced by Leopold Vietoris [50] as the basis of an early homology theory; they were later independently discovered by Eliyahu Rips and popularized by Gromov [32] (who coined the name 'Rips complex') as a tool for studying hyperbolic groups. Ghrist [29], Carlsson [8], and others have proposed Rips complexes as a lightweight representation of the topological structure of high-dimensional data. A recent example of this approach is the analysis by Carlsson and others [7, 9, 29] of a large set of nine-dimensional feature points extracted from digital photographs (the "Mumford data set"). Rips complexes of points in the plane have also been used to model coverage problems in sensor networks [30, 46, 45]; this particular use of Rips complexes was the original motivation for our work.

Our paper contains two main algorithmic results. The first is an efficient algorithm to determine whether a given cycle in a planar Rips complex is contractible (Section 5). Our algorithm requires $O(m \log n)$ time to preprocess a set of n points in the plane in which m pairs have distance at most 1. After preprocessing, we can determine whether a cycle of k Rips edges is contractible in $O(k)$ time. Our second algorithm (Section 6) computes the shortest non-contractible cycle in the Rips complex of a given planar point set, minimizing either number of edges or total Euclidean length, in $O(n^2 \log n + mn)$ time. The efficiency of our algorithms follows from special geometric properties of the

*Work on this paper was partially supported by NSF grant DMS-0528086. See http://www.cs.uiuc.edu/~jeffe/pubs/ripscon.html for the most recent version of this paper.

[†]Also partially supported by an NSF Graduate Research Fellowship.

Rips shadow, described in Sections 4 and 5.1, which we believe are of independent interest.

2. PRELIMINARIES

We begin by recalling some standard definitions. For further background on algebraic and computational topology, see Edelsbrunner [22], Hatcher [34], Stillwell [47], and Zomorodian [53].

A *simplicial complex* $\mathcal{X}$ is a collection of simplices (points, segments, triangles, etc.) with the following properties: (1) Any face of a simplex in $\mathcal{X}$ is another simplex in $\mathcal{X}$; (2) Any two simplices in $\mathcal{X}$ intersect in a common face. The k*-skeleton* of $\mathcal{X}$ is the subcomplex consisting of all simplices in $\mathcal{X}$ of dimension k or less. The *flag complex* $\mathcal{F}(G)$ of a graph G is the largest simplicial complex whose 1-skeleton is G; every $(k+1)$-clique in G defines a k-simplex in $\mathcal{F}(G)$.

Let P be a set of points in some metric space. The ***Vietoris-Rips complex*** $\mathcal{R}_\varepsilon(P)$ is the simplicial complex that contains a k-simplex for each subset of $k+1$ points with maximum pairwise distance at most ε. For simplicity, we will refer to $\mathcal{R}(P) = \mathcal{R}_1(P)$ as the *Rips complex* of P. Equivalently, the Rips complex of P is the flag complex of the *proximity graph* of P, whose edges are all pairs of points $p, q \in P$ such that $|pq| \leq 1$. See Figure 1 for an example. The closely related *Čech complex* $\check{\mathcal{C}}_\varepsilon(P)$ is the simplicial complex that contains a k-simplex for every subset of $k+1$ points contained in a ball of radius ε.

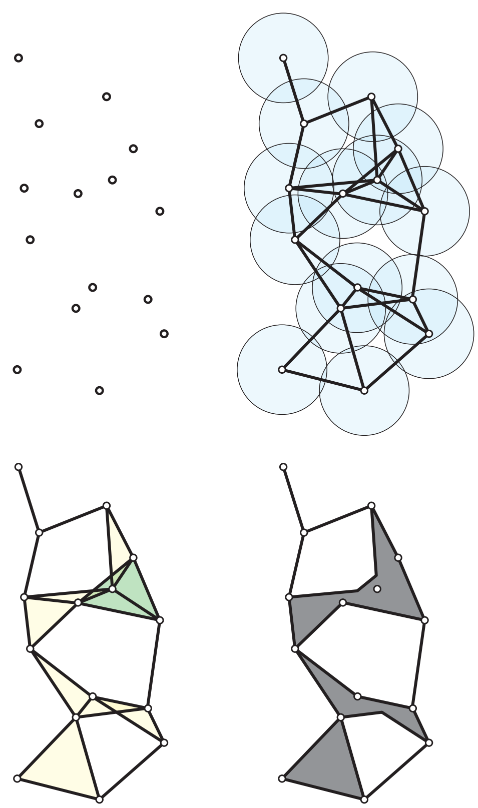

Figure 1. A set of points in the plane, its proximity graph (the intersection graph of circles of radius 1/2), its Rips complex, and its Rips shadow (see Section 3).

Given some topological space M, a *path* is a continuous function $p : [0,1] \to M$; a path whose endpoints coincide is called a *loop*. A *homotopy* between two paths p and p' with the same endpoints is a continuous function $h\colon [0,1]\times[0,1] \to M$ such that $H(0,t) = p(t)$ and $H(1,t) = p'(t)$ for all t, and $H(s,0) = p(0) = p'(0)$ and $H(s,1) = p(1) = p'(1)$ for all s. If M is a simplicial complex, a path is generally constrained to lie along the 1-skeleton of the complex, and homotopies are comprised of a series of elementary moves, each of which moves a portion of the path or cycle across a triangle. Two paths are *homotopic* if there is a homotopy from one to the other. A loop is *contractible* if it is homotopic to a point.

It is an easy exercise to verify that homotopy is an equivalence relation on the set of loops with any fixed basepoint. The *fundamental group* $\pi_1(X, x)$ of a space X with basepoint $x \in X$ is the group of homotopy classes of loops with endpoint x, with concatenation as the group operation and the set of contractible cycles through x as the identity element. If X is connected, then $\pi_1(X, x_1) \simeq \pi_1(X, x_2)$ for any x_1 and x_2 in X; consequently, we often simply write $\pi_1(X)$.

3. THE RIPS SHADOW

For any planar Rips complex $\mathcal{R}$ (indeed for any abstract simplicial complex whose vertices are points in the plane), there is a canonical projection map $p\colon \mathcal{R} \to \mathbb{R}^2$ that maps each simplex in $\mathcal{R}$ affinely onto the convex hull of its vertices in $\mathbb{R}^2$. The ***Rips shadow*** $\mathcal{S}(P)$ is the image of this canonical projection map, or equivalently, the union of the convex hulls of all subsets of P with diameter at most 1:

$$\mathcal{S}(P) := \bigcup \{\operatorname{conv}(Q) \mid Q \subseteq P \text{ and } \max_{p,q \in Q} |pq| \leq 1\}.$$

The Rips shadow is a planar region, possibly with holes, with a piecewise-linear boundary; intuitively, the shadow is a polygon with holes whose boundaries may touch themselves and/or each other. The boundary of the shadow can be decomposed into maximal line segments, which we call *(shadow) boundary edges*, meeting at *(shadow) boundary vertices*. The collection of boundary vertices and boundary edges comprise the *(shadow) boundary graph*. We define the *complexity* of the shadow to be the total number of boundary vertices and edges.

The canonical projection map p naturally induces a map $\pi_1(p)\colon \pi_1(\mathcal{R}(P)) \to \pi_1(\mathcal{S}(P))$ from the fundamental group of the Rips complex to the fundamental group of its shadow. Our results build on the following recent result of Chambers *et al.* [13]:

Theorem 3.1. *For any set P of points in the plane, the induced map $\pi_1(p)\colon \pi_1(\mathcal{R}(P)) \to \pi_1(\mathcal{S}(P))$ is an isomorphism.*

Equivalently, Theorem 3.1 states that a cycle γ in the Rips complex is contractible if and only if its projection $p(\gamma)$ is contractible in the Rips shadow. Theorem 3.1 immediately implies that the projection map p also induces a isomorphism between the first homology groups of $\mathcal{R}(P)$ and $\mathcal{S}(P)$.

The Rips shadow $\mathcal{S}(P)$ is homotopy equivalent to a set of loops with common basepoint, where each loop winds around a hole in the shadow exactly one time. An immediate but important consequence of Theorem 3.1 is that the fundamental group $\pi_1(\mathcal{R}(P))$ is a *free* group. Moreover, a minimal presentation of $\pi_1(\mathcal{R}(P))$ contains exactly one generator for each hole in the shadow, and thus can be computed quickly once the shadow is constructed. (The proof of Theorem 3.1

describes an efficient method to 'lift' any cycle γ in $\mathcal{S}(P)$—for example, the boundary of a hole—to a cycle $\hat{\gamma}$ in $\mathcal{R}(P)$ such that γ and $p(\hat{\gamma})$ are homotopic [13].)

Unlike Čech complexes and α-shapes [21], the Rips complex and its shadow are *not* homotopy equivalent in general. For example, for any positive integer n, if P is a set of $2n+2$ regularly-spaced points on a circle of radius strictly between 1 and $\sec(\pi/(2n+2)) < 1 + \pi/8n^2$, then $\mathcal{R}(P)$ is combinatorially isomorphic to an $(n+1)$-dimensional cross-polytope and therefore homeomorphic to $\mathbb{S}^n$, even though $\mathcal{S}(P)$ is a disk. Haussman [35] and Latschev [40] proved that the Rips complex of any sufficiently dense point set *near* a closed Riemannian manifold M is homotopy-equivalent to M; in contrast, Theorem 3.1 requires *no* sampling conditions.

Theorem 3.1 immediately implies a polynomial-time algorithm determine whether a cycle in $\mathcal{R}(P)$ is contractible. The shadow $\mathcal{S}(P)$ is defined as the union of $O(n^3)$ triangles, so it can be constructed in $O(n^6)$ time. After $O(N^{6+\varepsilon})$ additional preprocessing time, where $N = O(n^6)$ is the complexity of the shadow, we can determine whether any cycle of k Rips edges is contractible in $O(k\sqrt{N}\log n)$ time, using the algorithm of Cabello *et al.* [5]. In the next two sections, we derive several geometric properties of the Rips shadow that considerably improve the efficiency of this approach.

4. COMPUTING THE SHADOW

In this section we develop an efficient algorithm to compute the Rips shadow $\mathcal{S}(P)$ of a given set P of n points in the plane. Our algorithm relies on two structural results, which may be of independent interest. First, although the Rips complex $\mathcal{R}(P)$ can contain $\Theta(n^2)$ edges and $\Theta(n^3)$ triangles in the worst case, the Rips shadow $\mathcal{S}(P)$, which is the union of those segments and triangles, always has complexity $O(n)$. Second, there is a subset of $O(n)$ Rips edges and Rips triangles whose union is the entire the Rips shadow $\mathcal{S}(P)$.

4.1 Linear Complexity

We begin with two trivial observations about a single pair of intersecting Rips edges.

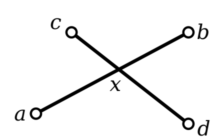

Figure 2. Intersecting Rips edges.

Lemma 4.1. *If Rips edges ab and cd intersect, then (1) either ac or bd is a Rips edge; (2) either ad or bc is a Rips edge; (3) at least one of abc, abd, acd, or bcd is a Rips triangle.*

Proof: Let $x = ab \cap cd$. The triangle inequality implies that $|ac| + |bd| \leq |ax| + |bx| + |cx| + |dx| = |ab| + |cd| \leq 2$. Thus, either $|ac| \leq 1$ or $|bd| \leq 1$, which implies (1). A symmetric argument implies (2). Finally, (3) follows immediately from (1) and (2). □

Lemma 4.2. *Let ab and cd be Rips edges that intersect at a point $x = ab \cap cd$, such that neither abc nor acd is a Rips triangle. Then $|ac| > 1$ and $\angle axc > \pi/3$.*

Proof: Lemma 4 implies that either ad or bc is a Rips edge. Thus, if ac were a Rips edge, then either abc or acd would be a Rips triangle.

We have both $|ax| \leq |ab| < 1$ and $|cx| \leq |cd| < 1$. Thus, ac is the unique longest side of triangle acx, so its opposite angle $\angle axc$ is the unique largest angle. □

Theorem 4.3. *The Rips shadow of n points in the plane has combinatorial complexity $O(n)$.*

Proof: Fix a set P of n points in the plane. We assume without loss of generality that $\mathcal{R}(P)$ and therefore $\mathcal{S}(P)$ are connected; if not, we can analyze each connected component independently. This assumption implies that each hole in $\mathcal{S}(P)$ has a single boundary cycle.

We bound the complexity of the Rips shadow by (over-)counting the number of boundary edges and vertices. The same boundary vertex or edge may appear multiple times on the same facial walk or on multiple walks; we count each occurrence separately. To simplify our presentation, we consider the two sides of any Rips edge or shadow boundary edge separately; for any edge uv, let $\overrightarrow{uv}$ and $\overrightarrow{vu}$ denote its two oriented *halfedges*. A facial walk now consists of a sequence of boundary halfedges, oriented with the hole on the left; two consecutive boundary halfedges $\overrightarrow{xy}$ and $\overrightarrow{yz}$ form a *boundary corner* at y. We prove that there are $O(n)$ boundary corners.

We say that a Rips halfedge $\overrightarrow{pq}$ is *uncovered* if there is no Rips triangle pqr with r to the left of the oriented line $\overrightarrow{pq}$. Every corner of the shadow boundary is located at the intersection of two uncovered halfedges, possibly at a common endpoint.

If $\overrightarrow{pq}$ and $\overrightarrow{pr}$ are two uncovered Rips halfedges with a common source point p, then $\angle qpr > \pi/3$ by Lemma 4.2. It follows that any point in P is the source of at most five uncovered halfedges, giving at most $5n$ uncovered edges total. In addition, there are at most five boundary corners at any point in P.

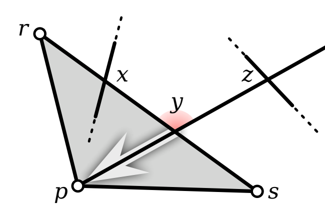

Figure 3. Charging a corner to one end of an uncovered halfedge.

Let $\overrightarrow{pq}$ and $\overrightarrow{rs}$ be uncovered Rips halfedges, with r to the left of $\overrightarrow{pq}$, whose interiors intersect at boundary vertex y. Suppose some pair of boundary halfedges $\overrightarrow{xy} \subset \overrightarrow{pq}$ and $\overrightarrow{yz} \subset \overrightarrow{rs}$ form a boundary corner at y. Lemma 4.1 implies that either prs or pqs is a Rips triangle, since either of the other two possible triangles would cover $\overrightarrow{pq}$ or $\overrightarrow{rs}$. If prs is a Rips triangle, segment py lies inside the shadow, so y is the closest boundary corner to p, among all boundary corners on $\overrightarrow{pq}$. See Figure 3. Similarly, if pqs is a Rips triangle, y is the boundary corner on $\overrightarrow{rs}$ that is closest to s.

Thus, every boundary corner that is not a point in P is either the first or last boundary corner on some uncovered halfedge. It follows that there are at most $10n$ boundary corners not at points in P, and thus at most $15n$ boundary corners overall. □

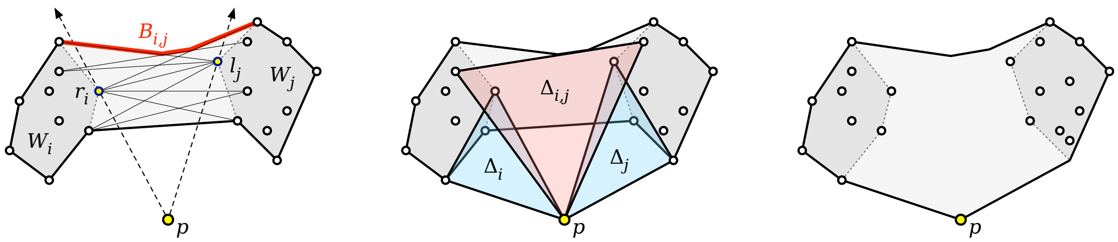

Figure 4. Left: $\mathcal{S}(W_i \cup W_j)$. Middle: Adding $\triangle_i$, $\triangle_{i,j}$, and $\triangle_j$. Right: $\mathcal{S}(W_i \cup W_j \cup \{p\})$

4.2 Linear Coverage

Theorem 4.4. *For any set P of n points in the plane, there is a set of $O(n)$ Rips edges and Rips triangles whose union is the Rips shadow $\mathcal{S}(P)$.*

Proof: Fix a set P of n points in the plane and an arbitrary point $p \in P$, and let $P' = P \setminus \{p\}$. As in the previous proof, we assume that $\mathcal{R}(P)$ and thus $\mathcal{S}(P)$ are connected; if not, we prove the theorem independently for each component. We prove that $\mathcal{S}(P)$ is the union of $\mathcal{S}(P')$ and a constant number of Rips edges and triangles, each of which have p as a vertex; the theorem then follows immediately by induction.

Let $Q' \subseteq P'$ denote the set of Rips neighbors of p, and let $Q = Q' \cup \{p\}$. $\mathcal{S}(P')$ is the union of all Rips edges and triangles that do not have p as a vertex, and $\mathcal{R}(Q)$ contains every edge and triangle in $\mathcal{R}(P)$ that is incident to p, so $\mathcal{S}(P) = \mathcal{S}(P') \cup \mathcal{S}(Q)$. Thus, it suffices to prove that $\mathcal{S}(Q)$ is the union of $\mathcal{S}(Q')$ and $O(1)$ Rips triangles incident to p.

We divide the plane into six congruent wedges by three lines through p. Let $W_0, \ldots, W_5$ be the (possibly empty) subsets of Q inside these wedges, indexed in order around p.

Each set W_i has diameter less than 1, which implies that $\operatorname{conv}(W_i \cup \{p\}) = \mathcal{S}(W_i) \subseteq \mathcal{S}(Q)$. For each non-empty set W_i, let ℓ_i and r_i denote the leftmost and rightmost points in W_i as seen from p, and let $\triangle_i = \operatorname{conv}\{p, \ell_i, r_i\}$. In particular, if W_i contains only one point, then $\ell_i = r_i$ and $\triangle_i$ is the segment pr_i. If W_i is empty, we define $\triangle_i$ to be the point p. We clearly have $\mathcal{S}(W_i \cup \{p\}) = \mathcal{S}(W_i) \cup \triangle_i$.

Consider two nonempty subsets W_i and W_j with $i < j$ and $i \neq j - 3$. (There are at most 12 such subset pairs.) We define a triangle $\triangle_{i,j}$ such that

$$\mathcal{S}(W_i \cup W_j) \cup \triangle_i \cup \triangle_{i,j} \cup \triangle_j \;=\; \mathcal{S}(W_i \cup W_j \cup \{p\}) \qquad (\star)$$

Let $B_{i,j}$ denote the concave chain of edges on the boundary of $\mathcal{S}_{i,j} = \mathcal{S}(W_i \cup W_j \cup \{p\})$, which starts at a vertex of $\operatorname{conv}(W_i)$, ends at a vertex of $\operatorname{conv}(W_j)$, and contains no other Rips vertices. Without loss of generality, suppose p lies 'below' $B_{i,j}$. Every boundary edge of $B_{i,j}$ is a subset of a Rips edge with one endpoint in W_i and the other in W_j.

We claim that for any Rips edge qr that touches $B_{i,j}$, the triangle $\triangle_{i,j} = pqr$ satisfies Equation ($\star$). See Figure 4. Consider another Rips edge st that touches $B_{i,j}$, such that s is above qr and r is above st. We easily verify that qrt and qst are Rips triangles. By considering all such edges st, we conclude that the region bounded by the chain $B_{i,j}$, the edge qr, and the rays $\overrightarrow{pr_i}$ and $\overrightarrow{pl_j}$ lies entirely within $\mathcal{S}(W_i \cup W_j \cup \{p\})$, which establishes our claim.

If non-empty subsets W_i and W_j lie in opposite wedges ($i = j \pm 3$), then we may need to define *two* triangles $\triangle_{i,j}$ and $\triangle_{j,i}$ on opposite sides of p, so that

$$\mathcal{S}(W_i \cup W_j) \cup \triangle_i \cup \triangle_{i,j} \cup \triangle_{j,i} \cup \triangle_j \;=\; \mathcal{S}(W_i \cup W_j \cup \{p\}).$$

It suffices to choose arbitrary triangles $\triangle_{i,j}$ and $\triangle_{j,i}$ incident to p that touch the two concave chains connecting W_i to W_j.

Altogether, we have chosen at most 24 Rips triangles and Rips edges—one for each of the 6 wedges, one for each of the 12 non-opposing wedge pairs, and two for each of the 3 opposing wedge pairs—whose union covers $\mathcal{S}(P) \setminus \mathcal{S}(P')$. $\square$

4.3 Construction Algorithm

Theorem 4.5. *Given a set P of n points in the plane, we can construct $\mathcal{S}(P)$ in $O((m+n)\log n)$ time, where m is the number of edges in the proximity graph of P.*

Proof: We first describe a simpler algorithm that runs in $O(n^2 \log n)$ time, and then describe a general reduction strategy that improves the running time to $O((m+n)\log n)$.

Our simple algorithm computes a set of $O(n)$ Rips triangles whose union is $\mathcal{S}(P)$, as follows. For each point $p \in P$, we execute the following subroutine. Fix a point $p \in P$. We can easily compute the six neighbor sets $W_1, \ldots, W_6$ of p from Theorem 4.4 in $O(n)$ time. For each set W_i, we compute the extreme points l_i and r_i in $O(n)$ time and add the triangle pr_il_i to the output list. For each of subsets W_i and W_j, we compute the triangle $\triangle_{i,j}$ as follows. We construct the Voronoi diagram of W_j, and then for each point $q \in W_i$, we compute its nearest neighbor $r \in W_j$ using a point-location query. Then, among all segments qr with length at most 1, we determine the one that intersects the ray $\overrightarrow{pr_i}$ furthest from p and add the resulting triangle pqr to the output list. The subroutine runs in $O(n \log n)$ time, so the total time to compute the $O(n)$ covering triangles is $O(n^2 \log n)$. Once we have the covering triangles, we can compute their union in $O(n^2)$ time with a standard sweepline algorithm.

To reduce the running time to $O((m+n)\log n)$, we impose a grid of $1/2 \times 1/2$ squares over the point set, and independently compute the intersection of $\mathcal{S}(P)$ with each grid square. Let $c_1, c_2, \ldots, c_N$ denote the grid cells that are within distance 3 of some point in P; clearly $N = O(n)$. For each i, let C_i denote the $5/2 \times 5/2$ square with the same center as grid cell c_i, let $P_i = P \cap C_i$, and let $N_i = |P_i|$. Observe that $\mathcal{S}(P) \cap c_i = \mathcal{S}(P_i) \cap c_i$, since all Rips edges have length at most 1. For each point $p \in P$, we determine the subsets P_i that contain it; this takes $O(n)$ time overall. Then for each index i, we compute $O(N_i)$ triangles whose union is $\mathcal{S}(P_i)$ using our earlier algorithm, intersect each of these triangles with c_i, and compute the union of the resulting polygons. Finally, we glue the resulting sub-shadows together along their common boundaries, in $O(n)$ time, to obtain $\mathcal{S}(P)$. The total running time of this algorithm is $\sum_i O(N_i^2 \log N_i)$.

To complete our analysis, we prove that $\sum_i N_i^2 = O(m+n)$. For each i, let $n_i = |P \cap c_i|$, and let us write $i \sim k$ to mean $c_i \subset C_k$, so that $N_k = \sum_{i\sim k} n_i$. We immediately have

$$\begin{aligned}\sum_k N_k^2 &= \sum_k \sum_{i\sim k} \sum_{j\sim k} n_i n_j \\ &\le \sum_k \sum_{i\sim k} \sum_{j\sim k} (n_i^2 + n_j^2)/2 \\ &= \sum_k \sum_{i\sim k} n_i^2 \\ &\le 25 \sum_i n_i^2.\end{aligned}$$

On the other hand, we also have $\sum_i \binom{n_i}{2} \le m$, because each cell c_i has diameter less than 1. We conclude that $\sum_i N_i^2 \le 50m + 25n$, as claimed. $\square$

5. TESTING CONTRACTIBILITY

In this section, we describe an efficient algorithm to determine whether a given cycle of Rips edges is contractible, or equivalently, whether two paths with common endpoints are homotopic, in the Rips complex $\mathcal{R}(P)$. Theorem 3.1 implies that a cycle γ is contractible in $\mathcal{R}(P)$ if and only if its projection $p(\gamma)$ is contractible in the Rips shadow $\mathcal{S}(P)$. Thus, testing contractibility is simply a matter of tracking how many times the projected cycle $p(\gamma)$ winds around each hole.

The fastest algorithm known for testing the contractibility of cycles in planar regions with holes is due to Cabello *et al.* [5]. Their algorithm tests whether a cycle of k edges is contractible in the plane minus n point obstacles in time $O(k\sqrt{n}\log n)$, after $O(n^{1+\varepsilon})$ preprocessing time, by applying a planar ray shooting data structure [37] to a spanning tree of the obstacles with stabbing number $O(\sqrt{n})$ [14, 43, 51]. We can directly apply their algorithm to Rips cycles by choosing an arbitrary point in each hole of the Rips shadow, plus one point outside the outer boundary; this requires $O(m \log n)$ additional preprocessing time.

However, the special geometric structure of Rips shadows allows us to develop a faster and simpler algorithm, with the same preprocessing time. Our algorithm constructs a spanning tree of the holes such that any *unit-length* line segment crosses a *constant* number of edges. (Such a spanning tree does not exist for arbitrary planar regions—consider a unit square containing a $\sqrt{n} \times \sqrt{n}$ grid of tiny holes.) The existence of this spanning tree follows from another structural result of independent interest: Although holes in the Rips shadow can have arbitrarily small area, they cannot have arbitrarily small diameter.

5.1 No Small Holes

Lemma 5.1. *Let $\mathcal{A}$ be a set of $n \ge 4$ arcs on a common circle C, each subtending an angle larger than $4\pi/5$, whose union is the entire circle. $\mathcal{A}$ contains four arcs that cover C.*

Proof: Let $|\alpha|$ denote the angular length of arc α.

Start with an arbitrary arc $\alpha_1 \in \mathcal{A}$. Let α_2 be the clockwise-most arc in $\mathcal{A}$ that overlaps α_1, and let α_3 be the clockwise-most arc in $\mathcal{A}$ that overlaps α_2. Arcs α_1 and α_3 must be disjoint, which implies that if $|\alpha_1 \cup \alpha_2| \le 6\pi/5$, then $|\alpha_2 \cup \alpha_3| \ge 6\pi/5$. Thus, we can assume with no loss of generality that $|\alpha_1 \cup \alpha_2| \ge 6\pi/5$.

Let α_1' be the counterclockwise-most arc in A that overlaps α_2, and let α_0' be the counterclockwise-most arc in $\mathcal{A}$ that overlaps α_1'. We still have $|\alpha_1' \cup \alpha_2| \ge 6\pi/5$, so $|C \setminus (\alpha_1' \cup \alpha_2)| \le 4\pi/5$. Thus, arcs α_0' and α_3 must overlap, which implies that $\alpha_0', \alpha_1', \alpha_2, \alpha_3$ cover the circle. $\square$

Theorem 5.2 (No small holes). *For any set P of n points in the plane, every hole in the Rips shadow $\mathcal{S}(P)$ has circumradius at least $(\sqrt{2}-1)/8\sqrt{3} \approx 0.029893$.*

Proof: Let H be a hole in the Rips shadow of a planar point set P. For any real $\rho > 0$, let D_ρ denote the open disk of radius ρ centered at the origin o, let $C_\rho = \partial D_\rho$ denote its boundary circle, and let $P_\rho = X \cap D_\rho$. Every pair of points in $P_{1/2}$ is connected by a Rips edge, so the Rips shadow of $P_{1/2}$ is equal to its convex hull.

Suppose H lies inside the disk D_ρ. Fix a real value $\sigma = (1+\sqrt{5})\rho < 1/2$. There are three cases to consider: (1) P_σ is empty; (2) P_σ is nonempty and H is convex; and (3) there is a point in P on the boundary of H. (It will become clear during the proof that these cases are exhaustive.) Each case implies a different lower bound on the radius ρ. The first case is by far the most involved.

Case 1: All points far from the hole. Suppose $P_\sigma = \varnothing$. In this case, H must be convex. Let $e_1, e_2, \ldots, e_r$ denote the Rips edges bounding H. For each edge e_i, let α_i denote the portion of C_σ separated from H by e_i. Because $\sigma = (1+\sqrt{5})\rho = \rho/\cos(2\pi/5)$, each arc α_i subtends an angle larger than $4\pi/5$, and these arcs cover the entire circle. Lemma 5.1 implies that four of these arcs $\alpha_i, \alpha_j, \alpha_k, \alpha_l$ also cover the circle. The corresponding Rips edges e_i, e_j, e_k, e_l bound a convex quadrilateral ***pseudohole*** $\tilde{H}$ that lies inside D_σ.

To simplify our notation, we relabel the endpoints and intersection points of the edges e_i, e_j, e_k, e_l, as illustrated in Figure 5. Let a, b, c, d, e, f, g, h denote the endpoints in clockwise order, so that (without loss of generality) $e_i = af$, $e_j = be$, $e_k = ch$, and $e_l = dg$. We also label the vertices of $\tilde{H}$ in clockwise order: $w = af \cap dg$, $x = af \cap be$, $y = be \cap ch$, and $z = ch \cap dg$.

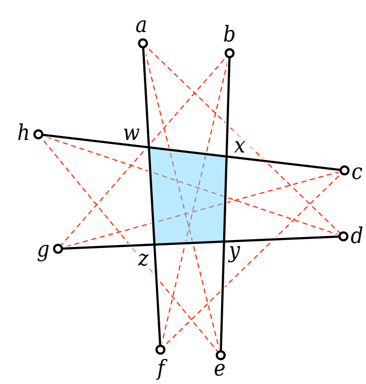

Figure 5. Four Rips edges bounding a quadrilateral pseudohole. The dashed segments are *not* Rips edges.

Lemma 4.2 implies that there are no other intersections among these segments; all eight endpoints are distinct. We prove that points a, b, e, f are in convex position as follows; a similar argument implies that c, d, g, h are also in convex position. Suppose b is inside triangle aef, and let $s = \overrightarrow{be} \cap af$. We immediately have $|sf| < |af| \leq 1$. Segment bf crosses through the hole H, so we must have $|bf| > 1$ and therefore $|ef| > 1$.

Point s is outside C_σ and segments se and sf intersect C_ρ, so $\angle esf < 2\arcsin(\rho/\sigma) \leq \pi/5$. Together with the inequalities $|sf| < 1$ and $|ef| > 1$, this implies that $|es| > (1+\sqrt{5})/2$, from which it follows that $|bs| > (\sqrt{5}-1)/2$ and therefore $|es| > (\sqrt{5}-1)/2$. We conclude that the minimum distance between sf and be is more than $(\sqrt{5}-1)/2$. On the other hand, sf and be both intersect C_ρ; it follows that $\rho > (\sqrt{5}-1)/4$, so $\sigma > 1$, which is impossible. Thus, b cannot be inside triangle aef, so (by symmetric arguments) a, b, e, f must be in convex position.

Because Rips edges af and dg intersect (at z), Lemma 4.1 implies that either ag or df is a Rips edge. The triangles adg and adf intersect the hole H and therefore must not be Rips triangles. It follows that ad is not a Rips edge; a similar argument implies that bg, cf, and eh are not Rips edges. Also, since segments ae, bf, cg and dh intersect the hole H, none of them can be Rips edges. We thus have the following inequalities:

$$|ae| > 1, \quad |bf| > 1, \quad |cg| > 1, \quad |dh| > 1,$$
$$|ad| > 1, \quad |bg| > 1, \quad |cf| > 1, \quad |eh| > 1.$$

Now the triangle inequality implies that

$$|gz| + |dz| \;=\; |dg| \;<\; 1 \;<\; |ad| \;<\; |az| + |dz|;$$

it follows immediately that $|gz| < |az|$. A symmetric argument implies that $|ey| < |gy|$; thus,

$$|ey| \;<\; |gy| \;=\; |gz| + |yz| \;<\; |az| + |yz|.$$

But the triangle inequality also implies that $|az|+|yz|+|ey| > |ae| > 1$, so we must have $|az| + |yz| > 1/2$. It follows that

$$\begin{aligned} |fw| &= |fz| + |wz| \\ &< 1 - |az| + |wz| \\ &< \frac{1}{2} + |yz| + |wz|. \end{aligned}$$

An analogous argument implies that $|cw| < 1/2 + |wx| + |xy|$.

The four interior angles of $\tilde{H}$ sum to 2π, so we can assume without loss of generality that $\angle cwf \leq \pi/2$, which implies that $|fw|^2 + |cw|^2 \geq |cf|^2 > 1$. Plugging in our upper bounds for $|cw|$ and $|fw|$ and simplifying, we find that

$$(|wx|+|xy|)+(|wx|+|xy|)^2+(|yz|+|wz|)+(|yz|+|wz|)^2 > \frac{1}{2}.$$

Thus, without loss of generality, we have

$$(|wx| + |xy|) + (|wx| + |xy|)^2 > 1/4,$$

which implies that $|wx| + |xy| > (\sqrt{2}-1)/2$.

On the other hand, each edge of $\tilde{H}$ intersects D_ρ, and by Lemma 5, each interior angle of $\tilde{H}$ is greater than $\pi/3$; these facts imply that each edge of $\tilde{H}$ has length less than $2\sqrt{3}\rho$. We conclude that $\boldsymbol{\rho > (\sqrt{2}-1)/8\sqrt{3} \approx 0.029893}$.

Case 2: Points nearby, but separated from the hole. Now suppose there is a point $p \in X$ such that $|op| < \sigma$, and there is a line ℓ separating $P_{1/2}$ from the interior of H. In this case, H must be convex. Without loss of generality, we assume the line ℓ is vertical and that H lies to its right. Let x be the rightmost vertex of H; this point is the intersection of two Rips edges ab and cd, both of which are uncovered on the left. In particular, abp and cdp are not Rips triangles, so we can assume without loss of generality that $|ap| > 1$ and $|cp| > 1$.

If $|ao| < 1/2$ and $|co| < 1/2$, then $|ac| < 1$, contradicting the fact that $\angle axc$ is an uncovered corner. Thus, at least one of the Rips edges through x has both endpoints outside $D_{1/2}$.

Suppose $|ao| > 1/2$ and $|bo| > 1/2$. Let y be the point on ab closest to the origin o; clearly $|oy| < \rho$. The triangle inequality implies (crudely!) that $|by| > 1/2 - \rho$ and therefore

$$\begin{aligned} 1 < |ap| &< |op| + |oy| + |ay| \\ &< \sigma + \rho + (1 - |by|) \\ &< 1/2 + \sigma + 2\rho. \end{aligned}$$

Thus, $\sigma + 2\rho > 1/2$. We conclude that $\boldsymbol{\rho > 1/2(3+\sqrt{5}) \approx 0.095492}$.

Case 3: Point on the hole boundary. Finally, suppose there is no line separating $P_{1/2}$ from the interior of H. In this case, the convex hulls of $P_{1/2}$ and H must have intersecting interiors, which implies that H is non-convex. In particular, some vertex p of conv $X_{1/2}$ lies in the interior of conv H and is therefore a concave vertex of H. Without loss of generality, suppose the vertical line through p has the rest of conv $X_{1/2}$ to its left.

The remainder of the argument is identical to the previous case, except now we have $|op| < \rho$. The rightmost vertex x of H is the intersection of two Rips edges ab and cd, both uncovered on the left. Without loss of generality, we have the following inequalities:

$$|ap| > 1, \quad |cp| > 1, \quad |ao| > 1/2, \quad |bo| > 1/2.$$

As in the previous case, the triangle inequality implies that $|by| > 1/2 - \rho$ and therefore

$$\begin{aligned} 1 < |ap| &< |op| + |oy| + |ay| \\ &< 2\rho + (1 - |by|) \\ &< 1/2 + 3\rho. \end{aligned}$$

We conclude that $\boldsymbol{\rho > 1/6 \approx 0.166666}$. □

5.2 Contractibility Algorithm

Our algorithm follows a standard strategy used by Cabello *et al.* [5] and many other authors [36, 19, 18, 15, 12, 16, 49] for encoding the homotopy class of paths and cycles in two-dimensional spaces. In a preprocessing phase, we find several line segments $\phi_1, \phi_2, \ldots, \phi_b$, which we call *fences*, that form a spanning tree of the holes; we assign each fence an arbitrary orientation. The *crossing word* of any cycle γ records the sequence of fences that γ crosses, along with the direction of each crossing. For example, the crossing word $12\overline{2}3$ indicates that γ first crosses ϕ_1 from left to right, then ϕ_2 from left to right, then ϕ_2 from right to left, and finally ϕ_3 from left to right. We can *reduce* any crossing word by removing any matching pairs of the form $x\overline{x}$ or $\overline{x}x$; each reduction corresponds to a continuous deformation of γ

that removes two fence crossings. Finally, γ is contractible if and only if its reduced crossing word is empty.

Our spanning tree construction is a straightforward consequence of Theorem 5.2.

Lemma 5.3. *Let P be a set of points whose Rips shadow has b holes. Given $\mathcal{S}(P)$, we can compute in $O(n)$ time a set of b disjoint line segments $\phi_1, \ldots, \phi_b$ such that (1) $\mathcal{S}(P) \setminus \bigcup_i \phi_i$ has no holes, and (2) any line segment of length 1 crosses $O(1)$ segments ϕ_i.*

Proof: Consider an axis-aligned grid of squares of width 1/100 that covers the shadow $\mathcal{S}(P)$. A simple packing argument implies that $\mathcal{S}(P)$ intersects at most $O(n)$ cells in this grid, and clearly, any unit-length segment intersects $O(1)$ grid cells.

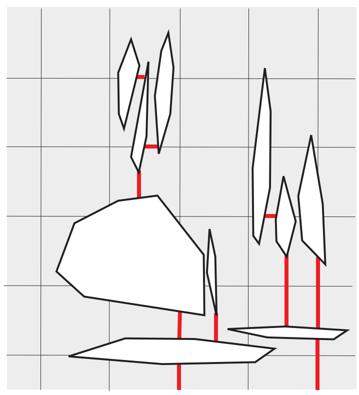

Figure 6. Fences in the Rips shadow.

We output three types of fences ϕ_i. Each *vertical* fences extends directly downward from the lowest intersection point of a hole with a vertical grid line, to either another hole or the outer boundary. Any hole that touches a vertical grid line is connected by a sequence of vertical fences and holes to the outer boundary; it remains only to connect *thin* holes that lie in *strips* between adjacent vertical grid lines. Each *horizontal* fence extends horizontally from the leftmost or rightmost intersection point of a thin hole with a horizontal grid line, to either another hole or the outer boundary, *without crossing any vertical fence.* After inserting as many horizontal fences as possible, some clusters of thin holes may still be isolated. Finally, for each cluster of thin holes, we extend a *cluster* fence from the lowest point in that cluster, either to the highest point in the next lower cluster in the same strip, or directly downward to some other boundary, whichever is closer. See Figure 6.

The horizontal and vertical fences are carried by the edges of the grid. By Theorem 5.2, each thin hole intersects at least two horizontal grid lines, so at most one cluster fence intersects any grid cell. It follows that any line segment crosses at most three fences within any grid cell, and thus $O(1)$ fences overall. $\square$

Theorem 5.4. *After $O(m \log n)$ preprocessing time, we can determine whether any given cycle of k Rips edges is contractible in $\mathcal{R}(P)$, either in $O(k \log n)$ time using $O(n)$ space, or in $O(k)$ time using $O(m)$ space.*

Proof: We can determine the crossing word for any unit-length line segment in $O(\log n)$ time in two different ways. The first method is to preprocess the degenerate simple polygon $\mathcal{S}' = \mathcal{S}(P) \setminus \bigcup_i \phi_i$ for ray-shooting, using the 'pedestrian' data structure of Hershberger and Suri [37]. The 'polygon' $\mathcal{S}'$ has complexity $O(n)$, so we can build the ray-shooting data structure in $O(n)$ time. The second, even simpler method is to store the grid cells that intersect any fence ϕ_i in a hash table. There are at most $O(n)$ such grid cells. For each such grid cell, we keep sorted arrays of the horizontal and vertical fences on its boundary, as well as the cluster fence in its interior (if any). Given any unit segment s, we can easily determine in constant time which grid cells it intersects. For each such grid cell $\square$, we can find the (at most two) fences on the boundary of $\square$ that s crosses using binary search, and then assemble the crossing word of $s \cap \square$ in $O(1)$ time.

To test whether a cycle is contractible, we consider the edges one by one in order, maintaining the reduced crossing word of the path traversed so far. Each new edge adds $O(1)$ symbols to the crossing word, so we can perform the necessary reductions in $O(1)$ time per edge. With no additional preprocessing, the total time is $O(k \log n)$. Alternately, if we precompute the crossing word of every Rips edge, we can process any cycle in constant time per edge. $\square$

6. FINDING THE SHORTEST NON-CONTRACTIBLE CYCLE

Finally, we describe how to find the shortest cycle in the Rips complex that is non-contractible. We assume that each edge pq in the proximity graph has a non-negative weight $w(pq)$; the length of a cycle is the sum of the weights of its edges. Our results hold for *any* non-negative edge weights. In particular, we can minimize either the number of edges or the total Euclidean length of the cycle; however, our algorithm does not require the weights to satisfy the triangle inequality.

For any point p and any Rips edge qr, let $C(p, qr)$ denote the cycle of Rips edges composed of the shortest path from p to q, the edge qr, and the shortest path from r back to p. The following characterization of shortest non-contractible cycles was first observed by Thomassen [48, 44] for graphs embedded on surfaces; see also [26].

Lemma 6.1. *For any point $p \in P$, the shortest non-contractible cycle in $\mathcal{R}(P)$ that passes through p is the cycle $C(p, qr)$ for some Rips edge qr.*

Proof: Let C be the shortest non-contractible cycle containing p. Let x be the point furthest from p on C; this point could be in the interior of a Rips edge. Points p and x partition C into two paths of equal length; call these paths α and β. Let γ be any other path from p to x. If γ is shorter than α and β, then the shorter loops $\alpha\overline{\gamma}$ and $\gamma\overline{\beta}$ must be contractible. But this is impossible, because the concatenation of those two loops is homotopic to $C = \alpha\overline{\beta}$, which is by definition non-contractible. We conclude that α and β are the *shortest* paths from p to x. Finally, let qr be any edge in C that contains the point x. $\square$

Theorem 6.2. *Given a set P of n points in the plane, we can compute the shortest non-contractible cycle in $\mathcal{R}(P)$ in $O(n^2 \log n + mn)$ expected time.*

Proof: As in the previous algorithm, we preprocess P for fast contractibility queries in $O(m \log n)$ time. We also construct the shortest path tree T_p by running Dijkstra's algorithm at each point $p \in P$, in total time $O(n^2 \log n + mn)$.

For each point p, we store the reduced crossing crossing words of the shortest paths from p to every other point in P in a *trie* [28], which we denote $\widetilde{T}_p$. We also store a pointer from each point q (as a node in T_p) to the corresponding node $\widetilde{q}$ in $\widetilde{T}_p$. Because the crossing word each edge of T_p has constant length, H_p has only $O(n)$ nodes. In a standard trie, each node would store an array of $O(n)$ child pointers, bringing the total size of $\widetilde{T}_p$ to $O(n^2)$. To avoid the extra space overhead, we use a *hashed trie*, in which each node with r children stores them in a hash table of size $O(r)$. Standard dynamic hashing techniques allow us to insert a new child at any node in $O(1)$ expected amortized time. Consequently, we can construct each hashed trie H_p in $O(n)$ expected time.

Finally, we preprocess H_p for constant-time least-common-ancestor queries in $O(n)$ time, using the algorithm of Bender and Farach-Colton [1]. The reduced crossing word $X_p(q,r)$ can be extracted from the edge labels on the unique path in $\widetilde{T}_p$ from $\widetilde{q}$ to $\widetilde{r}$; no further reductions are necessary. In particular, if this path is longer than $X(qr)$, then the cycle is non-contractible.

We now have all the necessary data structures to determine in $O(1)$ time, given any point p and any Rips edge qr, whether the cycle $C(p, qr)$ is non-contractible. First we find the least common ancestor z of $\widetilde{q}$ and $\widetilde{r}$. We then assemble the crossing word $X(p, qr)$ by walk up the trie from $\widetilde{q}$ to z, and then walking up the trie from $\widetilde{r}$ to z. If either walk is longer than $X(qr)$, we abort and report that the cycle is non-contractible. Otherwise, we report that the cycle is contractible if and only if $X_p(q,r) = X(qr)$. The total time to test all possible cycles $C(p, qr)$ is $O(mn)$. $\square$

7. CONCLUSION

We have presented efficient algorithms to answer two basic questions about contractibility in the Rips complex $\mathcal{R}(P)$ of a point set P in the plane. First, given a cycle in $\mathcal{R}(P)$, is it contractible? Second, given a point set P, find the shortest non-contractible cycle in its Rips complex $\mathcal{R}(P)$. Both algorithms exploit a natural isomorphism between the fundamental group of the Rips complex and the fundamental group of its shadow $\mathcal{S}(P)$. The efficiency of our algorithm relies on several new structural properties of the Rips shadow: it has combinatorial complexity $O(n)$, it is the union of only $O(n)$ Rips triangles, and it has no arbitrarily small holes.

The most obvious open question is to improve the running time of our algorithms even further. In particular, is there an algorithm to construct the Rips shadow of n points in $O(n \log n)$ time? There are two barriers to improving our current running time. First, for each input point p, we currently spend $O(n \log n)$ time finding a constant number of triangles that cover the neighborhood of p in the shadow. Second, we do not know how to compute the union of these $O(n)$ triangles in sub-quadratic time, even though the output has only linear complexity.

Another more difficult problem is generalizing our results to Rips complexes in higher dimensions. It is an open question whether Theorem 3.1 generalize to Rips complexes of points in $\mathbb{R}^3$; it does *not* generalize to more than three dimensions [13]. Nevertheless, the special geometric structure of Rips complexes may allow for more efficient algorithms for many topological questions than for general simplicial complexes. In particular, we could consider the Rips complex of a sample X of a smooth surface Σ in $\mathbb{R}^3$. Results of Haussman [35] and Latschev [40] imply that $\mathcal{R}(X)$ is homotopy-equivalent to X if X is sufficiently dense, but even if the sampling conditions for homotopy-equivalence are not met, Theorem 3.1 offers some hope that $\mathcal{R}(X)$ faithfully captures the topology of (the sampled portion of) Σ.

References

[1] M. A. Bender and M. Farach-Colton. The LCA problem revisited. *Proc 4th Latin American Symp. Theoret. Informatics (LATIN)*, 88–94, 2000.

[2] S. Bespamyatnikh. Computing homotopic shortest paths in the plane. *Proc. 14th Annu. ACM-SIAM Sympos. Discrete Algorithms*, 609–617, 2003.

[3] S. Bespamyatnikh. Encoding homotopy of paths in the plane. *Proc. LATIN 2004: Theoretical Infomatics*, 329–338, 2004. Lect. Notes Comput. Sci. 2976, Springer-Verlag.

[4] S. Cabello. Many distances in planar graphs. *Proc. 17th Ann. ACM-SIAM Symp. Discrete Algorithms*, 1213–1220, 2006.

[5] S. Cabello, Y. Liu, A. Mantler, and J. Snoeyink. Testing homotopy for paths in the plane. *Discrete Comput. Geom.* 31(1):61–81, 2004.

[6] S. Cabello and B. Mohar. Finding shortest non-separating and non-contractible cycles for topologically embedded graphs. *Discrete Comput. Geom.* 37(2):213–235, 2007.

[7] E. Carlsson, G. Carlsson, and V. de Silva. An algebraic topological method for feature identification. *Intl. J. Comput. Geom. Appl.* 16(4):291–314, 2006.

[8] G. Carlsson. Persistent homology and the analysis of high dimensional data. Presentation at the Symposium on the Geometry of Very Large Data Sets, Fields Institute for Research in Mathematical Sciences, February 24, 2005. ⟨http://comptop.stanford.edu/preprints/ottawa.pdf⟩.

[9] G. Carlsson, T. Ishkhanov, V. de Silva, and A. Zomorodian. On the local behaevior of spaces of natural images. Preprint, 2006.

[10] G. Carlsson, A. Zomorodian, A. Collins, and L. J. Guibas. Persistence barcodes for shapes. *Int. J. Shape Modeling* 11(2):149–187, 2005.

[11] E. W. Chambers and S. Cabello. Multiple source shortest paths in a genus g graph. *Proc. 18th Ann. ACM-SIAM Symp. Discrete Algorithms*, 89–97, 2007.

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

[13] E. W. Chambers, V. de Silva, J. Erickson, and R. Ghrist. Rips complexes of planar point sets. Preprint, 2007, ArXiv:0712.0395.

[14] B. Chazelle, M. Sharir, and E. Welzl. Quasi-optimal upper bounds for simplex range searching and new zone theorems. *Proc 6th Ann. Symp. Comput. Geom.*, 23–33, 1990.

[15] É. Colin de Verdière and J. Erickson. Tightening non-simple paths and cycles on surfaces. *Proc. 17th Ann. ACM-SIAM Symp. Discrete Algorithms*, 192–201, 2006.

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

[17] E. D. Demaine, M. Hajiaghayi, and B. Mohar. Approximation algorithms via contraction decomposition. *Proc. 18th Ann. ACM-SIAM Symp. Discrete Algorithms*, 278–287, 2007.

[18] T. K. Dey and S. Guha. Transforming curves on surfaces. *J. Comput. Sys. Sci.* 58(2):297–325, 1999.

[19] T. K. Dey and H. Schipper. A new technique to compute polygonal schema for 2-manifolds with application to null-homotopy detection. *Discrete Comput. Geom.* 14:93–110, 1995.

[20] C. A. Duncan, A. Efrat, S. G. Kobourov, and C. Wenk. Drawing with fat edges. *Int. J. Found. Comput. Sci.* 17(5):1143–1164, 2006.

[21] H. Edelsbrunner. The union of balls and its dual shape. *Discrete Comput. Geom.* 13:415–440, 1995.

[22] H. Edelsbrunner. *Geometry and Topology for Mesh Generation*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, New York, NY, USA, 2006.

[23] H. Edelsbrunner and J. Harer. Persistent homology: A survey. *Discrete & Computational Topology: Twenty Years Later*, to appear. AMS Press. ⟨http://www.cs.duke.edu/~edels/BookSurv/PersistenceSurvey.pdf⟩.

[24] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.* 28(4):511–533, 2002.

[25] A. Efrat, S. G. Kobourov, and A. Lubiw. Computing homotopic shortest paths efficiently. *Comput. Geom. Theory Appl.* 35(3):162–172, 2006.

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

[27] J. Erickson and K. Whittlesey. Greedy optimal homotopy and homology generators. *Proc. 16th Ann. ACM-SIAM Symp. Discrete Algorithms*, 1038–1046, 2005.

[28] E. Fredkin. Trie memory. *Commun. ACM* 3(9):490–499, 1960.

[29] R. Ghrist. Barcodes: The persistent topology of data. Preprint, 2007. ⟨http://www.math.uiuc.edu/~ghrist/preprints/barcodes.pdf⟩.

[30] R. Ghrist and A. Muhammad. Coverage and hole-detection in sensor networks via homology. *Proc. 4th Int. Symp. Information Processing in Sensor Networks*, p. 34, 2005. IEEE Press.

[31] D. Grigoriev and A. Slissenko. Polytime algorithm for the shortest path in a homotopy class amidst semi-algebraic obstacles in the plane. *Proc. Internat. Sympos. Symbolic and Algebraic Computation*, 17–24, 1998.

[32] M. Gromov. Hyperbolic groups. *Essays in Group Theory*, 75–265, 1987. MSRI Publications 8, Springer-Verlag.

[33] I. Guskov and Z. Wood. Topological noise removal. *Proc. Graphics Interface*, 19–26, 2001.

[34] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2001. ⟨http://www.math.cornell.edu/~hatcher/⟩.

[35] J.-C. Hausmann. On the Vietoris-Rips complexes and a cohomology theory for metric spaces. *Prospects in Topology: Proceedings of a Conference in Honor of William Browder*, 157–188, 1995. Ann. Math. Stud. 138, Princeton Univ. Press.

[36] J. Hershberger and J. Snoeyink. Computing minimum length paths of a given homotopy class. *Comput. Geom. Theory Appl.* 4(2):63–97, 1994.

[37] J. Hershberger and S. Suri. A pedestrian approach to ray shooting: Shoot a ray, take a walk. *J. Algorithms* 18(3):403–431, 1995.

[38] K. Kawarabayashi and B. Reed. Computing crossing number in linear time. *Proc. 39th ACM Symp. Theory Comput.*, 382–390, 2007.

[39] M. Kutz. Computing shortest non-trivial cycles on orientable surfaces of bounded genus in almost linear time. *Proc. 22nd Ann. Sympos. Comput. Geom.*, 430–438, 2006.

[40] J. Latschev. Vietoris-Rips complexes of metric spaces near a closed Riemannian manifold. *Archiv der Mathematik* 77(6):522–528, 2001.

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

[42] A. A. Markov. Insolubility of the problem of homeomorphy. *Proceedings of the International Congress of Mathematics*, 14–21, 1958. Cambridge University Press.

[43] J. Matoušek. Spanning trees with low crossing numbers. *Informatique Théorique et Applications* 6(25):103–123, 1991.

[44] B. Mohar and C. Thomassen. *Graphs on Surfaces*. Jons Hopkins Univ. Press, 2001.

[45] A. Muhammad and A. Jadbabaie. Asymptotic stability of switched higher order Laplacians and dynamic coverage. *Hybrid Systems: Computation and Control*, p. to appear, 2007. Lecture Notes Comput. Sci., Springer-Verlag. ⟨http://www.seas.upenn.edu/~jadbabai/papers/Muhammad_Jadbabaie_HSCC07_FinalSubmission.pdf⟩.

[46] A. Muhammad and A. Jadbabaie. Dynamic coverage verification in mobile sensor networks via switched higher order Laplacians. *Proceedings of Robotics: Science and Systems*, p. to appear, 2007. ⟨http://www.seas.upenn.edu/~jadbabai/papers/RSS_FINAL.pdf⟩.

[47] J. Stillwell. *Classical Topology and Combinatorial Group Theory*, 2nd edition. Graduate Texts in Mathematics 72. Springer-Verlag, 1993.

[48] C. Thomassen. Embeddings of graphs with no short noncontractible cycles. *J. Comb. Theory Ser. B* 48(2):155–177, 1990.

[49] É. C. de Verdière and F. Lazarus. Optimal pants decompositions and shortest homotopic cycles on an orientable surface. *Proc. 11th Int. Symp. Graph Drawing*, 478–490, 2003.

[50] L. Vietoris. Über den hoheren Zusammenhang kompakter Raume und eine Klasse von zusammenhangstreuen Abbildungen. *Math. Ann.* 97:454–472, 1927.

[51] E. Welzl. On spanning trees with low crossing numbers. *Data Structures and Efficient Algorithms, Final Report on the DFG Special Joint Initiative*, 233–249, 1992. Lecture Notes Comput. Sci. 594, Springer-Verlag.

[52] Z. Wood, H. Hoppe, M. Desbrun, and P. Schröder. Removing excess topology from isosurfaces. *ACM Trans. Graph.* 23(2):190–208. ACM Press, 2004.

[53] A. Zomorodian. *Topology for Computing*. Cambridge University Press, New York, NY, 2005.

[54] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete Comput. Geom.* 33(2):249–274, 2005.